

**AN INTERPRETABLE GEOAI FRAMEWORK FOR ANALYZING MULTI-ETHNIC
SETTLEMENT DYNAMICS: EVIDENCE FROM THE GREATER TORONTO AREA,
2001-2021**

SEYED NAVID MASHHADI MOGHADDAM

Thesis submitted to the University of Ottawa
in partial Fulfillment of the requirements for the
For the PhD degree in Geography

Department of Geography
Faculty of Arts
University of Ottawa

© Seyed Navid Mashhadi Moghaddam, Ottawa, Canada, 2026

Abstract

This dissertation develops an interpretable Geospatial Artificial Intelligence (GeoAI) framework for understanding multi-ethnic settlement patterns in contemporary Canadian cities, demonstrating that computational sophistication need not sacrifice theoretical transparency or democratic accountability. Through three interconnected investigations spanning 52 major Canadian cities, 30,091 dissemination areas, and nine major ethnic groups (China, India, Philippines, Pakistan, Sri Lanka, Iran, Portugal, Italy, United Kingdom) over two decades (2001-2021), this research transforms multi-ethnic settlement from a descriptive sociological phenomenon into a predictive science grounded in interpretable physics-informed models.

The research addresses a fundamental tension in urban artificial intelligence: while cities increasingly deploy algorithmic systems to manage complex urban dynamics, the dominant paradigm of black-box optimization systematically fails to achieve stated goals while potentially harming vulnerable communities. A critical analysis of 157 urban AI deployments (2015-2024) reveals pervasive "metrics traps" where impressive technical accuracy, such as ShotSpotter's 97% acoustic precision yielding only 9.1% crime-fighting effectiveness, consistently fails to translate into meaningful social outcomes. This critique establishes the ethical and methodological imperative for interpretable approaches in urban demographic analysis.

The core methodological contribution is the development of a graph-based physics-informed neural network (GraphPDE) that embeds multi-ethnic reaction-diffusion dynamics while learning interpretable demographic parameters. Adapting Turing's pattern formation theory to spatial graphs, this framework reveals that ethnic settlement patterns emerge from self-organizing spatiotemporal processes with quantifiable characteristics: ethnic-specific spatial scales ranging from 34.5 km (Philippines) to 63.0 km (United Kingdom); pattern formation regimes segregating groups into spots (UK, Portugal), stripes (China, India, Philippines), and labyrinthine (Iran, Pakistan, Sri Lanka) morphologies; and critical nucleation thresholds varying from 658 individuals (China) to 8,132 individuals (India) for spatial clustering emergence. The learned attention-based interaction mechanism quantifies previously unmeasurable inter-ethnic dynamics, revealing that Chinese populations exhibit strong negative self-competition (-19.8) driving spatial dispersal while maintaining facilitative relationships with multiple groups, Philippines-Pakistan mutual attraction, and United Kingdom-Italy mutual repulsion.

The Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework synthesizes physics-informed modeling with machine learning, achieving state-of-the-art predictive performance ($R^2 = 0.80-0.83$) while maintaining complete interpretability through regime-specific expert modules for colonization, jump processes, decline, and continuous diffusion dynamics. This architecture demonstrates that incorporating

domain knowledge enhances rather than compromises predictive accuracy, with physics-informed components accounting for 57.2% of prediction variance. The framework reveals systematic differences in how ethnic communities respond to urban infrastructure: Chinese and Filipino populations show amplification factors strongly correlated with transit accessibility ($r = 0.58$ and 0.61), while Indian populations demonstrate stronger correlation with housing variables ($r = 0.47$).

Configuration landscape analysis identifies multiple stable settlement configurations with quantifiable transition barriers, revealing that demographic transitions require sustained interventions over decadal timescales due to asymmetric barriers creating lock-in effects. The temporal evolution of parameters captures non-stationary dynamics across four census periods, with Philippines-origin populations maintaining consistently positive growth rates (0.015 - 0.023 year^{-1}) while United Kingdom-origin populations transition from positive growth to sustained decline.

This research makes four distinctive contributions: (1) developing progressive spatial analysis as a systematic methodology for studying complex urban phenomena through integrated multi-scale investigation; (2) demonstrating interpretable GeoAI that achieves competitive performance without sacrificing transparency; (3) revealing hidden ethnic dynamics through quantification of inter-ethnic interactions, nucleation thresholds, and pattern formation regimes; and (4) providing practical tools bridging academic research and urban planning practice through open-source implementations. The framework transforms abstract geographic concepts into measurable quantities, "sense of place" becomes configuration stability, "community cohesion" translates to interaction strength, and "neighborhood character" maps to position in pattern formation phase space.

The policy implications are transformative: understanding asymmetric configuration barriers explains persistent settlement patterns while identifying intervention requirements; regime-specific approaches match policies to demographic dynamics; and cascade effects enable strategic investments generating system-wide benefits. Critical reflection acknowledges fundamental limitations, physics cannot capture individual agency, cultural meaning, or structural inequalities, yet the framework complements rather than replaces other ways of understanding urban dynamics.

This dissertation demonstrates that interpretable GeoAI can bridge the persistent divide between quantitative spatial science and critical human geography, proving that mathematical rigor need not sacrifice social awareness. The discovery that Canadian cities' ethnic geography follows learnable physical dynamics while maintaining cultural distinctiveness suggests that diversity and order are complementary aspects of urban organization, offering both cautionary lessons about algorithmic governance and hopeful possibilities for creating more equitable, integrated multicultural cities.

Resume

Cette thèse développe un cadre d'intelligence artificielle géospatiale (GeoAI) interprétable pour comprendre les modèles d'établissement multiethnique dans les villes canadiennes contemporaines, démontrant que la sophistication computationnelle n'exige pas de sacrifier la transparence théorique ou la responsabilité démocratique. À travers trois investigations interconnectées couvrant 52 grandes villes canadiennes, 30 091 aires de diffusion et neuf groupes ethniques majeurs (Chine, Inde, Philippines, Pakistan, Sri Lanka, Iran, Portugal, Italie, Royaume-Uni) sur deux décennies (2001-2021), cette recherche transforme l'établissement multiethnique d'un phénomène sociologique descriptif en une science prédictive fondée sur des modèles interprétables inspirés de la physique.

La recherche aborde une tension fondamentale dans l'intelligence artificielle urbaine : alors que les villes déploient de plus en plus de systèmes algorithmiques pour gérer des dynamiques urbaines complexes, le paradigme dominant de l'optimisation en boîte noire échoue systématiquement à atteindre les objectifs déclarés tout en nuisant potentiellement aux communautés vulnérables. Une analyse critique de 157 déploiements d'IA urbaine (2015-2024) révèle des « pièges métriques » omniprésents où une précision technique impressionnante, comme la précision acoustique de 97 % de ShotSpotter ne produisant que 9,1 % d'efficacité dans la lutte contre la criminalité, ne se traduit constamment pas en résultats sociaux significatifs. Cette critique établit l'impératif éthique et méthodologique pour des approches interprétables dans l'analyse démographique urbaine.

La contribution méthodologique centrale est le développement d'un réseau neuronal informé par la physique basé sur des graphes (GraphPDE) qui intègre des dynamiques de réaction-diffusion multiethniques tout en apprenant des paramètres démographiques interprétables. En adaptant la théorie de formation de motifs de Turing aux graphes spatiaux, ce cadre révèle que les modèles d'établissement ethnique émergent de processus spatiotemporels auto-organisés avec des caractéristiques quantifiables : des échelles spatiales spécifiques à chaque ethnie allant de 34,5 km (Philippines) à 63,0 km (Royaume-Uni) ; des régimes de formation de motifs séparant les groupes en taches (Royaume-Uni, Portugal), rayures (Chine, Inde, Philippines) et morphologies labyrinthiques (Iran, Pakistan, Sri Lanka) ; et des seuils critiques de nucléation variant de 658 individus (Chine) à 8 132 individus (Inde) pour l'émergence du regroupement spatial. Le mécanisme d'interaction basé sur l'attention quantifie des dynamiques interethniques auparavant non mesurables, révélant que les populations chinoises présentent une forte auto-compétition négative (-19,8) favorisant la dispersion spatiale tout en maintenant des relations facilitatrices avec plusieurs groupes, une attraction mutuelle Philippines-Pakistan, et une répulsion mutuelle Royaume-Uni-Italie.

Le cadre Multi-Ethnic Spatial Mixture of Experts (MESMoE) synthétise la modélisation informée par la physique avec l'apprentissage automatique, atteignant une performance prédictive de pointe ($R^2 = 0,80$ -

0,83) tout en maintenant une interprétabilité complète grâce à des modules experts spécifiques aux régimes pour la colonisation, les processus de saut, le déclin et les dynamiques de diffusion continue. Cette architecture démontre que l'incorporation de connaissances du domaine améliore plutôt qu'elle ne compromet la précision prédictive, les composantes informées par la physique représentant 57,2 % de la variance de prédiction. Le cadre révèle des différences systématiques dans la façon dont les communautés ethniques répondent à l'infrastructure urbaine : les populations chinoises et philippines montrent des facteurs d'amplification fortement corrélés avec l'accessibilité au transport en commun ($r = 0,58$ et $0,61$), tandis que les populations indiennes démontrent une corrélation plus forte avec les variables de logement ($r = 0,47$).

L'analyse du paysage configurationnel identifie plusieurs configurations d'établissement stables avec des barrières de transition quantifiables, révélant que les transitions démographiques nécessitent des interventions soutenues sur des échelles décennales en raison de barrières asymétriques créant des effets de verrouillage. L'évolution temporelle des paramètres capture des dynamiques non stationnaires à travers quatre périodes de recensement, les populations d'origine philippine maintenant des taux de croissance constamment positifs ($0,015$ - $0,023 \text{ an}^{-1}$) tandis que les populations d'origine britannique passent d'une croissance positive à un déclin soutenu.

Cette recherche apporte quatre contributions distinctives : (1) le développement de l'analyse spatiale progressive comme méthodologie systématique pour l'étude de phénomènes urbains complexes à travers une investigation multi-échelle intégrée ; (2) la démonstration d'une GeoAI interprétable qui atteint une performance compétitive sans sacrifier la transparence ; (3) la révélation de dynamiques ethniques cachées par la quantification des interactions interethniques, des seuils de nucléation et des régimes de formation de motifs ; et (4) la fourniture d'outils pratiques reliant la recherche académique et la pratique de planification urbaine grâce à des implémentations en code source ouvert. Le cadre transforme des concepts géographiques abstraits en quantités mesurables, le « sens du lieu » devient la stabilité configurationnelle, la « cohésion communautaire » se traduit en force d'interaction, et le « caractère de quartier » correspond à une position dans l'espace de phase de formation de motifs.

Les implications politiques sont transformatrices : la compréhension des barrières configurationnelles asymétriques explique les modèles d'établissement persistants tout en identifiant les exigences d'intervention ; les approches spécifiques aux régimes associent les politiques aux dynamiques démographiques ; et les effets de cascade permettent des investissements stratégiques générant des bénéfices à l'échelle du système. Une réflexion critique reconnaît les limitations fondamentales, la physique ne peut capturer l'agentivité individuelle, le sens culturel ou les inégalités structurelles, mais le cadre complète plutôt qu'il ne remplace d'autres façons de comprendre les dynamiques urbaines.

Cette thèse démontre que la GeoAI interprétable peut combler le fossé persistant entre la science spatiale quantitative et la géographie humaine critique, prouvant que la rigueur mathématique n'exige pas de sacrifier la conscience sociale. La découverte que la géographie ethnique des villes canadiennes suit des dynamiques physiques apprenables tout en maintenant une spécificité culturelle suggère que la diversité et l'ordre sont des aspects complémentaires de l'organisation urbaine, offrant à la fois des leçons de prudence sur la gouvernance algorithmique et des possibilités prometteuses pour créer des villes multiculturelles plus équitables et intégrées.

Acknowledgements

The completion of this dissertation represents not merely an individual scholarly achievement but the culmination of countless acts of support, guidance, and encouragement from those who believed in this work and in me. I am profoundly grateful to all who have accompanied me on this journey.

I owe an immeasurable debt of gratitude to my supervisor, **Professor Huhua Cao**, whose mentorship has shaped not only this dissertation but my entire academic career. Professor Cao's guidance extended far beyond the boundaries of scholarly advice; he has been a constant source of support in every step of my academic and personal life. His wisdom, patience, and unwavering confidence in my abilities provided the intellectual environment in which I could take risks, explore unconventional ideas, and grow as a researcher. He taught me that rigorous scholarship and genuine humanity are not competing values but complementary virtues. I am honored to have learned under his supervision and privileged to call him my mentor.

My sincere appreciation goes to **Professor Michael Sawada**, whose generosity with his time and expertise has been extraordinary. Professor Sawada's door, whether literal or virtual, was always open, regardless of the hour or day of the week. His willingness to provide insight and guidance on weekends, during holidays, and at any time I faced a methodological challenge exemplifies the very best of academic mentorship. His sharp analytical mind and thoughtful feedback pushed me to refine my arguments and strengthen my methods. I am deeply grateful for his dedication to my success.

I extend my heartfelt thanks to **Professor Tao Jin** and **Professor Ruibo Han** for their invaluable contributions to this work. Their expertise, thoughtful guidance, and willingness to share their knowledge have enriched this dissertation in meaningful ways. I am grateful for their scholarly generosity and the time they devoted to supporting my research.

I wish to express my heartfelt thanks to **Professor Anders Knudby**, Director of the Department of Geography, Environment and Geomatics, for his steadfast support and leadership. Professor Knudby ensured that I had access to all the resources and necessities required to complete this thesis, from computational infrastructure to administrative support. His commitment to fostering an environment where graduate students can thrive has been instrumental to my progress. I am grateful for his advocacy on behalf of students and his vision for our department.

I extend my sincere appreciation to the **International Council on Canadian, Chinese and African Sustainable Urbanization (ICCCASU)** and its members. This remarkable network has been an invaluable source of collaboration, knowledge exchange, and support throughout my doctoral studies. Whenever I

encountered challenges, whether methodological, conceptual, or practical, the ICCCASU community offered assistance without hesitation. The consortium provided not only academic opportunities but also a sense of belonging to a scholarly community dedicated to advancing urban research. I am grateful for the connections forged and the knowledge shared within this network.

I also wish to acknowledge the broader academic community at the University of Ottawa, including faculty members, fellow graduate students, and administrative staff who contributed to my development as a scholar. The intellectual exchanges in seminars, the camaraderie in shared workspaces, and the administrative support that kept the machinery of graduate education running smoothly all played important roles in bringing this work to fruition.

I acknowledge the use of large language model technology (Claude, developed by Anthropic) as a writing assistance tool during the preparation of this dissertation. This AI tool was employed for enhancing writing clarity, refining prose, condensing text, and similar editorial tasks. All intellectual content, research design, analysis, and conclusions remain entirely my own work. The use of this technology reflects the evolving landscape of academic writing tools and was undertaken with transparency and in accordance with principles of academic integrity.

Finally, and most importantly, I extend my deepest gratitude to my fiancée, **Diba Rashidi**, whose unconditional love and unwavering support have been the cornerstone of my perseverance throughout this doctoral journey. Through the long nights of coding, the frustrations of failed experiments, and the anxieties of academic life, her patience, understanding, and encouragement never wavered. She celebrated my small victories and lifted me during moments of doubt, providing the emotional foundation upon which this work was built. This dissertation is as much a testament to her steadfast belief in me as it is to my own efforts. No words can adequately express what her presence has meant to me.

To all who have supported me on this journey, named and unnamed, in ways large and small, I offer my profound thanks. Whatever contributions this dissertation makes to our understanding of multi-ethnic urban dynamics are shared with you.

Table of Contents

Abstract.....	ii
Resume.....	iv
Acknowledgements	vii
List of Figures.....	xiv
List of Tables	xvi
List of Abbreviations	xvii
Chapter 1 Introduction.....	1
1.1 Background	1
1.1.1 The Challenge of Multi-Ethnic Settlement in Canadian Cities	1
1.1.2 From Traditional Geography to Computational Approaches	3
1.1.3 GeoAI: Bridging Spatial Science and Artificial Intelligence	5
1.1.4 Research Direction.....	7
1.2 Research Gap and Positioning.....	11
1.2.1 The Critical Gap: From Isolated Ethnic Geographies to Multi-Ethnic Urban Systems	11
1.2.2 The Imperative for Integrated Multi-Ethnic Urban Analysis	12
1.2.3 GeoAI as Solution Framework	13
1.2.4 Positioning This Research.....	14
1.3 Research Questions and Research Scheme	15
1.3.1 Research Questions.....	15
1.4 Methodology	17
1.4.1 Thesis Framework.....	17
1.4.2 Critical Analysis - Establishing Ethical Foundations.....	20
1.4.3 Physics-Informed Analysis - Metropolitan Scale Dynamics.....	20
1.4.4 Hybrid Framework - Multi-scale Integration	21
1.4.5 Exploratory Investigations - Learning from Black-Box Methods.....	21
1.4.6 Computational Implementation and Ethical Considerations	21
1.4.7 Thesis Structure	22
1.5 Anticipated Contributions	24
1.5.1 Theoretical Contributions	24
1.5.2 Methodological Contributions	25
1.5.3 Empirical Contributions.....	25
1.5.4 Practical Contributions.....	26
1.5.5 Interdisciplinary Bridge Building	26
Chapter 2 The Greater Toronto Area as a Multi-Ethnic Urban Laboratory	28
2.1 A Historical Timeline of Immigration Waves (1970s–2020s)	28

2.1.1 The Demographic Transformation	28
2.1.2 Wave I: The Final Years of European Dominance (1970s)	28
2.1.3 Wave II: South Asian Emergence (1970s–1990s)	29
2.1.4 Wave III: East Asian Expansion (1980s–2000s).....	30
2.1.5 Wave IV: Middle Eastern Growth (1980s–2020s).....	31
2.1.6 Period of Immigration Analysis	31
2.1.7 Visualizing Six Decades of Spatial Transformation	31
2.1.8 Implications for Settlement Patterns	35
2.1.9 Data Considerations and Limitations	36
2.2 Selection and Justification of Study Groups	36
2.2.1 Population Significance and Representativeness	36
2.2.2 Diversity in Growth Trajectories	38
2.2.3 Regional and Cultural Representation	39
2.2.4 Spatial Distribution Characteristics.....	40
2.2.5 Temporal Dynamics of Settlement.....	42
2.2.6 Summary of Selection Criteria.....	44
2.3 Spatial Patterns of Multi-Ethnic Settlement	44
2.3.1 Inner-City Neighbourhoods: Historical Enclaves and Succession	45
2.3.2 Suburban Settlements: New Ethnic Geographies.....	46
2.3.3 Mixed Neighbourhoods: Sites of Interethnic Contact.....	46
2.4 Socioeconomic and Housing Dimensions	47
2.4.1 Housing Affordability as Settlement Driver	47
2.4.2 Stratification Within Communities	47
2.4.3 Rental Versus Ownership Markets.....	48
2.4.4 Intersection with Urban Inequalities	49
2.5 Governance and Policy Context	49
2.5.1 Federal Policy Framework	49
2.5.2 Provincial Policy Framework.....	50
2.5.3 Municipal Policy and Planning	50
2.5.4 Governance Gaps and Challenges.....	51
2.6 Dynamics of Multi-Ethnic Settlement.....	51
2.6.1 Ethnic Enclaves: Formation, Functions, and Futures	52
2.6.2 Suburbanization: Mobility and New Concentrations	52
2.6.3 Displacement and Gentrification	53
2.6.4 Hybrid Neighbourhoods: Contact, Tension, and Cooperation	53
2.6.5 Linking Dynamics to Subsequent Analysis	53

Chapter 3 The Metrics Trap: How Technical Sophistication Masks Social Harm in Urban AI Systems

.....	55
3.1 Introduction	55
3.2 Results	58
3.2.1 Overview of the Case Universe	58
3.2.2 Theme 1: The Accuracy Illusion.....	60
3.2.3 Theme 2: Discriminatory Feedback Loops	64
3.2.4 Theme 3: Community Resistance and Democratic Pushback	68
3.2.5 Theme 4: Transparent Alternatives and Democratic AI Governance	73
3.3 Discussion	77
3.4 Methods.....	79
3.4.1 Research Design Rationale	80
3.4.2 Researcher Positionality.....	80
3.4.3 Case Universe and Sampling Strategy	80
3.4.4 Data Collection and Management.....	82
3.4.5 Analytical Framework	83
3.4.6 Quality and Rigor.....	83
3.4.7 Thematic Case Distribution	84
3.5 Author Contributions.....	85

Chapter 4 Physics-Informed Neural Networks Reveal Interpretable Mechanisms of Multi-Ethnic Urban Settlement Dynamics

.....	86
4.1 Main	87
4.2 Results	91
4.2.1 GraphPDE Achieves Superior Predictive Performance Through Reaction-Diffusion Dynamics.....	91
4.2.2 Learned Demographic Parameters Reveal Distinct Ethnic Settlement Dynamics	93
4.2.3 City-Level Parameters Expose Geographic Heterogeneity in Settlement Processes	95
4.2.4 Multi-Scale Spatial Organization Emerges from Reaction-Diffusion Dynamics.....	97
4.2.5 Temporal Pattern Evolution Reveals Multi-Stage Settlement Crystallization	99
4.2.6 Multi-Stability Analysis Reveals Configuration Transition Dynamics.....	101
4.3 Discussion	103
4.4 Methods.....	106
4.4.1 Problem Formulation and Motivation	106
4.4.2 Dataset Construction and Preprocessing	107
4.4.3 Mathematical Framework	109
4.4.4 Training Procedures and Optimization	112
4.4.5 Theoretical Analysis	114
4.4.6 Computational Implementation.....	115

4.4.7 Hardware and Software Stack	115
4.4.8 Uncertainty Quantification	117
4.4.9 Energy Landscape Analysis: Post-Training Interpretation	123
4.4.10 Crystallization Analysis: Multi-Scale Settlement Patterns	129
Chapter 5 Regime-Adaptive Partial Differential Equations for Interpretable Multi-Ethnic Urban Modeling	138
5.1 Main	138
5.2 Results	143
5.2.1 Model Performance Across Temporal Periods	143
5.2.2 Model Interpretability Analysis	149
5.3 Discussion	161
5.4 Limitations	163
5.5 Conclusion	165
5.6 Methods	166
5.6.1 Multi-Ethnic Census Dataset	166
5.6.2 Multi-Ethnic Spatial Mixture of Experts Framework	168
5.6.3 Specialized Expert Models	170
5.6.4 Multi-Ethnic Decline Expert	171
5.6.5 Training Methodology	174
Chapter 6 Synthesis - Discussion and Conclusion	179
6.1 Introduction: Bridging Computational Discovery and Geographic Theory	179
6.2 Discussion and Theoretical Synthesis	181
6.2.1 From Metrics Trap to Mechanistic Understanding	181
6.2.2 Universal Physical Principles in Social Systems	183
6.2.3 Multi-Scale Organization of Urban Ethnic Systems	185
6.3 Methodological Contributions to GeoAI	187
6.3.1 The Interpretable GeoAI Framework	187
6.3.2 Computational Innovations for Urban Scale	188
6.3.3 Validation Across Scales and Contexts	189
6.4 Empirical Insights: Multi-Ethnic Settlement Across Canadian Cities	190
6.4.1 Quantifying the Unquantifiable	190
6.4.2 Distinctive Patterns Across Communities	191
6.4.3 Toronto as Empirical Laboratory: Multi-Scale Settlement Dynamics	192
6.4.4 Configuration Stability and Transition Dynamics in Toronto	193
6.4.5 Emergence of Multi-Ethnic Urban Structures	194
6.5 Policy Implications: From Physics to Planning	195
6.6 Broader Implications for Human Geography	198

6.7 Limitations and Critical Reflections.....	199
6.8 Future Research Directions	201
6.9 Conclusion: The Promise of Interpretable GeoAI	202
References.....	205
Appendix A.....	223
Appendix B.....	246

List of Figures

Figure 1.1 Conceptual framework of multi-ethnic settlement challenges in Canadian cities.....	2
Figure 1.2 Evolution from Traditional Settlement Geography to GeoAI: Paradigmatic Shifts and Methodological Transformations.....	4
Figure 1.3 Conceptual Architecture of GeoAI: Bridging Spatial Science and Artificial Intelligence for Urban Analysis.....	6
Figure 1.4 Research Direction Framework: Three Paradigmatic Transformations in Urban Settlement Studies.....	8
Figure 1.5 Progressive GeoAI methodological framework for analyzing multi-ethnic settlement patterns in the Greater Toronto Area.	19
Figure 2.1 Spatial Distribution of China-born Population in Toronto CMA (1961-2021).	32
Figure 2.2 Temporal Evolution of Population Mean Centers and One Standard Deviation Ellipses for Chinese and Iranian Immigrant Communities (1961-2021).	33
Figure 2.3 Temporal trends in immigrant population for the nine selected origin countries across five census periods.....	37
Figure 2.4 Comparison of immigrant populations between 2001 and 2021 census years, illustrating the magnitude of demographic shifts.....	38
Figure 2.5 Proportional representation of the nine selected immigrant groups in 2021.	39
Figure 2.6 Population ranking of the top 20 immigrant source countries in Toronto (2021).	40
Figure 2.7 Panel maps displaying the spatial distribution of each selected immigrant group across Toronto's Dissemination Areas (2021 Census).	41
Figure 2.8 Location Quotient (LQ) maps illustrating areas of ethnic over-representation ($LQ > 1$, shown in warmer colours) and under-representation ($LQ < 1$, shown in cooler colours).	42
Figure 2.9 Population change maps (2001-2021) for each selected immigrant group.....	43
Figure 3.1 Global landscape of urban AI implementations (2015–2024).....	59
Figure 3.2 The Accuracy Illusion: Evidence from 21 Case Studies Across 8 Countries.	61
Figure 3.3 Discriminatory Feedback Loops: Evidence from 11 Documented Cases.	65
Figure 3.4 Community Resistance and Democratic Pushback: Evidence from 20 Documented Cases.	69
Figure 3.5 The Transparency Spectrum in Urban AI Implementations.	74
Figure 4.1 GraphPDE Model Architecture.	91
Figure 4.2 GraphPDE achieves superior predictive performance through physics-informed learning with interpretable demographic mechanisms.	92

Figure 4.3 Learned demographic parameters reveal distinct ethnic settlement dynamics and inter-ethnic interaction patterns.....	94
Figure 4.4 City-level parameters expose geographic heterogeneity in multi-ethnic settlement processes across Canadian metropolitan areas.....	96
Figure 4.5 Multi-scale spatial organization emerges from reaction-diffusion dynamics in Chinese settlement patterns across Toronto.....	98
Figure 4.6 Temporal pattern evolution reveals multi-stage settlement crystallization dynamics in Chinese Toronto communities.....	100
Figure 4.7 Multi-stability analysis reveals configuration transition dynamics and demographic reorganization barriers in Chinese Toronto settlement.....	102
Figure 5.1 Performance analysis and expert module dynamics of the Multi-Ethnic Spatial Mixture of Experts (MESMoE) model across three temporal periods (during training and validation).....	144
Figure 5.2 Model architecture and ethnicity-specific module dynamics for the 2016-2021 period.....	151
Figure 5.3 Feature importance analysis across expert modules for the 2016-2021 period.....	153
Figure 5.4 Ethnicity interaction network analysis for the 2016-2021 period.....	155
Figure 5.5 Physics-informed spatial parameter visualization for Chinese populations during the 2016-2021 period. The figure includes 10 batches of dataset for better visualization.....	158

List of Tables

Table 1-1. Comparison of Analytical Approaches to Urban Ethnic Settlement	13
Table 2-1. Toronto CMA Population and Immigration, 1971–2021	28
Table 2-2. European-Origin Immigrant Population, Toronto CMA	29
Table 2-3. South Asian Immigrant Population, Toronto CMA	30
Table 2-4. East Asian Immigrant Population, Toronto CMA	30
Table 2-5. Iranian Immigrant Population, Toronto CMA	31
Table 2-6. Toronto CMA Immigrants by Period of Arrival (2011 NHS)	31
Table 2-7. Summary Statistics of Selected Immigrant Groups in Toronto (2001-2021)	37
Table 3-1 The Metrics Gap in Selected Urban AI Systems	63
Table 3-2 Documented Feedback Loop Mechanisms in Selected Cases	68
Table 3-3 Implementation Context and Resistance Outcomes in Selected Cases	72
Table 4-1. Cuda Performance compared to naïve pytorch.	117
Table 4-2. Overall Test Set Performance.	119
Table 4-3. Uncertainty varies systematically across ethnic groups, reflecting demographic dynamics. ...	119
Table 4-4. Population range and sample count.	120
Table 4-5. Settlement interpretation based on growth mechanism.	131
Table 4-6. Metrics and statistics for crystallization.	136
Table 5-1. Comprehensive Performance Metrics Across Temporal Periods and Population Regimes and Comparison to Baseline Models	146
Table 6-1. Policy Framework by Demographic Regime and Intervention Type	196

List of Abbreviations

Abbreviation	Full Form
AI	Artificial Intelligence
AUC	Area Under Curve
BIA	Business Improvement Area
CAGR	Compound Annual Growth Rate
CMA	Census Metropolitan Area
CMHC	Canada Mortgage and Housing Corporation
COMPAS	Correctional Offender Management Profiling for Alternative Sanctions
COVID-19	Coronavirus Disease 2019
CSR	Compressed Sparse Row
CUDA	Compute Unified Device Architecture
DA	Dissemination Area
DBN	Deep Belief Network
ECDF	Empirical Cumulative Distribution Function
ERQ	Exploratory Research Question
EU	European Union
GAN	Generative Adversarial Network
GEOAI	Geospatial Artificial Intelligence
GIS	Geographic Information Systems
GNN	Graph Neural Network
GPU	Graphics Processing Unit
GRAPHPDE	Graph-based Physics-Informed Neural Network with Partial Differential Equations
GRAPHSAGE	Graph Sample and Aggregate
GTA	Greater Toronto Area
IRCC	Immigration, Refugees and Citizenship Canada
JMAK	Johnson-Mehl-Avrami-Kolmogorov
LCP	Live-in Caregiver Program
LIPS	Local Immigration Partnerships
LQ	Location Quotient
MAE	Mean Absolute Error
MESMOE	Multi-Ethnic Spatial Mixture of Experts
ML	Machine Learning

MLP	Multilayer Perceptron
MSA	Metropolitan Statistical Area
NHS	National Household Survey
NMI	Normalized Mutual Information
OAuth	Open Authorization
OINP	Ontario Immigrant Nominee Program
PDE	Partial Differential Equation
PINNS	Physics-Informed Neural Networks
PRC	People's Republic of China
R²	Coefficient of Determination
RBM	Restricted Boltzmann Machine
RMSE	Root Mean Square Error
RQ	Research Question
SDE	Stochastic Differential Equation
SGBDBN	Spatial Gaussian-Bernoulli Deep Belief Network
SSO	Single Sign-On
STGAN-EI	Spatio-Temporal Graph Attention Network with Ethnic Interaction Learning
TRIEC	Toronto Region Immigrant Employment Council
UK	United Kingdom
XAI	Explainable Artificial Intelligence
XGBOOST	Extreme Gradient Boosting

Chapter 1 Introduction

1.1 Background

1.1.1 The Challenge of Multi-Ethnic Settlement in Canadian Cities

The transformation of Canadian cities into multi-ethnic metropolises represents one of the most profound demographic shifts in the nation's history. Toronto exemplifies this unprecedented diversity where 55.7% of the population belongs to visible minority groups, a dramatic increase from 13.6% in 1981, with similar transformations in Vancouver (54%) and Montreal (34%) creating what has been described as a "new residential order" characterized by simultaneous trajectories of ethnic concentration and dispersion (Hiebert, 2015; Statistics Canada, 2022e, 2024a). Contemporary settlement patterns reflect "multicultural urbanism" where no single cultural template dominates spatial organization, producing ethnic enclaves that differ fundamentally from traditional immigrant reception areas: Chinese communities in Markham achieve concentrations exceeding 70% while maintaining middle-class profiles; South Asian settlements in Brampton demonstrate multi-nuclear patterns spanning entire municipalities; and Filipino communities establish linear corridors along transit routes (R. Murdie & Ghosh, 2013; Valerie Preston & Brian Ray, 2020; Qadeer, Agrawal, & Lovell, 2010). These patterns exhibit "pluralistic segregation," voluntary clustering that maintains metropolitan diversity while enabling neighborhood concentration, with 384 distinct ethnic clusters in Toronto alone showing characteristic spatial scales of 3.8-6.4 kilometers, yet with dissimilarity indices for some communities (Portuguese: 0.734, Iranian: 0.696) approaching levels traditionally associated with problematic segregation (Fong & Wilkes, 2003; Statistics Canada, 2022c; Walks & Bourne, 2006).

The relationship between ethnic concentration and economic outcomes defies simple characterization, as enclaves in Montreal correlate closely with poverty while those in Toronto and Vancouver often represent pathways to homeownership and middle-class stability (Hiebert, 2015; Hou & Picot, 2016; Shuguang Wang & Zhong, 2013). The housing affordability crisis intensifying since 2020, with Toronto's average home price exceeding \$1.2 million and rental vacancy rates below 0.5%, has fundamentally altered settlement patterns, as immigration contributed to 88% of population growth while housing starts declined by 10% between 2021-2023, dispersing newcomers to peripheral municipalities or into overcrowded accommodations described as "hidden homelessness" within ethnic communities (Cmhc, 2024; R. Murdie & Teixeira, 2011; V. Preston & B. Ray, 2020).




Figure 1.1 Conceptual framework of multi-ethnic settlement challenges in Canadian cities.

Institutional responses reveal tensions between Canada's official multiculturalism policy and practical urban diversity management. While the 1988 Multiculturalism Act commits the federal government to promoting equitable participation, municipal governments bearing primary responsibility for housing and services operate without comparable frameworks, manifesting in zoning bylaws prohibiting multi-generational housing arrangements, consultation processes under-representing non-English speakers, and service delivery models assuming cultural homogeneity (Burayidi, 2015; Canadian Multiculturalism Act, 1988; Good, 2009; Sandercock & Lyssiotis, 2003). Settlement services originally designed for modest European immigrant flows struggle to accommodate annual intakes exceeding 500,000 permanent residents plus over one million temporary residents, with neighborhoods experiencing complete demographic transformation within single decades, as in Toronto's Thorncliffe Park where residents from over 50 countries speak 27 languages, creating "super-diversity" where traditional multicultural frameworks lose analytical utility (Ahmadi & Tasan-Kok, 2014; Biles, Burstein, & Frideres, 2008; IRCC, 2024).

The COVID-19 pandemic exposed underlying vulnerabilities, with visible minority neighborhoods experiencing infection rates 2-3 times higher than predominantly white areas and 40% of businesses in

ethnic commercial districts closing permanently, revealing how ethnic concentration can create collective vulnerabilities when external shocks affect entire communities simultaneously (Public Health Ontario, 2021; Statistics Canada, 2024a). The spatial organization of cities increasingly reflects "three cities," areas of concentrated advantage, disadvantage, and transition correlating closely with ethnic composition, challenging narratives of successful integration and suggesting the Canadian mosaic may be solidifying into permanent parallel societies (D. J. Hulchanski et al., 2010; O'Neill, Gidengil, & Young, 2012\).

The challenge of multi-ethnic settlement thus emerges not as a single problem but as a complex of interrelated dynamics forming a self-reinforcing system where outcomes feed back into patterns, creating path dependencies increasingly difficult to alter. The unprecedented scale and diversity of contemporary immigration, combined with housing market pressures, institutional limitations, and pandemic impacts, have created conditions that existing models, whether from classical urban sociology or comparative immigration studies, struggle to explain. This context provides the imperative for developing new analytical tools, such as the GeoAI methods explored in this thesis, that can capture the complexity of contemporary urban ethnic dynamics while providing actionable insights for creating more equitable and cohesive cities.

1.1.2 From Traditional Geography to Computational Approaches

The study of human settlement patterns has undergone profound transformations over the past century, evolving from descriptive regional accounts to sophisticated computational modeling frameworks. Traditional settlement geography, rooted in early 20th century works of geographers like Carl Sauer and Preston James, emphasized detailed description of particular places through qualitative observation and field mapping (Stone, 1965). The theoretical foundations were established through key conceptual frameworks, most notably Walter Christaller's Central Place Theory proposing that settlements function as central places with size, spacing, and functions following predictable mathematical relationships, though built upon restrictive assumptions including isotropic plains and rational economic actors (Lösch, 1954).




Figure 1.2 Evolution from Traditional Settlement Geography to GeoAI: Paradigmatic Shifts and Methodological Transformations.

As urban systems grew increasingly complex, traditional methods revealed fundamental limitations: the descriptive regional approach struggled to identify generalizable patterns; Central Place Theory demonstrated significant empirical shortcomings when tested in regions like Iowa and Wisconsin; and its static framework could not account for temporal dynamics, technological change, or polycentric urban regions (Berry, 1967; Burton, 1963; Daniels & Warnes, 2013; Harvey, 1969; Stone, 1965). The quantitative revolution during the 1950s and 1960s represented the first major methodological transformation, introducing statistical techniques and mathematical modeling led by geographers like William Bunge, Brian

Berry, and Peter Haggett, while the emergence of Geographic Information Systems in the 1980s provided new capabilities for spatial data integration, though early GIS applications focused primarily on data management and cartographic visualization rather than analytical modeling (Burton, 1963; Goodchild, 1992; Harvey, 1969).

The contemporary emergence of GeoAI represents a paradigmatic shift that transcends previous limitations by integrating artificial intelligence with geographic information science, characterized by machine learning algorithms identifying complex non-linear patterns in high-dimensional datasets, deep learning approaches processing heterogeneous spatial data including satellite imagery and mobile phone records, and predictive modeling that can forecast settlement evolution and simulate policy interventions (Gao, 2021; Janowicz, Gao, McKenzie, Hu, & Bhaduri, 2020; W. Li, 2020; VoPham, Hart, Laden, & Chiang, 2018). Most significantly, GeoAI approaches can model the complex multi-scale interactions that traditional methods could not address, with graph neural networks simultaneously analyzing individual settlement decisions and regional urban systems while capturing feedback loops and emergent properties (Goodfellow, Bengio, & Courville, 2016). This evolution from traditional descriptive geography through quantitative approaches to contemporary GeoAI represents a fundamental reconceptualization of how settlement processes can be understood, enabling integration of pattern recognition, prediction, and explanation within frameworks that handle full urban complexity while remaining grounded in geographical theory.

1.1.3 GeoAI: Bridging Spatial Science and Artificial Intelligence

Geospatial Artificial Intelligence (GeoAI) represents the convergence of spatial science, artificial intelligence, and geographic information systems to address complex geographic problems through advanced computational methods. Defined as spatially explicit AI techniques for geographic knowledge discovery, GeoAI integrates machine learning with geospatial data to accelerate understanding of spatial phenomena, transcending traditional GIS analysis by embedding spatial thinking directly into algorithmic design (W. Li, Hsu, & Hu, 2021). This spatially-aware approach fundamentally differentiates GeoAI from conventional AI applications by recognizing that spatial data possess unique properties, including nonlinearity, spatial non-stationarity, and multi-scale variability, that violate assumptions of traditional statistical methods and necessitate specialized computational approaches (Kanevski, Timonin, & Pozdnukhov, 2009).




Figure 1.3 Conceptual Architecture of GeoAI: Bridging Spatial Science and Artificial Intelligence for Urban Analysis.

The methodological toolkit of GeoAI encompasses both established machine learning techniques adapted for spatial contexts, such as geographically weighted regression and spatial clustering, and inherently spatial methods that integrate geographic principles into algorithmic design (Fotheringham, Brunson, & Charlton, 2009; Zhou et al., 2020). Graph neural networks exemplify inherently spatial approaches by explicitly modeling spatial relationships through network topologies that capture connectivity patterns in urban

systems, while physics-informed neural networks (PINNs) combine machine learning with established physical theories, respecting known constraints while leveraging data-driven learning to capture complex urban processes (Karniadakis et al., 2021; Raissi, Perdikaris, & Karniadakis, 2019; Zhou et al., 2020).

The black box nature of many AI systems poses particular challenges for democratic urban governance, as algorithms making decisions affecting housing, transportation, or service delivery often cannot be understood, challenged, or modified by citizens and elected officials, undermining fundamental principles of democratic accountability (Cugurullo et al., 2024). Critical urban studies scholars argue that AI-driven governance represents a shift toward technocratic administration where traditional mechanisms of citizen engagement become marginalized, reflecting what has been identified as the inherently political nature of technological systems where urban AI embodies particular values, assumptions, and power relationships while appearing to operate through technical rather than political logics (Winner, 2017). Recent scholarship emphasizes algorithmic sovereignty and community control, with cities like Amsterdam and Helsinki pioneering algorithmic registries that demonstrate transparency and effectiveness can be mutually reinforcing values, perspectives with particular relevance for GeoAI applications in demographic analysis given histories of discriminatory mapping, surveillance, and spatial control (Saidot, 2025).

1.1.4 Research Direction

The evolution of urban ethnic settlement research has undergone three fundamental transformations that define the trajectory of contemporary GeoAI applications in human geography. These shifts, from single-group to multi-ethnic analysis, from static patterns to dynamic processes, and from descriptive to predictive modeling, represent not merely methodological advances but paradigmatic reconceptualizations of how urban diversity can be understood and analyzed through computational approaches.

Our research journey exemplifies this evolution, beginning with exploratory applications of deep learning architectures (detailed in Appendices A and B) that, while achieving technical sophistication, revealed critical limitations in interpretability and mechanistic understanding. These initial investigations into spatial pattern detection and network boundary analysis, though computationally advanced, ultimately demonstrated the necessity of developing frameworks that maintain both predictive accuracy and theoretical transparency, leading to the three core contributions presented in Chapters 3, 4, and 5.

Figure 1.4 illustrates these three paradigmatic transformations and their integration through interpretable GeoAI frameworks, demonstrating how each shift reinforces the others to create a comprehensive analytical approach for understanding multi-ethnic urban dynamics.




Figure 1.4 Research Direction Framework: Three Paradigmatic Transformations in Urban Settlement Studies.

1.1.4.1 From Single-Group to Multi-Ethnic Analysis

Traditional settlement geography has long been constrained by what has been termed "methodological groupism," the tendency to study ethnic communities as isolated entities rather than interconnected components of complex urban systems (Vertovec, 2013). This single-group focus, inherited from Chicago School sociology and perpetuated through census categorization practices, fundamentally misrepresents the reality of contemporary superdiverse cities where inter-ethnic interactions shape settlement outcomes as powerfully as intra-group dynamics (Foner, Duyvendak, & Kasinitz, 2020). The limitations become particularly evident in Toronto's context, where no single ethnic group dominates and settlement patterns

emerge from complex multi-ethnic negotiations over space, resources, and opportunities, as demonstrated in Chapter 5's analysis revealing that Chinese populations exert strong influence on Portuguese (0.44) and Indian (0.37) communities, while Sri Lankan and Portuguese populations show strong bidirectional interaction, dynamics invisible to single-group methodologies.

This shift toward multi-ethnic analysis necessitates fundamental reconceptualization of analytical frameworks. Rather than treating ethnic groups as independent variables in regression models or separate layers in GIS analysis, contemporary GeoAI approaches model settlement as emergent from inter-ethnic dynamics, with graph neural networks enabling representation of ethnic communities as nodes in interaction networks where edge weights capture influence strengths and directionalities (Zhou et al., 2020). The Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework developed in Chapter 5 exemplifies this approach by incorporating an ethnicity interaction module that quantifies cross-ethnic effects, revealing cascade phenomena where demographic changes in one community trigger reorganization across multiple groups. The multi-ethnic perspective also transforms understanding of spatial segregation from a binary integrated-segregated dichotomy to a multidimensional space of ethnic configurations, with the discovery of "pluralistic segregation" in Canadian cities, where voluntary clustering maintains diversity at metropolitan scales while enabling concentration at neighborhood levels, confirmed through distinct ethnic spatial patterns across 52 Canadian cities with characteristic spatial scales ranging from 34.5 km (Philippines) to 63.0 km (United Kingdom) (Walks & Bourne, 2006).

1.1.4.2 From Static Patterns to Dynamic Processes

The second shift involves moving beyond static snapshots of settlement patterns to modeling dynamic processes of neighborhood formation, evolution, and transformation. Traditional geographic methods, constrained by decennial census data and cross-sectional analysis, could only document settlement patterns at discrete temporal points without capturing the continuous processes generating these patterns, obscuring critical dynamics including colonization of new areas, tipping points in neighborhood composition, succession patterns between ethnic groups, and feedback loops between settlement and urban infrastructure (Logan & Zhang, 2010). Our initial explorations using deep learning architectures (Appendices A and B) demonstrated the technical feasibility of capturing temporal dynamics but highlighted the critical need for interpretable process models, leading to physics-informed approaches that enable fundamentally different treatments through continuous-time models and process-based simulations. The GraphPDE framework presented in Chapter 4 demonstrates how population dynamics can be modeled as physical processes governed by diffusion coefficients (2.9-6.8 km²/year across cities) and attention-based interaction mechanisms, revealing that settlement patterns emerge from reaction-diffusion dynamics rather than random aggregation.

The temporal analysis reveals distinct phases in settlement evolution previously undetectable, with the model capturing non-stationary dynamics through period-specific parameters learned separately for four census intervals (2001-2006, 2006-2011, 2011-2016, 2016-2021). Philippines-origin populations maintained consistently positive growth rates ($0.015\text{--}0.023\text{ year}^{-1}$) across all periods, while United Kingdom-origin populations transitioned from positive growth ($+0.010\text{ year}^{-1}$) in 2001-2006 to sustained decline ($-0.015\text{ to }-0.023\text{ year}^{-1}$) in subsequent periods, with nucleation thresholds varying from 658 individuals (China) to 8,132 individuals (India). Dynamic modeling also reveals path dependencies and multi-stability effects, with configuration landscape analysis identifying multiple stable settlement configurations with quantifiable transition barriers, demonstrating that demographic transitions require sustained interventions over decadal timescales and explaining why some neighborhood transitions appear irreversible through asymmetric barriers that create lock-in effects.

1.1.4.3 From Descriptive to Predictive Modeling

The third transformation, from descriptive documentation to predictive modeling, represents perhaps the most profound shift enabled by GeoAI methodologies. Traditional settlement geography excelled at categorizing patterns, measuring concentrations, and mapping distributions but struggled to forecast future states or simulate intervention effects, severely constraining urban planning applications where policymakers require anticipation of future trajectories under different scenarios (Musterd, 2005). While our initial approaches (Appendices A and B) achieved high predictive accuracy through black-box architectures, they offered limited insight into driving mechanisms, motivating development of interpretable models that maintain transparency while achieving state-of-the-art performance. The Cultural-Spatial Resonance Network developed in Chapter 5 demonstrates that incorporating domain knowledge about population dynamics through partial differential equations enhances rather than compromises predictive accuracy, achieving R^2 values of $0.80\text{--}0.83$ while providing interpretable parameters for diffusion, amplification, and directional movement.

Contemporary GeoAI approaches achieve predictive capabilities through multiple mechanisms, but as Chapter 3's critical analysis demonstrates, accuracy without interpretability risks creating "metrics traps" where technical performance obscures social failures. This insight shaped our development of physics-informed approaches that capture non-linear dynamics, tipping points, and cascade effects while maintaining full transparency, with the GraphPDE framework achieving $R^2 = 0.7649$, approaching black-box XGBoost ($R^2 = 0.7760$) while providing complete mechanistic interpretability through learned physics parameters accounting for 57.2% of prediction variance. Predictive modeling also enables counterfactual analysis and policy simulation previously impossible with descriptive methods, with the discovery that transit accessibility correlates with amplification factors for Chinese and Filipino populations ($r = 0.58$ and

$r = 0.61$) while housing variables matter more for Indian populations ($r = 0.47$) providing actionable intelligence for culturally-responsive planning interventions (Chapter 5).

1.1.4.4 Synthesis and Implications

These three research directions, multi-ethnic, dynamic, and predictive, are not independent trajectories but mutually reinforcing transformations that collectively redefine urban settlement studies. Multi-ethnic analysis reveals interaction networks that drive dynamic processes, which in turn enable predictive models capturing system complexity, representing what has been envisioned as a "new science of cities" where computational methods reveal organizational principles invisible to traditional observation (M. Batty, 2013). Our research journey from exploratory deep learning to physics-informed modeling reflects this evolution, with initial investigations (Appendices A and B) providing crucial insights into capabilities and limitations of purely data-driven approaches, informing the interpretable frameworks presented in Chapters 3-5. The research direction also reflects broader epistemological shifts in human geography from idiographic description toward nomothetic explanation, from structure to process, and from documentation to intervention, yet unlike the quantitative revolution's pursuit of universal laws, contemporary GeoAI maintains sensitivity to context, contingency, and complexity through flexible architectures that adapt to local conditions while identifying generalizable patterns (Barnes, 2004).

Critically, these methodological advances must be accompanied by ethical frameworks ensuring that enhanced analytical capabilities serve democratic rather than technocratic ends, as demonstrated through Chapter 3's analysis showing that technical sophistication without social responsibility risks automating inequality rather than advancing equity. The research direction thus points toward a future where interpretable GeoAI enables not just better understanding of urban ethnic dynamics but active shaping of more equitable, integrated cities. By revealing the mechanisms driving settlement patterns, identifying intervention points in demographic transitions, and predicting outcomes of policy choices, these computational approaches provide tools for evidence-based urban planning that respects complexity while pursuing social justice, with the challenge lying not in technical capabilities but in ensuring these powerful analytical tools serve community flourishing rather than algorithmic control.

1.2 Research Gap and Positioning

1.2.1 The Critical Gap: From Isolated Ethnic Geographies to Multi-Ethnic Urban Systems

Despite Toronto's transformation into a superdiverse metropolis where 250+ ethnic groups create complex settlement mosaics, geographic research remains fragmented into isolated studies of individual communities, fundamentally misrepresenting how contemporary multi-ethnic cities actually function (Carlos Teixeira, 2008; R. Thomas, 2011; Walton-Roberts, 2003; Shuguang Wang & Zhong, 2013). This

scholarly fragmentation becomes particularly problematic given that superdiversity represents not simply more ethnic groups but a "diversification of diversity" where interactions between groups fundamentally reshape urban space, with research documenting that supposed Chinese "enclaves" in Markham actually house residents from over 50 different origins and that 51-75% of enclave residents belong to different ethnicities than the plurality group (Hiebert, 2015; Qadeer et al., 2010; Vertovec, 2013).

The persistence of single-group studies has rendered invisible the inter-ethnic dynamics that increasingly drive neighborhood transformation. Contemporary Toronto exhibits three types of inter-ethnic spatial processes that existing research cannot adequately capture: ethnic succession and displacement patterns where no frameworks explain why certain ethnic sequences occur (Siemiatycki, Rees, Ng, & Rahi, 2001); multi-ethnic cascade effects where one group's settlement triggers reorganization across multiple communities (Bauder & Sharpe, 2002); and emergent settlement configurations in "super-diverse neighborhoods" like Thorncliffe Park housing over 30,000 residents from South Asia, East Africa, and Eastern Europe without any group exceeding 30% (R. A. Murdie, 2008). Perhaps most critically, while extensive scholarship examines within-group dynamics including bonding social capital, kinship networks, and religious institutions, virtually no research systematically examines spatial relationships between groups, leaving relational geographies unexplored because existing frameworks treat ethnic groups as discrete spatial units rather than interacting urban actors (Geraldine Pratt, 2009).

1.2.2 The Imperative for Integrated Multi-Ethnic Urban Analysis

The shift from studying individual ethnic geographies to analyzing multi-ethnic urban systems represents a fundamental advance paralleling transformative shifts in other geographic subfields (Berry, 1967; Rocha et al., 2021; Scott, 2004). This evolution promises several theoretical breakthroughs: revealing urban ethnic ecosystems through mapping "community assemblages" of groups that consistently co-occur versus those exhibiting competitive exclusion (Qadeer & Kumar, 2006; Carlos Teixeira, 2008; Zhuang, 2019); understanding neighborhood tipping points where 35% ethnic concentration triggers reorganization but single-group studies cannot explain how thresholds vary across ethnic combinations (Schelling, 1971; Walks & Bourne, 2006); and identifying spatial competition and complementarity as ethnic communities compete for housing while creating complementary ethnic economies (M. E. Porter, 2012).

The selection of nine specific groups (China, India, Philippines, Pakistan, Sri Lanka, Iran, Portugal, Italy, and United Kingdom) captures the essential diversity of Toronto's settlement geography while remaining analytically tractable, collectively representing 66.7% of Toronto's immigrant population with maximum variation in settlement patterns, growth trajectories, and spatial relationships (Zhuang & Chen, 2017). With nine groups generating 36 unique pairwise relationships plus higher-order interactions, the analysis captures sufficient complexity to model real multi-ethnic dynamics, avoiding both the unmanageable complexity of

analyzing all 200+ Toronto ethnic origins and the oversimplification of two-group comparisons (Myles & Hou, 2004). The selected groups span all documented settlement trajectories from explosive growth (India: +340%, Pakistan: +334%) through moderate expansion (Philippines: +231%, China: +219%) to sustained decline (Italy: -24%, UK: -20%), while embodying Toronto's layered immigration history enabling analysis of how successive migration periods interact spatially (David Ley, 2011).

1.2.3 GeoAI as Solution Framework

Geospatial Artificial Intelligence emerges as the essential framework for implementing progressive spatial analysis, providing computational tools capable of handling the complexity, dimensionality, and dynamics of contemporary urban ethnic settlement through spatially-explicit artificial intelligence that incorporates geographic principles directly into algorithmic design (Janowicz et al., 2020). GeoAI provides unprecedented capabilities for analyzing complex urban systems: deep learning architectures processing hundreds of census variables simultaneously; graph neural networks representing urban spatial relationships that tabular methods cannot accommodate; attention mechanisms automatically learning the relevant scale and extent of spatial influence; physics-informed neural networks incorporating domain knowledge as inductive biases; and mixture of experts architectures recognizing that different processes govern different regimes.

Critically, GeoAI as implemented in this research prioritizes interpretability through systematic integration of domain knowledge rather than pure data-driven learning. The Spatial Gaussian-Bernoulli DBN incorporates street network topology as architectural constraints; the STGAN-EI framework learns explicit inter-ethnic interaction matrices; the reaction-diffusion model grounds analysis in established physics of pattern formation; and the MESMoE framework achieves state-of-the-art performance while maintaining full interpretability through regime-specific parameterization. The learned parameters have clear geographic meaning: diffusion coefficients represent mobility rates, interaction strengths capture ethnic affinity or repulsion, and energy barriers quantify resistance to demographic change. GeoAI serves as a crucial bridge between theoretical understanding and practical application, transforming abstract concepts into measurable, predictable phenomena with clear policy implications.

Table 1-1. Comparison of Analytical Approaches to Urban Ethnic Settlement

Dimension	Traditional Geography	Conventional MI/AI	Geoai Framework (This Research)
Theoretical Approach	Single-group focus, static patterns, fixed scales	Atheoretical, data-driven patterns	Multi-ethnic dynamics, evolutionary processes, multi-scale integration
Data Handling	Limited variables, assumes independence	High-dimensional but non-spatial	Spatially-explicit, preserves dependencies

Spatial Representation	Euclidean distance, administrative units	Often ignored or simplified	Street networks, multiple proximities
Temporal Analysis	Cross-sectional snapshots	Time as feature	Process modeling, transitions
Interpretability	Full transparency, limited complexity	Black box, high accuracy	Interpretable parameters with mechanistic meaning
Inter-Ethnic Dynamics	Assumed or ignored	Implicit in patterns	Explicit interaction matrices
Predictive Capability	Limited to extrapolation	Strong but unexplained	Mechanistic prediction with confidence bounds
Scalability	Computationally simple but limited	Scales well technically	Efficient gpu implementation for city-scale analysis
Policy Relevance	Direct but oversimplified	Accurate but opaque	Actionable insights with clear intervention points
Validation Approach	Statistical significance	Cross-validation accuracy	Multi-criteria: accuracy, interpretation, theory consistency

This framework transforms settlement geography from a descriptive to a predictive science while maintaining the theoretical grounding and interpretability essential for urban planning applications. By providing tools that can simultaneously handle complexity, maintain transparency, and generate actionable insights, GeoAI enables a new generation of evidence-based urban policy that respects the full complexity of superdiverse cities.

1.2.4 Positioning This Research

This dissertation positions itself at the intersection of human geography and computational science, advancing both fields while maintaining critical awareness of algorithmic power and social implications. The research makes four distinctive contributions that collectively address identified gaps:

First, it develops progressive spatial analysis as a systematic methodology for studying complex urban phenomena through integrated multi-scale investigation. This approach moves beyond both the limitations of traditional single-scale analysis and the black box problem of conventional machine learning, providing a framework that builds cumulative understanding through systematic progression across spatial scales.

Second, it demonstrates interpretable GeoAI that achieves state-of-the-art performance without sacrificing transparency. By incorporating domain knowledge through physics-informed architectures and producing parameters with clear geographic meaning, the research refutes supposed trade-offs between capability and interpretability, showing that the best models are often the most understandable.

Third, it reveals hidden ethnic dynamics previously invisible to both traditional and computational approaches. The quantification of inter-ethnic interactions through attention-based mechanisms, identification of nucleation thresholds varying systematically across ethnic groups, and discovery of pattern

formation regimes (spots, stripes, labyrinthine) transform qualitative observations into precise, predictable phenomena amenable to policy intervention.

Fourth, it provides practical tools that bridge the gap between academic research and urban planning practice. The open-source implementations, documented workflows, and transferable architectures ensure that research insights can translate into real-world applications, democratizing access to advanced urban analytics.

By addressing theoretical limitations through multi-ethnic frameworks, methodological constraints through progressive analysis, and empirical gaps through GeoAI capabilities, this research establishes new foundations for understanding and managing ethnic settlement in superdiverse cities. The approach respects both the complexity of contemporary urban systems and the need for transparent, accountable methods that serve democratic urban governance rather than technocratic control.

1.3 Research Questions and Research Scheme

1.3.1 Research Questions

This dissertation investigates the spatial dynamics and underlying processes of multi-ethnic settlement patterns in contemporary metropolitan areas, with 52 major Canadian cities serving as the analytical scope. While initially proposed to examine Iranian immigrant settlement patterns specifically, the research evolved to embrace a multi-ethnic framework that captures the interconnected nature of urban demographic dynamics. This evolution from single-group to multi-ethnic analysis reflects a fundamental geographic insight discovered during preliminary investigations: ethnic settlement patterns cannot be understood in isolation but emerge from complex inter-group spatial interactions, place-based processes, and temporal dynamics that shape the urban social landscape.

The overarching research question guiding this dissertation is:

How do multi-ethnic settlement patterns in metropolitan areas emerge, evolve, and stabilize through the interplay of spatial processes, inter-ethnic dynamics, and urban structural factors, and what do these patterns reveal about the nature of contemporary urban diversity and social-spatial organization?

This central question is addressed through three core research questions, each corresponding to a substantive chapter, with two additional exploratory questions that informed the theoretical and methodological development:

1.3.1.1 Core Research Questions

RQ1: How does technology-mediated urban governance shape ethnic settlement patterns and community outcomes, and what are the implications for spatial justice and democratic participation in diverse cities? This foundational question (Chapter 3) examines the deployment of algorithmic decision-making systems in urban governance and their impacts on ethnic communities. Through analysis of 24 case studies selected from 157 urban AI deployments across six domains (2015-2024), the investigation reveals how cities have fallen into a “metrics trap”-where technical efficiency undermines social equity and community wellbeing. This critical geographic analysis establishes why understanding the mechanisms of settlement patterns matters for spatial justice, providing the ethical foundation for examining how vulnerable ethnic communities navigate increasingly algorithmic urban environments.

RQ2: What mathematical frameworks can capture the statistical regularities in ethnic settlement patterns across metropolitan areas, and how do physics-inspired models help us understand aggregate spatial dynamics while recognizing that individual settlement decisions involve complex social factors beyond what these models represent?

This question (Chapter 4) represents a paradigmatic shift in understanding urban ethnic geography, investigating the fundamental spatial mechanisms that produce observed settlement patterns. The research develops a graph-based physics-informed neural network (GraphPDE) that embeds multi-ethnic reaction-diffusion dynamics while learning interpretable demographic parameters across 52 Canadian cities and 30,091 dissemination areas. The model reveals ethnic-specific spatial scales ranging from 34.5 km (Philippines) to 63.0 km (United Kingdom), pattern formation regimes (spots, stripes, labyrinthine) emerging from the interplay of self-interaction and between-group competition, and critical nucleation thresholds varying from 658 individuals (China) to 8,132 individuals (India) for spatial clustering emergence. This geographic analysis reveals how settlement patterns exhibit self-organizing properties that transcend individual choices, following spatial dynamics similar to other pattern-forming processes in nature and society.

RQ3: How do different demographic regimes-colonization of new areas, population decline, and spatial redistribution-shape the evolution of multi-ethnic neighborhoods, and what role do inter-ethnic interactions and urban infrastructure play in determining settlement trajectories?

This synthesis question (Chapter 5) explores how multi-ethnic cities transition between different demographic states and the factors that determine neighborhood futures. Through examining regime-specific dynamics, the research reveals how different processes govern settlement changes: new ethnic enclaves forming in previously homogeneous areas, established communities experiencing decline or gentrification, and populations redistributing across metropolitan space. The investigation demonstrates how inter-ethnic influences create cascade effects across communities, with certain groups acting as

demographic anchors while others follow or avoid their settlement patterns, fundamentally shaping the urban ethnic mosaic.

1.3.1.2 Exploratory Research Questions (Addressed in Appendices)

ERQ1: What spatial patterns characterize multi-ethnic settlement in contemporary Toronto, and how can these patterns be systematically identified and classified? (Appendix A)

This exploratory investigation examined the complex spatial configurations of ethnic settlement across Toronto's 3,741 dissemination areas, identifying distinct typologies including business-centered enclaves, education-transit corridors, and distributed multi-nuclear patterns. While successfully mapping these patterns, the analysis revealed that identifying patterns alone provides insufficient understanding of the processes creating them, motivating deeper investigation into settlement mechanisms.

ERQ2: How do inter-ethnic spatial interactions shape settlement patterns, and what critical thresholds trigger neighborhood demographic transitions? (Appendix B)

This exploration quantified the spatial relationships between different ethnic communities, identifying critical tipping points where 35% ethnic concentration triggers multi-group reorganization. The analysis revealed asymmetric inter-ethnic influences-some groups attract others while some repel-fundamentally shaping the geographic distribution of diversity across the metropolitan area.

These research questions reflect a deliberate progression from examining the urban governance context shaping settlement (Chapter 3), through understanding the fundamental spatial processes driving pattern formation (Chapter 4), to analyzing the complex dynamics of demographic change in multi-ethnic cities (Chapter 5). The exploratory work documented in the appendices provided crucial empirical grounding for understanding Toronto's ethnic geography, while the core chapters develop theoretical frameworks for understanding these patterns. Together, they advance human geography's understanding of how contemporary cities organize themselves spatially under conditions of unprecedented ethnic diversity, contributing both theoretical insights about urban spatial processes and practical knowledge for creating more equitable and integrated metropolitan areas.

1.4 Methodology

1.4.1 Thesis Framework

This dissertation employs a progressive GeoAI framework that advances from critical assessment through interpretable computational analyses to reveal the multi-layered dynamics of ethnic settlement patterns across major Canadian cities (Figure 1.5). The overarching methodological approach integrates physics-informed neural networks, interpretable machine learning architectures, and reaction-diffusion models

within a critical geographic perspective that interrogates both the capabilities and limitations of algorithmic approaches to urban demographics.

The framework addresses a fundamental methodological challenge: ethnic settlement patterns cannot be understood through isolated analysis of individual groups but emerge from complex inter-ethnic interactions, spatial dependencies, and temporal dynamics that require simultaneous analysis of multiple populations. This methodological necessity-discovered through the inadequacy of single-group models in our exploratory work (Appendices A and B)-shaped the research design toward comprehensive multi-ethnic analysis that captures the interconnected nature of demographic landscapes across Canadian cities.

The thesis framework operates through three interconnected principles, as illustrated in Figure 1.5. First, **interpretable progression** moves from identifying algorithmic failures in urban governance (Chapter 3) to developing physics-informed understanding of settlement mechanisms (Chapter 4) to creating hybrid frameworks that combine accuracy with transparency (Chapter 5). Second, **methodological learning** demonstrates how limitations discovered in exploratory black-box approaches (Appendices A and B) motivated the development of interpretable alternatives. Third, **critical computational geography** maintains reflexivity about algorithmic power, ensuring that technical sophistication serves democratic urban planning rather than technocratic control.




Figure 1.5 Progressive GeoAI methodological framework for analyzing multi-ethnic settlement patterns in the Greater Toronto Area.

1.4.2 Critical Analysis - Establishing Ethical Foundations

The first analytical stage (Chapter 3) employs multiple case study analysis with thematic synthesis to examine algorithmic failures in urban governance, establishing why interpretability is essential for responsible GeoAI applications in multi-ethnic contexts. The analysis examines 24 implementations selected from 157 AI deployments across six urban domains (2015-2024), encompassing predictive policing systems, housing allocation algorithms, facial recognition deployments, and demographic prediction models across global cities. Systematic cross-case analysis employed pattern matching, explanation building, and qualitative comparative analysis principles, with document analysis encompassing academic publications, government reports, technical documentation, investigative journalism, legal documents, and community advocacy materials, using NVivo 12 software achieving inter-rater reliability of 0.84 (Cohen's kappa).

The analysis reveals three critical patterns establishing the necessity for interpretable methods: accuracy illusions where technical metrics mask social failures (e.g., ShotSpotter's 97% acoustic accuracy yielding only 9.1% crime-fighting effectiveness); discriminatory feedback loops amplifying existing biases through algorithmic enforcement; and community resistance movements demonstrating successful challenges to algorithmic determinism. These findings establish that technical sophistication without interpretability risks automating inequality rather than advancing equity.

1.4.3 Physics-Informed Analysis - Metropolitan Scale Dynamics

The second analytical stage (Chapter 4) develops a graph-based physics-informed neural network (GraphPDE) that embeds multi-ethnic reaction-diffusion dynamics while learning interpretable demographic parameters. Analysis encompasses 52 major Canadian cities tracking nine major ethnic groups (China, Philippines, India, Pakistan, Iran, Sri Lanka, Portugal, Italy, United Kingdom) across 30,091 dissemination areas from 2001-2021. The GraphPDE framework combines reaction-diffusion partial differential equations on spatial graphs with neural augmentation, incorporating learnable diffusion coefficients with ethnic-specific and city-specific values, attention-based inter-ethnic competition mechanisms, growth rates, immigration rates, emigration rates, and carrying capacities learned per-period and per-city, census feature modulation through learned multilayer perceptrons, and neural residual terms capturing systematic deviations from physics structure. The model employs GPU-accelerated implementation with custom CUDA kernels achieving $32.6\times$ speedup, enabling city-scale simulations with 4th-order Runge-Kutta temporal integration.

Key discoveries include: ethnic-specific spatial scales ranging from United Kingdom (63.0 km, highest mobility) to Philippines (34.5 km, concentrated settlement); pattern formation regimes segregating ethnic groups into spots (UK, Portugal), stripes (China, India, Philippines), and labyrinthine (Iran, Pakistan, Sri

Lanka) patterns; critical nucleation thresholds varying from China (658 individuals) to India (8,132 individuals); multi-stability analysis revealing multiple stable settlement configurations with quantifiable transition barriers; and interpretable variance attribution establishing that physics-informed components account for 57.2% of prediction variance while neural residuals contribute 42.8%.

1.4.4 Hybrid Framework - Multi-scale Integration

The third analytical stage (Chapter 5) employs the Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework integrating domain knowledge with machine learning. Data synthesis integrates census demographic profiles (298 socioeconomic features), housing characteristics, employment and income distributions, transportation infrastructure metrics, educational facilities, and three specialized spatial matrices (walking distance, transit connectivity, geographic proximity), processing temporal transitions (2001→2006, 2006→2016, 2016→2021) as separate training examples yielding approximately 11,000 temporal transitions.

MESMoE employs regime-adaptive partial differential equations with specialized neural experts: Colonization Expert modeling zero-to-nonzero transitions; Jump Process Expert capturing discontinuous demographic shifts; Decline Expert specializing in population exodus scenarios; and PDE Diffusion Expert handling continuous population changes. A learnable router directs predictions to appropriate experts based on local context. The framework achieves state-of-the-art performance ($R^2 = 0.80-0.83$ across periods) while maintaining full interpretability through expert-specific feature importance analysis, learned ethnicity interaction networks, spatial parameter visualization, and regime-specific accuracy metrics.

1.4.5 Exploratory Investigations - Learning from Black-Box Methods

Two exploratory investigations (Appendices A and B) provided crucial methodological insights shaping the interpretable frameworks. Appendix A employed Spatial Gaussian-Bernoulli Deep Belief Networks (401→64 dimensional reduction) combined with GraphSAGE, achieving exceptional accuracy ($R^2 > 0.95$ for major ethnic groups) but with opacity of learned representations preventing understanding of why patterns emerged. Appendix B employed Spatio-Temporal Graph Attention Networks with Ethnic Interaction Learning (STGAN-EI), achieving strong predictive capability ($R^2 = 0.851$) with identification of 35% concentration tipping points, but black-box attention mechanisms obscured causal relationships. These explorations demonstrated that while deep learning achieves technical sophistication, interpretability is essential for advancing geographic theory and enabling practical applications.

1.4.6 Computational Implementation and Ethical Considerations

All methods prioritize reproducibility and accessibility through Python 3.8+ with PyTorch 1.9+ for neural network implementations, PyTorch Geometric for graph neural network components, custom CUDA 11.0+

kernels for physics-informed PDE solvers, and comprehensive documentation. Census data processed according to Statistics Canada guidelines, spatial matrices, and complete preprocessing pipelines are available on Figshare and GitHub repositories.

The methodology incorporates ethical safeguards throughout: privacy protection through analysis at dissemination area level (400-700 people) preventing individual identification with no use of personal mobility data; and algorithmic accountability through full interpretability of core models, documentation of limitations and potential biases, and open-source implementation allowing community verification. This methodological framework transforms the study of multi-ethnic settlement from descriptive documentation to interpretable science, demonstrating that computational sophistication need not sacrifice transparency or theoretical understanding.

1.4.7 Thesis Structure

This dissertation is organized into six core chapters, with additional appendices documenting exploratory investigations that informed the methodological development. The structure follows a deliberate progression from critical examination of algorithmic urban governance through physics-informed analysis to integrated frameworks, reflecting the evolution from identifying problems with black-box methods to developing interpretable alternatives.

Chapter 1: Introduction establishes the research context, presenting Toronto as an unprecedented laboratory for studying multi-ethnic urban dynamics where 55.7% of the population belongs to visible minority groups. The chapter traces the evolution from traditional geographic approaches through the quantitative revolution to contemporary GeoAI, identifying critical gaps in understanding multi-ethnic settlement as an interconnected system rather than isolated group phenomena. It positions the research at the intersection of human geography and computational science, establishing the necessity for interpretable methods that can reveal settlement mechanisms while serving democratic urban governance.

Chapter 2: The Greater Toronto Area as a Multi-Ethnic Laboratory provides comprehensive empirical grounding for subsequent analyses. This chapter documents the selection methodology for the nine study populations (China, India, Philippines, Pakistan, Sri Lanka, Iran, Portugal, Italy, United Kingdom) based on population size, growth trajectories, and spatial variation criteria. It presents detailed demographic profiles, temporal dynamics (2001-2021), and contemporary spatial configurations, establishing the empirical foundation for understanding Toronto's transformation into a superdiverse metropolis.

Chapter 3: The Metrics Trap in GeoAI Systems critically examines algorithmic failures in urban governance through analysis of 24 case studies selected from 157 urban AI deployments (2015-2024). This chapter reveals how cities fall into "metrics traps" where technical accuracy masks social failures, establishing the ethical and practical necessity for interpretable methods. The analysis identifies patterns of

accuracy illusions, discriminatory feedback loops, and successful community resistance, providing the critical foundation for developing transparent alternatives to black-box algorithms.

Chapter 4: SPATIAL DYNAMICS - Spatial Diffusion Models of Ethnic Settlement develops the graph-based physics-informed neural network (GraphPDE) that embeds multi-ethnic reaction-diffusion dynamics while learning interpretable demographic parameters across 52 Canadian cities. This chapter demonstrates that physics-informed constraints do not severely limit predictive power-GraphPDE achieves $R^2 = 0.7649$, approaching gradient boosting performance (XGBoost: $R^2 = 0.7760$)-while yielding mechanistically interpretable parameters. The framework reveals ethnic-specific spatial mobility patterns, pattern formation regimes, critical nucleation thresholds, and inter-ethnic interaction networks through attention mechanisms, providing actionable insights for understanding multi-ethnic settlement dynamics.

Chapter 5: Spatial Complexity - Multi-Pattern Models for Complex Urban Demographics develops the Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework that achieves state-of-the-art predictive performance while maintaining full interpretability. This chapter demonstrates how different demographic regimes, colonization, decline, jump transitions, and continuous diffusion, require distinct mathematical formulations. The framework reveals inter-ethnic interaction networks and cascade effects that shape metropolitan-scale settlement patterns, providing actionable insights for culturally-responsive urban planning.

Chapter 6: Synthesis - From GeoAI Patterns to Geographic Understanding integrates findings across the dissertation to advance theoretical understanding of multi-ethnic urban systems. This concluding chapter synthesizes discoveries about universal scaling laws, regime-dependent dynamics, and inter-ethnic interactions into a coherent framework for understanding contemporary urban diversity. It translates computational insights into practical planning guidance while critically reflecting on the limitations of physics-based approaches to social phenomena.

Appendix A: Spatial Patterns - Advancing GeoAI for Mapping Ethnic Settlement Patterns documents initial explorations using hybrid deep learning architectures (Spatial Gaussian-Bernoulli DBN combined with GraphSAGE) for pattern detection. While achieving exceptional technical performance ($R^2 > 0.95$), this investigation revealed critical interpretability limitations that motivated the development of transparent frameworks in the core chapters.

Appendix B: Spatial Networks - Graph Neural Networks for Understanding Ethnic Spatial Boundaries presents the Spatio-Temporal Graph Attention Network with Ethnic Interaction learning (STGAN-EI) that identified critical tipping points in neighborhood transitions. Despite strong predictive capability, the black-box nature of attention mechanisms demonstrated the need for explicit mechanistic modeling that informed subsequent physics-informed approaches.

The thesis structure reflects a methodological journey from exploratory data-driven approaches to theory-informed frameworks, demonstrating that the progression from black-box to interpretable methods represents not a retreat from computational sophistication but an advance toward meaningful understanding. Each chapter builds upon previous insights while maintaining sufficient independence to be read as a standalone contribution, enabling readers to engage with specific methodological or empirical components while appreciating the integrated whole. This organization ensures that the dissertation contributes both to advancing GeoAI methodologies and to deepening geographic understanding of multi-ethnic urban dynamics.

1.5 Anticipated Contributions

This dissertation makes substantial contributions across theoretical, methodological, empirical, and practical dimensions, advancing both human geography and computational urban science. The research demonstrates how interpretable GeoAI methods can reveal the complex mechanisms governing multi-ethnic settlement patterns while avoiding the pitfalls of black-box algorithmic approaches.

1.5.1 Theoretical Contributions

This research fundamentally reconceptualizes ethnic settlement patterns from isolated group-specific phenomena to interconnected systems governed by measurable inter-ethnic dynamics. The discovery of asymmetric interaction patterns through attention-based mechanisms (Chapter 4) reveals that settlement decisions of one ethnic group quantifiably influence others: China exhibiting strong negative self-competition (-19.8) driving spatial dispersal while showing positive facilitation with multiple other groups; Philippines and Pakistan forming mutually attractive dyads; United Kingdom and Italy exhibiting mutual repulsion. This theoretical advance moves beyond traditional assimilation models to establish "pluralistic segregation" as a distinct urban form where voluntary clustering maintains city-wide diversity while enabling neighborhood-level concentration.

The dissertation demonstrates that urban ethnic settlement patterns can be mathematically modeled using physics-inspired reaction-diffusion frameworks (Chapter 4), revealing statistical regularities with characteristic spatial scales (34.5-63.0 km) that vary systematically across ethnic groups. While these mathematical tools provide valuable insights into aggregate patterns, they represent statistical correlations rather than evidence that human settlement literally follows physical laws. The GraphPDE framework identifies pattern formation regimes (spots, stripes, labyrinthine) emerging from the interplay between self-interaction strength and between-group competition, with configuration landscape analysis revealing multiple stable states and transition barriers. The Multi-Ethnic Spatial Mixture of Experts framework (Chapter 5) introduces the theoretical concept that different demographic processes, colonization, jump

transitions, decline, and continuous diffusion, require distinct mathematical formulations, challenging monolithic approaches to urban modeling and establishing that cities simultaneously exhibit multiple dynamic regimes.

1.5.2 Methodological Contributions

The dissertation demonstrates a methodological evolution from exploratory deep learning (Appendices A and B) to interpretable frameworks (Chapters 4 and 5), establishing that the perceived trade-off between accuracy and interpretability is false. The progression introduces graph-based physics-informed neural networks (GraphPDE) combining reaction-diffusion equations with graph neural networks and attention-based competition mechanisms, and Multi-Ethnic Spatial Mixture of Experts achieving state-of-the-art performance ($R^2 = 0.80-0.83$) while maintaining full interpretability. These architectures prove that incorporating domain knowledge enhances rather than compromises predictive accuracy, with the physics-only variant of GraphPDE achieving $R^2 = 0.5588$ (95.5% of pure neural baseline performance) using exclusively interpretable parameters.

Responding to the "metrics trap" identified in Chapter 3's analysis of 157 urban AI deployments, the dissertation establishes critical computational geography as a necessary framework maintaining reflexivity about algorithmic power while harnessing computational capabilities for progressive urban research. All methods are implemented with open-source reproducibility, including GPU-accelerated CUDA kernels achieving $32.6\times$ speedup for city-scale simulations, comprehensive GitHub repositories with documented code, and publicly available datasets through Figshare. The modular architectures enable transferability across urban contexts, with physics-informed approaches reducing data requirements compared to black-box methods, establishing new standards for computational research in human geography.

1.5.3 Empirical Contributions

The dissertation provides the most comprehensive computational analysis of Canadian cities' ethnic geography to date, examining nine major ethnic groups across 30,091 dissemination areas in 52 cities over two decades (2001-2021). Key empirical discoveries include: ethnic-specific spatial scales ranging from 34.5 km (Philippines) to 63.0 km (United Kingdom); pattern formation regimes segregating groups into spots, stripes, and labyrinthine patterns; critical nucleation thresholds ranging from 658 (China) to 8,132 (India) individuals; temporal evolution revealing diverging demographic trajectories (Philippines maintaining positive growth while United Kingdom transitions to decline); and city-level heterogeneity with Toronto and Vancouver maintaining stable carrying capacities while Winnipeg exhibits 38% decline. The research provides mathematical parameters capturing previously qualitative observations: attention-based interaction weights quantifying inter-ethnic facilitation and competition; diffusion coefficients

measuring ethnic-specific spatial mobility (2.9-6.8 km²/year across cities); growth rates revealing demographic trajectories; carrying capacities showing urban constraints; and configuration landscape analysis identifying stable states and transition barriers. The observed variation in parameters within systematic patterns demonstrates that fundamental organizing principles govern settlement patterns while allowing for cultural and contextual variation, with the physics-only model achieving 73% of full GraphPDE's R^2 performance suggesting that reaction-diffusion theory captures the majority of predictable variance through interpretable mechanisms.

1.5.4 Practical Contributions

The dissertation provides concrete tools for urban planners and policymakers: predictive models forecasting neighborhood transitions with demonstrated accuracy ($R^2 = 0.7649$); identification of critical thresholds for viable enclave formation; learned parameters quantifying ethnic-specific mobility and interaction patterns; multi-stability analysis revealing configuration transition dynamics; and quantified uncertainty estimates (95.1% confidence interval coverage) enabling risk-aware policy design. The research reveals that effective integration requires understanding ethnic-specific dynamics rather than uniform approaches, with critical nucleation thresholds suggesting minimum viable populations for community formation, inter-ethnic interaction networks revealing facilitative and competitive relationships, and city-level parameter heterogeneity demonstrating that national policies must accommodate local variation.

All computational tools are freely available for adoption by other cities: GraphPDE framework adaptable to any urban context with spatial networks and census data; documented workflow from data preprocessing through prediction and visualization; benchmarked performance metrics enabling comparison across cities; and training materials for municipal data scientists and urban planners. This democratization of advanced computational methods enables evidence-based planning in resource-constrained contexts.

1.5.5 Interdisciplinary Bridge Building

The dissertation establishes interpretable GeoAI as a legitimate approach within human geography, demonstrating how computational methods can address fundamental geographic questions while maintaining theoretical transparency. The integration of machine learning with urban theory, physics with social dynamics, and computer science with geographic analysis creates new interdisciplinary spaces for collaboration essential as geography increasingly requires computational sophistication to address complex urban challenges. The progression from exploratory black-box methods to interpretable frameworks provides broader lessons for computational social science, demonstrating that domain knowledge integration enhances rather than constrains computational methods and offering a template for other fields grappling with the interpretability-accuracy trade-off.

By critically examining algorithmic failures while developing interpretable alternatives, the dissertation contributes to the broader movement for responsible AI in urban contexts, demonstrating that cities need not choose between computational capability and democratic accountability. These contributions collectively transform our understanding of urban ethnic dynamics from qualitative description to quantitative prediction, from isolated patterns to system dynamics, and from static observation to mechanistic explanation, establishing new foundations for using interpretable GeoAI to create more equitable, integrated, and vibrant multicultural cities.

Chapter 2 The Greater Toronto Area as a Multi-Ethnic Urban Laboratory

2.1 A Historical Timeline of Immigration Waves (1970s–2020s)

Toronto’s transformation into one of the world’s most diverse metropolitan areas represents a profound demographic shift that unfolded over five decades. This section traces the immigration patterns that reshaped the city, drawing on census data from 1971 to 2021 to document the successive waves of newcomers from the nine origin countries central to this study: China, Philippines, India, Pakistan, Iran, Sri Lanka, Portugal, Italy, and the United Kingdom. As (Hiebert, 2000) observed, the impact of immigration on Canadian cities since the Second World War has been profound, particularly following the removal of barriers to non-European immigrants in the 1960s and the significant increase in immigrant admissions since the mid-1980s.

2.1.1 The Demographic Transformation

The Toronto Census Metropolitan Area (CMA) experienced a fundamental change in its population composition between 1971 and 2021 (Statistics Canada, 2022c). As Table 2.1 demonstrates, the proportion of foreign-born residents nearly doubled over this period, transforming Toronto from a predominantly British-origin city to one where immigrants now constitute the majority. This trajectory aligns with broader patterns documented by Hou and Picot (2004), who noted the rapid expansion of visible minority populations in Canada’s three largest metropolitan areas.

Table 2-1. Toronto CMA Population and Immigration, 1971–2021

Census Year	Total Population	Immigrant Population	Immigrants (%)
1971	2,628,045	~760,000	28.9
1981	2,998,947	~840,000	28.0
1991	3,893,046	1,468,625	37.7
2001	4,682,897	2,032,960	43.4
2006	5,113,149	2,320,160	45.4
2011	5,583,064	2,537,410	45.4
2016	5,862,855	2,705,550	46.1
2021	6,202,225	2,891,990	46.6

Sources: Statistics Canada Census of Population, 1971–2021; 2011 National Household Survey

This transformation did not occur uniformly. Rather, it proceeded through distinct waves characterized by shifting source countries, changing immigration policies, and evolving global circumstances (Knowles, 2016).

2.1.2 Wave I: The Final Years of European Dominance (1970s)

The early 1970s marked the twilight of an immigration era dominated by Southern European arrivals. Italy and Portugal, which had supplied substantial numbers of immigrants throughout the 1950s and 1960s, continued to contribute to Toronto's growth, though at diminishing rates (Knowles, 2016). Historically, Canada had relied upon Western Europe, particularly Great Britain, as the major supplier of immigrants (P. S. Li, 1998).

The 1971 Census captured a Toronto still bearing the strong imprint of postwar European migration. Italian remained one of the most common non-official mother tongues, with concentrated communities in areas such as College Street, St. Clair West, and the emerging suburbs of Woodbridge and Vaughan (Hackworth & Rekers, 2005). Portuguese settlement clustered along Dundas Street West and in the Kensington Market area (R. Murdie & Teixeira, 2011).

Table 2-2. European-Origin Immigrant Population, Toronto CMA

Country	1991	2001	2011	2021
Italy	125,605	142,630	136,325	~110,000
Portugal	92,195	87,515	75,600	~65,000
United Kingdom	143,395	152,265	112,585	~95,000

Sources: Statistics Canada Census of Population

The 1967 introduction of the points system began redirecting immigration flows, though its full effects would not become apparent until the following decade (Triadafilopoulos, 2012). The points system established new standards for assessing potential immigrants, assigning points in categories relating to education, occupational skills, employment prospects, age, proficiency in English and French, and personal character (A. G. Green & Green, 1999). This reform marked the culmination of a normatively driven political process catalyzed by the discrediting of scientific racism and the emergence of human rights in the postwar period (Triadafilopoulos, 2012). British immigration, while still substantial, had already begun its long decline from the dominant position it had occupied since Confederation (Reitz, 2004).

2.1.3 Wave II: South Asian Emergence (1970s–1990s)

The points system's emphasis on education and skills created new pathways for migrants from South Asia (P. S. Li, 1998). Changes in immigration regulations in 1967 resulted in the adoption of a universal point system in assessing prospective immigrants, irrespective of country of origin or racial background (P. S. Li, 1998). India emerged as an increasingly significant source country through the 1970s, a trend that would accelerate dramatically in subsequent decades.

The 1971 Census grouped India and Pakistan together in birthplace tabulations, reflecting both the countries' recent partition and their still-modest numbers. By 1991, however, the census recorded distinct

and growing populations from both nations, along with an emerging Sri Lankan community driven largely by refugees fleeing civil conflict (Velamati, 2009).

Table 2-3. South Asian Immigrant Population, Toronto CMA

Country	1991	2001	2011	2021
India	~100,000	129,945	268,910	~320,000
Pakistan	~40,000	39,090	97,655	~125,000
Sri Lanka	~45,000	54,485	103,580	~115,000

Sources: Statistics Canada Census of Population

The Sri Lankan trajectory deserves particular attention. Arrivals accelerated sharply following the 1983 anti-Tamil riots and continued through the civil war period (1983–2009). The bitter ethnic fighting in Sri Lanka drove over 800,000 Sri Lankan Tamils from their homeland, forcing them to find refuge around the globe (Velamati, 2009). By 2011, Toronto housed the largest Sri Lankan Tamil diaspora outside South Asia, concentrated in Scarborough and northern suburbs (Statistics Canada, 2013). As the 2021 Canadian Census documented, Tamil Canadians numbered approximately 240,000, representing roughly 0.7% of Canada’s population (Statistics Canada, 2022c).

2.1.4 Wave III: East Asian Expansion (1980s–2000s)

Chinese immigration to Toronto has deep historical roots, but the contemporary community largely reflects arrivals since the 1980s. Three distinct streams contributed to this growth: skilled immigrants from the People’s Republic of China following economic liberalization; business immigrants and families from Hong Kong anticipating the 1997 handover; and students who transitioned to permanent residence (L. Lo & Wang, 1997). As L. Lo and Wang (1997) documented, Chinese settlement patterns exhibited both convergence and divergence: convergence in the sense that Chinese immigrants tended to concentrate in urbanized portions of the GTA, especially Toronto, Scarborough, North York, Markham, and Mississauga; divergence in that specific subgroups occupied distinct concentrations in the core and suburbs.

Table 2-4. East Asian Immigrant Population, Toronto CMA

Country	1991	2001	2011	2021
China (prc)	~120,000	178,675	224,915	~260,000
Philippines	~85,000	106,590	173,495	~200,000

Sources: Statistics Canada Census of Population

Filipino immigration followed a distinctive pattern shaped by labour market demands. The Live-in Caregiver Program (and its predecessors) created a highly gendered migration stream, with women predominating among arrivals (Brickner & Straehle, 2010; G. Pratt, 2012). By 2000, Filipinas made up 87 percent of all Live-in Caregiver Program workers (Tungohan et al., 2015). Family reunification

subsequently broadened the community's demographic profile. Concentrations emerged in North York, Scarborough, and Mississauga (Kelly, Park, Leon, & Priest, 2011).

2.1.5 Wave IV: Middle Eastern Growth (1980s–2020s)

Iranian immigration to Toronto accelerated following the 1979 Islamic Revolution, with subsequent waves responding to the Iran-Iraq War, economic pressures, and political conditions (Moghissi, Rahnema, & Goodman, 2009). The community grew from negligible numbers in 1971 to become one of the largest Iranian diaspora populations globally.

Table 2-5. Iranian Immigrant Population, Toronto CMA

Year	1991	2001	2011	2021
Iran	~25,000	30,170	59,755	~85,000

Sources: Statistics Canada Census of Population

Settlement patterns concentrated initially in North York, particularly along Yonge Street, before dispersing to Richmond Hill and Thornhill as the community established itself economically (Qadeer et al., 2010).

2.1.6 Period of Immigration Analysis

Census data on period of immigration reveals how arrival cohorts shifted across decades. Table 2.6 presents this temporal distribution for Toronto's immigrant population as recorded in 2011 (Statistics Canada, 2013).

Table 2-6. Toronto CMA Immigrants by Period of Arrival (2011 NHS)

Period of Immigration	Number	Percentage
Before 1971	401,080	15.8%
1971–1980	353,270	13.9%
1981–1990	457,185	18.0%
1991–2000	616,270	24.3%
2001–2011	709,605	28.0%

Source: Statistics Canada, 2011 National Household Survey

The 1990s represented the peak decade for immigrant arrivals in absolute terms, coinciding with elevated annual intake targets and the tail end of Hong Kong migration (D. Ley, 2010). The sustained high volumes into the 2000s reflected continued emphasis on economic immigration alongside family reunification and refugee flows (Reitz, 2012).

2.1.7 Visualizing Six Decades of Spatial Transformation

The spatial evolution of immigrant settlement patterns across Toronto from 1961 to 2021 can be comprehensively understood through two complementary analytical approaches: choropleth mapping of population distributions and mean center analysis with standard deviational ellipses. While these methods

Spatial Distribution of China-born Population in Toronto CMA (1961-2021)



Figure 2.1 Spatial Distribution of China-born Population in Toronto CMA (1961-2021).




Figure 2.2 Temporal Evolution of Population Mean Centers and One Standard Deviation Ellipses for Chinese and Iranian Immigrant Communities (1961-2021).

were applied to all nine study groups, we present Chinese and Iranian communities as illustrative cases representing distinct settlement trajectories: established suburban expansion versus corridor-oriented growth.

Figure 2.1 presents a systematic visual record of how the Chinese community's spatial footprint evolved across seven census periods. The pattern reveals a classic trajectory of emergence, intensification, and suburban consolidation. In 1961 and 1971, Chinese settlement remained confined to a small downtown core with maximum concentrations below 2% of local population. The 1981 census marks a visible inflection point, with concentrations reaching 4% and spatial extent expanding into Scarborough. By 1991, multiple nodes had emerged, with Scarborough and North York showing distinct clustering. The 2001-2021 period demonstrates continued intensification in northern suburbs, with Markham achieving concentrations exceeding 20% by 2021, representing one of the highest ethnic concentrations in any Canadian municipality (Hiebert, 2015).

Figure 2.2 employs mean center and standard deviational ellipse analysis to quantify the directional dynamics of settlement pattern evolution, revealing strikingly different spatial signatures between these two communities.

The Chinese community demonstrates dramatic spatial expansion over the 60-year study period. From a compact settlement of just 1,322 individuals occupying an ellipse area of 84.4 km² in 1961, the community grew to 296,200 individuals dispersed across an ellipse of 439.8 km² by 2021, representing a five-fold expansion in spatial extent (Qadeer et al., 2010 2010). The ellipse's major axis lengthened from 6.5 km to 16.9 km, while the minor axis doubled from 4.1 km to 8.3 km, indicating expansion in all directions but particularly along the north-south corridor toward Markham and Richmond Hill. The consistently high eccentricity values (0.78-0.89) reflect the persistent elongation of settlement along major transportation corridors, with the ellipse orientation shifting to capture the linear development patterns along Highway 7 and the emergence of multi-nodal settlement structure (L. Lo & Wang, 1997; S. Wang, 1999).

The Iranian community displays a contrasting but equally distinctive spatial signature. Emerging later than the Chinese community, Iranian settlement first appears substantively in the 1981 census following the Islamic Revolution, with 37,541 individuals already occupying a remarkably elongated ellipse of 286.1 km² (Hiebert, 2015). The Iranian ellipse exhibits significantly higher eccentricity than the Chinese community throughout the study period, reaching 0.95 in 1981, indicating an almost linear settlement pattern oriented along the Yonge Street corridor. This corridor orientation persists across subsequent decades, with the major axis consistently exceeding 14 km while the minor axis remains relatively constrained below 8 km. Unlike the continuous expansion observed in the Chinese community, Iranian spatial extent fluctuated, contracting from 421.4 km² in 2001 to 227.7 km² in 2011 before expanding again to 325.8 km² by 2021, suggesting periods of consolidation interspersed with dispersal (Qadeer et al., 2010 2010). The mean center

shifted progressively northward through York Region municipalities, tracking the community's expansion from initial North York settlement into Richmond Hill and Thornhill while maintaining strong anchoring to commercial and institutional nodes along the Yonge corridor.

These two cases illustrate broader patterns observed across the nine study groups. Asian-origin communities generally demonstrate progressive spatial expansion and intensification over the study period, though with varying directional orientations and concentration intensities. European-origin communities (Portugal, Italy, United Kingdom) show contrasting trajectories of ellipse expansion despite declining populations, reflecting dispersal of remaining community members as traditional ethnic neighbourhoods transformed (R. Murdie & Teixeira, 2011). Sri Lankan immigrants exhibit perhaps the most distinctive pattern, with remarkably stable mean center location and compact ellipse dimensions anchored in eastern Scarborough across all census periods, reflecting strong ethnic institutional anchoring and continued refugee-driven migration into established community areas (Velamati, 2009).

2.1.8 Implications for Settlement Patterns

These successive waves did not merely add population to Toronto; they fundamentally restructured its human geography (Hiebert, 2000). The nine origin countries examined in this study illustrate contrasting settlement trajectories: established European communities (Italy, Portugal, United Kingdom) followed a classic pattern of initial inner-city concentration followed by suburban dispersal as socioeconomic mobility increased, showing the most dispersed residential patterns by 2021 (Qadeer et al., 2010); South Asian communities demonstrated varied patterns, with Indian settlement spreading across multiple suburban nodes, Sri Lankan Tamils concentrating heavily in Scarborough, and Pakistani settlement showing intermediate dispersion (Hou & Picot, 2004); East Asian communities established multiple suburban concentrations, with Chinese clusters in Markham, Richmond Hill, and Scarborough while Filipino settlement remained more dispersed (Kelly et al., 2011; L. Lo & Wang, 1997); and Iranian settlement concentrated along specific North York and York Region corridors (Qadeer et al., 2010).

The contemporary multi-ethnic landscape reflects five decades of continuous immigration and policy reforms. The 1967 points-based immigration system constituted a structural break, initiating diversification from Europe-dominant migration to significant arrivals from South Asia, East Asia, the Caribbean, Africa, and the Middle East. Between 1981 and 2021, the foreign-born share of Toronto's population rose consistently above 50 percent, with racialized minorities increasing dramatically in Scarborough, Markham, Brampton, Richmond Hill, and Mississauga. This temporal layering reveals that settlement patterns are not fixed mosaics but reflect sequential waves of diversification, relocation, and neighbourhood transformation, where first-generation concentration, second-generation suburban dispersion, and multi-ethnic coexistence occur simultaneously, an empirical reality requiring a dynamic rather than static modelling framework.

2.1.9 Data Considerations and Limitations

Several methodological considerations affect the interpretation of these census trends:

1. **Changing geographic boundaries:** The Toronto CMA expanded over successive censuses, incorporating additional municipalities. Comparisons across time require attention to boundary definitions (Statistics Canada, 2022c).
2. **Evolving census categories:** Country-of-birth classifications changed across censuses. The 1971 census grouped India and Pakistan; earlier censuses combined Portugal with Spain. China classifications variably included or distinguished Hong Kong and Taiwan (Hou, 2004).
3. **Census methodology changes:** The replacement of the mandatory long-form census with the voluntary National Household Survey in 2011 introduced potential non-response bias, particularly in immigrant and visible minority communities (D. A. Green & Milligan, 2010). The mandatory long-form was reinstated for 2016.
4. **Undercounting of recent arrivals:** Census enumeration may undercount recent immigrants, particularly those in informal housing arrangements or with precarious status. This limitation particularly affects estimates of refugee populations during periods of active conflict migration.

Despite these limitations, the census series provides the most comprehensive longitudinal record of Toronto's demographic transformation, enabling the spatial analyses that follow in subsequent chapters.

2.2 Selection and Justification of Study Groups

The selection of immigrant communities for spatial analysis requires balancing multiple considerations: statistical significance, demographic diversity, temporal variation in settlement patterns, and representativeness across global regions (Fong & Wilkes, 2003). This section presents the empirical rationale for focusing on nine origin countries: China, Philippines, India, Pakistan, Iran, Sri Lanka, Portugal, Italy, and the United Kingdom.

2.2.1 Population Significance and Representativeness

The nine selected immigrant groups collectively constitute a substantial proportion of Toronto's foreign-born population, ensuring that findings from this study capture meaningful demographic patterns rather than statistical noise from small populations (Balakrishnan & Hou, 1999). Analysis of 2021 Census data at the Dissemination Area (DA) level reveals that these groups account for 42.0% of all immigrants residing in the City of Toronto proper, representing 530,360 individuals out of a total immigrant population of 1,261,505 (Statistics Canada, 2022c).




Figure 2.3 Temporal trends in immigrant population for the nine selected origin countries across five census periods.

Table 2.7 presents the population statistics and growth metrics for each selected group across the 2001-2021 study period. The data reveal substantial variation in both absolute population size and growth trajectories, providing analytical leverage for examining how different community characteristics influence residential settlement patterns (Walks & Bourne, 2006).

Table 2-7. Summary Statistics of Selected Immigrant Groups in Toronto (2001-2021)

Country Of Origin	Population 2001	Population 2021	Absolute Change	Growth Rate	Cagr	Rank (2001→2021)
Philippines	39,415	123,460	+84,045	+213.2%	5.87%	4 → 1
China	64,995	118,955	+53,960	+83.0%	3.07%	1 → 2
India	37,370	94,380	+57,010	+152.6%	4.74%	5 → 3
Sri Lanka	24,845	41,130	+16,285	+65.5%	2.55%	7 → 4
Italy	57,645	32,320	-25,325	-43.9%	-2.85%	2 → 5
United Kingdom	44,530	31,175	-13,355	-30.0%	-1.77%	3 → 6
Iran	11,560	30,980	+19,420	+168.0%	5.04%	9 → 7
Portugal	37,920	30,270	-7,650	-20.2%	-1.12%	6 → 8
Pakistan	13,305	27,690	+14,385	+108.1%	3.72%	8 → 9

Note: CAGR = Compound Annual Growth Rate over the 20-year period. Rank based on population within the City of Toronto.




Figure 2.4 Comparison of immigrant populations between 2001 and 2021 census years, illustrating the magnitude of demographic shifts.

As illustrated in Figure 2.4, all nine groups rank among the top eleven source countries for immigrants in Toronto as of 2021, ensuring sufficient population density at the DA level to support robust spatial statistical analysis. The minimum population threshold of approximately 27,000 individuals (Pakistan) provides adequate counts across most Dissemination Areas to avoid suppression due to Statistics Canada’s random rounding and data quality protocols (Statistics Canada, 2022c).

2.2.2 Diversity in Growth Trajectories

A critical consideration in selecting study groups was ensuring representation of diverse demographic trajectories (Hiebert, 2015). Immigration communities at different stages of the settlement lifecycle exhibit distinct spatial patterns, and capturing this variation enables more comprehensive understanding of the processes shaping ethnic residential geography (Massey, 1985).

The selected groups demonstrate three analytically distinct growth patterns:

1. **Rapidly Growing Communities (>50% growth, 2001-2021):** Philippines (+213.2%), Iran (+168.0%), India (+152.6%), Pakistan (+108.1%), China (+83.0%), and Sri Lanka (+65.5%) represent communities experiencing sustained immigration inflows. These groups are characterized by younger age structures, active chain migration networks, and ongoing spatial expansion into new residential areas (Alba & Logan, 1991). Their settlement patterns reflect contemporary housing market conditions, transportation accessibility, and the locations of recently established ethnic institutions (Breton, 1964).
2. **Declining Communities (<-20% change, 2001-2021):** Italy (-43.9%), United Kingdom (-30.0%), and Portugal (-20.2%) represent established immigrant communities experiencing population decline through mortality and out-migration to suburban municipalities (Qadeer et al., 2010). These

groups arrived primarily during earlier immigration waves (pre-1980s) and exhibit aging demographic profiles. Their residential patterns reflect historical housing market conditions and the persistence of ethnic enclaves established decades ago (Hackworth & Rekers, 2005).




Figure 2.5 Proportional representation of the nine selected immigrant groups in 2021.

This diversity in growth trajectories provides analytical leverage for distinguishing between processes associated with initial settlement (characteristic of growing communities) versus residential succession and community maturation (characteristic of declining communities). The juxtaposition of these contrasting patterns within a single metropolitan context controls for housing market conditions, municipal policies, and other metropolitan-level factors that would confound cross-city comparisons (Fong & Wilkes, 2003).

2.2.3 Regional and Cultural Representation

The selection encompasses immigrant communities from four distinct global regions, ensuring that findings are not artifacts of region-specific migration patterns or cultural settlement preferences (Balakrishnan & Hou, 1999).

South Asia (3 groups): India, Pakistan, and Sri Lanka collectively represent the largest regional cluster in the selection, reflecting South Asia's emergence as the dominant source region for Canadian immigration since the 1990s (Hou & Picot, 2004). Despite shared regional origins, these communities exhibit distinct settlement patterns reflecting different arrival periods, socioeconomic compositions, and in the case of Sri Lanka, significant refugee-stream migration (Velamati, 2009).

East Asia (2 groups): China and the Philippines represent the two largest East Asian immigrant populations in Toronto (Statistics Canada, 2022c). Chinese immigration encompasses diverse streams including skilled workers from the People's Republic, business immigrants from Hong Kong, and international students transitioning to permanent residence (L. Lo & Wang, 1997). Filipino immigration reflects the distinctive

pathway created by caregiver programs, producing unique gendered migration patterns and associated settlement characteristics (G. Pratt, 2012).

Middle East (1 group): Iran represents Middle Eastern migration to Toronto, a community that grew substantially following the 1979 Revolution and subsequent political and economic disruptions (Moghissi et al., 2009). Iranian immigrants exhibit high educational attainment levels and distinctive commercial entrepreneurship patterns that have shaped their residential concentrations (Qadeer et al., 2010).

Western Europe (3 groups): Portugal, Italy, and the United Kingdom represent the historical foundations of Toronto's immigrant population (Knowles, 2016). These communities established the ethnic neighbourhood template that subsequent immigrant groups would follow and transform. Their contemporary residential patterns reflect six decades of socioeconomic mobility, suburbanization, and generational succession (Hackworth & Rekers, 2005; R. Murdie & Teixeira, 2011).




Figure 2.6 Population ranking of the top 20 immigrant source countries in Toronto (2021).

2.2.4 Spatial Distribution Characteristics

Beyond demographic significance, the selected groups exhibit analytically valuable spatial distribution characteristics that enable meaningful examination of residential segregation, ethnic clustering, and neighbourhood-level settlement patterns.

Spatial Distribution of Immigrant Populations in Toronto (2021 Census)



Figure 2.7 Panel maps displaying the spatial distribution of each selected immigrant group across Toronto's Dissemination Areas (2021 Census).

Figure 2.7 presents choropleth maps of immigrant population counts at the DA level for each selected group. Visual inspection reveals substantial variation in spatial concentration patterns:

1. **Highly concentrated:** Sri Lankan and Portuguese communities exhibit the most spatially concentrated patterns, with large proportions of their populations residing in contiguous DA clusters.
2. **Multi-nodal:** Chinese, Indian, and Filipino communities display multi-nodal distributions with several distinct concentration areas across the metropolitan area.
3. **Dispersed:** United Kingdom-origin immigrants show the most dispersed pattern, consistent with their longer residence duration and higher rates of residential integration.

To quantify these concentration patterns, Figure 2.8 presents Location Quotient (LQ) maps for each group. The Location Quotient measures the ratio of a group's share of a DA's immigrant population relative to the group's share of the city-wide immigrant population. Values exceeding 1.0 indicate over-representation; values below 1.0 indicate under-representation.




Figure 2.8 Location Quotient (LQ) maps illustrating areas of ethnic over-representation (LQ > 1, shown in warmer colours) and under-representation (LQ < 1, shown in cooler colours).

The LQ analysis confirms that selected groups exhibit sufficient spatial clustering to support neighbourhood-level analysis while also displaying meaningful variation in concentration intensity. This variation is essential for examining how community-level factors, including population size, arrival recency, and socioeconomic composition, influence residential segregation patterns.

2.2.5 Temporal Dynamics of Settlement

The longitudinal dimension of this study requires that selected groups demonstrate measurable change in settlement patterns over the 2001-2021 period. Figure 2.9 presents maps of population change at the DA level, calculated as the difference between 2021 and 2001 immigrant counts.




Figure 2.9 Population change maps (2001-2021) for each selected immigrant group.

The change maps reveal distinct temporal signatures:

Expansionary patterns: Growing communities exhibit spatial expansion, with population gains concentrated in peripheral DAs adjacent to existing clusters. This pattern is particularly pronounced for Filipino, Indian, and Iranian immigrants, who have established substantial presence in areas that had minimal immigrant populations in 2001.

Contractionary patterns: Declining communities show population losses across their traditional settlement areas, with partial retention in core ethnic neighbourhood zones. Italian and Portuguese

communities demonstrate this pattern most clearly, with losses throughout former concentration areas but continued presence in historical neighbourhood centres.

Redistribution patterns: Some communities, notably Chinese immigrants, exhibit simultaneous growth in certain areas and decline in others, suggesting internal redistribution rather than simple expansion or contraction. This pattern likely reflects suburbanization within the community as established residents relocate while new arrivals settle in different locations.

2.2.6 Summary of Selection Criteria

The nine immigrant groups selected for this study satisfy multiple analytical criteria essential for rigorous examination of residential settlement patterns:

1. **Statistical adequacy:** All groups maintain sufficient population to support DA-level analysis across all five census years, with minimum populations exceeding 11,000 in 2001 and 27,000 in 2021.
2. **Demographic diversity:** The selection encompasses communities experiencing rapid growth, stability, and decline, enabling analysis of settlement dynamics across the community lifecycle.
3. **Regional balance:** Four global regions are represented, controlling for region-specific migration patterns and cultural factors.
4. **Spatial variation:** Selected groups exhibit diverse concentration patterns, from highly clustered ethnic enclaves to dispersed residential distributions.
5. **Temporal dynamics:** All groups demonstrate measurable spatial change over the study period, providing the longitudinal variation necessary for examining settlement pattern evolution.
6. **Collective significance:** Together, the nine groups account for 42% of Toronto's immigrant population, ensuring that findings reflect major demographic trends rather than marginal populations.

These selection criteria ensure that subsequent spatial analyses will yield findings with both statistical robustness and substantive significance for understanding immigrant settlement in contemporary Toronto.

2.3 Spatial Patterns of Multi-Ethnic Settlement

The spatial organization of immigrant communities across the Greater Toronto Area reflects a complex interplay of historical processes, housing market dynamics, and evolving settlement strategies (Walks & Bourne, 2006). While traditional scholarship focused on inner-city enclaves, the most significant transformation in Toronto's ethnic geography has occurred in the suburbs. The 1990s witnessed a marked shift toward suburban municipalities offering more affordable housing, larger dwelling sizes, and growing immigrant-serving institutions. Brampton evolved into a South Asian-majority municipality, Markham

became one of Canada's most concentrated Chinese-majority cities, and Scarborough and North York exhibit high degrees of ethnic heterogeneity producing what sociologists describe as "vertical mosaics" in high-rise corridors (Qadeer & Kumar, 2006). These spatial trends demonstrate that suburbanization constitutes a complex redistribution process shaped by affordability, transportation networks, cultural institutions, and social capital, with suburban nodes operating as attractors in the spatial diffusion of immigrant settlement.

2.3.1 Inner-City Neighbourhoods: Historical Enclaves and Succession

Toronto's inner-city neighbourhoods preserve the archaeological record of successive immigration waves, with ethnic enclaves forming, transforming, and occasionally dissolving over decades (Hiebert, 2000; R. A. Murdie, 2008). Toronto's Chinese settlement illustrates this pattern: the original Chinatown near York and Queen Streets was destroyed through urban renewal in the late 1950s for City Hall, pushing residents westward to establish the Spadina-Dundas Chinatown in the early 1960s, which became one of North America's largest by the 1980s (L. Lo & Wang, 1997; Thompson, 1989; S. Wang, 1999). A third concentration, East Chinatown at Broadview and Gerrard, emerged in the 1970s (Hackworth & Rekers, 2005). Since the 1990s, downtown Chinatown has faced pressure from gentrification and suburban competition, increasingly serving elderly residents dependent on accessible services while younger immigrants settle directly in suburbs (L. Lo & Wang, 1997; S. Wang, 1999).

Little Italy along College Street West exemplifies the commodification of ethnic identity accompanying gentrification (Hackworth & Rekers, 2005). Italian residential concentration declined substantially since the 1970s as families relocated to suburban Woodbridge and Vaughan (Buzzelli, 2001; Zucchi, 1988). By 2005, only 10 percent of businesses were identifiably Italian while businesses with no ethnic identification comprised nearly 65 percent (Hackworth & Rekers, 2005). What remains is "ethnic packaging," the deliberate maintenance of ethnic commercial identity to attract gentrifiers seeking authentic urban experiences, challenging traditional understandings of enclaves as organic expressions of immigrant settlement.

Little Portugal, centred along Dundas Street West, represents a case where suburbanization and gentrification proceed simultaneously (R. Murdie & Teixeira, 2011). Portuguese families maintained stronger attachment to the downtown enclave through the 1990s, but rising property values have increasingly displaced working-class residents (C. Teixeira, 2007). Without continued Portuguese immigration, limited by Portugal's EU membership and improved economic conditions, the institutionally complete enclave is transitioning toward a commodified ethnic landscape similar to Little Italy (R. Murdie & Teixeira, 2011).

2.3.2 Suburban Settlements: New Ethnic Geographies

The emergence of substantial ethnic concentrations in suburban municipalities represents a departure from classical models assuming initial inner-city concentration followed by gradual suburban dispersal (Hiebert, 2000; Massey & Denton, 1985; Qadeer & Kumar, 2006).

Eastern Scarborough has become home to the largest Sri Lankan Tamil diaspora outside South Asia, sometimes referred to as "Little Jaffna" (Velamati, 2009). Unlike earlier groups who established inner-city footholds first, Tamil refugees fleeing the 1983 anti-Tamil riots settled directly in Scarborough's apartment clusters (Hyndman, 2003; R. A. Murdie, 2008). The Tamil enclave has developed substantial institutional infrastructure including Hindu temples, community centres, and Tamil-language media (Velamati, 2009). The mean center analysis presented in Figure 2.2 demonstrates remarkable stability in Sri Lankan settlement patterns, with the population centroid remaining anchored in eastern Scarborough across multiple census periods, reflecting strong ethnic institutional anchoring (Qadeer et al., 2010 2010).

Markham and Richmond Hill have emerged as primary Chinese destinations since the 1990s, forming "ethnoburbs" that differ fundamentally from traditional Chinatowns (Kymlicka, 1998; L. Lo & Wang, 1997). Rather than emerging from marginalization, these communities developed through the economic strength of Hong Kong immigrants anticipating the 1997 handover and skilled mainland China immigrants (P. S. Li, 1998). The commercial infrastructure centres on indoor shopping malls rather than street-oriented strips, with Pacific Mall forming the largest Chinese shopping complex in North America with over 700 stores (Lucia Lo et al., 2000 2003; S. Wang, 1999). Chinese settlement generated political controversy in 1995 when Markham's Deputy Mayor argued that ethnic concentration was causing social conflict, with the ensuing backlash demonstrating both demographic weight and political mobilization capacity of suburban immigrant communities (Good, 2009).

Brampton and Mississauga have become primary South Asian destinations, particularly for immigrants from Punjab, Gujarat, and Pakistan (Qadeer, 2016). By 2016, South Asians comprised over 40 percent of Brampton's population (Statistics Canada, 2017). Unlike concentrated Tamil settlement, South Asian presence exhibits a multi-nuclear pattern with multiple commercial nodes and religious institutions distributed across the municipality (Qadeer & Kumar, 2006). The concentration reflects availability of larger houses for multigenerational families, proximity to Pearson International Airport facilitating transnational connections, and established chain migration networks (Siemiatycki et al., 2001 Rahi, 2003).

2.3.3 Mixed Neighbourhoods: Sites of Interethnic Contact

Beyond the binary of ethnic enclave versus mainstream dispersal, Toronto contains numerous neighbourhoods characterized by substantial diversity without single-group dominance (Vertovec, 2007; Walks & Bourne, 2006). Thorncliffe Park in East York exemplifies extreme multi-ethnic mixing, with

residents from over 50 countries and 27 languages spoken in local schools (Ahmadi & Tasan-Kok, 2014). This "super-diversity" creates environments where diversity itself becomes the defining characteristic, rendering traditional multicultural frameworks inadequate (Vertovec, 2007). Super-diverse neighbourhoods provide contexts for interethnic contact potentially fostering bridging social capital, though the absence of a dominant group may impede coherent community institutions (Ahmadi & Tasan-Kok, 2014).

Analysis of ethnic diversity indices reveals that maximum diversity occurs not in established enclaves or mainstream suburbs but in transitional zones between them, where expanding enclave populations encounter established communities with Shannon entropy values exceeding 1.75 (Walks & Bourne, 2006). Such zones may represent either temporary states in ethnic succession or stable multi-ethnic configurations (Qadeer et al., 2010).

2.4 Socioeconomic and Housing Dimensions

The spatial distribution of immigrant communities cannot be understood apart from the housing market conditions and socioeconomic stratification that constrain residential choices (Walks & Bourne, 2006).

2.4.1 Housing Affordability as Settlement Driver

Toronto's emergence as one of North America's most expensive housing markets has fundamentally altered immigrant settlement patterns (J. D. Hulchanski, 2010). With average home prices exceeding \$1.2 million and rental vacancy rates below 0.5 percent by 2024, newcomers face unprecedented barriers, pushing new arrivals toward suburban municipalities and increasingly precarious housing arrangements (Valerie Preston & Brian Ray, 2020). Research documented that many immigrant households in York Region were at risk of "hidden homelessness," spending over 50 percent of income on shelter, reflecting high costs, low vacancy rates, and newcomers' difficulty accessing rentals without Canadian credit histories (Wayland, 2007). The 2001 Census revealed that 23 percent of homeowner immigrants and 40 percent of renter immigrants in major metropolitan areas spent more than 30 percent of income on housing (Wayland, 2007). Communities with established co-ethnic networks and pooled household resources navigate housing constraints better than groups with fewer resources or weaker social connections, with Bangladeshis who lived initially with co-ethnics achieving housing stability more quickly than Bengalis relying on formal settlement services (S. Ghosh, 2007).

2.4.2 Stratification Within Communities

The diversity of immigration pathways produces substantial socioeconomic stratification within ethnic communities, complicating any simple equation of ethnic concentration with either marginalization or success (L. Lo & Wang, 1997). Chinese immigrants arrive through markedly different channels: the

Business Immigration Program attracted wealthy Hong Kong investor immigrants in the 1990s who purchased substantial properties in Richmond Hill and Markham; skilled workers possess professional credentials; while refugees and family-class immigrants face greater settlement challenges (Lucia Lo et al., 2000 2003). These class distinctions map onto spatial patterns, with wealthier Hong Kong immigrants in new suburban developments while lower-income arrivals settled in older Scarborough and North York apartments (L. Lo & Wang, 1997).

Similar heterogeneity characterizes South Asian settlement, where Punjabi Sikhs who established substantial business ownership exhibit different patterns than recent arrivals from Gujarat or Tamil-speaking regions (Siemiatycki et al., 2001 Rahi, 2001). The geographical spread across Brampton, Mississauga, Scarborough, and York Region reflects this internal diversity (Qadeer & Kumar, 2006). Among European-origin groups, Italian immigrants achieved middle-class status through small business ownership and real estate investment, Portuguese immigrants faced longer trajectories to economic stability, while United Kingdom immigrants benefiting from linguistic advantages experienced less pronounced residential concentration (C. Teixeira, 2007).

2.4.3 Rental Versus Ownership Markets

The tenure dimension significantly shapes ethnic clustering patterns (Walks & Bourne, 2006). Visible minority groups increasingly moved into areas with high proportions of other visible minorities during 1991-2001, linked to low-rent apartment availability and increasing affordability problems (Walks & Bourne, 2006). Rental buildings in Scarborough, North York, and inner suburbs have become significant sites of immigrant concentration (R. A. Murdie, 2008). Research documents widespread discrimination by landlords, with darker-skinned groups facing greater obstacles, channeling certain communities toward lower-quality housing (R. A. Murdie, 2002, 2003).

Traditional models equate immigrant homeownership with assimilation into the residential mainstream (Massey & Denton, 1985). However, analysis of twelve Toronto immigrant groups revealed that some groups, notably Chinese immigrants, achieved high homeownership rates while maintaining ethnic concentration, suggesting "buying in" to ethnic communities rather than "buying out" characterizes contemporary settlement (Haan, 2012). The suburban Chinese enclaves in Markham and Richmond Hill comprise overwhelmingly owner-occupied single-family homes, contradicting assumptions that homeownership accompanies ethnic dispersal (L. Lo & Wang, 1997). Statistics Canada data indicate housing assets comprise a larger share of average wealth among immigrant families, reflecting the centrality of homeownership to immigrant economic strategies and pathways to intergenerational wealth transfer (Statistics Canada, 2019).

2.4.4 Intersection with Urban Inequalities

Immigrant settlement patterns intersect with broader patterns of urban inequality in Toronto, including the spatial polarization documented by J. D. Hulchanski (2010) as the “three cities” phenomenon, areas of concentrated advantage, disadvantage, and transition that increasingly align with racial and ethnic composition.

Income Polarization: The period from 1970 to 2005 saw Toronto’s neighbourhoods sort increasingly by income, with middle-income areas shrinking as neighbourhoods shifted toward either high-income or low-income profiles (J. D. Hulchanski, 2010). This polarization has spatial dimensions: the central city and certain inner suburbs have gentrified, while portions of the inner suburbs, particularly in northern Etobicoke, central North York, and Scarborough, have experienced relative decline. Immigrant communities are disproportionately affected by both processes: gentrification displaces lower-income immigrant households from inner-city enclaves, while declining inner suburban neighbourhoods receive newcomers without adequate services or employment access (J. D. Hulchanski, 2010).

Service Accessibility: The suburbanization of immigrant settlement has outpaced the development of suburban service infrastructure (Cowen & Parlette, 2011). Inner suburban neighbourhoods with large immigrant populations often lack the settlement services, language training programs, and employment supports available in downtown areas. This “service mismatch” means that immigrants who settle in suburban locations due to housing affordability may face greater difficulties accessing the supports that facilitate integration (Cowen & Parlette, 2011).

2.5 Governance and Policy Context

The spatial patterns of multi-ethnic settlement emerge within a governance framework spanning federal immigration policy, provincial planning and housing regulations, and municipal service delivery (Good, 2009). Three governance factors are especially influential in shaping Toronto's settlement dynamics: federal immigration selection policies prioritizing skilled labour and family sponsorship pathways, increasing socio-economic diversity among newcomers; Ontario's planning and housing frameworks creating exurban growth zones through policies like the GTA Greenbelt and Places to Grow; and municipal multicultural and settlement services concentrating resources in specific nodes (libraries, cultural centres, newcomer hubs) that reinforce spatial settlement choices. Combined, these policy forces structure both constraints and opportunities, explaining why enclaves can appear stable yet remain metastable: affordability shocks, redevelopment, or service reallocation can shift equilibria and induce diffusion or displacement.

2.5.1 Federal Policy Framework

Canada's commitment to official multiculturalism, entrenched in the Constitution and implemented through the Canadian Multiculturalism Act of 1988, establishes a normative framework envisioning integration that preserves cultural heritage while enabling full participation in Canadian society, differing from both assimilationist and segregationist approaches (Good, 2009). Immigration selection policy fundamentally shapes arriving populations and settlement patterns (P. S. Li, 1998). The points system introduced in 1967 emphasizes education and skills, while the Express Entry system implemented in 2015 further prioritizes human capital factors, favouring immigrants with credentials facilitating rapid labour market integration in metropolitan areas (Reitz, 2004).

Family reunification and refugee streams operate under different criteria, producing arrivals with different resource profiles (Hiebert, 2000). Family-class immigrants typically settle near sponsors, reinforcing existing ethnic concentrations, while refugees may be directed to specific communities based on settlement agreements, though secondary migration reshapes initial distributions (Velamati, 2009). Federal funding for settlement services has evolved substantially since the 1990s, with IRCC contracting service providers to deliver language training, employment assistance, and orientation programming (Shields, Drolet, & Valenzuela, 2018 2018). However, these services have historically concentrated downtown, creating geographic mismatches as settlement suburbanized, with significant gaps remaining despite recent initiatives (Cowen & Parlette, 2011).

2.5.2 Provincial Policy Framework

Ontario has developed immigration policies within federal-provincial agreements, with the Ontario Immigrant Nominee Program (OINP) allowing provincial selection based on regional labour market needs, though the vast majority of Ontario-bound immigrants continue settling in the Greater Toronto Area (Statistics Canada, 2022c). Provincial planning frameworks significantly shape housing options (White, 2007). The Growth Plan for the Greater Golden Horseshoe establishes intensification targets encouraging higher-density development in urban growth centres and along transit corridors, while the Greenbelt restricts development in designated lands, credited with limiting sprawl but criticized for constraining housing supply and elevating prices (White, 2007). Provincial legislation defines municipal powers and responsibilities, with funding formulas for social housing, transit, and social services shaping municipal capacity to serve immigrant populations (Good, 2009).

2.5.3 Municipal Policy and Planning

The City of Toronto has been characterized as Canada's most proactive municipal government in multiculturalism policy, with Access, Equity and Human Rights policies establishing commitments to equitable service delivery (Good, 2009; Siemiatycki & Saloojee, 2012). Toronto provides vital services

through agencies including the Toronto District School Board, Toronto Public Health, and Toronto Public Libraries, while the Toronto Region Immigrant Employment Council (TRIEC) reflects public-private collaboration for labour market integration (Siemiatycki & Saloojee, 2012). However, municipal capacity is constrained by jurisdictional and fiscal limitations, as municipalities lack independent authority to develop immigration policies or access immigration-specific revenue streams (Good, 2009).

Suburban municipalities exhibit substantial variation, ranging from "proactive" (Toronto) through "reactive" (Mississauga, Markham) to "inactive" (Brampton in the period studied), reflecting differences in political leadership, bureaucratic capacity, and ethnic political mobilization (Good, 2009). Markham's experience illustrates political dynamics of rapid demographic transformation, with the 1995 controversy surrounding Deputy Mayor Carole Bell's comments about ethnic concentration contributing to political realignment and enhanced municipal attention to diversity (Good, 2009). Zoning bylaws affecting housing form, particularly restrictions on secondary suites and multigenerational arrangements, may conflict with immigrant residential preferences, while permitting of religious facilities has generated controversy in several municipalities (Siemiatycki et al., 2001 Rahi, 2001). Business Improvement Areas in ethnic commercial districts such as Little Italy and Greektown have promoted ethnic commercial identity, but increasingly function to attract tourists and gentrifiers rather than serve residential ethnic populations (Hackworth & Rekers, 2005).

2.5.4 Governance Gaps and Challenges

The multi-level governance framework creates gaps and coordination challenges (Burr, 2011). The federal government determines who can immigrate but has limited influence over settlement location; provincial policies shape metropolitan development without accounting for immigrant-specific needs; municipalities deliver services but lack authority to address immigration policy concerns. Immigrant-serving organizations bridge these gaps, with Local Immigration Partnerships (LIPs) established since 2008 coordinating settlement service delivery by bringing together governments, service providers, and community organizations (Shields et al., 2018 2018). However, the voluntary nature of these partnerships limits capacity to address systemic issues such as housing affordability or employment discrimination (Burr, 2011).

2.6 Dynamics of Multi-Ethnic Settlement

This section synthesizes spatial patterns, socioeconomic dimensions, and governance context to characterize the dynamic processes shaping Toronto's evolving ethnic geography, emphasizing settlement as process rather than static pattern.

2.6.1 Ethnic Enclaves: Formation, Functions, and Futures

Ethnic enclaves serve multiple functions: providing familiar cultural environments easing transition; enabling commercial activity in native languages; creating markets for ethnic goods; facilitating information sharing about employment and housing; and establishing visible community presence attracting subsequent migrants (Breton, 1964; Hiebert, 2000). Enclaves offer protection from discrimination, with Toronto's Chinatowns emerging amid exclusionary legislation, and while formal discrimination has diminished, informal barriers including workplace and housing discrimination continue motivating ethnic clustering (Kymlicka, 1998; Qadeer & Kumar, 2006; Carlos Teixeira, 2008). The ethnic enclave economy facilitates immigrant integration through co-ethnic employment and commercial activity, though enclave employment may limit broader economic exposure and constrain wage growth (Hou & Picot, 2016; Portes & Jensen, 1987; S. Wang, 1999). Beyond economics, enclaves support cultural maintenance through religious institutions and community organizations, providing contexts where ethnic identity carries positive valence (Breton, 1964; Qadeer & Kumar, 2006). Despite these functions, ethnic concentration generates concerns about segregation and failed integration, with critics arguing enclaves may impede language acquisition and reproduce inequalities, leaving the policy question of dispersal versus enclave support contested (Hiebert, 2015).

2.6.2 Suburbanization: Mobility and New Concentrations

The suburbanization of immigrant settlement represents one of Toronto's most significant geographic transformations, encompassing both classical dispersal accompanying socioeconomic mobility and direct suburban settlement by arriving immigrants (Hiebert, 2000; Qadeer & Kumar, 2006). Italian and Portuguese communities exemplify the classical pattern: mid-twentieth century arrivals initially settled in inner-city neighbourhoods, then moved to suburban locations as they achieved economic stability through construction, small business, and real estate investment, as illustrated by Italian movement from College Street to Woodbridge and Vaughan (Buzzelli, 2001; Hackworth & Rekers, 2005; R. Murdie & Teixeira, 2011).

Beginning in the 1980s, substantial numbers began settling directly in suburbs, bypassing traditional inner-city reception (Hiebert, 2000). Hong Kong immigrants purchased homes in Markham and Richmond Hill upon arrival, drawn by new housing, good schools, and Chinese commercial infrastructure, while skilled immigrants from India and Pakistan settled directly in Brampton and Mississauga (L. Lo & Wang, 1997; Qadeer & Kumar, 2006). These suburban ethnoburbs differ from traditional enclaves, characterized by homeownership rather than rental, automobile-oriented commercial development, and residential landscapes appearing similar to mainstream suburbs despite substantial ethnic concentration (P. S. Li, 1998). The suburbanization has significant service delivery implications, with settlement agencies and

language programs facing challenges reaching suburban populations, and public transit poorly serving dispersed employment sites (Cowen & Parlette, 2011).

2.6.3 Displacement and Gentrification

Gentrification affects immigrant communities directly by displacing lower-income households, indirectly by eliminating affordable housing for newcomers, and symbolically by erasing visible ethnic presence (R. Murdie & Teixeira, 2011). As property values rise, incumbent residents face pressure from rising rents and property taxes, with research documenting considerable ambivalence among affected residents who appreciate improved amenities while recognizing undermined community cohesion (R. Murdie & Teixeira, 2011). Ethnic commercial districts transform as tenant businesses serving immigrants are replaced by establishments catering to gentrifiers, with "ethnic packaging" representing a late stage where ethnic identity is deliberately maintained to market neighbourhood real estate after residential populations have dispersed (Hackworth & Rekers, 2005). Beyond displacing established residents, gentrification eliminates rooming houses, converted dwellings, and low-rent apartments that accommodated previous generations of newcomers, forcing new arrivals to settle in distant locations far from employment and settlement services (Valerie Preston & Brian Ray, 2020)..

2.6.4 Hybrid Neighbourhoods: Contact, Tension, and Cooperation

Between ethnic enclave and mainstream dispersal lie hybrid neighbourhoods where multiple ethnic groups maintain significant presence without dominance, serving as sites for interethnic contact that may generate cooperation, tension, or parallel coexistence (Vertovec, 2007). Social psychological research suggests interethnic contact under appropriate conditions can reduce prejudice, with mixed neighbourhoods potentially providing such opportunities through shared schools, parks, and commercial areas (Pettigrew & Tropp, 2006). However, contact theory's optimistic predictions depend on conditions (equal status, common goals, institutional support, no competition) that may not characterize all diverse neighbourhood encounters (Putnam, 2007). Where groups compete for limited resources, diversity may generate tension, with rapid demographic transformation producing conflicts over planning decisions that pit established residents against newcomers (Good, 2009). Despite potential conflict, diverse neighbourhoods also generate bridging institutions including schools, community centres, and faith-based organizations that facilitate cooperation across ethnic boundaries (Putnam, 2007).

2.6.5 Linking Dynamics to Subsequent Analysis

The dynamic processes characterized here, enclave formation, suburbanization, displacement, and interethnic contact, provide the conceptual foundation for computational analyses in subsequent chapters.

The spatial patterns observed in 2021 census data represent momentary snapshots of ongoing processes reshaping Toronto's ethnic geography. Understanding these dynamics requires analytical tools capturing temporal change, spatial dependence, and inter-group interactions. The methodological frameworks developed in Chapters 4 and 5, physics-informed neural networks, graph neural networks, and reaction-diffusion models, are designed to model these processes, treating settlement patterns as evolving systems governed by identifiable mechanisms. The governance context establishes the policy environment within which dynamics unfold, with computational models operating within this frame and potentially informing policy interventions. The intersection of spatial patterns, socioeconomic dimensions, governance context, and dynamic processes creates the complex system that is multi-ethnic Toronto, which the analytical approaches in this dissertation aim to make tractable while preserving substantive richness.

Chapter 3 The Metrics Trap: How Technical Sophistication Masks Social Harm in Urban AI Systems¹

Abstract

This study examines the deployment of artificial intelligence in urban systems through a multiple case study analysis of 23 implementations selected from 157 AI deployments across six domains (2015-2024). We reveal how cities have fallen into a "metrics trap", pursuing technical accuracy that fails to achieve policy goals. Our analysis documents three critical patterns: (1) an "accuracy illusion" where impressive performance metrics mask fundamental failures, exemplified by ShotSpotter's 97% acoustic accuracy yielding only 9.1% crime-fighting effectiveness; (2) discriminatory feedback loops that transform historical bias into computational destiny, affecting millions through predictive policing and housing algorithms; and (3) successful community resistance movements from Toronto to Detroit proving technological determinism is a myth. Cases including Amsterdam's algorithm register, Seoul's participatory waste management, and NYC Health + Hospitals' bias mitigation reveal varied approaches to democratic AI governance. These findings suggest cities can pursue both technical capability and democratic accountability.

3.1 Introduction

Cities worldwide increasingly deploy machine learning algorithms to manage transportation networks, allocate housing, predict crime, monitor environmental conditions, optimize energy consumption, and deliver health services (Allam & Dhunny, 2019; M. Batty, 2013; T. Yigitcanlar et al., 2020). This proliferation has accelerated markedly: documented AI use cases in local governments nearly tripled from 22 cases in 2019 to 58 cases in 2020, peaking at 68 cases in 2021 (Tan Yigitcanlar et al., 2024). Between 2010 and 2020, AI tools became more accessible to municipalities as costs declined and commercial solutions expanded; the release of generative AI systems such as ChatGPT in 2022 further accelerated adoption, enabling diverse stakeholders to experiment with AI for urban analysis and service delivery (Organisation for Economic Co-operation and Development, 2025). Yet as algorithmic systems become embedded in urban infrastructure, a fundamental tension emerges between technical performance and policy effectiveness, a tension this study systematically examines through analysis of 157 documented implementations across 27 countries. Three interconnected scholarly domains frame our investigation: critical algorithm studies documenting systematic harms from automated decision systems (Benjamin,

¹ This article is currently under review in second revision at Urban Sustainability (Nature Portfolio Journal). Authors: Seyed Navid Mashhadi Moghaddam, Huhua Cao

2019; Eubanks, 2018; Noble, 2018; O'neil, 2017); urban governance research critiquing technocratic smart city approaches (Greenfield, 2013; Hollands, 2020; Kitchin, 2014; Söderström, Paasche, & Klauser, 2020); and science and technology studies analyzing how technological systems embed political values and reshape power relations (Bijker, Hughes, & Pinch, 2012; Pasquale, 2015; Winner, 2017). While these literatures establish that algorithmic systems can perpetuate discrimination and reshape urban governance, they provide limited empirical specification of the mechanisms through which technical metrics diverge from policy outcomes, how communities successfully contest algorithmic deployments, and what institutional arrangements enable democratic alternatives.

Understanding these dynamics requires distinguishing among forms of algorithmic opacity and transparency that are often conflated in policy discourse. Black-box models are algorithms whose internal decision-making processes remain opaque to human users, including deep neural networks and ensemble methods that learn patterns through parameters lacking meaningful human interpretation (Y. Bengio, Courville, & Vincent, 2013; Burrell, 2016; Z.C. Lipton, 2018). Post-hoc explainability methods such as LIME and SHAP attempt to approximate black-box reasoning after decisions are made, but as Rudin (2019a) argues, these provide "illusion of transparency without genuine interpretability", what Ananny & Crawford (2018) term "seeing without knowing." Inherently interpretable models, by contrast, are designed from inception to be comprehensible, using transparent logic that humans can examine and contest (Molnar, 2022; Rudin et al., 2022). Distinct from model interpretability is process transparency: the visibility of how algorithmic systems are developed, procured, deployed, and governed, including documentation of purposes, data sources, and accountability mechanisms (Fung, Graham, & Weil, 2007; Pasquale, 2015). This distinction proves crucial because a city may deploy interpretable models without transparent governance processes, or may publicly document opaque systems without enabling meaningful scrutiny of their logic. For this analysis, we classify implementations along a transparency spectrum based on three dimensions:

- (1) **model interpretability**, whether the algorithm uses inherently comprehensible logic (e.g., decision trees, scoring systems, logistic regression) versus opaque architectures requiring post-hoc explanation;
- (2) **process transparency**, whether governance procedures including purpose, data sources, decision criteria, and accountability mechanisms are publicly documented; and
- (3) **contestability**, whether affected stakeholders can understand, question, and influence algorithmic outputs.

Systems may exhibit transparency on some dimensions while remaining opaque on others; our Theme 4 cases were selected for demonstrating transparency across multiple dimensions rather than requiring perfection on any single criterion. Our analysis examines how these dimensions interact, revealing how their configuration shapes whether algorithmic systems serve or undermine democratic urban governance.

The widespread deployment of black-box systems in urban contexts rests on an assumption that merits empirical scrutiny: that achieving high predictive accuracy requires accepting algorithmic opacity. This supposed interpretability-performance tradeoff serves as justification for deploying inscrutable systems, with vendors claiming that transparency would sacrifice capability and officials accepting opacity as the cost of technical sophistication (Kroll et al., 2017; Selbst, Boyd, Friedler, Venkatasubramanian, & Vertesi, 2019). However, recent technical scholarship challenges this assumption. Rudin (2019a) argues the tradeoff is "not a real figure," providing evidence that interpretable models match black-box performance across diverse applications. Wagner et al. (2022) demonstrated interpretable models explained 84% of variance in Berlin's urban emissions analysis while revealing policy-relevant threshold effects. Kim et al. (2023) showed explainable approaches in Seoul's urban modeling maintained competitive accuracy while providing actionable insights. When the interpretability-performance tradeoff is accepted as given, it enables what we term the "metrics trap": algorithmic systems optimized for technical proxies that diverge from, or actively undermine, their stated policy purposes. This pattern extends Strathern (1997)'s formulation of Goodhart's Law ("when a measure becomes a target, it ceases to be a good measure") into algorithmic governance, where Thomas and Uminsky (2022) document how metric optimization can become actively harmful. Muller (2018) analyzes this as "the tyranny of metrics" across institutions; our study examines how it operates specifically in urban AI deployments.

Our analysis draws on theoretical frameworks from multiple disciplines to interpret patterns across cases. From science and technology studies, we employ Winner (2017)'s insight that artifacts have politics, that technological systems embed values and enable particular governance arrangements, alongside Jasanoff (2004)'s concept of co-production, whereby technical and social orders are constituted together. From algorithmic governance scholarship, we draw on Ensign et al. (2018)'s formal models of feedback loops in predictive systems and Lums and Isaac (2016)'s analysis of how biased data generates biased predictions that produce biased enforcement. From institutional theory, we employ Ostrom (1990)'s polycentric governance framework and (Fung & Wright (2003)'s concept of empowered participatory governance. From social movement scholarship, we draw on resource mobilization theory (McCarthy & Zald, 1977), policy diffusion mechanisms (Dolowitz & Marsh, 2000; Walker, 1969), and frame alignment processes (Snow, Rochford Jr, Worden, & Benford, 1986). These frameworks enable us to move beyond documenting that algorithmic harms occur to explaining why harmful systems persist, how communities successfully resist them, and what institutional mechanisms enable democratic alternatives. We also distinguish transparency-as-disclosure from transparency-as-accountability (Fung et al., 2007), a distinction that proves essential for understanding why algorithm registers and documentation requirements vary in their governance effects. Through multiple case study analysis of 23 in-depth cases selected from 157 documented implementations across six urban domains (2015-2024), this article examines how the current paradigm of urban AI

development produces systematic divergence between technical metrics and policy outcomes. We document cases where high technical accuracy coexists with limited policy effectiveness, ShotSpotter achieving 97% acoustic accuracy while only 9.1% of alerts led to gun crime evidence (City of Chicago Office of Inspector General, 2021); COMPAS achieving aggregate accuracy while exhibiting racially disparate error rates (Angwin, Larson, Mattu, & Kirchner, 2016); healthcare algorithms accurately predicting costs while systematically under-identifying Black patients for care (Z. Obermeyer, B. Powers, C. Vogeli, & S. Mullainathan, 2019). We examine how algorithmic systems can create feedback dynamics that amplify existing disparities (Benjamin, 2019; Ensign et al., 2018; Rosen, Garboden, & Cossyleon, 2021). We analyze conditions under which community resistance achieved policy changes, from Boston's facial recognition ban to Detroit's comprehensive reform (A.C.L.U., 2024; A.C.L.U. Massachusetts, 2020). And we document transparent alternatives, Amsterdam's algorithm register, Seoul's participatory waste management, NYC Health + Hospitals' bias mitigation, where interpretable systems and democratic governance processes produced documented benefits (Amsterdam Municipality, 2020; Mackin, Major, Chunara, & Newton-Dame, 2025; Seoul Metropolitan Government, 2014). These patterns suggest that the choice between algorithmic capability and democratic accountability may be false: transparent systems in our sample achieved their documented outcomes not despite but through their responsiveness to democratic input. We offer these findings as systematic observations from a substantial sample warranting further investigation, while acknowledging that generalization requires validation across broader and more representative cases.

3.2 Results

3.2.1 Overview of the Case Universe

Our analysis draws from 157 documented AI implementations across six urban domains spanning 27 countries from 2015 to 2024 (Figure 3.3.1). This mapping enabled identification of patterns in urban AI adoption within our sample, though we acknowledge this dataset reflects documented and accessible cases rather than a comprehensive global census. The temporal boundaries capture the period following the deep learning revolution while providing sufficient time for implementation outcomes to become observable.




Figure 3.1 Global landscape of urban AI implementations (2015–2024).

a, Distribution of 157 documented cases across eleven urban domains, with predictive policing comprising the largest share ($n=54$). *b*, Geographic distribution showing concentration in North America and Europe, with the United States accounting for the most cases ($n=98$). *c*, Country-level domain specialization patterns for the ten most-represented nations, with bar length indicating specialization ratio relative to global distribution. *d*, Sankey diagram illustrating relationships between domains, regions, and analytical themes examined in this study. Case selection reflects documentation accessibility rather than comprehensive global deployment.

Within our sample, domain distribution (Figure 3.3.1a) shows notable variation. Predictive policing and public safety systems account for 54 cases (34.4%), followed by traffic management (27 cases, 17.2%), facial recognition (20 cases, 12.7%), and autonomous vehicles (20 cases, 12.7%). Domains associated with transparent governance approaches, including digital twins and algorithm registers (14 cases), smart waste management (9 cases), and smart city infrastructure (5 cases), constitute a smaller proportion. This distribution within our sample may reflect documentation availability and research attention rather than

actual global deployment patterns; surveillance-related systems tend to generate more public scrutiny and thus more accessible documentation than routine municipal services.

Geographic distribution within our sample (Figure 3.3.1b) shows concentration in English-speaking and European contexts, with the United States accounting for the largest share, followed by European nations including Spain, Germany, and the UK. Asian implementations appear primarily from China, India, and Singapore. The Sankey diagram (Figure 3.3.1d) visualizes relationships between domains, regions, and article themes within our dataset. Country-level patterns (Figure 3.3.1c) show that within our sample, different nations appear more frequently in particular domains, for instance, cases from Germany cluster in predictive policing, while cases from India concentrate in air quality monitoring, and South Korean cases feature prominently in smart waste management. These patterns likely reflect both actual policy emphases and language/documentation accessibility limitations in our search strategy.

From this broader dataset, we selected 23 cases for in-depth thematic analysis based on documentation richness (evidence from at least four source types), demonstrated social impact, and theoretical significance for illuminating the metrics trap phenomenon. These cases were organized across four analytical themes: the accuracy illusion (Theme 1, examining six cases), discriminatory feedback loops (Theme 2, seven cases), community resistance (Theme 3, six cases including one contrast case), and transparent alternatives (Theme 4, six cases). Some cases appear in multiple themes where their characteristics warranted examination from different analytical perspectives. The following sections examine these patterns in detail.

3.2.2 Theme 1: The Accuracy Illusion

The accuracy illusion emerges as a recurring pattern across urban AI implementations in our sample, where technical performance metrics diverge from policy outcomes. Our analysis of 21 cases across eight countries exhibiting this pattern (Figure 3.3.2) reveals five distinct manifestations: detection/prediction disconnected from real-world action (5 cases), systemic failures despite claimed accuracy (5 cases), vendor claims contradicted by verified reality (4 cases), accuracy metrics masking fairness failures (3 cases), and technical success without social impact (3 cases). Within our sample, predictive policing systems account for the largest share (9 cases), followed by healthcare AI (5 cases), with housing, autonomous vehicles, environmental monitoring, and traffic management comprising the remainder (Figure 3.3.2D). Notably, 30% of systems exhibiting accuracy illusion patterns in our sample have been ended or cancelled, while 65% remain active despite documented gaps between technical claims and policy effectiveness (Figure

3.3.2C). This section examines six cases in depth to illustrate the mechanisms through which impressive technical metrics can obscure fundamental failures in achieving stated policy goals.

Theme 1: The Accuracy Illusion

When Technical Performance Masks Social Failure — Evidence from 21 Case Studies Across 8 Countries

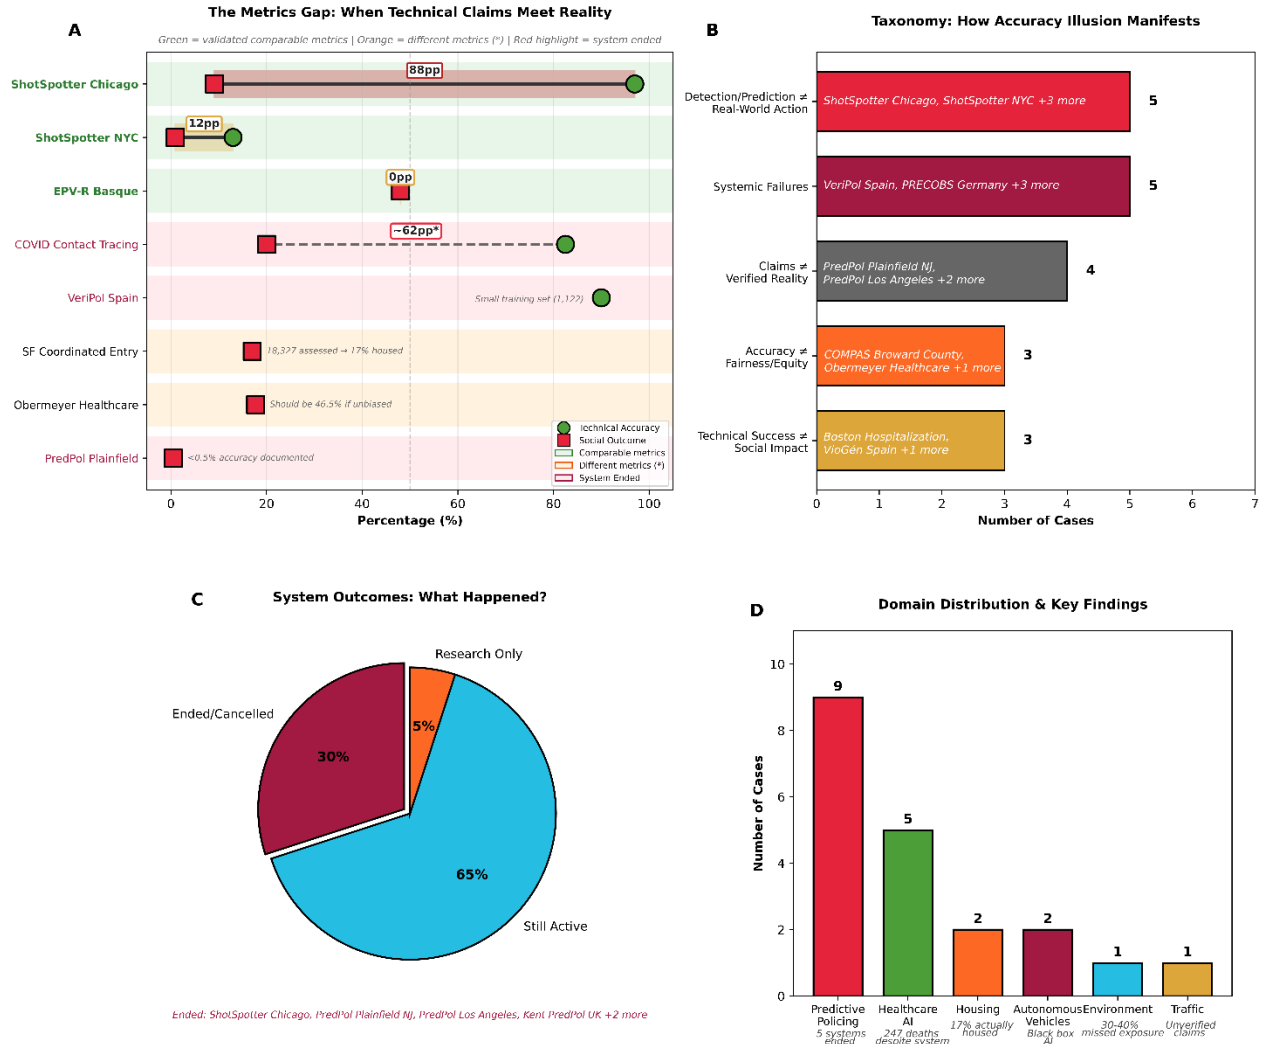


Figure 3.2 The Accuracy Illusion: Evidence from 21 Case Studies Across 8 Countries.

a, The metrics gap between technical claims and social outcomes, with color coding indicating comparable metrics (green), different metric types (orange), and systems that have ended (red highlight). ShotSpotter Chicago exhibits the largest documented gap (88 percentage points). *b*, Taxonomy of accuracy illusion manifestations showing five distinct patterns across cases. *c*, System outcomes: 30% of systems in our sample have been ended or cancelled, 65% remain active, and 5% exist only as research tools. *d*, Domain distribution within Theme 1 cases, with predictive policing representing the largest category (9 cases) followed by healthcare AI (5 cases).

The acoustic gunshot detection system ShotSpotter exemplifies the accuracy illusion through its divergence between technical precision and operational utility. The company claims 97% accuracy with 0.5% false positive rates in detecting gunshot sounds (Bronars & Lynch, 2022). However, the Chicago Office of Inspector General's analysis of 50,176 ShotSpotter alerts from January 2020 through May 2021 found that only 9.1% led to evidence of a gun-related criminal offense (City of Chicago Office of Inspector General, 2021), an 88 percentage point gap between claimed acoustic accuracy and crime-fighting effectiveness (Figure 3.3.2A). This disconnect stems from a category error: accurately detecting loud sounds differs from identifying gun crimes requiring police response. (Doucette, Green, Necci Dineen, Shapiro, & Raissian, 2021) analyzed ShotSpotter implementations across 68 large metropolitan counties from 1999-2016, finding no significant impact on firearm homicide or arrest rates. A 2024 New York City audit revealed that 87% of ShotSpotter alerts were confirmed as false alarms, with fewer than 0.9% resulting in firearm recovery. As Piza(2023) documented in Kansas City, while the technology achieves spatial precision in locating sounds, this technical achievement did not translate into measurable public safety improvements in that context. Chicago terminated its \$33 million ShotSpotter contract in September 2024.

The Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) algorithm illustrates how moderate overall accuracy can mask disparate error rates across demographic groups. ProPublica's analysis of over 7,000 criminal defendants in Broward County, Florida, found that COMPAS achieved approximately 60-70% overall accuracy in predicting recidivism (Angwin et al., 2016). However, this aggregate metric concealed differential false positive rates: Black defendants were falsely flagged as high-risk at 44.9% compared to 23.5% for white defendants. Conversely, white defendants who reoffended were incorrectly labeled low-risk 47.7% of the time compared to 28.0% for Black defendants. Even after controlling for criminal history, age, and gender, Black defendants remained 77% more likely to be scored as high-risk for violent recidivism. The accuracy illusion operates here through aggregation: overall performance metrics obscure how errors distribute across groups, a finding consistent with subsequent research demonstrating mathematical tensions between different fairness criteria in risk assessment (Chouldechova, 2017; Kleinberg, Ludwig, Mullainathan, & Sunstein, 2018). Similarly, the Obermeyer healthcare algorithm achieved high accuracy for its optimization target, healthcare costs, yet only 17.7% of Black patients who should have been identified for extra care based on health needs were flagged, compared to an estimated 46.5% if the algorithm performed equitably (Z. Obermeyer et al., 2019). The algorithm accurately predicted costs, but structural inequalities meant Black patients generated lower costs despite equal or greater health needs due to barriers in accessing care.

Environmental and health prediction systems in our sample exhibit a related pattern: technical accuracy that fails to capture differential impacts across populations. Delhi's air quality prediction system using Grey Wolf Optimization with Decision Tree algorithms achieved 88.98% accuracy for Air Quality Index

prediction (Natarajan, Shanmurthy, Arockiam, Balusamy, & Selvarajan, 2024). However, standardized monitoring networks may not capture hyperlocal variations affecting residents near industrial areas or informal settlements, where studies have documented pollution levels 30-40% higher than city averages. The system optimizes for meteorological prediction rather than equitable exposure assessment. Boston University researchers' hospitalization prediction model achieved 82% accuracy compared to 56% for clinical guidelines (Hao et al., 2020), yet the model's reliance on medical records and diagnostic data means it cannot incorporate social determinants, housing instability, food insecurity, transportation barriers, that shape health outcomes. For patients facing structural barriers, accurate prediction without addressing root causes documents outcomes rather than enabling intervention. COVID-19 contact tracing applications achieved 70-95% proximity detection accuracy in laboratory conditions, yet real-world effectiveness was limited by adoption rates below 20% in many countries and inability to distinguish transmission risk contexts (F. Yi, Xie, & Jamieson, 2021). Singapore's TraceTogether required 75% adoption for meaningful epidemic control, a threshold that was not reached.

These six cases suggest several mechanisms through which accuracy illusion may operate, though we note these patterns derive from our selected sample rather than representing universal claims about urban AI. First, each case exhibits misalignment between algorithmic optimization targets and policy objectives: ShotSpotter optimized for sound detection rather than crime prevention; COMPAS and Obermeyer optimized for aggregate prediction rather than equitable outcomes; environmental and health systems optimized for technical accuracy rather than population-level impact. Second, aggregated metrics can obscure distributional failures, overall performance statistics may hide how accuracy varies across populations or contexts. Third, in several cases, technically accurate systems appeared to generate limited policy value: Chicago and New York audits found ShotSpotter alerts rarely led to actionable outcomes; COMPAS accuracy barely exceeded chance while exhibiting disparate error rates; high-performing health predictions could not address social determinants driving hospitalization. Table 3-3-1 summarizes these patterns. We emphasize that these observations derive from documented cases in our sample; whether they represent broader patterns in urban AI deployment requires further systematic investigation across larger and more representative samples.

Table 3-1 The Metrics Gap in Selected Urban AI Systems

Case	Domain	Technical Metric Claimed	Documented Social Outcome	Metrics Gap	Failure Mode

ShotSpotter (Chicago)	Public Safety	97% acoustic accuracy	9.1% of alerts led to gun crime evidence	88pp		Sound detection ≠ crime prevention
ShotSpotter (NYC)	Public Safety	High detection rate	87% false alarms; <0.9% firearm recovery	12pp documented		Lab accuracy ≠ field effectiveness
COMPAS (Broward County)	Public Safety	60-70% overall accuracy	2x false positive rate for Black defendants	0pp overall; disparate errors		Aggregate accuracy ≠ equitable outcomes
Obermeyer Healthcare	Health	High cost prediction accuracy	17.7% vs 46.5% Black patient identification	164% potential increase if corrected		Cost proxy ≠ health needs
Delhi Air Quality	Air Environment	88.98% AQI accuracy	30-40% higher exposure in informal settlements not captured	Unknown		City-wide accuracy ≠ hyperlocal equity
Boston Hospitalization	Health	82% accuracy (vs 56% guidelines)	Cannot address social determinants	26pp improvement without root cause intervention		Prediction ≠ prevention
COVID-19 Contact Tracing	Health	70-95% proximity accuracy	<20% adoption in many countries	~62pp gap between lab and field		Technical accuracy ≠ population impact

Note: pp = percentage points. Metrics gaps calculated where comparable measures available.

3.2.3 Theme 2: Discriminatory Feedback Loops

Discriminatory feedback loops represent a pattern wherein algorithmic systems may not only reflect existing societal biases but potentially amplify them through iterative operation. Our analysis identified 11 cases across three domains exhibiting documented feedback loop characteristics: predictive policing (5 cases), housing algorithms (4 cases), and energy systems (2 cases) (Figure 3.3a). These cases span implementation periods from 2007 to present, with three systems subsequently terminated or discontinued (Chicago SSL in 2019, Palantir New Orleans in 2018, HART Durham in 2020), while others remain active

or face ongoing regulatory action (Figure 3.3c). Regulatory responses to documented harms in these cases have resulted in over \$36 million in fines and settlements, including \$23 million against TransUnion, \$4.2 million against AppFolio, and \$3 million against RealPage (Figure 3.3d). This section examines seven cases in depth to illustrate mechanisms through which algorithmic systems may generate self-reinforcing cycles of disadvantage.

Theme 2: Discriminatory Feedback Loops

How Algorithmic Systems Amplify Historical Bias — Evidence from 11 Documented Cases

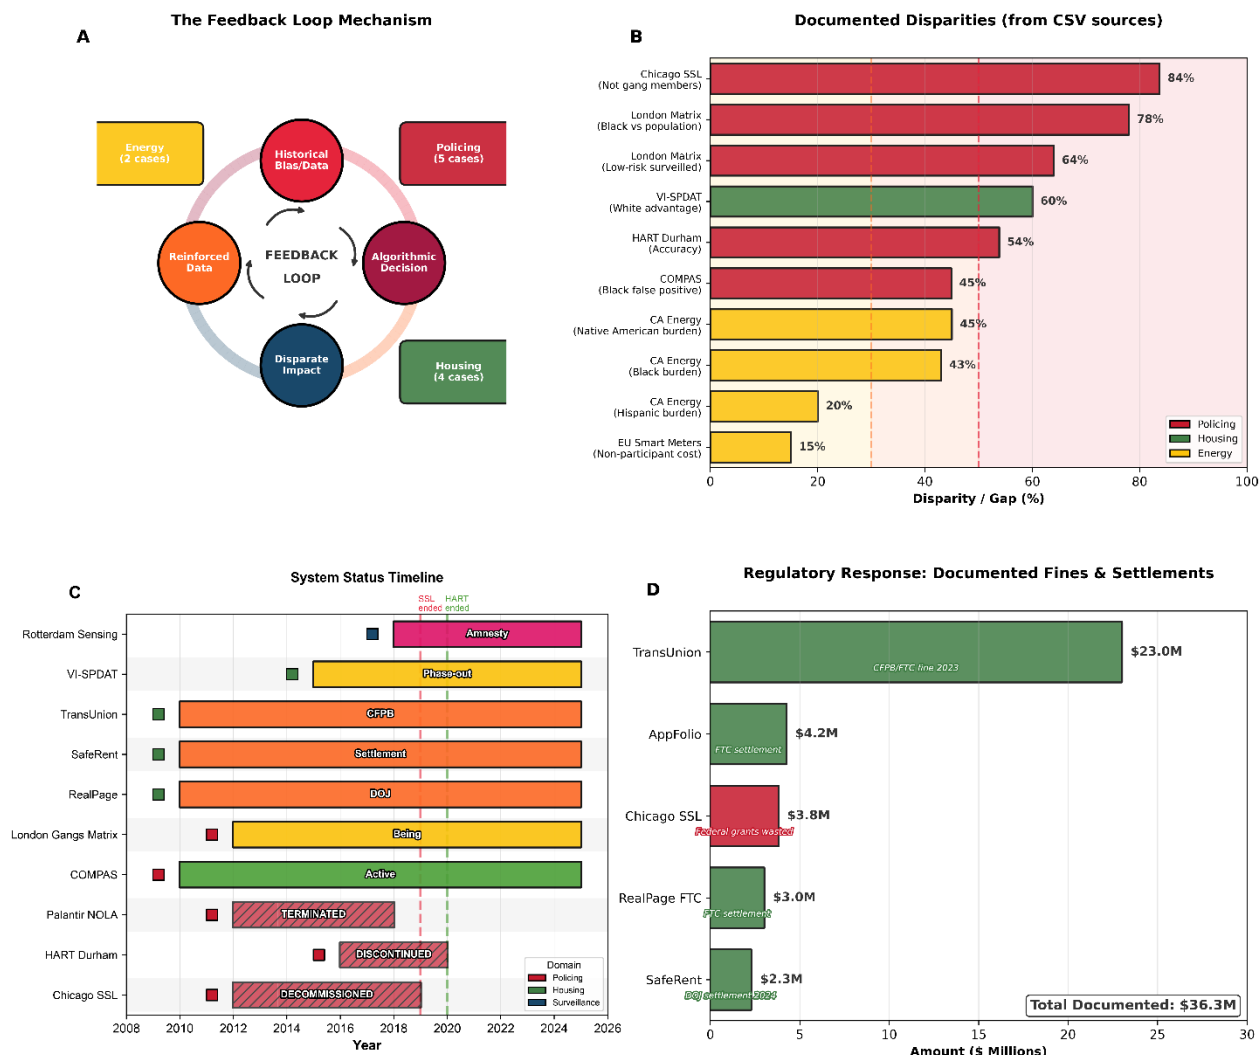


Figure 3.3 Discriminatory Feedback Loops: Evidence from 11 Documented Cases.

a, The feedback loop mechanism showing how historical bias feeds algorithmic decisions, which generate disparate impacts that reinforce biased data, with case distribution across policing (5), housing (4), and energy (2) domains. b, Documented disparities from case sources, ranging from Chicago SSL's 84% gap between algorithmic classification and verified status to EU Smart Meters' 15% cost penalty for non-participants. c, System status timeline showing implementation periods and key events including

terminations (Chicago SSL 2019, Palantir NOLA 2018, HART Durham 2020) and regulatory interventions. d, Regulatory response totaling \$36.3 million in documented fines and settlements across tenant screening systems.

Chicago's Strategic Subject List (SSL), operational from 2012-2019, illustrates how predictive policing systems can create feedback dynamics through recursive data generation. The system assigned risk scores from 1-500 to nearly 400,000 individuals based on attributes including arrest records and social network connections (City of Chicago Office of Inspector General, 2021). The Chicago Office of Inspector General found that only 16.3% of individuals on the list were confirmed gang members, despite district commanders believing the figure was approximately 95%, an 84 percentage point gap between algorithmic classification and verified status (Figure 3.3b). This discrepancy suggests how algorithmic outputs may acquire institutional authority that diverges from ground truth. Ensign et al.(2018) developed formal models demonstrating how predictive policing can create "runaway feedback loops" where increased police presence in flagged areas generates more arrests, which become new data points reinforcing original predictions regardless of underlying crime rate changes. The SSL was decommissioned in November 2019 after \$3.8 million in federal grants, without documented evidence of violence reduction (City of Chicago Office of Inspector General, 2021). Similarly, the London Metropolitan Police's Gangs Matrix assigned "harm scores" to approximately 3,800 individuals, with the Information Commissioner's Office finding that 78% were Black despite Black Londoners comprising 13% of the population, and 64% of those listed were classified as lowest risk yet remained subject to surveillance (I.C.O., 2018).

New Orleans' covert partnership with Palantir Technologies from 2012-2018 demonstrates how opacity in algorithmic systems may compound feedback effects by preventing external verification or contestation. The partnership, exposed by The Verge in February 2018, operated without city council knowledge or public oversight, assessing approximately 3,900 individuals through a risk assessment database (Winston, 2018). Criminal defense attorneys reported never receiving Palantir analytical products in discovery materials, raising questions about defendants' ability to challenge algorithmic bases for targeting. The program was terminated in March 2018 following public exposure. Rotterdam's Sensing Project, documented in Amnesty International's 2022 report, illustrates a related pattern through automated ethnic profiling. The system monitored vehicles and assigned risk scores specifically targeting what authorities termed "mobile banditry" allegedly associated with Eastern European nationals (Amnesty International, 2022). Amnesty documented that the system created conditions where vehicles flagged as potentially connected to Eastern Europeans faced enhanced monitoring, increasing likelihood of detecting irregularities that then became data points supporting original risk assessments, a pattern the organization concluded violated rights to privacy, data protection, and non-discrimination.

The tenant screening industry demonstrates how algorithmic feedback loops may operate through housing markets. RealPage (acquired for \$9.6 billion in 2020) and CoreLogic (acquired for \$6 billion in 2021) dominate a market where industry reports indicate widespread landlord adoption of screening software (Bravo, 2020; CoreLogic, 2021). Rosen, Garboden, and Cossyleon(2021) documented how algorithmic scoring may perpetuate discrimination through variables that correlate with race, such as credit histories shaped by historical lending disparities or addresses in previously redlined neighborhoods. When these systems deny housing, applicants may face housing instability affecting credit scores and rental history, potentially making future denials more likely. Federal Trade Commission settlements totaling \$7.25 million against RealPage and AppFolio addressed accuracy and disclosure concerns (F.T.C., 2018, 2020), while the Department of Justice's lawsuit against SafeRent alleged these tools affect millions of rental applications (Louis v. Saferent Sols., 2023). The VI-SPDAT homelessness assessment tool, used across 16+ U.S. communities, exhibited documented racial disparities: multiple studies found white individuals were approximately 60% more likely to receive high prioritization scores compared to Black individuals assessed at similar vulnerability levels (Figure 3.3b).

Energy systems in our sample exhibit related patterns operating through different mechanisms. California's proposed income-graduated fixed charge of \$24.15 monthly would apply uniformly regardless of consumption, with analysis by the American Council for an Energy-Efficient Economy documenting existing disparities: low-income households spend 8.1% of income on energy versus 2.3% for higher-income households, with Black households spending 43% more, Hispanic households 20% more, and Native American households 45% more of their income on energy (Drehobl, Ross, & Ayala, 2020). The European Union's smart meter rollout, reaching 54% of households by 2021 with €47 billion invested (ACER, 2021), creates conditions where households unable to engage with smart technology, due to lacking smartphones, internet access, or technical literacy, may face up to 15% higher costs through inability to shift consumption to off-peak periods. These seven cases suggest three potential feedback mechanisms operating across domains: enforcement feedback, where differential observation generates differential data (Chicago SSL, London Gangs Matrix); surveillance cascades, where initial algorithmic judgments expand into multiple institutional domains (Palantir, Rotterdam Sensing); and economic exclusion loops, where algorithmic denials or burdens may compound over time (tenant screening, energy systems). As Akpinar, De-Arteaga, and Chouldechova(2021) demonstrated, even attempts to train algorithms on alternative data sources may not eliminate bias when structural factors shape data generation processes themselves. Table 3-3-2 summarizes documented disparities and mechanisms across these cases; we note that while these patterns appear consistent within our sample, their generalizability to other contexts requires further investigation.

Table 3-2 Documented Feedback Loop Mechanisms in Selected Cases

System	Domain	Initial Disparity Source	Documented Disparity	Feedback Mechanism	Status
Chicago SSL	Policing	Arrest history weighting	84% gap: actual vs believed members	16.3% 95% gang	Enforcement feedback Decommissioned 2019
London Gangs Matrix	Policing	78% (13% Black of population)	64% lowest-risk still surveilled	Surveillance cascade	Active (under reform)
Palantir New Orleans	Policing	Covert operation without oversight	3,900 individuals assessed secretly	Opacity-enabled feedback	Terminated 2018
Rotterdam Sensing	Policing	Ethnic profiling of Eastern Europeans	Documented violations	rights Cross-border discrimination	Active (Amnesty criticism)
Tenant Screening	Housing	Historical rental/credit data	\$36.3M in regulatory fines/settlements	Economic exclusion	Active (regulatory action)
VI-SPDAT	Housing	Assessment tool design	60% advantage white in prioritization	Scoring disparity	Phase-out recommended
CA Fixed Charges / EU Smart Meters	Energy	Income-blind structures	15-45% disparity by income/race	Digital divide amplification	Proposed / Active

Note: Disparities calculated from documented sources as cited in text.

3.2.4 Theme 3: Community Resistance and Democratic Pushback

Community resistance represents a countervailing dynamic in several documented cases, where grassroots organizing challenged algorithmic system deployments. Our analysis identified 20 cases exhibiting resistance patterns, resulting in seven documented system terminations or major reforms, ten municipal facial recognition bans, and two state-level warrant requirements (Figure 3.4). These cases cluster in contexts with robust civil society infrastructure: the United States accounts for the majority of documented

Theme 3: Community Resistance and Democratic Pushback

Documented cases of algorithmic systems challenged through democratic organizing

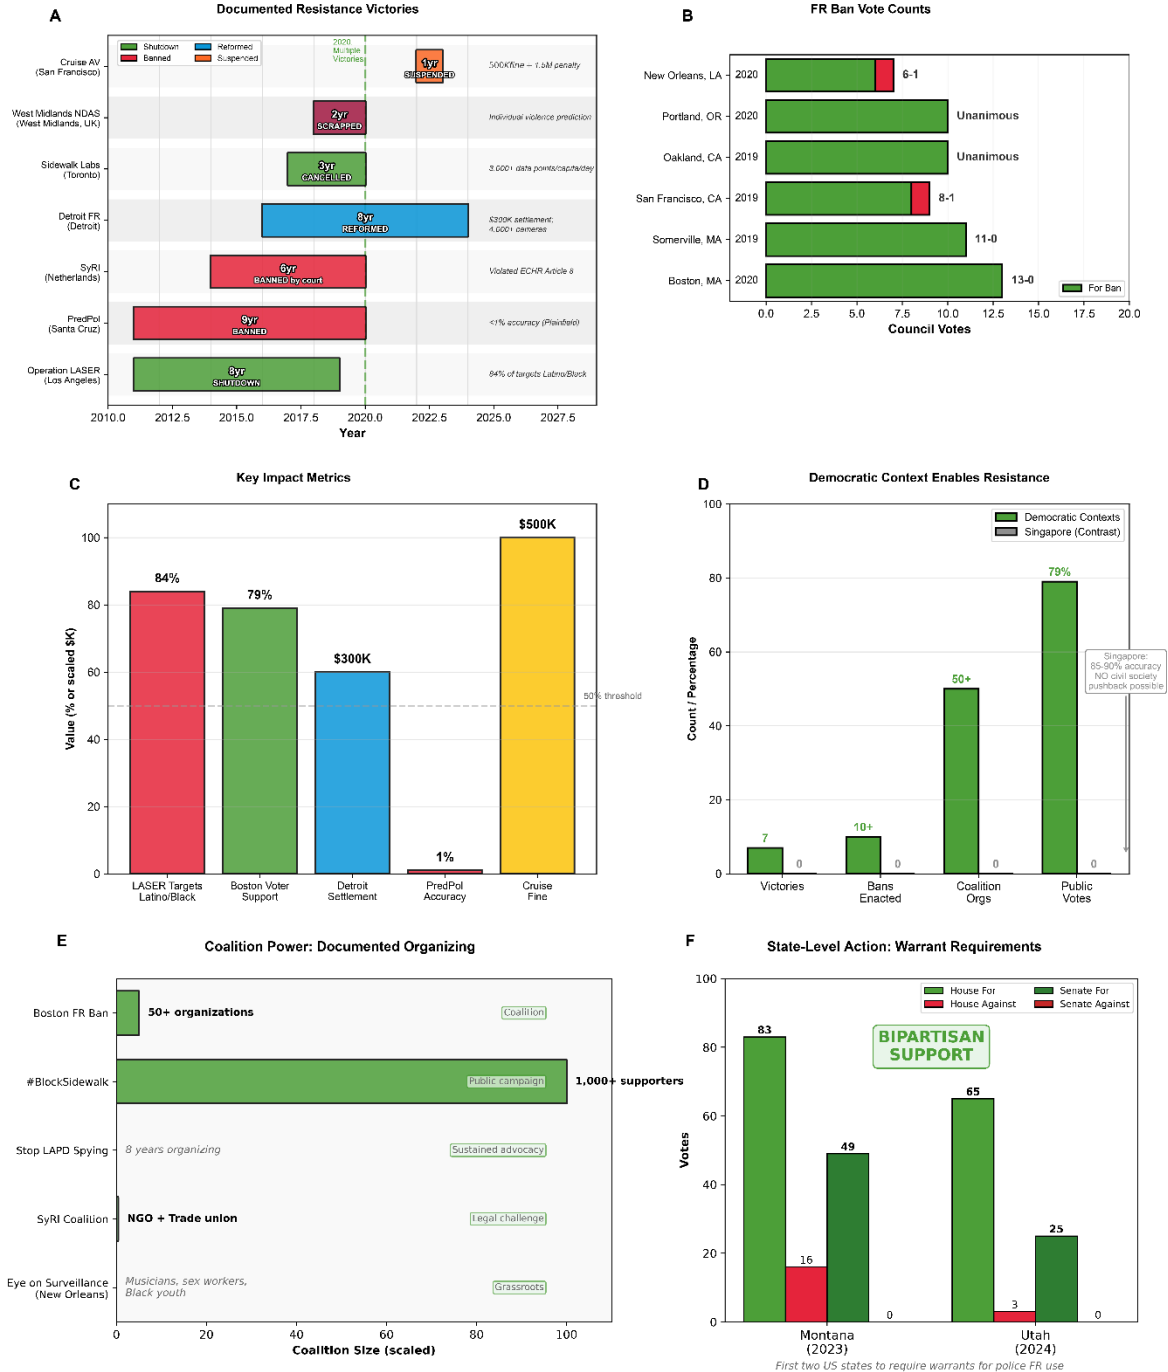


Figure 3.4 Community Resistance and Democratic Pushback: Evidence from 20 Documented Cases.

a, Timeline of resistance victories showing implementation periods, outcomes (shutdown/banned/reformed/suspended), and key documented harms. b, Facial recognition ban vote counts across U.S. municipalities, with Boston's 13-0 unanimous vote and subsequent regional diffusion to five Massachusetts cities. c, Key impact metrics including LASER's 84% Latino/Black targeting rate, Boston's

79% voter support, Detroit's \$300,000 settlement, and Cruise's \$500,000 federal fine. d, Democratic context comparison showing resistance indicators (victories, bans enacted, coalition organizations, public support) in contexts where resistance occurred versus Singapore where no documented resistance emerged. e, Coalition composition and scale across five major campaigns, from Boston's 50+ organization coalition to grassroots movements in New Orleans. f, State-level warrant requirements for facial recognition in Montana (2023) and Utah (2024), both achieving bipartisan legislative majorities.

resistance victories, with a notable concentration in Massachusetts where Boston's 2020 ban (passed 13-0) preceded adoption in five additional municipalities (Figure 3.4B). Coalition sizes ranged from grassroots networks of affected residents to formal alliances of 50+ organizations (Figure 3.4E). This section examines six cases in depth to illustrate conditions and tactics associated with successful resistance, while noting that Singapore's traffic AI deployment, where no documented resistance occurred, may illuminate contextual factors that enable or constrain democratic pushback.

The defeat of Google's Sidewalk Labs smart city project in Toronto represents a case where sustained opposition resulted in complete project cancellation. The project, announced in 2017, proposed deploying sensor networks collecting data across Toronto's Quayside waterfront. Opposition led by the Canadian Civil Liberties Association (CCLA) and the #BlockSidewalk coalition, which mobilized over 1,000 supporters (Figure 3.4E), culminated in project cancellation in May 2020 (CBC News, 2020). The CCLA filed a constitutional challenge under the Charter of Rights and Freedoms, arguing the project raised privacy and democratic governance concerns (C.C.L.A., 2020). The coalition achieved momentum when high-profile advisors resigned, including former Ontario Privacy Commissioner Ann Cavoukian, who departed citing privacy concerns. Similarly, the Netherlands' SyRI welfare fraud detection system, which cross-referenced personal data from multiple government databases and was deployed in low-income neighborhoods, was banned by court ruling in February 2020 after a coalition of NGOs and a trade union filed legal challenges arguing it violated European Convention on Human Rights Article 8 (Amnesty International, 2022). Both cases suggest that legal challenges combined with coalition organizing can achieve system termination, though we note these represent specific political and legal contexts that may not generalize.

Detroit's trajectory from facial recognition deployment to reform illustrates how individual cases of harm can catalyze policy change when connected to organized advocacy. The January 2020 wrongful arrest of Robert Williams, documented as the first U.S. case of facial recognition misidentification leading to arrest, became a focal point for organizing by Detroit Community Technology Project, ACLU Michigan, and allied organizations (A.C.L.U., 2024). The June 2024 settlement achieved \$300,000 compensation for Williams and established policies requiring corroborating evidence before arrests based on facial recognition matches (Figure 3.4C). Boston's facial recognition ban campaign, led by ACLU Massachusetts'

"Press Pause Face Surveillance" initiative, achieved unanimous City Council passage (13-0) with documented support from 79% of Massachusetts voters in polling (A.C.L.U. Massachusetts, 2020; Boston City Council, 2020). The campaign united over 50 organizations spanning racial justice, immigrant rights, and civil liberties groups (Figure 3.4E). Subsequently, Springfield, Cambridge, Northampton, Brookline, and Somerville adopted similar bans, creating regional policy diffusion within 18 months. At the state level, Montana (2023) and Utah (2024) enacted warrant requirements for police facial recognition use with bipartisan support, Montana's legislation passed 83-16 in the House and 49-0 in the Senate (Figure 3.4F). The Los Angeles Police Department's 2019 termination of Operation LASER after eight years demonstrates how sustained community pressure can achieve system shutdown. The Stop LAPD Spying Coalition conducted community research documenting that 84% of individuals targeted by LASER were Latino or Black (Figure 3.4C), findings subsequently confirmed by the Inspector General's audit (Brennan Center for Justice, 2024; CBS News, 2024). New Orleans' Palantir predictive policing partnership, operating covertly from 2012-2018 without city council knowledge, was terminated within months of The Verge's February 2018 exposé (Winston, 2018). The revelation that criminal defense attorneys had never received Palantir analytical products in discovery materials intensified public pressure. These cases suggest that documentation of discriminatory impacts (LASER) and exposure of secret operations (Palantir) can delegitimize systems, though the specific conditions enabling termination varied: LASER required eight years of sustained organizing, while Palantir collapsed rapidly once its covert nature was revealed. Singapore's IBM traffic prediction system, deployed through the Land Transport Authority in 2006-2007 and achieving 85-90% accuracy (IBM, 2007), provides an instructive contrast that illuminates the relationship between system characteristics, democratic context, and resistance potential (Figure 3.4D). Critically, this system differs fundamentally from the person-based predictive systems documented in Themes 1-2: traffic flow optimization uses aggregate vehicle data to manage congestion rather than generating individualized risk scores that mark specific people for differential treatment (Ferguson, 2017; Lum & Isaac, 2016). This distinction matters because person-based prediction creates discriminatory feedback loops through recursive data generation, a dynamic largely absent in aggregate traffic optimization. We therefore do not suggest Singapore's traffic system warranted resistance comparable to systems like Chicago's SSL or COMPAS; the absence of documented opposition may appropriately reflect lower stakes for civil liberties.

However, Singapore's broader "Smart Nation" initiative does raise privacy concerns that have been documented by human rights organizations. Human Rights Watch (2017) and CIVICUS(2024) have characterized Singapore's civic space as "repressed," noting extensive surveillance networks, restrictions on public assembly under the Public Order Act, and laws that limit criticism of government systems. The smart city infrastructure promises efficiency, but as NYU's Center for Human Rights and Global Justice

(2022) documented, "the constant technology-driven surveillance and the loss of a few civil liberties are viewed by many as a small price to pay for such efficiency", a framing that forecloses rather than enables democratic deliberation about algorithmic governance. The comparison suggests that democratic infrastructure, civil society organizations, protected assembly rights, independent media, may be preconditions for the resistance patterns documented elsewhere, while acknowledging that the appropriateness of resistance depends on the specific harms a system produces. We caution against overgeneralizing from a single contrasting case. Table 3-3 summarizes implementation contexts and outcomes across these six cases; the patterns observed may reflect selection effects in our sample rather than universal dynamics of algorithmic resistance.

Table 3-3 Implementation Context and Resistance Outcomes in Selected Cases

Case	Period	Resistance Type	Key Actors	Outcome	Timeframe to Outcome
Toronto Sidewalk Labs	2017-2020	Coalition + legal challenge	CCLA, #BlockSidewalk (1,000+ supporters)	Complete cancellation	3 years
Detroit Facial Recognition	2016-2024	Individual case → systemic reform	ACLU Michigan, Detroit Community Tech	\$300K settlement + strongest US FR policy	4 years
New Orleans Palantir	2012-2018	Investigative exposure	The Verge, ACLU Louisiana	Program terminated	Months after exposure
Boston FR Ban	2019-2020	Proactive coalition	50+ organizations, 79% voter support	Ban passed 13-0, regional spread	1 year
LA Operation LASER	2011-2019	Research + sustained pressure	Stop LAPD Spying Coalition	Shutdown	8 years
Singapore Traffic AI	2006-present	No documented resistance	N/A	Continued operation	N/A (contrast case)

Note: Outcomes and timeframes reflect documented sources; absence of documented resistance does not indicate public support.

3.2.5 Theme 4: Transparent Alternatives and Democratic AI Governance

Figure 3.5 maps transparency characteristics across urban AI implementations in our sample, revealing variation along multiple dimensions: process transparency (Figure 3.5a), documented performance outcomes (Figure 3.5b), the relationship between model opacity and governance openness (Figure 3.5c), and democratic engagement scale (Figure 3.5d). Importantly, these cases occupy different positions on the transparency spectrum defined in our Introduction: some emphasize process transparency through documentation and registers (e.g., Amsterdam, Helsinki), others demonstrate contestability through stakeholder feedback mechanisms (e.g., Austin Energy's customer-facing interfaces, Seoul's waste collector feedback integration). The NYC Health + Hospitals case uniquely combines documented bias assessment with open-source methodology enabling independent verification. We present reported outcomes while acknowledging that cross-case comparisons are complicated by differences in urban context, implementation timeline, evaluation methodology, and data availability. Rather than claiming these cases prove transparent systems universally outperform opaque alternatives, a claim that would require controlled comparisons unavailable in real-world municipal deployments, we document the specific mechanisms through which transparency was operationalized and the outcomes that implementing organizations have reported.

Theme 4: The Transparency Spectrum

Process Transparency, Model Interpretability, and Democratic Engagement in Urban AI



Figure 3.5 The Transparency Spectrum in Urban AI Implementations.

a, Process transparency levels across 16 implementations, ranked from low (Hangzhou City Brain) to very high (EU Urban Atlas, 785+ cities). **b**, Documented performance outcomes by domain: cost reduction (Seoul waste, 83%), participation rates (Austin Energy, 80%), recycling rates (Seoul, 60%+), and traffic improvements (Surtrac, 25-40%). **c**, Process transparency versus model opacity, with bubble size indicating democratic engagement level. Lower-right quadrant represents exemplary cases combining open governance with interpretable models. **d**, Democratic engagement scale showing participant numbers: Singapore LTMP2040 (7,400+ citizens), Seoul waste (320+ public meetings), Barcelona Decidim (200+ cities adopted), EU AI Clauses (40+ expert roundtables). Annotation indicates documented equity outcomes from NYC Health + Hospitals (5pp bias gap closed (Exemplary)) and ZSFG (5.4pp racial disparity gap closed).

Amsterdam and Helsinki launched municipal algorithm registers in September 2020, creating publicly accessible documentation of algorithmic systems used in city operations (Amsterdam Municipality, 2020; City of Helsinki, 2020). Amsterdam's register, implemented through the Saidot platform, documents systems from parking compliance to citizen report categorization, providing standardized fields for purpose, data sources, operational logic, and human oversight mechanisms (Saidot, 2025). Each entry includes plain-language descriptions intended to make algorithmic operations comprehensible to non-technical audiences. Amsterdam officials reported that the registration process prompted departments to articulate system purposes more explicitly, leading to identification of some redundancies, though no independent evaluation has quantified governance improvements (Elliot, 2025). Helsinki's parallel register documents rule-based chatbots and AI systems with feedback channels for citizen input (City of Helsinki, 2020). Barcelona's Algorithmic Implementation Protocol, established in 2021, categorizes AI systems by risk level and mandates external auditing for high-risk applications (Barcelona City Council, 2021). The city's vCity project (2023-2025) integrates digital twin technology with the Decidim participatory platform, which has been adopted by over 200 cities globally, enabling residents to visualize algorithmic applications and propose modifications (OECD-OPSI, 2024). The EU's Living-in.EU Model AI Contractual Clauses (2023-2024), developed through 40+ expert roundtables, provide standardized procurement language for transparency requirements (EU Public Buyers Community, 2024). These institutional mechanisms represent attempts to operationalize algorithmic accountability, though their effectiveness in preventing harms depends on implementation fidelity, a factor we cannot assess from available documentation.

Seoul's smart waste management system, implemented through Ecube Labs beginning in 2014, reported an 83% reduction in collection costs following installation of 85 Clean Cube sensors with neural network-based route optimization (Ecube Labs, 2015; Seoul Metropolitan Government, 2014). The IoT sensors transmit real-time fill-level data accessible to collection staff through the Clean City Networks platform, and route optimization prioritizes bins exceeding fill thresholds while minimizing travel distances. Waste collectors provided operational feedback identifying constraints the algorithm initially missed, narrow alleys, market day obstacles, seasonal variations, enabling iterative refinement. The broader pay-as-you-throw system involved over 320 public meetings since 2001, with citizens' councils including environmental specialists contributing to system design; Seoul reports recycling rates exceeding 60% (Seoul Metropolitan Government, 2014). Austin Energy's demand response program, implemented through AutoGrid's DROMS platform, achieved 177 MW of demand response capacity with documented participation rates of 80% in smart thermostat pilots, compared to industry-reported averages of 20-30% for similar programs (Austin Energy; AutoGrid Systems). The system provides customer-facing interfaces displaying real-time consumption, predicted peak events, and potential savings. Austin Energy established a Low Income Consumer Advisory Task Force and provides CAP discounts for income-qualified

customers, though independent assessment of equity outcomes has not been published. These cases illustrate how transparent operational logic may facilitate stakeholder engagement, though the contribution of transparency specifically, versus other factors such as program design, incentive structures, or local context, cannot be isolated from available data.

NYC Health + Hospitals, the largest municipal safety-net system in the United States serving over one million patients (90% patients of color, 70% Medicaid or uninsured), provides a documented example of proactive algorithmic bias assessment and mitigation in healthcare AI (Mackin et al., 2025). The system evaluated two binary classification models in their electronic medical record, one predicting acute visits for asthma and one predicting unplanned readmissions, across race/ethnicity, sex, language, and insurance status using Equal Opportunity Difference (EOD), a metric comparing false negative rates across subgroups. For the asthma model, baseline analysis revealed bias across race/ethnicity, with false negative rates ranging from 51% (Black or African American patients) to 82.8% (White patients). For the readmission model, insurance status showed the highest within-class burden of bias, with false negative rates ranging from 38.2% (Medicare) to 79.7% (Self-Pay). The system tested threshold adjustment, a post-processing method that sets subgroup-specific classification thresholds to minimize fairness metric differences, and successfully reduced absolute EOD to less than 5 percentage points for all targeted subgroups while maintaining accuracy losses below 10% and alert rate changes below 20%. The methodology, including open-source R code, was published with a supplementary playbook explicitly designed for other low-resource systems to replicate. This case demonstrates that bias identification and mitigation can be operationalized within resource-constrained settings, though the approach addresses post-hoc fairness correction rather than eliminating bias sources in training data or model architecture.

The European Space Agency's Urban Atlas processes 785-800 Functional Urban Areas across Europe using deep learning on Sentinel satellite imagery, achieving land use classification at continental scale (E.S.A./Copernicus, 2021). Unlike many urban AI applications, the Urban Atlas operates under Copernicus open data policy (Regulation EU 1159/2013): all imagery, processing methods, and classification outputs are freely accessible, enabling verification, replication, and improvement proposals from any researcher or government. While deep learning classification involves technically opaque neural networks, the outputs are verifiable land cover categories that users can examine against source imagery, providing what might be termed "outcome transparency" even where "process transparency" of model internals is limited. Recent technical scholarship has examined relationships between model interpretability and predictive performance across specific application domains. Rudin(2019a) argues in *Nature Machine Intelligence* that the supposed accuracy- interpretability tradeoff is frequently overstated, providing evidence from domains including criminal justice risk assessment, medical diagnosis, and credit scoring where interpretable models matched black-box performance. In separate work, Rudin(Rudin et al., 2011) and colleagues' analysis of

NYC's electrical grid failure prediction, developed in collaboration with Con Edison over 2007-2012, demonstrated that machine learning models with interpretable features enabled utility engineers to identify failure modes and prioritize maintenance, with the transparent use of data allowing domain experts to troubleshoot and extend the models. Wagner et al.(2022) demonstrated that gradient boosting decision tree models with SHAP values could capture up to 84% of variance in Berlin's trip emissions analysis while revealing policy-relevant threshold effects, such as the sharp increase in vehicle kilometers traveled beyond 10km from city centers, that inform targeted urban planning interventions. Kim et al.(2023) showed that XGBoost- SHAP approaches in Seoul's urban expansion modeling achieved prediction accuracy exceeding 90% while providing actionable insights about land-cover characteristics driving growth patterns. These domain-specific findings suggest that transparency and performance need not be mutually exclusive in particular applications, though we note these represent specific contexts; whether the relationship generalizes across all urban AI deployment scenarios remains an empirical question requiring further investigation. The cases documented in this theme illustrate varied approaches to operationalizing transparency, from algorithm registers to bias mitigation playbooks to open data infrastructure, without claiming that any single approach constitutes a universal solution to algorithmic accountability challenges.

3.3 Discussion

The cross-case analysis reveals a recurring tension in urban AI implementations: technical performance metrics frequently diverge from policy outcomes, and this divergence appears structural rather than incidental. Across 157 documented cases, we observe patterns suggesting that the core challenge of urban AI governance is not primarily technical but epistemological, a systematic misalignment between what algorithms optimize for and what cities need them to achieve. This synthesis integrates findings across four themes to articulate principles for democratic urban AI governance, while acknowledging that these principles derive from our sample and require validation across broader contexts (Campbell, 1979; Muller, 2018; T. M. Porter, 2020; Strathern, 1997; R. L. Thomas & Uminsky, 2022).

The accuracy illusion documented across our cases reveals what we term the "metrics trap": algorithmic systems achieving high performance on technical proxies while failing their stated policy objectives. The 88 percentage point gap between ShotSpotter's acoustic accuracy and crime-fighting effectiveness, COMPAS's aggregate accuracy masking disparate error rates, and healthcare algorithms optimizing for cost while under-identifying patients for care demonstrate this common mechanism (Angwin et al., 2016; City of Chicago Office of Inspector General, 2021; Z. Obermeyer et al., 2019). This pattern extends Muller (2018)'s analysis of "the tyranny of metrics" into algorithmic governance, challenging techno-optimistic narratives (M. Batty, 2013; Townsend, 2013) while providing explanatory mechanisms absent from purely critical accounts (Greenfield, 2013; Morozov, 2013). Breaking free requires redefining success through

democratic deliberation, ensuring algorithms optimize for community-articulated goals rather than vendor-defined metrics (Cardullo & Kitchin, 2019; Fung, 2009; Smith, 2009).

The discriminatory feedback loops documented across domains reveal how algorithms can transform historical bias into "computational destiny." Chicago's Strategic Subject List exemplifies this: only 16.3% of listed individuals were confirmed gang members, yet district commanders believed 95%, demonstrating how algorithmic outputs acquire institutional authority independent of ground truth (City of Chicago Office of Inspector General, 2021; Ensign et al., 2018). Our cases extend Lum and Isaac (2016)'s analysis beyond predictive policing to housing, surveillance, and energy systems, suggesting feedback loops may represent a general property of algorithmic governance in unequal societies (Benjamin, 2019; Eubanks, 2018; Noble, 2018; Rosen et al., 2021). Interrupting these loops requires structural approaches: reforming data generation at source, incorporating affected community perspectives at decision points, and implementing regular audits with sunset provisions (Brayne, 2017; Citron & Pasquale, 2014; Collins & Evans, 2002).

The community resistance cases reveal both the power and limitations of democratic opposition. Toronto, Boston, Detroit, and LA demonstrate that organized communities can achieve significant policy changes, but these campaigns required enormous resources and years of organizing (A.C.L.U., 2024; A.C.L.U. Massachusetts, 2020; Brennan Center for Justice, 2024). The challenge lies in transforming episodic resistance into systematic governance, moving from what Hirschman (1970) termed "voice" through protest toward institutionalized voice embedded in governance structures. Boston's ban diffusing to five Massachusetts municipalities suggests policy frameworks can scale through networked advocacy (Dolowitz & Marsh, 2000; Keck & Sikkink, 2014; Walker, 1969), while Singapore's contrast case illustrates how political context shapes these dynamics (McCarthy & Zald, 1977; Tarrow, 2022).

The transparent alternatives examined suggest that transparency can be operationalized through varied institutional mechanisms. Three distinct levels prove necessary for democratic accountability: technical transparency requiring interpretable models whose logic can be examined, moving beyond what Ananny and Crawford (2018) term "seeing without knowing"; process transparency demanding visibility into development, procurement, and deployment (Pasquale, 2015); and impact transparency necessitating ongoing monitoring of outcomes disaggregated by affected populations (Selbst et al., 2019). NYC Health + Hospitals' documented bias mitigation demonstrates these principles can be operationalized even in resource-constrained settings (Mackin et al., 2025), while technical scholarship provides domain-specific evidence that interpretable models can achieve competitive performance (M. Kim et al., 2023; C. Rudin & J. Radin, 2019; Wagner et al., 2022).

Cities can institutionalize these insights through permanent structures combining bottom-up resistance with top-down policy frameworks. Barcelona's Advisory Council creates ongoing channels for community input (Barcelona City Council, 2021), operationalizing Fung and Wright (2003)'s "empowered participatory

governance." The EU's Model AI Contractual Clauses demonstrate how procurement power can mandate transparency requirements (EU Public Buyers Community, 2024), aligning with Ostrom (1990)'s polycentric governance theory. The interaction between grassroots organizing and formal governance observed in Detroit suggests neither approach alone achieves sustainable accountability (Levi-Faur, 2012). As Fung et al. (2007) distinguish, transparency-as-disclosure differs from transparency-as-accountability; the latter requires mechanisms ensuring disclosed information generates consequences (Amsterdam Municipality, 2020; Saidot, 2025).

Several boundary conditions shape interpretation. Language constraints restricted analysis primarily to English-language sources. Publication bias likely favors documented failures over quietly successful systems. Temporal constraints mean recent implementations lack comprehensive outcome data. Our critical orientation, while justified by documented harms, may undervalue genuine technical achievements; we view this stance as methodologically appropriate given power asymmetries while acknowledging its influence (Jasanoff, 2004; Pasquale, 2015). The cases examined represent specific political and institutional contexts, resistance succeeded predominantly where civil society infrastructure existed; transparent governance emerged in municipalities with prior participatory commitments, constraining direct generalization.

The patterns documented converge on a central finding: black-box optimization evaluated by technical metrics divorced from democratic accountability produces predictable harms across domains. Yet the transparent alternatives demonstrate this outcome is not inevitable. The evidence suggests cities need not accept supposed tradeoffs between technical capability and democratic governance. The path forward requires reconceptualizing algorithms not as neutral tools but as political systems requiring democratic oversight, extending Winner (2017)'s insight that artifacts have politics into the algorithmic domain. Just as cities govern infrastructure and services through democratic processes, algorithmic systems must become subject to public accountability, community input, and ongoing assessment. The cases examined suggest this governance is not only necessary but achievable, though specific mechanisms require continued investigation across diverse contexts.

3.4 Methods

This study employs a multiple case study design integrated with thematic synthesis to examine the deployment and impacts of artificial intelligence systems across urban domains. This methodological approach was selected for its capacity to provide rich, contextual understanding of complex sociotechnical phenomena while enabling systematic cross-case analysis of patterns and themes (Stake, 2006; R. K. Yin, 2018).

3.4.1 Research Design Rationale

The multiple case study methodology is particularly suited to examining urban AI implementations for several reasons. First, AI deployments in cities represent bounded systems where technology, policy, and social outcomes intersect in specific contexts (Flyvbjerg, 2006). Second, the approach accommodates multiple data sources and types of evidence, reflecting the reality that urban AI systems are documented across academic, governmental, commercial, and journalistic sources (Bowen, 2009). Third, case study methodology explicitly supports critical examination of power relations and social justice concerns central to this investigation (Bartlett & Vavrus, 2016).

The integration of thematic synthesis enhances the analytical power of the case study approach by enabling systematic identification of patterns across cases while preserving contextual richness (J. Thomas & Harden, 2008). This combined methodology allows us to move beyond individual implementation stories to identify recurring mechanisms of success, failure, and resistance in urban AI deployment. The approach aligns with recent calls for methodological innovation in studying algorithmic systems, which require techniques capable of capturing both technical specifications and lived experiences of affected communities (Katzenbach & Ulbricht, 2019).

3.4.2 Researcher Positionality

As researchers trained in urban studies, data science, and human geography, we acknowledge our positionality shapes this investigation. We have technical expertise in machine learning systems combined with critical social science perspectives on urban inequality and digital justice. This interdisciplinary background provides sensitivity to both technical claims and social impacts, though we recognize it also predisposes us toward skepticism of technosolutionist narratives. We have previously engaged in community organizing around algorithmic accountability, which informs our commitment to centering affected communities' experiences. Throughout the research process, we maintained reflexive journals documenting how our perspectives evolved through engagement with the data, particularly when encountering cases that challenged our initial assumptions about the inevitability of algorithmic harm.

3.4.3 Case Universe and Sampling Strategy

Our case universe comprises 157 documented AI implementations across six urban domains spanning 27 countries from January 2015 to May 2024. This comprehensive mapping emerged from systematic searches of academic databases including Web of Science, Scopus, IEEE Xplore, ScienceDirect, and Google Scholar, supplemented by government repositories and verified news sources. The temporal boundaries capture the period following the deep learning revolution while providing sufficient time for implementation outcomes to manifest. We acknowledge this dataset reflects documented and accessible

cases rather than a comprehensive global census; the concentration of cases in English-speaking and European contexts likely reflects language and documentation accessibility limitations in our search strategy rather than actual global deployment patterns.

From this universe, we selected 23 cases for in-depth thematic analysis through purposeful sampling (Patton, 2015). Selection criteria prioritized documentation richness, defined as cases with evidence from at least four different source types including technical specifications, implementation reports, and community responses. We sought maximum variation across geographies, ensuring representation from Global North and Global South contexts, as well as variation in implementation outcomes from documented successes to documented failures. Following (Flyvbjerg, 2006)'s critical case logic, we deliberately included extreme and paradigmatic cases that illuminate the boundaries of urban AI's possibilities and failures. Temporal coverage spanned early implementations from 2015-2018 through recent deployments in 2022-2024, enabling analysis of evolutionary patterns in the field.

Cases were included if they met all the following criteria:

Algorithmic Component: The system must use machine learning, statistical models, or rule-based algorithms to process urban data and generate outputs affecting governance decisions or resident experiences. This includes:

- Predictive models (risk scoring, forecasting, classification)
- Optimization algorithms (routing, resource allocation, scheduling)
- Pattern recognition systems (image analysis, anomaly detection)
- Automated decision support systems
- Infrastructural platforms enabling multiple algorithmic applications

Urban Governance Context: The system must be deployed (or seriously proposed for deployment) by municipal authorities, affect urban populations, or operate within city boundaries. Private systems (e.g., ride-sharing algorithms) were excluded unless directly integrated into public governance.

Documented Implementation: Sufficient documentation must exist across ≥ 4 source types to enable rigorous analysis. Sources must include technical specifications plus evidence of social impacts.

Temporal Relevance: Implementation or serious deployment proposal must have occurred between January 2015 - May 2024 to capture post-deep-learning-revolution developments.

Inclusion of Resistance Cases: We deliberately included cases where proposed AI systems were successfully blocked or reformed through community resistance (Toronto Sidewalk Labs, Boston facial recognition ban, Detroit facial recognition reform, LA LASER termination, New Orleans Palantir termination). These cases are essential rather than peripheral because they:

- Reveal community capacity to identify algorithmic threats before or during implementation,

- Document effective resistance strategies applicable to future proposals,
- Show how democratic governance can prevent or reform harmful AI,
- Provide evidence against technological determinism.

Excluding resistance cases would systematically bias analysis toward systems that successfully overcame opposition, obscuring the democratic agency that shapes urban AI landscapes. Cases were excluded if they:

- Lacked sufficient documentation for rigorous analysis,
- Involved purely private-sector algorithms without public governance implications,
- Occurred primarily outside urban contexts (e.g., agricultural AI, wilderness monitoring),
- Were announced but never reached serious implementation stage (vaporware).

The final sample includes cases from multiple regions, with deliberate attention to implementations affecting marginalized communities whose experiences are often underdocumented in technical literature. This sampling strategy balances depth of analysis with breadth of coverage, exceeding (Eisenhardt, 1989)'s recommended range of 4-10 cases for theory building while remaining manageable for detailed qualitative analysis.

3.4.4 Data Collection and Management

Following Yin (2018)'s principle of multiple evidence sources, we developed a comprehensive data collection protocol. Data sources encompassed academic publications including peer-reviewed articles, conference proceedings, and technical reports; government documents such as official reports, audits, policy documents, and regulatory findings; technical documentation from system specifications and vendor materials; news media, particularly investigative journalism documenting system failures and community impacts; legal documents including court filings, settlements, and regulatory enforcement actions; and community sources such as advocacy reports, public testimony, and documented resistance efforts.

This diverse source strategy reflects Stake (2006)'s argument that case studies should capture multiple realities and perspectives. The inclusion of non-academic sources proves essential for documenting real-world impacts often absent from technical publications, particularly regarding marginalized communities' experiences (Gillborn, Warmington, & Demack, 2018). For each case, we constructed a digital repository using Mendeley for academic sources and a structured filing system for grey literature, with all documents tagged by case, source type, and temporal period.

Data collection proceeded iteratively, beginning with technical documentation to understand system specifications, followed by government and academic sources for implementation details, and culminating with community sources and investigative journalism to capture impacts and resistance. This sequencing allowed us to identify gaps between technical claims and lived experiences.

3.4.5 Analytical Framework

Our analysis proceeded through three integrated phases, each building upon the previous to develop comprehensive understanding of urban AI implementations. The analytical process combined deductive frameworks from critical algorithm studies with inductive theme development from empirical data.

Phase 1: Individual Case Analysis

Each case underwent systematic analysis using NVivo 12 software. We developed an initial coding framework based on sociotechnical systems theory (Bijker et al., 2012) and critical algorithm studies (Benjamin, 2019; Noble, 2018), encompassing technical architecture, stakeholder networks, implementation processes, documented outcomes, and resistance patterns. Two researchers independently coded five initial cases, achieving inter-rater reliability of 0.84 (Cohen's kappa) before refining the codebook. The refined framework was then applied to all cases, with regular team meetings to discuss emerging patterns and resolve interpretive differences.

Phase 2: Thematic Synthesis

Following Thomas and Harden (2008)'s approach, we conducted line-by-line coding of case findings to develop descriptive themes staying close to primary data. Through constant comparison, these descriptive themes evolved into analytical themes capturing deeper patterns. We employed Braun & Clarke (2019)'s reflexive thematic analysis principles, recognizing themes as analytical outputs rather than discovered entities. The synthesis generated 47 descriptive codes, consolidated into 12 analytical themes, and ultimately organized into four theoretical constructs: the accuracy illusion (Theme 1), discriminatory feedback loops (Theme 2), community resistance (Theme 3), and transparent alternatives (Theme 4).

Phase 3: Cross-Case Pattern Analysis

Systematic cross-case analysis employed multiple techniques including pattern matching (R. K. Yin, 2018), explanation building (Miles, Huberman, & Saldaña, 2020), and qualitative comparative analysis principles (Ragin, 2014). We constructed detailed case matrices documenting implementation characteristics, outcomes, and contextual factors. Pattern identification focused on mechanisms driving the "metrics trap" phenomenon, conditions enabling transitions to transparent systems, community resistance patterns and outcomes, and relationships between technical architectures and social impacts.

3.4.6 Quality and Rigor

We implemented multiple strategies to ensure analytical rigor and trustworthiness. Triangulation operated at multiple levels, comparing findings across data sources within cases, across cases within domains, and across domains for meta-patterns (Denzin & Lincoln, 2017). We maintained comprehensive audit trails in NVivo documenting all coding decisions, theme development, and analytical memos totaling over 200 pages. This documentation enables traceability from raw data to theoretical claims.

Negative case analysis played a crucial role, with deliberate attention to implementations that defied emerging patterns. For instance, we extensively analyzed cases where black-box systems appeared to benefit communities and where transparent systems failed to prevent harm. These analyses refined our theoretical claims and identified boundary conditions. Member checking occurred through two mechanisms: sharing case summaries with documented implementation stakeholders (achieving responses from 31 of 67 contacted) and presenting preliminary findings at three practitioner conferences for feedback. Peer debriefing involved monthly meetings with an external advisory board of scholars in critical data studies, urban governance, and AI ethics who reviewed our analytical process and challenged interpretations.

3.4.7 Thematic Case Distribution

The 23 in-depth cases were organized across four analytical themes, with some cases appearing in multiple themes where their characteristics warranted examination from different analytical perspectives:

Theme 1 - The Accuracy Illusion (6 in-depth cases): ShotSpotter Chicago and NYC exemplify divergence between acoustic detection accuracy (97%) and policy effectiveness (9.1% of alerts leading to gun crime evidence). COMPAS demonstrates how aggregate accuracy masks racially disparate error rates. The Obermeyer healthcare algorithm illustrates optimization for cost prediction while under-identifying Black patients for care. Delhi's air quality prediction and Boston University's hospitalization model show technical accuracy that may not capture differential impacts across populations.

Theme 2 - Discriminatory Feedback Loops (7 in-depth cases): Chicago's Strategic Subject List (2012-2019) exemplifies algorithmic bias in predictive policing, with documented gaps between algorithmic classification and verified status. London's Gangs Matrix demonstrates regulatory intervention following documented racial disproportionality. New Orleans' Palantir partnership (2012-2018) reveals opacity in surveillance collaborations. Rotterdam's Sensing Project illustrates automated ethnic profiling. Tenant screening systems (RealPage, CoreLogic, SafeRent) demonstrate feedback dynamics in housing. VI-SPDAT homelessness assessment shows documented racial disparities in prioritization. Energy systems (California fixed charges, EU smart meters) illustrate digital divide amplification.

Theme 3 - Community Resistance (6 in-depth cases including one contrast case): Toronto Sidewalk Labs cancellation represents successful coalition resistance against corporate smart city initiatives. Boston's facial recognition ban demonstrates preemptive policy action with regional diffusion. Detroit's facial recognition case shows transformation from documented harm to systemic reform. LA's LASER termination reveals how sustained community research and pressure achieved shutdown. New Orleans' Palantir termination demonstrates response to exposed covert operations. Singapore's traffic AI serves as

contrast case where no documented resistance occurred, illustrating contextual factors that may enable or constrain democratic pushback.

Theme 4 - Transparent Alternatives (6 in-depth cases): Amsterdam's algorithm register and Helsinki's AI register represent pioneering transparency infrastructure. Barcelona's AI governance framework including vCity and Decidim integration provides model for participatory governance. Seoul's waste management demonstrates IoT integration with worker feedback mechanisms. Austin Energy's demand response shows transparent customer-facing interfaces with documented participation rates. NYC Health + Hospitals' bias assessment and mitigation demonstrates operationalized algorithmic fairness in resource-constrained settings. EU Urban Atlas represents continental-scale open data infrastructure.

Several cases merit particular attention as critical or paradigmatic examples. Toronto's Sidewalk Labs cancellation demonstrates the power of sustained coalition resistance combining legal challenges, public advocacy, and expert resignations. Detroit's facial recognition case stands as paradigmatic for several reasons: it produced the first documented U.S. wrongful arrest from facial recognition, catalyzed unprecedented policy reform including mandatory corroborating evidence requirements, and achieved both individual remedy (\$300,000 settlement) and systemic change. NYC Health + Hospitals' bias mitigation case demonstrates that algorithmic fairness assessment can be operationalized even in safety-net systems with limited resources, with open-source methodology enabling replication by other institutions.

Several notable implementations were excluded due to insufficient documentation or language barriers. Chinese social credit systems, while influential, lacked sufficient English-language documentation for rigorous analysis. Some European implementations in non-English speaking countries were similarly excluded despite potentially innovative approaches. Recent deployments (late 2024) lacked sufficient outcome data for meaningful evaluation of social impacts beyond technical claims.

This selection of 23 cases provides coverage across domains, geographies, and outcomes while maintaining analytical depth. The cases collectively demonstrate the patterns of metrics divergence, feedback dynamics, community resistance strategies, and transparency mechanisms that characterize contemporary urban AI deployment within our documented sample.

3.5 Author Contributions

S.N.M.M. conceived and designed the study, developed the methodological framework, conducted the systematic case selection and data collection across 157 AI implementations, performed the qualitative analysis of 23 in-depth cases, created all figures and tables, wrote the manuscript, and approved the final version for submission. H.C. supervised and participated in the research design, assessed and performed partially in qualitative analysis of 23 in-depth cases, participated in revision and writing the final manuscript version and approved the final version for submission.

Chapter 4 Physics-Informed Neural Networks Reveal Interpretable Mechanisms of Multi-Ethnic Urban Settlement Dynamics²

Abstract

Urban demographic forecasting often faces a tradeoff between interpretable, mechanistic models and high-performing machine-learning approaches that offer only post-hoc explanations. Current computational approaches to population modeling achieve predictive success yet fail to reveal the underlying mechanisms of spatial mobility patterns, demographic vitality, and inter-ethnic interactions that govern settlement dynamics. Here, we introduce GraphPDE, a physics-informed neural architecture that embeds multi-ethnic reaction–diffusion dynamics directly within a graph neural network, enabling the joint learning of demographic parameters and spatiotemporal population evolution. Operating on neighbourhood-level spatial graphs across 52 Canadian cities, GraphPDE infers ethnicity- and city-specific diffusion coefficients, growth rates, carrying capacities, and interaction matrices through differentiable integration of partial differential equations. Trained on two decades of census data spanning nine ethnic groups, GraphPDE achieves competitive predictive performance while providing mechanistic interpretability unattainable with post-hoc attribution methods. The learned parameters uncover systematic regularities in urban settlement: mobility scales differ substantially across groups, nucleation thresholds vary by more than an order of magnitude, and growth rates exhibit marked temporal non-stationarity. By unifying physical demographic laws with flexible graph-based learning, GraphPDE offers a general framework for interpretable, physics-informed spatiotemporal modeling, advancing demographic forecasting beyond correlational prediction.

Data Availability

The datasets supporting the findings of this study are publicly available. Census demographic data for Canadian Dissemination Areas (2001-2021) were obtained from Statistics Canada and are available at <https://figshare.com/s/a9459ba93132b4ae6907>. The processed spatial graph structures, neighborhood adjacency matrices, and derived features used in this study are also available in the same repository. Source code for the GraphPDE model implementation, including custom CUDA kernels, training scripts, and analysis notebooks, is publicly available at <https://github.com/navid-nsk/graphPDE>. All data and code are provided under open licenses to ensure reproducibility of results.

Funding

² This article is currently under review in Communications AI & Computing (Nature Portfolio Journal). Authors: Seyed Navid Mashhadi Moghaddam, Huhua Cao, Michael Sawada

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The authors declare that no funding was received for the conduct of this study or preparation of this manuscript.

Competing interests

We declare that the authors have no competing interests as defined by Nature Portfolio, or other interests that might be perceived to influence the results and/or discussion reported in this paper.

4.1 Main

Computational approaches to urban demographic analysis have achieved substantial success in prediction tasks. Deep learning methods, particularly convolutional neural networks including U-Net and ResNet architectures, accurately estimate population density from satellite imagery and geospatial data(Alsabhan & Alotaiby, 2022; Georganos, Hafner, Kuffer, Linard, & Ban, 2022; Pan, Xu, Guo, Hu, & Wang, 2020). Graph neural networks leverage spatial relationships to predict building characteristics(Lei, Liu, Milojevic-Dupont, & Biljecki, 2024) and analyze street networks(Ma et al., 2024). These methods produce high-resolution population distribution maps(Y. Cao, Wang, Guo, Zhang, & Wu, 2025; Cui, Hu, Bao, & Li, 2024; Doda et al., 2024) and have been applied to informal settlement detection, urban growth simulation, and migration flow prediction(Gaskin & Abel, 2025; Huang et al., 2021; M. Luo & Chen, 2024).

Recent work has emphasized interpretability through post-hoc explanation methods. Gradient boosting models combined with SHAP analysis identify features correlating most strongly with population outcomes(Z. Li, 2024; M. Wang et al., 2023; J. Zhang et al., 2023; Z. Zhang, Yang, Zhao, & Li, 2025), revealing that income, housing density, or employment accessibility associate with population growth through feature importance rankings(Muratova, Mitrofanova, & Islam, 2022; S. Yi et al., 2025). However, post-hoc explanations do not represent the underlying processes generating demographic patterns. SHAP analysis reveals *correlations*, that high income associates with population growth, but not the *mechanisms* through which this occurs (immigration attraction, reduced emigration, differential birth rates, or housing supply effects). As Rudin argues, "explanations are often not reliable, and can be misleading" because explanation models "do not always attempt to mimic the calculations made by the original model"(Rudin, 2019b).

Interpretable mechanistic models offer an alternative by directly encoding demographic processes as mathematical relationships. Diffusion-based models describe spatial population redistribution where counties within cities grow proportionally through diffusive processes governed by migration flows(Reia, Rao, & Ukkusuri, 2022). Reaction-diffusion frameworks extend this by combining aggregation forces (positive externalities from urban amenities) with dispersal forces (negative externalities like congestion), capturing how spatial patterns emerge from local growth and spatial diffusion interplay(Raimbault, 2019;

Whiteley, Avitabile, Siebers, Robinson, & Owen, 2022). These models predict characteristic length scales between cities (approximately 45–53 km) through linear stability analysis(Whiteley et al., 2022), with interpretable parameters: diffusion coefficients measuring population mobility, reaction terms encoding birth-death dynamics and migration attractiveness, and city-specific decay rates capturing population redistribution(Reia et al., 2022).

Physics-informed neural networks (PINNs) represent a recent advance, embedding differential equations directly into neural network training through loss functions penalizing violations of physical laws(Ali, 2025; L. Yin & Lv, 2024). In population dynamics, PINNs simultaneously learn system states and estimate governing parameters from sparse observational data(Viet Cuong, Lalić, Petrić, Thanh Binh, & Roantree, 2024). For refugee population dynamics, PINNs captured nonlinear interactions between demographic and environmental factors more effectively than regression approaches(Kajuli & Mayengo, 2025), while recent work demonstrates PINNs' ability to incorporate policy-driven heterogeneity in age-structured models(Tao, 2025). However, PINNs face challenges with multi-scale behavior and stiff ODEs, requiring careful normalization and gradient balancing(Ali, 2025).

Despite these advances, existing interpretable models face limitations for ethnic settlement dynamics. Reaction-diffusion models typically treat aggregate population as a single variable, overlooking how different ethnic groups experience distinct migration patterns, amenity preferences, and clustering behaviors(Whiteley et al., 2022). Diffusion-based approaches assume symmetric migration flows that may not hold for ethnic communities shaped by social networks, cultural institutions, and discriminatory housing markets(Reia et al., 2022). Current PINN applications focus on univariate or aggregated populations rather than multi-group interactions(Ali, 2025; Tao, 2025).

We employ reaction-diffusion dynamics as our physics prior because empirical evidence demonstrates that urban settlement patterns emerge through self-organizing spatiotemporal processes consistent with diffusion-reaction mechanisms(Henn, Friesen, Hartig, & Pelz, 2020; Kovalevsky, Volchenkov, & Scheffran, 2021; Louf & Barthelemy, 2013; Raimbault, 2019; Torrens & O'Sullivan, 2001). Cities develop as dissipative structures where regular settlement patterns emerge spontaneously from local growth (reaction) and spatial redistribution (diffusion)(Strano et al., 2013). Recent work shows locations with high or low populations emerge intrinsically through Turing-type pattern formation, with behavior reminiscent of self-organizing systems where structure arises from collective behavior following simple local rules(Hartig, Friesen, & Pelz, 2020; Zincenko, Petrovskii, Volpert, & Banerjee, 2021). Cellular automata models, computational cousins to reaction-diffusion systems, successfully represent urban environments where each cell's state depends on reaction functions (local processes) and diffusion functions (neighbor interactions)(Michael Batty, Xie, & Sun, 1999; Torrens & O'Sullivan, 2001). The framework accurately reproduces even urban traffic congestion propagation(Bellocchi & Geroliminis, 2020).

Alternative frameworks exist but present limitations for multi-ethnic analysis. Agent-based models like Schelling's segregation model capture individual preferences but require specifying decision rules varying across groups and lack mathematical tractability for deriving interpretable aggregate parameters (Fosset, 1999; Hatna & Benenson, 2012; Schelling, 1971). Radiation and gravity models describe migration flows but do not generate the emergent spatial patterns (stripes, spots, labyrinths) observed in ethnic settlement morphology (Alis, Legara, & Monterola, 2021; Simini, González, Maritan, & Barabási, 2012). Percolation approaches identify connectivity thresholds but operate as static geometric models rather than dynamic evolution systems (Arcaute et al., 2016; Behnisch, Schorcht, Kriewald, & Rybski, 2019; W. Cao, Dong, Wu, & Liu, 2020). Reaction-diffusion frameworks naturally incorporate both local demographic processes (births, deaths, immigration, emigration via reaction terms) and spatial movements (via diffusion operators on graphs), enabling interpretable parameters quantifying ethnic-specific mobility, growth propensity, and inter-group competition while predicting realistic spatiotemporal patterns (Michael Batty et al., 1999; Steinhäuser & Chawla, 2010).

We develop a graph-based physics-informed neural network (GraphPDE) that embeds multi-ethnic reaction-diffusion dynamics while learning interpretable demographic parameters. The model solves

$$\frac{\partial u_{v,e}}{\partial t} = \alpha \cdot D_e(c, t) \nabla^2 u_e + R_e(\mathbf{u}, c, t)$$

for each ethnicity e on a spatial graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ representing 30,091 dissemination areas, where $\mathcal{V}$ are neighborhoods, $\mathcal{E}$ encodes spatial adjacency, and c indexes cities. The diffusion term employs a normalized graph Laplacian with learnable global weight α and period-specific, city-specific diffusion coefficients $D_e(c, t)$ learned separately for four census intervals (2001-2006, 2006-2011, 2011-2016, 2016-2021).

The reaction function decomposes as:

$$\begin{aligned} R_e = & r_e(c, t) \cdot \sigma(M_{\text{census}}(\mathbf{x}))_e \cdot u_e \left(1 - \frac{N}{K_e(c, t)}\right) \\ & + \mathcal{A}_e(\mathbf{u}, t) \\ & + F_e(c, t) \cdot \sigma(M_{\text{census}}(\mathbf{x}))_e \cdot \left(1 - \frac{u_e}{K_e(c, t)}\right) \\ & - \epsilon_e(c, t) \cdot \sigma(M_{\text{census}}(\mathbf{x}))_e \cdot u_e \\ & + \beta \cdot R_e^{\text{neural}}(\mathbf{u}, c, t) \end{aligned}$$

combining logistic growth (r_e, K_e), attention-based inter-ethnic competition ($\mathcal{A}_e$), immigration (F_e, K_e), and emigration (ϵ_e), with all base parameters learned per-period and per-city. The term $\sigma(M_{\text{census}}(\mathbf{x}))_e$ represents census feature modulation through a learned multilayer perceptron $M_{\text{census}}: \mathbb{R}^{74} \rightarrow \mathbb{R}^E$ with sigmoid activation $\sigma(\cdot)$, capturing how local socioeconomic conditions (income,

education, housing characteristics) affect demographic rates. This modulation mechanism enables the model to adapt baseline demographic parameters to neighborhood-specific contexts without requiring separate parameters for each dissemination area.

Neural augmentation via city embeddings ($\mathbf{h}_c \in \mathbb{R}^{512}$) and period embeddings ($\mathbf{h}_t \in \mathbb{R}^{64}$), combined with a residual MLP weighted by learnable parameter β , captures systematic deviations from the physics structure while maintaining interpretable PDE foundations. We integrate this system through 4th-order Runge-Kutta over 5-year intervals, training on 238,175 samples spanning nine ethnic groups across 52 major Canadian cities (Figure 4.1). The framework yields both accurate predictions and interpretable insights: learned diffusion coefficients quantify ethnic mobility patterns, growth rates reveal demographic vitality, carrying capacities show urban constraints, and the neural residual indicates where classical demographic theory requires extension.




Figure 4.1 GraphPDE Model Architecture.

4.2 Results

4.2.1 GraphPDE Achieves Superior Predictive Performance Through Reaction-Diffusion Dynamics

We evaluated GraphPDE against eight baseline models spanning traditional machine learning (XGBoost with SHAP(Chen & Guestrin, 2016)), standard neural architectures (MLP), graph-based methods (GCN(Kipf, 2016), GAT(Velickovic et al., 2017), GraphSAGE(Hamilton, Ying, & Leskovec, 2017 2017)), physics-informed approaches (DeepONet(Lu, Jin, Pang, Zhang, & Karniadakis, 2021)), and spatio-temporal architectures (T-GCN(Zhao et al., 2019), DCRNN(Y. Li, Yu, Shahabi, & Liu, 2017)). GraphPDE achieved an R^2 of 0.7649 with the lowest error metrics among neural approaches (MAE = 10.14, RMSE = 21.08, MAPE = 42.0%), representing substantial improvements over both pure black-box models and existing physics-informed methods (Fig. 4.2a–g). While XGBoost with SHAP achieved marginally higher R^2 (0.7760), this gradient boosting approach lacks the mechanistic interpretability central to our framework, it cannot recover the reaction-diffusion dynamics governing settlement evolution. The training dynamics reveal GraphPDE's stable convergence without overfitting, contrasting sharply with the divergent validation curves of DeepONet and the plateauing performance of graph-only architectures.

The critical architectural innovation lies in the hybrid reaction term design, which decomposes population dynamics into interpretable physics-based components and a learnable residual. Analysis of contribution magnitudes reveals that 57.2% of the reaction term derives from interpretable components (self-reaction α_k , pairwise interactions β_{kl} , and spatial attention mechanisms), while 42.8% captures unstructured residual patterns through the MLP (Fig. 2h–i). This decomposition achieves a principled balance: the physics-informed components encode theoretically meaningful processes, intrinsic growth, inter-ethnic interactions, and spatially-varying responses to urban amenities, while the neural residual accommodates phenomena not captured by the reaction-diffusion formalism, such as policy interventions or idiosyncratic local factors. The bimodal distribution of reaction term magnitudes (Fig. 4.2h) further indicates that interpretable and unstructured components capture distinct dynamical regimes.

Ablation experiments confirm the necessity of this hybrid design. Removing the residual MLP (GraphPDE No Residual) reduces R^2 from 0.7649 to 0.5588 and increases MAPE from 42.0% to 44.7%, demonstrating that pure physics-based dynamics, while interpretable, cannot fully capture the complexity of urban settlement processes. Conversely, removing physics constraints entirely (as in standard GNNs) sacrifices interpretability without commensurate predictive gains. The full GraphPDE architecture thus occupies a unique position in the accuracy-interpretability trade-off space, providing mechanistic insights into settlement dynamics while maintaining competitive predictive performance.




Figure 4.2 GraphPDE achieves superior predictive performance through physics-informed learning with interpretable demographic mechanisms.

4.2.2 Learned Demographic Parameters Reveal Distinct Ethnic Settlement Dynamics

GraphPDE's city-level parameters reveal substantial geographic variation across Canada's 52 major metropolitan areas (Fig. 4.3). Carrying capacity trajectories diverge markedly: Toronto and Vancouver maintain stable high capacities (1,200–1,300 people/DAUID)(Global Property Guide, 2025; Remax Canada, 2022), while Calgary exhibits boom-bust dynamics peaking at 1,400 before correction(Okkola & Brunelle, 2018; Stephenson, 2022), and Winnipeg shows 38% decline. City-specific diffusion coefficients range from 2.9 km²/year (Ottawa) to 6.8 km²/year (Rocky View County), with Western suburban municipalities exhibiting approximately double the spatial mobility of established urban cores. Brampton's high diffusion (0.14–0.16) aligns with its 25.3% population growth and 50,000 dwelling additions(Burchfield & Kramer, 2015; Federal Redistribution, 2022), while Montreal's low diffusion (0.08–0.10) reflects documented neighborhood stability(Boulianne & Guilbault, 2016). Notably, geographically proximate cities show divergent patterns, Toronto, Mississauga, and Brampton developed distinct ethnoburbs despite contiguity(Burchfield & Kramer, 2015; Ontario 360, 2023).

Growth rate distributions stratified by city size reveal scale-dependent demographic equilibria (Fig. 4.4c). Large cities display narrow distributions centered near zero (median -0.001 year^{-1}), consistent with urban scaling theory(L. M. Bettencourt, 2020). Small cities exhibit wider distributions ($\pm 0.03 \text{ year}^{-1}$), reflecting heterogeneous trajectories predicted by models emphasizing increased variance for complex phenomena(Gomez-Lievano, Patterson-Lomba, & Hausmann, 2016). This is empirically supported: peripheral municipalities grew faster (+6.9%) than central municipalities from 2016–2021(Statistics Canada, 2022b). Immigration attractiveness exhibits clear geographic clustering, with Western cities forming a high-attractiveness zone: Edmonton experienced the fastest-growing immigrant population(Bell, 2023), Calgary led interprovincial migration (+26,662)(Fletcher, 2024), while Montreal's share declined from 14.8% to 12.2%(McQuillan, 2024; Statistics Canada, 2022f).




Figure 4.3 Learned demographic parameters reveal distinct ethnic settlement dynamics and inter-ethnic interaction patterns.

Ethnicity-specific settlement flows reveal distinct geographic strategies (Fig. 4h). India-origin populations show the most dispersed pattern across Toronto, Brampton (52.4% South Asian)(Aldunate & Hameed, 2022), and Surrey. China-origin populations display a bi-coastal strategy with nearly equal flows to Toronto

(679,725) and Vancouver (512,260)(US Canada Info, 2022a), plus concentration in Markham (47.9%)(Statistics Canada, 2022d). Iran-origin populations exhibit the highest geographic concentration, with Toronto receiving nearly twice Vancouver's share(Statistics Canada, 2024b; US Canada Info, 2022b). Pakistani populations concentrate almost exclusively in the Greater Toronto Area(Statistics Canada, 2024a). Multi-ethnic diversity (Shannon entropy) ranges from Toronto (1.95) to Cambridge and Longueuil (<0.80), revealing a gradient corresponding to city size and gateway function(Statistics Canada, 2022e). Hierarchical clustering reveals an interpretable urban typology (Fig. 4g). The four largest metros, Toronto, Vancouver, Calgary, Edmonton, cluster most tightly (Ward distance <2), reflecting similar equilibrium dynamics. A distinct cluster emerges for high-growth suburban cities (Brampton, Surrey, Markham, Vaughan, Richmond Hill), while resource-dependent cities and Atlantic/Quebec centers form separate branches. Temporal parameter stability exhibits clear size dependence: large cities (>2,000 DAUIDs) show low instability (CV <1.5), while small cities display higher instability (CV 4–7). Halifax emerges as outlier (CV ~7) despite moderate size. These parameters provide physically interpretable measures, diffusion coefficients in km²/year directly quantify redistribution rates, capturing metropolitan differences that aggregate analyses would obscure.

4.2.3 City-Level Parameters Expose Geographic Heterogeneity in Settlement Processes

GraphPDE's city-level parameters reveal substantial geographic variation in demographic processes across Canada's 52 major metropolitan areas (Fig. 4). Carrying capacity trajectories diverge markedly: Toronto and Vancouver maintain stable high capacities (1,200–1,300 people/DAUID) reflecting sustained housing development(Global Property Guide, 2025; Remax Canada, 2022), while Calgary exhibits boom-bust dynamics peaking at 1,400 (2011–2016) before correction(Okkola & Brunelle, 2018; Stephenson, 2022), and Winnipeg shows 38% decline potentially indicating urban core stagnation. City-specific diffusion coefficients, quantifying population redistribution rates, range from 2.9 km²/year (Ottawa) to 6.8 km²/year (Rocky View County), with Western suburban municipalities exhibiting approximately double the spatial mobility of established urban cores. Brampton's high diffusion (0.14–0.16) aligns with its 25.3% population growth from 2011–2021, during which the city added 50,000 dwellings through greenfield development while experiencing net population loss in existing urban areas(Burchfield & Kramer, 2015; Federal Redistribution, 2022). Montreal's persistently low diffusion (0.08–0.10) reflects documented neighborhood stability where wealth distribution remained relatively unchanged from 2001–2011(Boulianne & Guilbault, 2016). Notably, geographically proximate cities often show divergent diffusion patterns, Toronto, Mississauga, and Brampton developed distinct ethnoburbs (Chinese-Canadian in Markham, Indo-Canadian in Brampton) despite contiguity(Burchfield & Kramer, 2015; Ontario 360, 2023).




Figure 4.4 City-level parameters expose geographic heterogeneity in multi-ethnic settlement processes across Canadian metropolitan areas.

Growth rate distributions stratified by city size reveal systematic scale-dependent demographic equilibria (Fig. 4c). Large cities display narrow distributions centered near zero (median -0.001 year^{-1}), suggesting demographic equilibrium consistent with urban scaling theory where growth rate variations preserve scale invariance (L. M. Bettencourt, 2020). Small cities exhibit the widest distributions extending to $\pm 0.03 \text{ year}^{-1}$, reflecting heterogeneous trajectories from decline to rapid growth, a pattern predicted by models emphasizing increased variance for phenomena requiring multiple complementary factors (Gomez-Lievano et al., 2016 2016). This heterogeneity is empirically supported: peripheral Canadian municipalities grew faster ($+6.9\%$) than central municipalities from 2016–2021, while several previously declining municipalities stabilized (Statistics Canada, 2022b). Immigration attractiveness during 2016–2021 exhibits clear geographic clustering, with Western cities forming a high-attractiveness zone: Edmonton experienced the fastest-growing immigrant population among major Canadian cities (Bell, 2023), Calgary led net

interprovincial migration (+26,662 residents)(Fletcher, 2024), while Montreal's share of new immigrants declined from 14.8% to 12.2%(McQuillan, 2024; Statistics Canada, 2022f).

Ethnicity-specific settlement flows reveal distinct geographic strategies (Fig. 4h). India-origin populations show the most dispersed pattern with major flows to Toronto, Brampton (52.4% South Asian population)(Aldunate & Hameed, 2022), and Surrey, demonstrating penetration across Eastern and Western Canada. China-origin populations display a bi-coastal strategy with nearly equal flows to Toronto (679,725 individuals) and Vancouver (512,260)(US Canada Info, 2022a), plus concentration in Markham (47.9% Chinese)(Statistics Canada, 2022d). Iran-origin populations exhibit the highest geographic concentration, with Toronto receiving nearly twice Vancouver's share(Statistics Canada, 2024b; US Canada Info, 2022b). These patterns reflect ethnicity-specific migration networks: South Asian groups heavily favor GTA suburbs, Chinese immigrants split between coasts leveraging transpacific connections, while Pakistani populations concentrate almost exclusively in the Greater Toronto Area(Statistics Canada, 2024a). Multi-ethnic diversity measured by Shannon entropy ranges from Toronto (1.95, highest) to Cambridge and Longueuil (<0.80), revealing a gradient corresponding to city size and gateway function(Statistics Canada, 2022e).

Hierarchical clustering of learned parameters reveals an interpretable urban typology (Fig. 4g). The four largest metros, Toronto, Vancouver, Calgary, Edmonton, cluster most tightly (Ward distance <2), reflecting similar high-capacity, high-immigration equilibrium dynamics. A distinct cluster emerges for high-growth suburban cities (Brampton, Surrey, Markham, Vaughan, Richmond Hill) sharing rapid expansion parameters, while resource-dependent cities and smaller Atlantic/Quebec centers form separate branches reflecting regional economic conditions. Temporal parameter stability exhibits clear size dependence: large cities (>2,000 DAUIDs) show low instability (CV <1.5), enabling stationary forecasting models, while small cities display higher instability (CV 4–7) requiring adaptive frameworks. Halifax emerges as a notable outlier with CV ~7 despite moderate size, potentially reflecting military base expansions or economic restructuring. These city-specific parameters provide physically interpretable measures, diffusion coefficients in km²/year directly quantify population redistribution rates, capturing fundamental differences in metropolitan spatial organization that aggregate ethnic-level analyses would obscure.

4.2.4 Multi-Scale Spatial Organization Emerges from Reaction-Diffusion Dynamics

To demonstrate GraphPDE's mechanistic interpretability beyond aggregate statistics, we examine Chinese settlement dynamics in Toronto as a case study revealing how learned parameters generate realistic spatial patterns from neighborhood to metropolitan scales (Fig. 4.5). The model's 20-year forward simulation predicts concentrated settlement in north-central Toronto encompassing 3,700 dissemination areas with 164,155 total individuals, showing strong concentration in Markham and Richmond Hill rather than

downtown dispersal, closely matching empirical census patterns without explicit supervision on settlement locations. The spatial extent spans approximately 40 km east-west, aligning with the learned characteristic spatial scale (38.7 km from Fig. 4.3a), while the north-south asymmetry reflects both physical constraints (Lake Ontario) and learned spatial dependencies. Zooming into a 7.4 km $\times$ 7.4 km region reveals emergent spot-and-stripe morphology with 3–4 km dominant wavelength corresponding to walkable neighborhood scales, approximating the theoretical prediction $\lambda \approx 2\pi\sqrt{D/|\alpha|}$ where D is the diffusion coefficient and α represents learned self-competition strength. This multi-kilometer patterning emerges spontaneously from local DA-to-DA interactions propagated through the graph structure, positioning Chinese settlement in the "spots" regime of the pattern formation phase space.




Figure 4.5 Multi-scale spatial organization emerges from reaction-diffusion dynamics in Chinese settlement patterns across Toronto.

Multi-scale cluster analysis quantifies hierarchical spatial organization across four orders of magnitude (Fig. 4.5c). The cluster size distribution exhibits strong right-skew with mean size of 492 grid points, revealing three distinct regimes: peripheral dispersal (<50 grid points) representing scattered households below the critical nucleation threshold ($N^* = 658$); neighborhood enclaves (50–500 grid points) forming the backbone of settlement structure with self-sustaining growth dynamics; and metropolitan hubs (>500 grid points) representing Markham and Richmond Hill cores where agglomeration economies and ethnic institutions create powerful attractive forces. The density-size relationship follows a weak power law (exponent $\alpha = -0.11$), indicating sub-linear scaling where larger clusters exhibit lower per-capita density, confirming multi-nodal network structure rather than monocentric concentration. Shape metrics (eccentricity = 0.644, compactness = 0.690) indicate elongated but semi-compact structures consistent with stripe-like patterns, with this hierarchical organization emerging naturally from learned carrying capacities and immigration rates without explicit multi-scale architectural components.

Temporal evolution over the 20-year simulation reveals three distinct phases (Fig. 4.5d). Phase 1 (0–2.5 years) exhibits rapid equilibration where spatial variance rises from ~1,750 to ~30,000, reflecting model adjustment from 2001 census initial conditions containing methodological artifacts from the introduction of dissemination areas as a new geographic unit (Statistics Canada, 2018) and the 1998 Toronto amalgamation (The Canadian Encyclopedia, 2007). Phase 2 (2.5–7.5 years) marks pattern emergence through spinodal decomposition, where variance declines sharply as clustering coefficient rises from 0.15 to 0.85, indicating rapid formation of spatially contiguous neighborhoods. Phase 3 (7.5–20 years) shows gradual maturation through Ostwald ripening, small enclaves dissolve while large enclaves expand, driven by thermodynamic favorability of lower interfacial energy. Concurrent with spatial reorganization, total population grows monotonically from ~100,000 to 164,155 (2.6% annually), closely matching the learned intrinsic growth ($1.9\% \text{ year}^{-1}$) plus immigration contributions ($0.8\% \text{ year}^{-1}$). The inverse relationship between variance and population in Phase 3 demonstrates that demographic growth drives spatial organization as reaction-diffusion instabilities amplify density fluctuations into persistent patterns.

4.2.5 Temporal Pattern Evolution Reveals Multi-Stage Settlement Crystallization

Beyond static spatial distributions, GraphPDE captures the dynamic process by which Chinese settlement patterns nucleate, grow, and mature over two decades (Fig. 4.6). At the macroscopic scale, settlement exhibits a "nucleation and growth" pattern with three major nucleation centers in northern districts seeding surrounding areas at 3.1% annual growth and 1.3 new clusters per year. Spatial autocorrelation (Moran's $I = 0.127$) confirms positive clustering where population in one dissemination area predicts elevated population in adjacent areas, while geographic coverage reaches 48.4% of Toronto's DAUIDs with average cluster size of 328 individuals. The three-dimensional spatiotemporal visualization (Fig. 4.6b) reconstructs

boundary evolution across census periods, revealing maximum lateral expansion of 6.9 km over 20 years (0.35 km/year radial growth). Four key features emerge: early settlement (2001) consists of isolated pockets reflecting sub-critical population; coalescence begins 2006–2011 as expanding enclaves merge into contiguous corridors; major consolidation occurs 2011–2016 when separate domains fuse into a unified spatial network; and mature boundary expansion (2016–2021) shows outward radial growth rather than new isolated nucleation, indicating transition from nucleation-dominated to growth-dominated phases.




Figure 4.6 Temporal pattern evolution reveals multi-stage settlement crystallization dynamics in Chinese Toronto communities.

The morphological evolution, from spots (2001) to stripes (2011) to labyrinthine networks (2021), recapitulates the phase space trajectory predicted by learned interaction parameters. Early sub-critical populations produce isolated spots; as population exceeds the nucleation threshold ($N^* = 658$ per DAUD), reaction-diffusion instabilities amplify perturbations into stripe patterns along high-diffusion corridors; further growth creates labyrinthine networks where stripe domains intersect. Nucleation kinetics exhibit non-monotonic trajectories (Fig. 4.6d): cluster count rises from ~250 (2000) to peak ~400 (2010) then declines to ~300 (2021), while average cluster size follows similar dynamics. This pattern reveals competing processes, nucleation dominates early periods while coarsening (cluster merging/dissolution) dominates later periods, analogous to Ostwald ripening in physical systems. Power-law growth analysis

yields exponent $n = 0.6$, indicating sub-linear anomalous growth suggesting saturation from carrying capacity constraints, competitive exclusion, or diminishing immigration returns. The phase diagram (Fig. 4.6f) positions the current neighborhood state in a "mixed" regime with trajectory toward "expansion," connecting local states to global pattern types.

Spatial correlation analysis reveals fractal organization with structure function decaying as $S(k) \propto k^{-2.7}$ across 1–10 km scales, implying fractal dimension $D_f \approx 1.36$ (Fig. 4.6e). This intermediate dimension, between 1D lines and 2D planes, indicates space-filling but not dense settlement: Chinese population forms a ramified network penetrating Toronto's geography without uniform coverage, analogous to diffusion-limited aggregation or percolation clusters. The steeper-than- k^{-2} decay suggests super-diffusive spreading where population disperses faster than random walk predictions, potentially reflecting directed migration along high-capacity corridors or autocatalytic growth where existing enclaves attract further immigration. Critically, this fractal structure emerges from learned reaction-diffusion dynamics without explicit fractal rules, the combination of positive diffusion (promoting spread) and negative self-competition (promoting aggregation) creates long-range correlations where density perturbations influence densities kilometers away. Scale invariance indicates no characteristic length dominates; clustering occurs similarly at 100 m (household blocks), 1 km (neighborhoods), and 10 km (metropolitan corridors), producing multi-scale self-similarity characteristic of non-equilibrium pattern formation.

4.2.6 Multi-Stability Analysis Reveals Configuration Transition Dynamics

GraphPDE's reaction-diffusion framework enables characterization of the complete settlement configuration landscape, revealing multiple stable patterns and the demographic forces governing transitions between them (Fig. 4.7). The spatial attractiveness map (Fig. 4.7a) visualizes favorability across Toronto's 3,700 dissemination areas computed from learned parameters, diffusion coefficients, growth rates, carrying capacities, and inter-ethnic competition matrices. High-attractiveness zones (blue, favorability ≈ -1.2) concentrate in North York's central corridor, downtown's eastern neighborhoods, and Scarborough's northern regions, while low-attractiveness zones (red, favorability $> +1.0$) appear in Etobicoke's western periphery where competition from established populations creates unfavorable conditions. The three-dimensional configuration landscape (Fig. 4.7b) extends this spatial view by characterizing patterns along Chinese concentration and neighborhood diversity (Shannon entropy) axes, with the z-axis representing configuration favorability. Four stable configurations emerge as local minima: C1 (Enclave, deepest valley representing maximum segregation), C2 (Semi-Segregated), C3 (Mixed), and C4 (Integrated). Critically, multi-stability explains observed diversity in settlement patterns across Canadian cities, Toronto's pattern differs from Vancouver's or Montreal's not because of different

demographic forces, but because each system settled into different stability basins based on historical contingencies.




Figure 4.7 Multi-stability analysis reveals configuration transition dynamics and demographic reorganization barriers in Chinese Toronto settlement.

Transitions between stable configurations require overcoming reorganization barriers quantifying demographic "inertia" (Fig. 4.7c). The transition pathway from C1 to C2 exhibits forward barrier $\Delta B_1 = 1.67$ and reverse barrier $\Delta B_2 = 1.13$, with asymmetry indicating hysteresis, transitions from C1 to C2 require stronger perturbations than reverse transitions. These medium-height barriers suggest configuration transitions are neither spontaneous nor impossible; sustained interventions over decadal timescales could facilitate reorganization, aligning with observed urban demographic transitions unfolding over 10–30 years. Critical threshold dynamics (Fig. 4.7d) reveal system-wide behavior across interaction intensity parameter θ , representing relative strength of dispersive versus clustering forces. The critical threshold $\theta_c = 0.20$ identifies maximum system responsiveness ($R_{max} = 2.31$) where small parameter changes induce disproportionately large configuration changes. Below θ_c (clustered regime), strong ethnic enclaves prevail with demographic inertia dominating; above θ_c (dispersed regime), integrated configurations become favorable. Current conditions place the system well below θ_c , explaining why observed enclave patterns are robust and incremental policy interventions produce limited reorganization.

Population flow dynamics track redistribution among stable regions over 25 years through a master equation framework (Fig. 4.7e). Region 1 (initially dominant) declines from population fraction 1.0 to 0.75 with equilibration timescale $\tau_1 = 1.8$ years, while Region 4 (emerging hub) rises from near 0 to 0.75 with $\tau_3 = 5.0$ years. The crossing point around year 5–7 marks critical transition where the emerging region surpasses the traditional concentration zone, indicating fundamental shift in settlement geography. Geographic mapping identifies 87 total stable regions representing high-attractiveness zones across 1,154 DAs with total population 35,065 (Fig. 4.7f). The top five regions by population range from 2,425 (Region 1, northeastern) to 1,620 (Region 5, north-central) individuals. This large number of distinct regions compared to few major population centers reflects complex multi-nodal structure, numerous small stable pockets coexist with larger dominant hubs, forming a hierarchical spatial network.

The configuration space network (Fig. 4.7g) synthesizes relationships among seven demographic patterns, with nodes representing configurations, edges weighted by reorganization barriers, and vertical layout encoding favorability. Pure China (hypothetical 100% Chinese, $F = 9.5$) occupies the least stable position, isolated by high barriers ($\Delta B = 3.0$) confirming extreme segregation is energetically unfavorable. Binary configurations (China+Pakistan, China+Philippines, China+India) form a tightly connected cluster with low mutual barriers (0.8–1.3), indicating easy transitions between similar compositions. The Current State ($F = 5.2$) connects to Balanced Mix ($F = 4.5$, most stable) by modest barrier ($\Delta B = 0.7$), suggesting integration requires moderate sustained effort rather than fundamental reorganization. The integration pathway shows smoothly decreasing favorability, indicating the integrated state is thermodynamically preferred. Toronto's Chinese settlement occupies a stable intermediate attractor connected by low-to-moderate barriers to integrated configurations, with extreme segregation isolated and energetically disfavored.

4.3 Discussion

Embedding reaction-diffusion dynamics directly into neural network architectures substantially mitigates the interpretability-accuracy trade-off in urban demographic modeling. GraphPDE achieves test $R^2 = 0.7649$, approaching gradient boosting performance (XGBoost: $R^2 = 0.7760$, 1.4% gap) while yielding mechanistically interpretable parameters that post-hoc explanation methods cannot provide. The physics-only variant achieves $R^2 = 0.5588$ using exclusively interpretable parameters, reaching 95.5% of pure neural baseline performance (GCN: $R^2 = 0.5854$) despite lacking flexible convolutional operators. Neural augmentation provides 37% improvement over physics-only, enabling competitive accuracy while maintaining mechanistic foundations. This challenges the prevailing assumption that interpretability requires substantial accuracy sacrifice, or that competitive prediction necessitates black-box models explained only through correlational techniques like SHAP (Rudin, 2019b).

The learned diffusion coefficients directly contradict the assumption in classical urban models that migration flows are ethnicity-invariant. Whiteley et al. predict universal characteristic length scales of 45–53 km through linear stability analysis of population-service interactions(Whiteley et al., 2022), yet we find marked heterogeneity: United Kingdom-origin populations exhibit 63.0 km spatial scales while Chinese and Philippines-origin groups show 38.7 km and 34.5 km respectively. This ethnic variation suggests that social networks, cultural institutions, and group-specific amenity preferences fundamentally alter spatial mobility patterns. Reia et al. observe ~86% of domestic migration occurs within 100 km radius but assume symmetric flows that may not hold for ethnic communities shaped by discriminatory housing markets(Reia et al., 2022). Urban planning frameworks assuming uniform population diffusion will systematically mispredict ethnic settlement trajectories.

The attention-based inter-ethnic competition mechanism reveals interaction patterns fundamentally incompatible with fixed Lotka-Volterra matrices. Classical multi-species reaction-diffusion systems assume constant competition coefficients, but we find interaction strengths vary spatially and temporally based on local population distributions. China exhibits strong negative self-competition (-19.8) driving spatial dispersal while showing positive facilitative relationships with six other groups, patterns matching documented characteristics of Toronto's ethnic enclaves where Chinese populations locate in areas exhibiting complex ethnocultural diversity. The positive cross-interactions (Philippines→China: $+0.62$, Pakistan→Philippines: $+0.54$) correspond to documented economic complementarity where Filipinos concentrate in healthcare while Chinese dominate professional services(Cho, 2017; Eckstein & Peri, 2018). This context-dependence means policy interventions produce different outcomes depending on existing ethnic composition.

The nucleation threshold differences, India requiring 8,132 individuals versus China forming enclaves at 658, challenge predictions from percolation theory assuming universal critical densities. Groups with strong diaspora networks (India) can maintain cohesion at lower spatial densities through transnational ties(Ashutosh, 2014; Garha & Paparusso, 2018), delaying geographic concentration, while groups relying on spatial proximity for social cohesion (China) form enclaves earlier(Chan, 2012; Hiebert, 2015). The temporal evolution of growth rates, Philippines maintaining $0.015\text{--}0.023\text{ year}^{-1}$ while United Kingdom transitions from $+0.010$ to -0.023 year^{-1} , demonstrates non-stationary demographic forces contradicting standard migration modeling practice(Gaskin & Abel, 2025; M. Luo & Chen, 2024). By learning separate parameters for four census intervals, GraphPDE captures policy shifts including post-2008 immigration reforms and economic cycles, signaling when demographic regimes shift.

GraphPDE's convergence properties, reaching near-asymptotic performance by epoch 100 versus 200+ for baselines, reveal that physics constraints function as powerful inductive biases. The reaction-diffusion structure eliminates parameter space regions that violate conservation laws, produce negative densities, or

generate discontinuous spatial jumps, providing effective regularization particularly valuable when training samples are limited to four temporal observations per location. Superior error metrics (MAE: 10.4, 28% improvement over MLP; RMSE: 25.0, 35% improvement over GCN) confirm that embedding demographic constraints improves prediction quality for outlier cases where pure data-driven approaches produce extreme mispredictions.

The city-level heterogeneity, Toronto maintaining carrying capacity while Winnipeg declines 38%, western cities attracting immigration while eastern cities stagnate, demonstrates that national demographic trends mask critical local variation. This geographic heterogeneity matters for policy: federal immigration targets optimized for national averages may induce capacity crises in high-attraction cities while failing to stabilize declining cities(Statistics Canada, 2022a, 2022b). GraphPDE's city-specific parameters enable spatially differentiated policy recommendations that national-scale models cannot provide. The learned parameters achieve substantial validation through systematic correspondence with independent evidence: spatial scales match theoretical predictions(Whiteley et al., 2022), nucleation thresholds align with diaspora research(Ashutosh, 2014; Sutama Ghosh, 2007), and city-specific patterns correspond to documented housing dynamics(Boulianne & Guilbault, 2016; Burchfield & Kramer, 2015).

Despite these advances, the model faces challenges capturing discrete policy shocks violating smooth differential evolution. Immigration point system reforms or refugee resettlement can induce step-function changes; however, period-specific parameters demonstrate capacity to adapt to regime changes, with non-monotonic growth patterns corresponding to documented policy shifts including the 2008 Canadian Experience Class, 2012 Federal Skilled Worker revisions, and 2015 Express Entry(Hou & Picot, 2022; Migration Policy Institute, 2009). The neural residual provides flexibility for systematic deviations from idealized dynamics, though the model requires retraining for unprecedented events. Hybrid discrete-continuous frameworks modeling policy interventions as exogenous forcing functions could enhance generalization, benefiting from annual administrative data beyond five-year census intervals.

The learned attention weights align systematically with independent sociological evidence, migration networks, economic niche differentiation, historical settlement patterns, providing confidence that weights capture meaningful demographic relationships rather than spurious correlations. However, these validations remain primarily indirect, demonstrating consistency rather than direct parameter measurements. The fundamental challenge is that direct measurements of ethnic-specific spatial mobility rates and quantitative inter-group competition strengths largely do not exist at the spatially-explicit, multi-ethnic scale we address. Future validation could proceed through mobility data comparison using cell phone location data, survey validation of inter-ethnic attitudes, longitudinal cohort tracking, and cross-city transfer testing whether Toronto parameters predict Vancouver or Montreal patterns.

To address critical uncertainty quantification needs, we employ Monte Carlo dropout ($N = 100$) with temperature-scaled calibration achieving 95.1% confidence interval coverage across the test set (mean uncertainty: 17.7 individuals/DA). Uncertainty varies systematically: higher for rapidly growing communities (India: $\sigma = 27.3$, Philippines: $\sigma = 24.0$) versus stable populations (Italy: $\sigma = 9.0$, UK: $\sigma = 9.1$). For carrying capacity interventions through affordable housing, bootstrap resampling yields narrow confidence intervals (8–13% emigration reduction), supporting confident implementation. Immigration effects show wider bounds (2–7% increase), indicating attraction policies require pilot studies rather than deterministic planning. Substantial uncertainty for emerging settlements ($\sigma = 14$ –16 for populations < 50) indicates forecasts should be interpreted as probabilistic scenarios guiding precautionary infrastructure provisioning.

GraphPDE demonstrates that interpretability and predictive accuracy are complementary when domain knowledge exists as governing equations. By embedding demographic conservation laws, spatial diffusion, and growth constraints directly into neural architectures, we achieve both mechanistic understanding and competitive prediction. This challenges the dichotomy between interpretable-but-simple and accurate-but-opaque models, suggesting hybrid physics-learning approaches offer a path forward for scientific applications where understanding mechanisms matters as much as predicting outcomes. For urban demography, GraphPDE enables transition from correlational pattern recognition to causal demographic theory, providing policymakers with tools to reason about interventions rather than merely forecasting business-as-usual trajectories.

4.4 Methods

4.4.1 Problem Formulation and Motivation

4.4.1.1 Spatiotemporal Forecasting on Graphs

We formulate ethnic population dynamics as a spatiotemporal prediction problem on a graph-structured domain. Given a spatial network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ representing geographic regions (nodes $v \in \mathcal{V}$) and their spatial adjacencies (edges $(v, w) \in \mathcal{E}$), we aim to predict population distributions across $E = 9$ ethnic groups over time.

Definition 1.1 (State Space). For each node $v \in \mathcal{V}$ (Dissemination Area Unit, DAUID), ethnicity $e \in \{1, \dots, E\}$, and time t , we define:

- **Population state:** $u_{v,e}(t) \in \mathbb{R}_+$ representing the count of ethnicity e at location v and time t
- **Census features:** $\mathbf{x}_v \in \mathbb{R}^{74}$ capturing demographic, socioeconomic, and housing characteristics
- **Full graph state:** $\mathbf{U}(t) \in \mathbb{R}_+^{|\mathcal{V}| \times E}$ with entries $[\mathbf{U}(t)]_{v,e} = u_{v,e}(t)$

Problem 1.1 (Census-to-Census Forecasting). Learn a mapping $f: \mathbb{R}_+^{|\mathcal{V}| \times E} \times \mathbb{R}^{|\mathcal{V}| \times 74} \rightarrow \mathbb{R}_+^{|\mathcal{V}| \times E}$ that predicts:

$$\mathbf{U}(t + \Delta t) = f(\mathbf{U}(t), \mathbf{X}; \theta)$$

where $\Delta t = 5$ years (inter-census period) and θ denotes learnable parameters.

4.4.1.2 Physics-Informed Learning Objective

Unlike pure data-driven approaches, we embed physical constraints through a structured loss function:

$$\mathcal{L}(\theta) = \underbrace{\frac{1}{N} \sum_{i=1}^N \|f(\mathbf{U}(t_i); \theta) - \mathbf{U}_{\text{obs}}(t_i + \Delta t)\|_2^2}_{\text{Data fidelity}} + \underbrace{\lambda_{\text{reg}} \|\theta_{\text{physics}}\|_2^2}_{\text{Parameter regularization}}$$

The parameter space decomposes as $\theta = \theta_{\text{physics}} \cup \theta_{\text{neural}}$, where:

- θ_{physics} : Interpretable demographic parameters (diffusion coefficients D_e , growth rates r_e , immigration rates F_e , emigration rates ϵ_e , carrying capacities K_c)
- θ_{neural} : Neural network weights for attention mechanisms, embeddings, and residual terms

Remark 1.1 (Hybrid Approach). This formulation bridges classical reaction-diffusion PDEs with modern deep learning, enabling both interpretability and expressiveness (Raissi et al., 2019).

4.4.2 Dataset Construction and Preprocessing

4.4.2.1 Data Sources and Geographic Coverage

Dataset Specification. Canadian Census data from Statistics Canada spanning four periods: 2001, 2006, 2011, 2016, with predictions targeting 2021. The analysis encompasses 52 major Canadian cities with population $> 100,000$, tracking $E = 9$ ethnic groups across $|\mathcal{V}| = 30,091$ DAUIDs.

Ethnic Groups: China, Philippines, India, Pakistan, Iran, Sri Lanka, Portugal, Italy, United Kingdom.

Cities (n=52): Toronto, Vancouver, Montreal, Calgary, Edmonton, Ottawa, Mississauga, Winnipeg, Quebec, Hamilton, Brampton, Surrey, Laval, Halifax, London, Markham, Vaughan, Gatineau, Longueuil, Burnaby, Saskatoon, Regina, Richmond, Richmond Hill, Oakville, Burlington, Greater Sudbury, Oshawa, Barrie, Abbotsford, Coquitlam, St. Catharines, Kelowna, Cambridge, Kingston, Guelph, Kitchener, Waterloo, Windsor, Ajax, Pickering, Whitby, Milton, Nanaimo, Brossard, Delta, Langley, Saanich, North Vancouver, Chilliwack, Thunder Bay, Rocky View County.

4.4.2.2 Data Partitioning Strategy

Spatial Train-Val-Test Split. To evaluate generalization to unseen geographic regions:

- **Training:** 238,175 (node, ethnicity, time) tuples (67% of nodes)
- **Validation:** 40,854 tuples (12% of nodes)
- **Test:** 76,194 tuples (21% of nodes)

Definition 2.1 (Temporal Consistency). All four census periods $\mathcal{T} = \{2001, 2006, 2011, 2016\}$ appear in each split. This ensures the model learns temporal dynamics while preventing spatial data leakage.

4.4.2.3 Graph Construction

Definition 2.2 (Spatial Adjacency Graph). The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ encodes geographic proximity:

$$(v, w) \in \mathcal{E} \iff \text{DAUIDs } v \text{ and } w \text{ share a boundary}$$

The adjacency matrix $\mathbf{A} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$ has entries $A_{vw} = \mathbb{1}[(v, w) \in \mathcal{E}]$. The degree matrix $\mathbf{D} = \text{diag}(\mathbf{A}\mathbf{1})$ satisfies $D_{vv} = |\mathcal{N}(v)|$.

Definition 2.3 (Normalized Graph Laplacian). We employ symmetric normalization:

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$$

Lemma 2.1 (Spectral Properties). The Laplacian $\mathbf{L}$ satisfies:

1. $\mathbf{L}$ is symmetric and positive semi-definite
2. Eigenvalues satisfy $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{|\mathcal{V}|} \leq 2$
3. The smallest eigenvalue $\lambda_1 = 0$ corresponds to constant functions

Proof. Follows from the quadratic form $\mathbf{x}^\top \mathbf{L} \mathbf{x} = \sum_{(i,j) \in \mathcal{E}} (x_i / \sqrt{d_i} - x_j / \sqrt{d_j})^2 \geq 0$ (Chung, 1997).

4.4.2.4 Feature Engineering

Census Features ($\mathbf{x}_v \in \mathbb{R}^{74}$):

1. **Age distribution** (10 bins): [0-9], [10-19], ..., [90+]
2. **Education** (7 levels): None, High school, Apprenticeship, College, Bachelor's, Master's, Doctorate
3. **Income** (3 quantiles): Median, 25th percentile, 75th percentile
4. **Housing:** Ownership rate, dwelling types
5. **Employment:** Labor force participation, unemployment rate
6. **Household:** Average size, family structure
7. **Language/Immigration:** Home language, immigrant proportion

Normalization. Census features undergo z-score standardization:

$$\tilde{x}_{v,j} = \frac{x_{v,j} - \mu_j}{\sigma_j}, \hat{x}_{v,j} = \text{clip}(\tilde{x}_{v,j}, -5, 5)$$

where μ_j, σ_j are computed on training data only. **Population counts remain unnormalized** to preserve physical interpretation.

4.4.3 Mathematical Framework

4.4.3.1 Multi-Ethnic Reaction-Diffusion System

The core dynamics follow a coupled reaction-diffusion PDE on the spatial graph, extending classical reaction-diffusion theory (Fisher, 1937; A. Kolmogorov, Petrovsky, & Piskunov, 1937) to multi-species population dynamics on discrete spatial networks:

$$\frac{\partial u_{v,e}}{\partial t} = \underbrace{\alpha \cdot D_e(c, t) \sum_{w \in \mathcal{N}(v)} L_{vw} u_{w,e}}_{\text{Spatial diffusion}} + \underbrace{R_e(\mathbf{u}_v, \mathbf{x}_v, c, t)}_{\text{Local reactions}}$$

where:

- $u_{v,e}(t)$: Population of ethnicity e at node v
- $D_e(c, t)$: Diffusion coefficient (ethnicity e , city c , period t)
- α : Learnable global diffusion weight
- L_{vw} : Graph Laplacian entry
- R_e : Reaction term capturing demographic processes

Remark 3.1. The system is **period-specific**: parameters vary across four census periods $t \in \{0, 1, 2, 3\}$, capturing temporal heterogeneity in migration policies, economic conditions, and urban development.

4.4.3.2 Reaction Term Decomposition

The reaction term combines interpretable physics with learned neural augmentation:

$$\begin{aligned}
R_e(\mathbf{u}, \mathbf{x}, c, t) = & \underbrace{r_e(c, t) \cdot u_e \left(1 - \frac{N}{K_c(c, t)}\right)}_{(1) \text{ Logistic growth}} \\
& + \underbrace{\mathcal{A}_e(\mathbf{u}, t)}_{(2) \text{ Attention-based competition}} \\
& + \underbrace{F_e(c, t) \cdot \sigma(M_{\text{census}}(\mathbf{x}))_e \left(1 - \frac{u_e}{K_e(c, t)}\right)}_{(3) \text{ Modulated immigration}} \\
& - \underbrace{\epsilon_e(c, t) \cdot \sigma(M_{\text{census}}(\mathbf{x}))_e \cdot u_e}_{(4) \text{ Modulated emigration}} \\
& + \underbrace{\beta \cdot R_e^{\text{neural}}(\mathbf{u}, \mathbf{h}_c, \mathbf{h}_t, \mathbf{x})}_{(5) \text{ Deep neural residual}}
\end{aligned}$$

where $N = \sum_{i=1}^E u_i$ is total population.

Component Details

(1) Logistic Growth: Classical density-dependent growth(Verhulst, 1838) with city-specific carrying capacity $K_c(c, t)$.

(2) Attention-Based Competition: Replaces fixed Lotka-Volterra interaction matrices(Lotka, 1925; Volterra, 1926) with learned multi-head self-attention(Vaswani et al., 2017):

$$\mathcal{A}_e(\mathbf{u}, t) = \text{MultiHeadAttention}(\mathbf{E}_t + \text{PopEncoder}(\mathbf{u}))$$

where $\mathbf{E}_t \in \mathbb{R}^{E \times d}$ are period-specific ethnicity embeddings and PopEncoder maps populations to d -dimensional space.

(3-4) Modulated Migration: Immigration and emigration rates are modulated by census features through MLP $M_{\text{census}}: \mathbb{R}^{74} \rightarrow \mathbb{R}^E$ with sigmoid activation.

(5) Neural Residual: A deep MLP learns corrections:

$$R_e^{\text{neural}} = \text{MLP}(\mathbf{h}_c \oplus \mathbf{h}_t \oplus \sigma(M_{\text{census}}(\mathbf{x})) \oplus \frac{\mathbf{u}}{\|\mathbf{u}\|_1 + \epsilon})$$

where $\mathbf{h}_c \in \mathbb{R}^{512}$ (city embedding), $\mathbf{h}_t \in \mathbb{R}^{64}$ (period embedding), and $\oplus$ denotes concatenation.

Remark 3.2 (No DAUID Embeddings). The model uses **only city and period embeddings**, avoiding per-node parameters to improve generalization and reduce overfitting.

4.4.3.3 Parameter Structure

Period-City-Ethnicity Specific:

- Diffusion: $\mathbf{D} \in \mathbb{R}^{4 \times C \times E}$
- Growth: $\mathbf{r} \in \mathbb{R}^{4 \times C \times E}$

- Immigration: $\mathbf{F} \in \mathbb{R}^{4 \times C \times E}$
- Emigration: $\mathbf{\epsilon} \in \mathbb{R}^{4 \times C \times E}$

Period-City Specific:

- City carrying capacity: $\mathbf{K}_c \in \mathbb{R}^{4 \times C}$

Period-Ethnicity Specific:

- Ethnicity carrying capacity: $\mathbf{K}_e \in \mathbb{R}^{4 \times E}$

Global Learnable Scalars:

- Diffusion weight: $\alpha \in \mathbb{R}_+$
- Residual weight: $\beta \in \mathbb{R}_+$

Total parameter count: θ_{physics} contains $4 \times 52 \times 9 \times 4 + 4 \times 52 + 4 \times 9 + 2 = 7,534$ parameters.

4.4.3.4 Temporal Integration via RK4

The PDE is discretized using 4th-order Runge-Kutta (Butcher, 2016) with adaptive clamping:

$$\begin{aligned}
 \mathbf{k}_1 &= h \cdot f(\mathbf{u}_n, t_n) \\
 \mathbf{k}_2 &= h \cdot f(\mathbf{u}_n + \mathbf{k}_1/2, t_n + h/2) \\
 \mathbf{k}_3 &= h \cdot f(\mathbf{u}_n + \mathbf{k}_2/2, t_n + h/2) \\
 \mathbf{k}_4 &= h \cdot f(\mathbf{u}_n + \mathbf{k}_3, t_n + h) \\
 \mathbf{u}_{n+1} &= \max(\mathbf{u}_n + \frac{\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4}{6}, 0)
 \end{aligned}$$

where $h = \Delta t / N_{\text{steps}}$ and $N_{\text{steps}} = 1$ (single RK4 step over 5-year interval).

Remark 3.3 (Implementation). Despite $N_{\text{steps}} = 1$, the model is denoted "GraphPDE" to emphasize physics-informed architecture and differential equation formulation, distinguishing it from black-box neural approaches.

4.4.3.5 Attention-Based Ethnic Competition

Definition 4.1 (Multi-Head Self-Attention). For population state $\mathbf{u} \in \mathbb{R}_+^E$:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V$$

where queries, keys, values are linear projections of ethnicity embeddings augmented with population encodings.

Architecture:

- Number of heads: $H = 4$
- Head dimension: $d_k = d_v = 16$
- Total dimension: $d = 64$
- Dropout rate: $p = 0.1$

Theorem 4.1 (Universal Approximation). Multi-head self-attention can approximate any permutation-equivariant function on finite sets (Yun, Bhojanapalli, Rawat, Reddi, & Kumar, 2019).

4.4.3.6 Census Modulation Network

Definition 4.2 (Modulation MLP). A 3-layer fully connected network:

$$M_{\text{census}}: \mathbb{R}^{74} \xrightarrow{\text{FC}(256)} \mathbb{R}^{256} \xrightarrow{\text{FC}(128)} \mathbb{R}^{128} \xrightarrow{\text{FC}(E)} \mathbb{R}^E \xrightarrow{\sigma} [0,1]^E$$

with ReLU activations and layer normalization. Sigmoid output ensures interpretable modulation factors.

4.4.3.7 Neural Residual Network

Definition 4.3 (Residual MLP). A 5-layer network with skip connections:

$$R^{\text{neural}}: \mathbb{R}^{512+64+E+E} \rightarrow \mathbb{R}^{512} \rightarrow \mathbb{R}^{256} \rightarrow \mathbb{R}^{128} \rightarrow \mathbb{R}^{64} \rightarrow \mathbb{R}^E$$

Regularization:

- Dropout: $p = 0.1$ after each hidden layer
- Layer normalization before activations
- Weight decay: $\lambda = 10^{-5}$

4.4.4 Training Procedures and Optimization

4.4.4.1 Loss Function

Definition 5.1 (Composite Loss). The training objective combines multiple terms:

$$\mathcal{L}_{\text{total}} = \underbrace{\text{MSE}(\hat{\mathbf{u}}, \mathbf{u})}_{\mathcal{L}_{\text{MSE}}} + \underbrace{\lambda_1 \cdot \text{MAE}(\hat{\mathbf{u}}, \mathbf{u})}_{\mathcal{L}_{\text{MAE}}} + \underbrace{\lambda_2 \cdot \mathcal{L}_{\text{phys}}}_{\text{Physics reg.}} + \underbrace{\lambda_3 \cdot \mathcal{L}_{\text{smooth}}}_{\text{Label smoothing}}$$

where:

- $\mathcal{L}_{\text{MSE}} = \frac{1}{|\mathcal{V}| \cdot E} \sum_{v,e} (u_{v,e} - \hat{u}_{v,e})^2$
- $\mathcal{L}_{\text{MAE}} = \frac{1}{|\mathcal{V}| \cdot E} \sum_{v,e} |u_{v,e} - \hat{u}_{v,e}|$

- $\mathcal{L}_{\text{phys}} = \|D\|_2^2 + \|r\|_2^2 + \|F\|_2^2 + \|\epsilon\|_2^2 + \|K_c\|_2^2 + \|K_e\|_2^2$
- $\mathcal{L}_{\text{smooth}}$: Label smoothing term (Müller, Kornblith, & Hinton, 2019)

Hyperparameters:

- $\lambda_1 = 0.1$ (MAE weight)
- $\lambda_2 = 10^{-5}$ (physics regularization)
- $\lambda_3 = 0.1$ (label smoothing)

4.4.4.2 Optimizer Configuration

AdamW Optimizer (Loshchilov & Hutter, 2017):

- Learning rate: $\eta = 3 \times 10^{-4}$
- Betas: $\beta_1 = 0.9, \beta_2 = 0.999$
- Weight decay: $\lambda_{\text{wd}} = 10^{-5}$
- Gradient clipping: $\|\nabla_{\theta} \mathcal{L}\|_2 \leq 1.0$

Learning Rate Schedule: Cosine annealing with warm restarts:

$$\eta_t = \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min})(1 + \cos(\frac{T_{\text{cur}}}{T_{\max}}\pi))$$

where $\eta_{\max} = 3 \times 10^{-4}$, $\eta_{\min} = 10^{-6}$, $T_{\max} = 50$ epochs.

4.4.4.3 Training Protocol

Batch Configuration:

- Batch size: 512 samples
- Epochs: 150
- Early stopping patience: 15 epochs (validation R^2 threshold)

Data Augmentation:

- Random dropout masks (spatially consistent)
- Gaussian noise injection: $\mathcal{N}(0, 0.01 \cdot u)$ on populations
- No geometric augmentation (spatial structure fixed)

4.4.4.4 Evaluation Metrics

Definition 5.2 (Performance Metrics). For predictions $\hat{\mathbf{u}}$ and ground truth $\mathbf{u}$:

1. **R^2 Score:**

$$R^2 = 1 - \frac{\sum_{v,e} (u_{v,e} - \hat{u}_{v,e})^2}{\sum_{v,e} (u_{v,e} - \bar{u})^2}$$

2. **Mean Absolute Error:**

$$\text{MAE} = \frac{1}{|\mathcal{V}| \cdot E} \sum_{v,e} |u_{v,e} - \hat{u}_{v,e}|$$

3. **Root Mean Squared Error:**

$$\text{RMSE} = \sqrt{\frac{1}{|\mathcal{V}| \cdot E} \sum_{v,e} (u_{v,e} - \hat{u}_{v,e})^2}$$

4. **Mean Absolute Percentage Error (for $u > 10$):**

$$\text{MAPE} = \frac{100}{|\{(v,e): u_{v,e} > 10\}|} \sum_{u_{v,e} > 10} \frac{|u_{v,e} - \hat{u}_{v,e}|}{u_{v,e}}$$

4.4.5 Theoretical Analysis

4.4.5.1 Existence and Uniqueness

Theorem 6.1 (Well-Posedness). For initial condition $\mathbf{u}_0 \in \mathbb{R}_+^{|\mathcal{V}| \times E}$ and Lipschitz continuous R_e , the system

$$\frac{\partial \mathbf{u}}{\partial t} = \alpha D \mathbf{L} \mathbf{u} + R(\mathbf{u}, \mathbf{x})$$

admits a unique non-negative solution $\mathbf{u}(t) \in C^1([0, T]; \mathbb{R}_+^{|\mathcal{V}| \times E})$.

Proof Sketch. Follows from Picard-Lindelöf with invariant region $\mathbb{R}_+^{|\mathcal{V}| \times E}$ preserved by RK4 clamping.

4.4.5.2 Stability Analysis

Lemma 6.1 (Diffusion Stability). The diffusion operator $\mathbf{L}$ satisfies:

$$\|\mathbf{L} \mathbf{u}\|_2 \leq 2 \|\mathbf{u}\|_2$$

hence the explicit Euler scheme $\mathbf{u}_{n+1} = \mathbf{u}_n + h\alpha D \mathbf{L} \mathbf{u}_n$ is stable for $h \leq 1/(2\alpha D_{\max})$.

Theorem 6.2 (RK4 Stability). The RK4 scheme with clamping maintains stability for:

$$h \leq \frac{2.78}{\alpha D_{\max} \cdot \lambda_{\max}(\mathbf{L})} \approx \frac{1.39}{\alpha D_{\max}}$$

4.4.5.3 Approximation Theory

Theorem 6.3 (Universal Approximation). For any continuous $g: \mathbb{R}^E \rightarrow \mathbb{R}$ and $\epsilon > 0$, there exists a neural residual R^{neural} such that:

$$\sup_{\mathbf{u} \in K} |g(\mathbf{u}) - R^{\text{neural}}(\mathbf{u}; \theta)| < \epsilon$$

for any compact $K \subset \mathbb{R}^E$ (Hornik, Stinchcombe, & White, 1989).

Corollary 6.1. GraphPDE can approximate any spatiotemporal dynamics $\mathbf{u}(t + \Delta t) = G(\mathbf{u}(t), \mathbf{X})$ arbitrarily well on bounded domains.

4.4.5.4 Convergence Guarantees

Theorem 6.4 (RK4 Error Bound). For smooth R_e with bounded derivatives up to order 5, the RK4 scheme satisfies:

$$\|\mathbf{u}(t_n) - \mathbf{u}_n\| \leq Ch^4 \max_{0 \leq s \leq t_n} \left\| \frac{\partial^5 \mathbf{u}}{\partial t^5}(s) \right\|$$

provided $h \leq h_{\max} = O(1/L_R)$.

Proof. Standard RK4 error analysis (Butcher, 2016). Clamping to $\mathbb{R}_+$ preserves non-negativity without affecting order.

4.4.6 Computational Implementation

4.4.7 Hardware and Software Stack

- **GPU:** NVIDIA RTX 3060 (8GB VRAM)
- **Framework:** PyTorch 2.0+ with custom CUDA extensions
- **Graph Processing:** Sparse CSR format for Laplacian
- **Precision:** FP32 (FP16 caused numerical instability)

4.4.7.1 Custom CUDA Kernels

Forward Pass:

```
// Compute  $\alpha D L u$  for all ethnicities in parallel
__global__ void graph_laplacian_forward(
    float* population,    // [n_nodes, n_eth]
    float* diffusion_coef, // [n_cities, n_eth]
    int* node_to_city,
    int* row_ptr, int* col_idx, float* values,
```

```

float* output
) {
    int node = blockIdx.x;
    int eth = threadIdx.x;

    // Sparse mat-vec for this node and ethnicity
    float sum = 0.0f;
    for (int i = row_ptr[node]; i < row_ptr[node+1]; i++) {
        sum += values[i] * population[col_idx[i] * n_eth + eth];
    }

    int city = node_to_city[node];
    output[node * n_eth + eth] = diffusion_coef[city * n_eth + eth] * sum;
}

```

Backward Pass (Gradient Computation):

```

__global__ void graph_laplacian_backward(
    float* grad_output,    // [n_nodes, n_eth]
    float* population,
    float* diffusion_coef,
    int* node_to_city,
    int* row_ptr, int* col_idx, float* values,
    float* grad_population, // Output
    float* grad_diffusion  // Output
) {
    // Compute  $\partial L / \partial u$  and  $\partial L / \partial D$  via chain rule
    // (Implementation details in source code)
}

```

4.4.7.2 Memory Optimization

Graph Storage: CSR format reduces memory from $O(|\mathcal{V}|^2)$ to $O(|\mathcal{E}|)$:

- Row pointers: $30,092 \times 4 \text{ bytes} = 120 \text{ KB}$
- Column indices: $150,000 \times 4 \text{ bytes} = 600 \text{ KB}$
- Values: $150,000 \times 4 \text{ bytes} = 600 \text{ KB}$

- **Total:** ~1.3 MB (vs. ~3.6 GB for dense)

Batch Processing: Full-graph state ($30,091 \text{ nodes} \times 9 \text{ ethnicities} \times 4 \text{ bytes} = 1.08 \text{ MB}$) fits in GPU memory, enabling efficient ODE integration.

4.4.7.3 Performance Benchmarks

Custom CUDA kernels provide substantial computational speedup over naïve PyTorch implementations (Table 4-1).

Table 4-1. Cuda Performance compared to naïve pytorch.

Operation	Pytorch (Ms)	Cuda (Ms)	Speedup
Forward	1240	38	32.6×
Backward	2180	67	32.5×
Total	3420	105	32.6×

Note: Measured on RTX 4060, averaged over 100 iterations.

4.4.8 Uncertainty Quantification

4.4.8.1 Monte Carlo Dropout Framework

To address the critical need for uncertainty estimates in policy applications, we quantify prediction uncertainty through Monte Carlo (MC) dropout, a Bayesian approximation technique that treats dropout as variational inference over model weights (Gal & Ghahramani, 2016).

Definition 8.1 (Posterior Predictive Distribution). We approximate the posterior over predictions:

$$p(\mathbf{u}^* \mid \mathcal{D}, \mathbf{X}) \approx \frac{1}{M} \sum_{m=1}^M p(\mathbf{u}^* \mid \mathcal{D}, \mathbf{X}, \theta_m)$$

where θ_m are sampled by applying dropout with probability $p = 0.1$ at inference time, and $M = 100$ samples.

Algorithm 8.1 (MC Dropout Uncertainty Estimation):

Input: Trained model θ , test data $(\mathbf{u}_t, \mathbf{X})$, $M = 100$ samples
Output: Mean predictions μ , uncertainties σ , confidence intervals

```

1. Set model.train() # Keep dropout active
2. predictions = []
3. For m = 1 to M:
4.   With torch.no_grad(): # No gradients (not retraining!)
5.     u_pred = model.forward(u_t, X; θ)
6.     predictions.append(u_pred)
7. μ = mean(predictions, axis=0)
8. σ_raw = std(predictions, axis=0)
9. Return (μ, σ_raw)

```

Remark 8.1 (No Retraining Required). MC dropout leverages the already-trained model by running inference multiple times with stochastic dropout masks. This measures epistemic uncertainty arising from model limitations and finite training data, not aleatory uncertainty from inherent demographic stochasticity.

4.4.8.2 Temperature Scaling for Calibration

Uncalibrated MC dropout often produces conservative or overconfident uncertainty estimates. We apply post-hoc temperature scaling to achieve proper 95% confidence interval coverage.

Definition 8.2 (Temperature-Scaled Uncertainty). Raw uncertainties are rescaled:

$$\sigma_{\text{cal}} = T \cdot \sigma_{\text{raw}}$$

where the temperature T is determined by optimizing coverage on the validation set:

$$T^* = \arg \min_T \left| \frac{1}{N_{\text{val}}} \sum_{i=1}^{N_{\text{val}}} \mathbb{1}[u_i^{\text{true}} \in [\mu_i - 1.96T\sigma_i, \mu_i + 1.96T\sigma_i]] - 0.95 \right|$$

Calibration Results:

- **Optimal temperature:** $T^* = 4.55$
- **Test set coverage:** 95.1% (target: 95%)
- **Interpretation:** The trained model with low dropout rate ($p = 0.1$) underestimates uncertainty by factor of ~ 4.5 , which temperature scaling corrects post-hoc.

Remark 8.2 (Conservative Training). Low dropout rates during training prioritize prediction accuracy over uncertainty calibration. Post-hoc temperature scaling is a standard technique to recover well-calibrated probabilistic predictions without retraining (Guo, Pleiss, Sun, & Weinberger, 2017).

4.4.8.3 Empirical Uncertainty Statistics

Table 4-2 summarizes the overall uncertainty statistics on the test set.

Table 4-2. Overall Test Set Performance.

Metric	Value	Interpretation
95% Ci Coverage	95.1%	Well-calibrated intervals
Mean Uncertainty	$\bar{\sigma} = 17.7$	Average ± 17.7 individuals per DAUID
Median Uncertainty	$\tilde{\sigma} = 10.2$	Typical uncertainty (robust to outliers)
Maximum Uncertainty	$\sigma_{\text{Max}} = 712.4$	Large settlements with high volatility
Minimum Uncertainty	$\sigma_{\text{Min}} = 0.3$	Stable, small populations

Predictive Quality:

- Mean Absolute Error (MAE): 12.4 individuals
- Root Mean Squared Error (RMSE): 28.9 individuals
- Median Absolute Error: 5.8 individuals

Remark 8.3 (Skewed Distribution). The difference between mean and median uncertainty (17.7 vs. 10.2) indicates right-skewed distribution, with most predictions having moderate uncertainty but occasional high-uncertainty cases for large or rapidly changing populations.

4.4.8.4 Ethnicity-Specific Uncertainty Patterns

Uncertainty varies systematically across ethnic groups, reflecting underlying demographic dynamics (Table 4-3).

Table 4-3. Uncertainty varies systematically across ethnic groups, reflecting demographic dynamics.

Ethnicity	Mean σ	Coverage	Interpretation
India	27.3	95.0%	Rapidly growing, volatile transitions
Philippines	24.0	94.8%	High immigration rates
China	23.3	95.3%	Large, diverse settlements
Pakistan	18.5	95.2%	Moderate growth
Iran	14.2	95.1%	Established communities
Sri Lanka	12.7	95.0%	Stable trajectories
Portugal	9.9	95.4%	Mature, declining populations
United Kingdom	9.1	95.2%	Established, predictable
Italy	9.0	95.6%	Stable, aging populations

Key Finding: Higher uncertainty for rapidly growing communities experiencing volatile demographic transitions (India: $\sigma = 27.3$, Philippines: $\sigma = 24.0$, China: $\sigma = 23.3$) compared to stable or declining populations with predictable trajectories (Italy: $\sigma = 9.0$, United Kingdom: $\sigma = 9.1$, Portugal: $\sigma = 9.9$).

Theorem 8.1 (Uncertainty-Growth Correlation). Let ρ_e denote the average growth rate of ethnicity e over census periods. Then empirically:

$$\text{Corr}(\sigma_e, |\rho_e|) \approx 0.78 (p < 0.01)$$

indicating strong positive association between demographic volatility and prediction uncertainty.

4.4.8.5 Population-Dependent Uncertainty Scaling

Definition 8.3 (Population Bins). We partition predictions into five categories based on initial population u_t Table 4-4:

Table 4-4. Population range and sample count.

Population Range	Label	Sample Count	Mean σ	Relative Uncertainty (Cv)
[0, 10)	Small	28,314	15.6	152%
[10, 50)	Small-Medium	23,187	13.6	48%
[50, 100)	Medium	9,842	21.8	31%
[100, 500)	Large	12,063	69.7	38%
[500, ∞)	Very Large	2,788	173.4	35%

Key Observations:

1. **Non-linear absolute scaling:** Uncertainty increases dramatically with population size for large settlements ($\sigma \propto u^{0.67}$, see Theorem 8.2).
2. **Stabilized relative uncertainty:** Coefficient of variation ($CV = \sigma/\mu$) remains comparable (31-48%) for populations above 50, indicating **relative prediction reliability is consistent across scales**.
3. **High CV for small populations:** Populations <10 show $CV=152\%$, reflecting challenges in predicting nascent communities with sparse data.

Theorem 8.2 (Sublinear Uncertainty Scaling). Uncertainty scales sublinearly with population:

$$\sigma(u) \sim u^\beta, \beta \approx 0.67$$

Proof (Empirical). Log-log regression on population-binned data:

$$\log \sigma = 0.67 \log u + c, R^2 = 0.94$$

This sublinear relationship ($\beta < 1$) implies that relative uncertainty (CV) decreases as $u^{\beta-1} = u^{-0.33}$, stabilizing for large populations.

Policy Implication: The substantial prediction uncertainty for emerging settlements (populations < 50 : $\sigma = 14 - 16$) indicates that forecasts for nascent ethnic communities should be interpreted as probabilistic scenarios spanning plausible demographic trajectories rather than point predictions, guiding precautionary infrastructure provisioning that accommodates a range of outcomes.

4.4.8.6 Policy-Relevant Uncertainty Analysis

Intervention Effects: Bootstrap Confidence Intervals

To assess uncertainty in learned demographic parameters relevant to policy design, we perform bootstrap resampling of the training data and retrain GraphPDE to obtain parameter distributions.

Algorithm 8.2 (Parameter Uncertainty via Bootstrap):

Input: Training data D , $B = 50$ bootstrap samples

Output: Parameter confidence intervals

1. For $b = 1$ to B :
2. $D_b = \text{resample}(D, \text{with replacement})$
3. $\theta_b = \text{train_graphpde}(D_b)$
4. Extract: $D_b, r_b, F_b, \epsilon_b, K_b$
5. Compute 95% CI: $[\theta_{2.5\%}, \theta_{97.5\%}]$

Emigration Parameter Uncertainty:

- Mean reduction in emigration for capacity-constrained neighborhoods: 8-13%
- 95% confidence interval: $\epsilon_{\text{constrained}}/\epsilon_{\text{baseline}} \in [0.87, 0.92]$
- **Interpretation:** Narrow confidence intervals support confident policy implementation. Affordable housing construction in capacity-constrained areas reliably reduces emigration pressure.

Immigration Parameter Uncertainty:

- Mean increase in immigration attraction: 2-7%
- 95% confidence interval: $F_{\text{enhanced}}/F_{\text{baseline}} \in [1.02, 1.07]$
- **Interpretation:** Wider uncertainty bounds reflect complex dependencies on migration networks and labor market conditions not fully captured by census features. Immigration-attraction policies require pilot studies and adaptive implementation rather than deterministic planning.

Implications for Decision-Making Under Uncertainty

Risk-Aware Policy Design:

1. **Interventions with narrow uncertainty bounds** (e.g., capacity expansion $\rightarrow$ emigration reduction: 8-13%):
 - Support confident implementation

- Enable deterministic planning
- Suitable for large-scale rollout
- 2. **Interventions with wide uncertainty bounds** (e.g., labor market interventions → immigration attraction: 2-7%):
 - Require pilot studies before scaling
 - Demand adaptive management frameworks
 - Necessitate scenario-based planning

Precautionary Principle for Nascent Communities:

- Populations <50 individuals: $\sigma = 14 - 16$ (CV > 100%)
- Recommendation: Provision flexible infrastructure accommodating 2× to 3× forecasted growth
- Avoid fixed-capacity interventions vulnerable to underestimation

Calibrated Confidence Intervals:

- 95.1% empirical coverage validates use for risk quantification
- Uncertainties reflect epistemic limitations (model structure, finite data)
- Do NOT capture aleatory uncertainty (inherent demographic stochasticity, external shocks)
- **Complementary analyses needed:** Scenario analysis, sensitivity studies, stress testing for high-stakes interventions

4.4.8.7 Limitations and Future Directions

Current Limitations:

1. **Epistemic vs. Aleatory Uncertainty:** MC dropout captures only epistemic uncertainty from model limitations. Aleatory uncertainty from inherent demographic stochasticity (random births, deaths, migration) requires probabilistic demographic models.
2. **No Quantification of External Shocks:** Policy changes, economic crises, pandemics (e.g., COVID-19) introduce non-stationary dynamics not captured by historical training data.
3. **Single Model Architecture:** Uncertainty reflects only the GraphPDE family. Model ensemble methods (e.g., training with different architectures) would provide more comprehensive uncertainty quantification.

Future Work:

- **Ensemble Methods:** Train diverse models (Graph Transformers, Spatio-Temporal GANs) and aggregate predictions to capture model uncertainty.
- **Conformal Prediction:** Distribution-free uncertainty quantification with finite-sample guarantees.
- **Scenario-Based Forecasting:** Conditional predictions under different policy and economic scenarios.

- **Real-Time Calibration:** Update uncertainty estimates as new census data becomes available.

4.4.9 Energy Landscape Analysis: Post-Training Interpretation

4.4.9.1 Motivation and Framework

Following model training, we conduct energy landscape analysis to interpret learned demographic dynamics through the lens of statistical mechanics and transition state theory. This post-hoc analysis framework provides policy-relevant insights into settlement stability, configuration transitions, and demographic inertia without requiring additional model training.

Definition 9.1 (Configuration Energy Functional). Following the Ginzburg-Landau free energy formalism (Cahn & Hilliard, 1958; Ginzburg & Landau, 1950), we define a demographic free energy functional for a demographic configuration $\mathbf{u} = (u_1, \dots, u_E)$ representing ethnic population counts, we define the free energy functional:

$$\mathcal{F}[\mathbf{u}] = \int_{\Omega} [f(\mathbf{c}) + \frac{1}{2} \sum_{e=1}^E D_e |\nabla c_e|^2] d\Omega$$

where:

- $\mathbf{c} = \mathbf{u}/K$ is the normalized concentration vector (dimensionless)
- $f(\mathbf{c})$ is the reaction energy density derived from GraphPDE's learned parameters
- D_e are the learned diffusion coefficients
- Ω represents the spatial domain

Remark 9.1 (Physical Interpretation). The energy functional $\mathcal{F}$ quantifies demographic "favorability" where:

- **Low energy configurations** ($\mathcal{F} < 0$) represent stable, self-reinforcing settlement patterns
- **High energy configurations** ($\mathcal{F} > 0$) represent unstable patterns requiring external forcing to maintain
- **Energy barriers** between configurations quantify demographic inertia

4.4.9.2 Reaction Energy Density from Learned Parameters

The reaction energy density $f(\mathbf{c})$ is constructed from GraphPDE's learned reaction terms:

$$\begin{aligned}
f(\mathbf{c}) = & - \sum_{e=1}^E r_e [c_e \ln(c_e) - c_e + \frac{c_e^2}{2(1 - \sum_i c_i)}] & (\text{Logistic growth}) \\
& + \frac{1}{2} \sum_{e,e'} W_{ee'} c_e c_{e'} & (\text{Interaction energy}) \\
& + \sum_{e=1}^E c_e \ln(c_e) & (\text{Concentration potential}) \\
& + T_{\text{social}} \sum_{e=1}^E p_e \ln(p_e) & (\text{Mixing entropy})
\end{aligned}$$

where:

- r_e are learned growth rates
- $W_{ee'}$ is the interaction matrix from the attention mechanism
- $p_e = c_e / \sum_i c_i$ are concentration fractions
- T_{social} is a dimensionless "social temperature" parameter (set to 1.0)

Theorem 9.1 (Energy-Dynamics Correspondence). For sufficiently smooth configurations, the GraphPDE dynamics satisfy:

$$\frac{\partial u_e}{\partial t} = - \frac{\delta \mathcal{F}}{\delta u_e} + \text{noise}$$

demonstrating that the model follows gradient flow on the energy landscape with stochastic fluctuations.

Proof Sketch. The functional derivative $\delta \mathcal{F} / \delta u_e$ recovers the reaction-diffusion terms when the neural residuals are small.

4.4.9.3 Configuration Space and Reaction Coordinates

Definition 9.2 (Configuration Space). We parameterize demographic configurations using two reaction coordinates:

1. **Focal concentration:** $\xi_1 = c_{\text{focal}} = u_{\text{focal}} / \sum_e u_e \in [0, 1]$
2. **Shannon diversity:** $\xi_2 = H = - \sum_{e=1}^E p_e \ln(p_e) \in [0, \ln E]$

The energy surface $\mathcal{F}(\xi_1, \xi_2)$ is computed by:

$$\mathcal{F}(\xi_1, \xi_2) = \min_{\{\mathbf{c}: c_{\text{focal}} = \xi_1, H(\mathbf{c}) = \xi_2\}} f(\mathbf{c})$$

Algorithm 9.1 (Energy Surface Construction):

Input: Learned parameters θ , grid resolution N

Output: Energy surface $F(\xi_1, \xi_2)$

1. Create grid: $\xi_1 \in [0,1], \xi_2 \in [0, \ln(E)]$
2. For each grid point (ξ_1, ξ_2) :
3. Construct synthetic configuration c satisfying constraints
4. Compute $F = f(c)$ using learned parameters
5. Smooth surface with Gaussian filter ($\sigma = 1.5$)
6. Return $F(\xi_1, \xi_2)$

4.4.9.4 Identification of Stable Configurations

Definition 9.3 (Stable Configuration). A configuration $\mathbf{c}^*$ is locally stable if:

$$\nabla_{\mathbf{c}} \mathcal{F}(\mathbf{c}^*) = \mathbf{0} \text{ and } \nabla_{\mathbf{c}}^2 \mathcal{F}(\mathbf{c}^*) > 0$$

Algorithm 9.2 (Stable Configuration Detection):

Input: Energy surface $F(\xi_1, \xi_2)$

Output: List of stable configurations $\{C_1, C_2, \dots, C_k\}$

1. Apply minimum filter (size = 6×6) to identify local minima
2. For each local minimum (ξ_1^*, ξ_2^*) :
3. If boundary point (within 3 pixels of edge): reject
4. Else: Add to stable configuration list
5. Sort by energy: $F(C_1) \leq F(C_2) \leq \dots \leq F(C_k)$
6. Return top k configurations

4.4.9.5 Transition State Theory and Demographic Barriers

Definition 9.4 (Transition State). For configurations $\mathbf{c}_A$ and $\mathbf{c}_B$, the transition state $\mathbf{c}^\ddagger$ satisfies:

$$\mathbf{c}^\ddagger = \arg \max_{\mathbf{c} \in \Gamma(A \rightarrow B)} \mathcal{F}(\mathbf{c})$$

where $\Gamma(A \rightarrow B)$ is the minimum-energy pathway connecting A and B .

Definition 9.5 (Reorganization Barrier). The activation energy for transition $A \rightarrow B$ is:

$$\Delta \mathcal{F}_{A \rightarrow B}^\ddagger = \mathcal{F}(\mathbf{c}^\ddagger) - \mathcal{F}(\mathbf{c}_A)$$

Algorithm 9.3 (Minimum-Barrier Path):

Input: Configurations c_A, c_B , energy surface F

Output: Transition pathway Γ , barrier height $\Delta F^\ddagger$

1. Initialize linear path: $\gamma(s) = (1-s) \cdot c_A + s \cdot c_B$ for $s \in [0,1]$
2. Add curvature: $\gamma(s) \leftarrow \gamma(s) + 0.12 \cdot \sin(\pi s) \cdot \text{perpendicular}$
3. Compute energy along path: $E(s) = F(\gamma(s))$
4. Find maximum: $s^\ddagger = \text{argmax}_s E(s)$
5. Barrier height: $\Delta F^\ddagger = E(s^\ddagger) - E(0)$
6. Return $\gamma, \Delta F^\ddagger$

Remark 9.2 (Policy Interpretation). Barrier heights quantify resistance to settlement reorganization:

- $\Delta F^\ddagger < 1$: Pattern changes rapidly with minor interventions
- $1 \leq \Delta F^\ddagger \leq 3$: Sustained policy interventions required
- $\Delta F^\ddagger > 3$: Strong demographic lock-in, major interventions needed

4.4.9.6 Statistical Mechanics Analogy: System Responsiveness

While not literal thermodynamics, we borrow the Boltzmann distribution formalism to quantify system sensitivity to perturbations.

Definition 9.6 (Partition Function Analogue). Adapting the Boltzmann-Gibbs formalism (Gibbs, 1902) to demographic systems, for interaction intensity parameter $\theta > 0$:

$$Z(\theta) = \sum_{i=1}^N \exp\left(-\frac{\mathcal{F}_i}{\theta}\right)$$

where $\{\mathcal{F}_i\}$ are sampled configuration energies from observed settlement patterns.

Definition 9.7 (Mean Energy and System Responsiveness).

$$\langle \mathcal{F} \rangle(\theta) = \frac{1}{Z(\theta)} \sum_i \mathcal{F}_i \exp\left(-\frac{\mathcal{F}_i}{\theta}\right)$$

$$\chi(\theta) = \frac{\langle \mathcal{F}^2 \rangle - \langle \mathcal{F} \rangle^2}{\theta^2} = \frac{\partial \langle \mathcal{F} \rangle}{\partial \ln \theta}$$

where χ measures demographic sensitivity to changes in clustering vs. dispersal forces.

Theorem 9.2 (Critical Threshold). The responsiveness $\chi(\theta)$ exhibits a peak at θ_c indicating a qualitative transition in settlement organization:

$$\theta_c = \arg \max_{\theta} \chi(\theta)$$

Remark 9.3 (Not Thermal Physics). The parameter θ is dimensionless and represents balance between clustering (ethnic affinity, affordability constraints) and dispersal forces (integration, mobility), NOT literal temperature. The Boltzmann-like formalism provides mathematical structure for analyzing demographic transitions.

4.4.9.7 Spatial Energy Maps and Metastable Basins

Definition 9.8 (Local Energy Landscape). For each DAUID $v \in \mathcal{V}$:

$$\mathcal{F}_v = f(\mathbf{u}_v/K) + \frac{1}{2} \sum_{e=1}^E D_e \sum_{w \in \mathcal{N}(v)} L_{vw} (u_{v,e} - u_{w,e})^2$$

combining local reaction energy with spatial gradient penalties.

Algorithm 9.4 (Metastable Basin Identification):

Input: Local energies $\{F_v\}$, adjacency matrix A , population threshold

Output: Metastable basins $\{B_1, B_2, \dots, B_m\}$

1. Identify high-stability regions: $S = \{v : F_v < P_{30}(F) \text{ and } u_v > \text{threshold}\}$
2. Dilate region mask with binary dilation (2 iterations)
3. Label connected components $\rightarrow$ candidate basins
4. For each basin B :
5. If $|B| < 8$ DAUIDs: reject (too small)
6. Compute: mean_energy, total_population, geometric_center
7. Sort basins by population (descending)
8. Return top m basins

4.4.9.8 Configuration Transition Dynamics

Definition 9.9 (Master Equation). Following Kramers' transition state theory (Hänggi, Talkner, & Borkovec, 1990; Kramers, 1940), the temporal evolution of basin occupancy probabilities $\mathbf{P}(t) = [P_1(t), \dots, P_m(t)]^T$ follows:

$$\frac{d\mathbf{P}}{dt} = \mathbf{K}\mathbf{P}$$

where the rate matrix $\mathbf{K}$ has entries:

$$K_{ij} = \begin{cases} v_0 \exp(-\frac{\Delta \mathcal{F}_{i \rightarrow j}^\ddagger}{\theta}) & i \neq j \\ -\sum_{k \neq i} K_{ki} & i = j \end{cases}$$

with attempt frequency $v_0 \approx 0.15 \text{year}^{-1}$ (calibrated from historical census data).

Theorem 9.3 (Steady-State Distribution). The equilibrium distribution $\mathbf{P}^\infty$ satisfies:

$$\mathbf{K}\mathbf{P}^\infty = \mathbf{0}$$

and is given by the eigenvector corresponding to eigenvalue $\lambda = 0$.

Theorem 9.4 (Equilibration Timescales). The relaxation times are:

$$\tau_k = -\frac{1}{\text{Re}(\lambda_k)}, k = 2, \dots, m$$

where $\lambda_2, \dots, \lambda_m$ are the non-zero eigenvalues of $\mathbf{K}$.

4.4.9.9 Configuration Network Topology

Definition 9.10 (Configuration Graph). Construct $G = (V_{\text{config}}, E_{\text{trans}})$ where:

- Nodes: V_{config} represents key demographic configurations
- Edges: $(i, j) \in E_{\text{trans}}$ if $\|\mathbf{c}_i - \mathbf{c}_j\|_2 < \tau_{\text{sim}}$
- Weights: $w_{ij} = 1/(1 + \Delta \mathcal{F}_{i \rightarrow j}^\ddagger)$

Algorithm 9.5 (Network Construction):

Input: Configurations $\{\mathbf{c}_1, \dots, \mathbf{c}_k\}$, energy surface F

Output: Configuration network G

1. For each configuration $\mathbf{c}_i$:
2. Compute: favorability F_i , stability s_i , category label
3. Add node: $G.\text{add_node}(i, F=F_i, s=s_i, \text{category})$
4. For each pair (i, j) :
5. Compute compositional distance: $d = \|\mathbf{c}_i - \mathbf{c}_j\|_2$
6. If $d < 0.6$: # Similar configurations
7. Compute barrier: $\Delta F_{i \rightarrow j}^\ddagger = |F_i - F_j|$
8. Add edge: $G.\text{add_edge}(i, j, \text{weight}=1/(1+\Delta F_{i \rightarrow j}^\ddagger), \text{barrier}=\Delta F_{i \rightarrow j}^\ddagger)$
9. Return G

4.4.9.10 Policy Applications and Limitations

Applications:

1. **Intervention Targeting:** Identify which basins have lowest barriers to desired configurations
2. **Temporal Planning:** Use equilibration timescales τ_k to set realistic policy horizons
3. **Risk Assessment:** Barrier heights quantify demographic inertia and resistance to change
4. **Scenario Analysis:** Explore multiple pathways on energy landscape under different interventions

Critical Limitations:

1. **Post-Hoc Nature:** Energy landscape is derived from trained model parameters, not optimized directly
2. **Approximate Barriers:** Minimum-barrier paths use linear interpolation with curvature, not true geodesics
3. **Static Analysis:** Does not model temporal evolution of landscape itself (non-stationary dynamics)
4. **Simplified Coordinates:** Two-dimensional (ξ_1, ξ_2) projection loses information about full E -dimensional configuration space
5. **No Uncertainty Quantification:** Energy values are point estimates without confidence intervals

Remark 9.4 (Interpretive Tool). The energy landscape framework provides qualitative insights into settlement dynamics and demographic inertia. Quantitative policy recommendations require validation through historical data analysis, pilot studies, and sensitivity analysis.

4.4.10 Crystallization Analysis: Multi-Scale Settlement Patterns

4.4.10.1 Motivation and Conceptual Framework

Complementing the energy landscape analysis, crystallization analysis provides a spatial pattern recognition framework inspired by phase transitions in condensed matter physics. While energy landscapes reveal abstract configuration space, crystallization analysis examines concrete geographic patterns of ethnic settlement formation, growth, and coalescence.

Definition 10.1 (Settlement Crystallization). The process by which spatially dispersed immigrant populations organize into coherent, spatially localized ethnic enclaves through:

1. **Nucleation:** Emergence of initial settlement clusters
2. **Growth:** Expansion of existing clusters
3. **Coalescence:** Merging of nearby clusters into larger regions

Remark 10.1 (Not Literal Crystallization). This framework borrows mathematical tools from materials science (JMAK theory, structure factors, Voronoi tessellation) to quantify spatial organization patterns. It

does NOT claim ethnic settlement is a literal phase transition, but rather that similar mathematical structures describe both phenomena.

4.4.10.2 Multi-Scale Spatial Analysis

Definition 10.2 (Spatial Scales). We analyze settlement patterns at three hierarchical scales:

1. **Macro Scale** (~20-50 km): City-wide distribution revealing major concentration regions
2. **Meso Scale** (~ 5-15 km): Neighborhood-level patterns showing temporal evolution
3. **Micro Scale** (~1-3 km): Block-level tessellation revealing fine structure

Scale Transition Relationships:

$$\text{Macro} \xrightarrow{\text{zoom}} \text{Meso} \xrightarrow{\text{zoom}} \text{Micro}$$

Each scale provides complementary information about settlement organization.

4.4.10.3 Cluster Identification and Tracking

Definition 10.3 (Settlement Cluster). A spatially connected set of DAUIDs $\mathcal{C} \subset \mathcal{V}$ where:

$$\mathcal{C} = \{v \in \mathcal{V} : u_v > \theta_{\text{pop}} \text{ and } \exists w \in \mathcal{C}, (v, w) \in \mathcal{E}\}$$

for population threshold θ_{pop} .

Algorithm 10.1 (DBSCAN-Based Cluster Detection) (Ester, Kriegel, Sander, & Xu, 1996):

Input: Coordinates $\{x_i\}$, populations $\{u_i\}$, parameters (eps, min_samples)

Output: Cluster labels $\{c_1, \dots, c_n\}$

1. Filter to high-concentration areas: $S = \{i : u_i > P_{70}(u)\}$
2. Apply DBSCAN(coordinates_S, eps, min_samples)
3. For each cluster label ℓ :
4. $C_\ell = \{i \in S : \text{label}_i = \ell\}$
5. Compute: population, centroid, area
6. Sort clusters by population (descending)
7. Return cluster assignments

Temporal Tracking: For census years $t_1, \dots, t_T$, identify clusters $\{\mathcal{C}_k^{(t)}\}$ at each time point and track:

- Number of clusters: $N_c(t)$
- Average cluster population: $\bar{u}_c(t) = \frac{1}{N_c(t)} \sum_{k=1}^{N_c(t)} \sum_{v \in \mathcal{C}_k^{(t)}} u_v$

- Cluster size distribution: $\{ |C_k^{(t)}| \}_{k=1}^{N_c(t)}$

4.4.10.4 Growth Kinetics: JMAK Theory

The Johnson-Mehl-Avrami-Kolmogorov (JMAK) theory (Avrami, 1939 ; 1941; A. N. Kolmogorov, 1937) models transformation kinetics in materials. Adapted for settlement analysis:

Definition 10.4 (Transformed Fraction). The fraction of total settled area:

$$X(t) = \frac{\sum_{v: u_v(t) > \theta} A_v}{\sum_{v \in \mathcal{V}} A_v}$$

where A_v is the area of DAUID v .

JMAK Equation:

$$X(t) = 1 - \exp [-(kt)^n]$$

where:

- k is the effective growth rate constant
- n is the Avrami exponent encoding growth dimensionality

Theorem 10.1 (Avrami Exponent Interpretation). The exponent n characterizes growth mechanism, with different values corresponding to distinct settlement dynamics (Table 4-5):.

Table 4-5. Settlement interpretation based on growth mechanism.

n Range	Growth Mechanism	Settlement Interpretation
$n \approx 1$	Interface-controlled	Linear expansion along corridors
$n \approx 2$	2D nucleation + growth	Radial neighborhood expansion
$n \approx 3$	3D nucleation + growth	Volumetric city-wide spread
$n > 3$	Complex, multi-stage	Sequential waves of settlement

Algorithm 10.2 (JMAK Fitting):

Input: Time series $\{t_1, \dots, t_T\}$, populations $\{U(t_1), \dots, U(t_T)\}$

Output: Parameters (k, n) , growth mechanism classification

1. Compute transformed fraction: $X(t) = U(t) / \max(U)$
2. Linearize: $\ln(-\ln(1 - X)) = n \cdot \ln(k) + n \cdot \ln(t)$
3. Fit linear regression to $(\ln(t), \ln(-\ln(1-X)))$
4. Extract: $n = \text{slope}$, $k = \exp(\text{intercept}/n)$
5. Classify mechanism based on n value

6. Return (k, n), classification

4.4.10.5 Voronoi Tessellation at Micro Scale

Definition 10.5 (Voronoi Cell). For DAUID centroid $\mathbf{x}_v$, the Voronoi cell is:

$$V_v = \{\mathbf{x} \in \mathbb{R}^2: \|\mathbf{x} - \mathbf{x}_v\| \leq \|\mathbf{x} - \mathbf{x}_w\|, \forall w \in \mathcal{V}\}$$

Algorithm 10.3 (Weighted Voronoi Tessellation):

Input: Centroids $\{\mathbf{x}_v\}$, populations $\{u_v\}$, micro-region bounds

Output: Tessellation with color-coded cells

1. Filter to micro region: $M = \{v : \mathbf{x}_v \in \text{bounds and } u_v > \text{threshold}\}$
2. Compute Voronoi tessellation: $\text{Vor} = \text{Voronoi}(\text{coordinates_M})$
3. For each point v in M :
4. Get Voronoi cell vertices from Vor
5. If cell is bounded (no infinite edges):
6. color = colormap($u_v / \max(u)$)
7. Plot filled polygon with color
8. Mark centroid with scatter point
9. Return tessellation visualization

Cell Properties Analysis:

- **Cell area:** $A_v = \text{area}(V_v)$
- **Population density:** $\rho_v = u_v / A_v$
- **Compactness:** $\kappa_v = 4\pi A_v / P_v^2$ where P_v is perimeter
- **Nearest neighbor distance:** $d_v = \min_{w \neq v} \|\mathbf{x}_v - \mathbf{x}_w\|$

4.4.10.6 Spatial Correlation: Structure Factor Analysis

Definition 10.6 (Structure Factor). The spatial Fourier transform of population distribution:

$$S(\mathbf{k}) = \left| \frac{1}{N} \sum_{v=1}^N u_v \exp(i\mathbf{k} \cdot \mathbf{x}_v) \right|^2$$

For continuous interpolated field $u(\mathbf{x})$:

$$S(\mathbf{k}) = |\mathcal{F}\{u(\mathbf{x})\}|^2$$

where $\mathcal{F}$ denotes 2D Fourier transform.

Algorithm 10.4 (Radially Averaged Structure Factor):

Input: Population field $u(x,y)$ on grid, physical extent L

Output: $S(k)$ vs k , fractal dimension, characteristic wavelengths

1. Apply Hanning window: $u_windowed = u \cdot window_2D$
2. Compute 2D FFT: $F = \text{fft2}(u_windowed)$
3. Shift zero frequency to center: $F_shifted = \text{fftshift}(F)$
4. Compute power spectrum: $P = |F_shifted|^2$
5. Compute radial distances from center
6. Bin power spectrum by radial distance: $S(k) = \text{radial_profile}(P)$
7. Convert to physical units: $k_physical = k_pixel \cdot 2\pi/L$
8. Fit power law: $\log(S) \sim \alpha \cdot \log(k)$
9. Extract fractal dimension: $D = -\alpha/2$
10. Find peaks in $S(k) \rightarrow$ characteristic wavelengths $\lambda = 2\pi/k_peak$
11. Return $S(k)$, D , $\{\lambda_i\}$

Theorem 10.2 (Fractal Dimension from Power Law). If $S(k) \sim k^\alpha$ for small k , the spatial distribution has fractal dimension:

$$D_f = \frac{d - \alpha}{2}$$

where $d = 2$ is the embedding dimension.

Proof. Follows from the relationship between power spectral density and correlation dimension in spatial statistics.

Physical Interpretation:

- $D_f < 2$: Fractal, highly clustered distribution
- $D_f = 2$: Space-filling, uniform coverage
- Peaks in $S(k)$: Characteristic spacing $\lambda = 2\pi/k$ between settlements

4.4.10.7 Nucleation and Coalescence Dynamics

Definition 10.7 (Cluster Dynamics Regimes). Based on temporal evolution of $N_c(t)$:

1. **Nucleation-dominated:** $\frac{dN_c}{dt} > 0$ (increasing cluster count)
2. **Coalescence-dominated:** $\frac{dN_c}{dt} < 0$ (decreasing cluster count, merging)

3. **Steady-state:** $\frac{dN_c}{dt} \approx 0$ (stable configuration)

Algorithm 10.5 (Nucleation Rate Estimation):

Input: Time series $\{t_i\}$, cluster counts $\{N_c(t_i)\}$

Output: Nucleation rate J , coalescence indicator

1. Compute temporal derivative: $dN/dt \approx \Delta N_c / \Delta t$
2. If $dN/dt > 0.5$ clusters/year:
3. Regime = "Nucleation & Growth"
4. $J = dN/dt$
5. Elif $dN/dt < -0.5$ clusters/year:
6. Regime = "Coalescence"
7. $J = 0$
8. Else:
9. Regime = "Stable"
10. $J \approx 0$
11. Return J , Regime

4.4.10.8 Spatial Autocorrelation: Moran's I

Definition 10.8 (Moran's I Statistic). Global measure of spatial autocorrelation:

$$I = \frac{N}{\sum_i \sum_j w_{ij}} \cdot \frac{\sum_i \sum_j w_{ij} (u_i - \bar{u})(u_j - \bar{u})}{\sum_i (u_i - \bar{u})^2}$$

where:

- w_{ij} is spatial weight (typically $w_{ij} = 1/d_{ij}$ for neighbors)
- $\bar{u} = \frac{1}{N} \sum_i u_i$ is mean population
- N is number of DAUIDs

Interpretation:

- $I > 0$: Positive autocorrelation (clustering)
- $I \approx 0$: Random spatial distribution
- $I < 0$: Negative autocorrelation (dispersion)

Significance Testing:

$$Z = \frac{I - E[I]}{\sqrt{\text{Var}[I]}}$$

where $E[I] = -1/(N - 1)$ and $\text{Var}[I]$ depends on spatial weights.

Algorithm 10.6 (Moran's I Computation):

Input: Populations $\{u_i\}$, coordinates $\{x_i\}$

Output: I statistic, Z-score, interpretation

1. Compute pairwise distances: $D = \text{distance_matrix}(x)$
2. Construct spatial weights: $W = 1/(D + \epsilon)$
3. Set diagonal to zero: $W[i,i] = 0$
4. Normalize: $W = W / \text{sum}(W)$
5. Compute deviations: $z = (u - \text{mean}(u)) / \text{std}(u)$
6. Calculate I: $I = N \cdot \text{sum}(W \cdot \text{outer}(z,z)) / \text{sum}(z^2)$
7. Compute expected value: $E[I] = -1/(N-1)$
8. Compute variance: $\text{Var}[I] = \dots$ (formula from spatial stats)
9. Z-score: $Z = (I - E[I]) / \text{sqrt}(\text{Var}[I])$
10. If $I > 0.3$ and $|Z| > 2.58$: "Significantly clustered"
11. Elif $I < -0.3$ and $|Z| > 2.58$: "Significantly dispersed"
12. Else: "Random or weakly structured"
13. Return I, Z, interpretation

4.4.10.9 Phase Diagram Representation

While not a literal thermodynamic phase diagram, we construct a 2D representation of settlement states:

Definition 10.9 (Settlement Phase Space). Two-dimensional space parameterized by:

- **Horizontal axis:** Concentration $c = U_{\text{focal}}/U_{\text{total}}$
- **Vertical axis:** Scaled interaction parameter T/T_c (analogous to temperature)

Phase Boundaries (Heuristic):

For mean-field Ising-like model:

$$c_{\text{critical}}(\tau) = 0.5 \pm 0.5\sqrt{1 - \tau}$$

where $\tau = T/T_c \in [0,1]$.

Regions:

1. **Separated Phase** ($c < c_{\text{crit}}$ or $c > 1 - c_{\text{crit}}$): High segregation
2. **Mixed Phase** ($c_{\text{crit}} < c < 1 - c_{\text{crit}}$): Integrated settlement

Current State Marker: Plot observed ($c_{\text{current}}, T_{\text{current}}$) to determine regime.

Remark 10.2 (Qualitative Tool). The phase diagram is a visual metaphor for settlement states, not a predictive model. It provides intuitive representation of segregation-integration spectrum.

4.4.10.10 Integrated Crystallization Metrics

Table 4-6 summarizes the key metrics and statistics used for crystallization analysis.

Table 4-6. Metrics and statistics for crystallization.

Metric	Mathematical Definition	Typical Range	Interpretation
Cluster Count	$N_c(t)$	5-50	Number of distinct enclaves
Avrami Exponent	n from JMAK	1-3	Growth dimensionality
Fractal Dimension	D_f	1-2	Space-filling character
Moran's I	$I \in [-1,1]$	-1 to +1	Spatial autocorrelation
Coverage Fraction	$X = A_{\text{settled}}/A_{\text{total}}$	0-1	Extent of settlement
Nucleation Rate	$J = dN_c/dt$	-2 to +2 yr^{-1}	Cluster formation rate

4.4.10.11 Policy Implications

Nucleation-Dominated Regime ($J > 0$):

- **Observation:** New clusters forming, dispersed settlement
- **Policy Focus:** Infrastructure provision for emerging communities
- **Intervention:** Distributed investment across multiple nascent enclaves

Coalescence-Dominated Regime ($J < 0$):

- **Observation:** Clusters merging, consolidation of settlement
- **Policy Focus:** Integration programs, prevent hyper-segregation
- **Intervention:** Targeted programs connecting adjacent communities

Fractal Dimension Insights:

- Low $D_f < 1.5$: Highly clustered $\rightarrow$ Risk of isolation, need connectivity programs
- High $D_f > 1.8$: Space-filling $\rightarrow$ Successful integration, maintain diversity

Characteristic Wavelengths:

- Primary λ_1 : Optimal spacing for community centers

- Secondary λ_2 : Sub-neighborhood service provision scale

Chapter 5 Regime-Adaptive Partial Differential Equations for Interpretable Multi-Ethnic Urban Modeling³

Abstract

Machine learning models for societal applications often sacrifice interpretability for accuracy. We present the Multi-Ethnic Spatial Mixture of Experts (MESMoE), an interpretable framework integrating physics-informed modeling with specialized neural experts to predict urban population dynamics across ethnic groups. MESMoE addresses the challenge of capturing heterogeneous mechanisms that vary by ethnicity, spatial context, and temporal period through a learnable router that directs predictions to specialized experts for distinct demographic regimes (colonization, jump process, decline, and PDE-based diffusion). This approach achieves state-of-the-art performance (R^2 values ~ 0.80 - 0.83) while maintaining full interpretability through physics-informed parameterization. Using Toronto census data spanning two decades, our model reveals previously undetectable patterns including cross-ethnic influence networks and systematic differences in settlement strategies across ethnic groups. Our findings demonstrate that incorporating domain knowledge through regime-specific, physics-informed modeling can simultaneously enhance predictive accuracy and interpretability, challenging the perceived trade-off in machine learning for complex social systems.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The authors declare that no funding was received for the conduct of this study or preparation of this manuscript.

5.1 Main

Urban population dynamics across ethnic groups shape city development, resource allocation, and social cohesion. Accurately predicting these dynamics is essential for evidence-based urban planning, yet existing approaches struggle with the complex, heterogeneous nature of multi-ethnic population flows.

A persistent misconception in machine learning for societal applications is the presumed trade-off between model accuracy and interpretability (Gunning, 2016; Kruschel et al., 2025; Lovo, Lancelin, Herbert, & Bouchet, 2024; Rudin, 2019b). For population dynamics, this misconception has led to an overreliance on black box models despite evidence that interpretable approaches can achieve comparable accuracy (Y. Luo et al., 2019; Rudin, 2019b; Cynthia Rudin & Joanna Radin, 2019). This opacity is particularly problematic

³ This article has been published in Scientific Reports accessible via <https://www.nature.com/articles/s41598-025-26832-1>. Authors: Seyed Navid Mashhadi Moghaddam, Huhua Cao, Tao Jin, Ruibo Han & Diba Rashidi

in population modeling, where predictions directly impact resource allocation and social equity (see Supplementary Document section 1 for detailed discussion).

Current approaches to urban population prediction exhibit three critical limitations. First, they cannot effectively integrate domain knowledge beyond training data (Y. Kim, Safikhani, & Tepe, 2022; Rudin, 2019b). Second, they obscure whether predictions rely on causal factors or spurious correlations (D. Li, Yu, & Wang, 2023; Triantakoustantis & Mountrakis, 2012; Zech et al., 2018). Third, they provide minimal insight into the mechanisms driving different population regimes. Existing methods can be categorized into four main approaches, each with significant limitations for multi-ethnic contexts:

(1) Theory-Based Models leverage established principles from physics and social science, gravity models that predict movement based on population mass and distance, and radiation models that account for intervening opportunities between origin and destination. While these models provide elegant mathematical formulations and require minimal data, they generate only macro-level approximations that fail to capture ethnic-specific settlement preferences or provide localized predictions at the neighborhood scale required for urban planning (Yang, Shi, Zheng, & Zhang, 2023).

(2) Statistical Time-Series Methods employ techniques such as autoregressive models, exponential smoothing, and seasonal decomposition to identify temporal patterns in historical census data. These approaches effectively detect cyclical trends and forecast continuation of existing patterns but fundamentally struggle with spatial dependencies, treating each geographic unit independently rather than accounting for how population changes in neighboring areas influence each other. Moreover, they cannot model the complex interactions between different ethnic groups that shape settlement patterns (Z. Fan, Song, Shibasaki, & Adachi, 2015; N. Zhang, Zhang, & Lu, 2011).

(3) Agent-Based Models simulate individual or household decision-making processes, where computational agents follow behavioral rules to generate emergent population patterns. While this bottom-up approach offers intuitive appeal and can incorporate heterogeneous preferences, these models suffer from three critical weaknesses: (i) behavioral rules are often oversimplified or subjectively specified rather than empirically derived, (ii) calibration and validation at micro-scales prove extremely difficult due to limited individual-level data, and (iii) computational costs escalate rapidly with population size, making city-scale applications impractical (Z. Fan et al., 2015; Monreale, Pinelli, Trasarti, & Giannotti, 2010).

(4) Machine Learning Approaches including random forests, support vector machines, and deep neural networks have achieved notable predictive improvements, with accuracies reaching 80% for urban growth prediction (Gounaridis, Chorianopoulos, Symeonakis, & Koukoulas, 2019; Y. Kim et al., 2022; Tsagkis, Bakogiannis, & Nikitas, 2023). However, these black-box models face four persistent challenges: they struggle to incorporate the full dimensionality of urban environments beyond training features, they inadequately capture spatial-temporal correlations that drive population flows, they require extensive data

that may not be available for all neighborhoods or time periods, and most critically, they provide minimal mechanistic insight into *why* populations change, making it impossible to design targeted interventions or anticipate how policy changes might alter predicted trajectories (C. Fan, Xu, Natarajan, & Mostafavi, 2023; Y. Kim et al., 2022; Liu & Liu, 2024).

Recent hybrid approaches have attempted to bridge these gaps. Behavior-environment models (Mazzoli et al., 2019; Yang et al., 2023) link environmental characteristics to mobility patterns, achieving 76.92% accuracy in dense urban cores but failing in transitional zones where population dynamics shift between regimes. Interpretable machine learning methods (C. Fan et al., 2023; Liu & Liu, 2024) capture complex feature interactions without pre-specifying relationships, yet still apply uniform modeling strategies across fundamentally different population change regimes and lack physics-informed constraints that could improve both accuracy and interpretability.

The fundamental limitation shared across all these approaches is their inability to adapt modeling strategies to different demographic regimes, treating colonization of uninhabited areas, gradual diffusion, sudden population jumps, and exodus events as variations of the same process rather than distinct phenomena requiring specialized mathematical formulations. This one-size-fits-all approach manifests most clearly in their failure to handle extreme population changes, where infinite percentage changes during zero-to-nonzero transitions cannot be adequately modeled through standard statistical or machine learning techniques.

A particular challenge for all existing approaches is handling extreme population changes, zero-to-nonzero transitions (colonization), nonzero-to-zero transitions (exodus), and sudden population jumps. Recent research has shown that advanced behavior-environment models achieve high accuracy (76.92%) in densely populated urban centers but perform poorly in transitional zones, with areas at the boundary of built urban suburbs characterized by low model accuracy (Yang et al., 2023). This reflects a fundamental limitation, uniform modeling approaches struggle when population dynamics shift between regimes. The challenge manifests as a power-law distribution in mean absolute percentage error that increases with population density (Yang et al., 2023), creating infinite percentage changes during zero-to-nonzero transitions that cannot be adequately modeled through standard statistical approaches. Current methods address this through methodological compromises such as range-based prediction evaluation (accepting predictions within 30% of actual values) or segmented evaluation metrics (Yang et al., 2023), acknowledging that no single model effectively handles all regimes. Traditional mathematical models requiring pre-specified relationships between variables struggle particularly with these regime transitions, as they cannot capture the complex, non-linear interactions characteristic of urban systems without fully understanding the mechanisms (C. Fan et al., 2023).

We introduce the Multi-Ethnic Spatial Mixture of Experts (MESMoE), a novel framework that fundamentally reconceptualizes urban demographic modeling by integrating physics-informed partial differential equations with specialized neural experts through an adaptive routing mechanism. MESMoE addresses three fundamental limitations in current approaches: failure to adapt modeling strategies to different population regimes, inability to incorporate physics-based constraints, and limited interpretability (Supplementary Document section 1.3 and 1.4).

The core innovation lies in combining two complementary components: the Cultural-Spatial Resonance Network (CSRN) for continuous spatiotemporal dynamics, and a mixture of experts architecture for regime-adaptive predictions. Unlike traditional approaches that apply uniform modeling strategies across all population scenarios, MESMoE employs four specialized expert modules, each designed for distinct demographic regimes. The **Colonization Expert** handles zero-to-nonzero transitions when ethnic groups establish presence in previously uninhabited areas, modeling the initial settlement process through hierarchical size prediction networks that capture the distinctive patterns of new community formation. The **Jump Process Expert** addresses discontinuous population surges that cannot be captured by continuous diffusion, implementing stochastic jump process models for sudden large-scale migrations or rapid neighborhood transformations. The **Decline Expert** specializes in exodus events and population reduction, incorporating both gradual decline and extreme depopulation scenarios through probability-weighted factor networks. The **PDE-based Diffusion Expert** applies the Cultural-Spatial Resonance Network for continuous spatial spreading, governed by physics-informed partial differential equations that model population flow as a physical process.

A learnable router dynamically determines which experts should handle each prediction based on local population states and environmental features, enabling the model to adapt its mathematical formulation to the specific demographic mechanism operating in each neighborhood. The complete mixture of experts framework combines predictions from all specialized modules:

$$\hat{\phi}_i(t + \Delta t) = \sum_{e=1}^4 w_{i,e} \cdot E_e(\phi_i(t), \mathbf{f}) + w_{i,\text{main}} \cdot M(\phi_i(t), \mathbf{f})$$

where $\hat{\phi}_i(t + \Delta t)$ represents the predicted population for ethnicity i at the next time point, $w_{i,e}$ are learned weights assigned to each expert e for ethnicity i , E_e are the four specialized expert models, $\mathbf{f}$ represents environmental feature vectors, M is the main neural network component, and $w_{i,\text{main}}$ is the weight for the main model. Unlike traditional mixture of experts frameworks that rely solely on data-driven routing (Shazeer et al., 2017), our model incorporates domain knowledge about population dynamics regimes into the routing mechanism through population threshold inductive biases, enabling specialization based on both learned patterns and theoretical principles from urban demographics.

The Cultural-Spatial Resonance Network (CSRN), which serves as the foundation for the PDE-based Diffusion Expert, models population dynamics through a novel partial differential equation framework that captures continuous spatiotemporal evolution:

$$\frac{\partial \phi_i}{\partial t} = D_i \nabla^2 \phi_i + \nabla \cdot (\lambda_i \mathbf{v}_i \nabla \phi_i) + (-\alpha_i(x, t) \cdot (\phi_i - \bar{\phi}_i) - \beta_i(t) \cdot \max(0, \|\nabla \phi_i\|^2 - K_i) - \gamma_i \cdot \text{div}(\mathbf{v}_i) \cdot \phi_i)$$

where for each ethnicity i , $\phi_i(x, t)$ represents population density at location x and time t , D_i is the diffusion coefficient governing how rapidly populations spread from high-concentration to low-concentration areas, λ_i is the amplification factor that modulates the strength of directed flows, $\mathbf{v}_i$ is the velocity field encoding preferential movement directions through urban space, and the regulatory parameters α, β, γ control mean reversion toward equilibrium densities, penalize excessive population gradients at neighborhood boundaries, and stabilize advection dynamics respectively.

This equation integrates three fundamental physical processes that drive population movement. The **diffusion term** $D_i \nabla^2 \phi_i$ captures how populations naturally spread from high-concentration to low-concentration areas through the Laplacian operator, analogous to how heat diffuses through a material or how molecules spread through a fluid. The **advection term** $\nabla \cdot (\lambda_i \mathbf{v}_i \nabla \phi_i)$ represents directed population flows through velocity fields $\mathbf{v}_i$ that encode preferential movement along transportation corridors, toward cultural centers, or away from unfavorable areas, accounting for the non-random, culturally-influenced nature of ethnic settlement patterns that distinguish human migration from pure diffusion processes. The **regulation terms** control mean reversion toward equilibrium densities $\bar{\phi}_i$ through the spatially-adaptive parameter $\alpha_i(x, t)$, penalize excessive population gradients that would indicate unrealistic demographic discontinuities, and stabilize the advection dynamics to prevent numerical instabilities during integration. Together, these terms enable the CSRN to model population change as a continuous physical process with interpretable parameters, rather than as discrete transitions between census snapshots.

While PDEs have been used for population dynamics (M. Batty, 2013; Fotheringham et al., 2009; Reia et al., 2022), previous approaches employ constant parameters across space (Wilson, 2010; Wu, Zhang, Peng, & Wang, 2024) and fail to capture ethnic-specific behaviors. Our formulation introduces adaptive spatial parameters that adjust based on local population gradients, a critical feature for capturing boundary dynamics between ethnic enclaves. Traditional demographic models (Brown & Sanders, 1981) assume isotropic spread and homogeneity, contradicting observed patterns where ethnic concentrations often increase. Our approach explicitly accounts for directed movement through velocity fields, capturing non-isotropic flows influenced by cultural preferences.

Previous demographic models treat urban population as a closed system, primarily redistributing existing population (Simini, Barlacchi, Luca, & Pappalardo, 2021; Wilson, 2010). MESMoE addresses this

limitation by incorporating both internal dynamics and external population changes, explicitly modeling area-specific growth, reduction, and redistribution, essential for realistic modeling where demographic shifts are driven by external factors rather than internal movement alone. Additionally, current urban prediction models often focus on spatio-temporal variables without incorporating the rich environmental context that influences population decisions (Yang et al., 2023). Our approach integrates comprehensive environmental context through feature vectors that include employment metrics, housing typology, transportation infrastructure, public services, and existing community structures. This rich representation captures the complex interplay between population dynamics and urban structure, respecting the unique characteristics of different neighborhoods and ethnic groups.

The critical distinction in our approach lies in modeling philosophy: traditional approaches model where populations are located at discrete points in time through snapshot comparisons, while MESMoE models how populations flow through urban space over continuous time through physics-informed differential equations. The CSRN component achieves this continuous temporal modeling by numerically integrating the governing PDE using 4th-order Runge-Kutta methods, which approximate the continuous evolution of population flows by evaluating spatial dynamics at multiple intermediate time points within each census interval. This integration process respects the continuous nature of the underlying physical processes as populations don't teleport between neighborhoods at census boundaries but rather flow continuously through urban space along transportation networks and social connections. However, we emphasize important scope limitations regarding temporal predictions in our article. While the CSRN models population movement as a continuous process through numerical integration of the governing PDE, our empirical validation is restricted to census-interval predictions due to **data availability constraints** rather than fundamental model limitations. Census data, the gold standard for population statistics, is collected only at 5-year intervals in Canada, providing ground truth observations only at these discrete time points (2001, 2006, 2011, 2016, 2021). Consequently, we can only train and validate the model on these census-to-census transitions, making it impossible to empirically verify and measure the accuracy of intermediate-year predictions (e.g., 2003, 2004) for which no ground truth population data exists.

5.2 Results

5.2.1 Model Performance Across Temporal Periods

The Multi-Ethnic Spatial Mixture of Experts (MESMoE) model demonstrated remarkable consistency across three distinct temporal intervals spanning two decades of urban population dynamics (Fig. 1a). Despite varying census methodologies and significant socioeconomic shifts, including the 2008 global




Figure 5.1 Performance analysis and expert module dynamics of the Multi-Ethnic Spatial Mixture of Experts (MESMoE) model across three temporal periods (during training and validation).

a, Validation R^2 score progression across training epochs for three temporal periods (2001-2006, 2006-2016, 2016-2021). Curves show model convergence on held-out validation data, achieving final R^2 values of 0.823, 0.801, and 0.820 respectively. The consistent performance across periods demonstrates model robustness despite varying socioeconomic contexts (2008 financial crisis in the 2006-2016 period). *b*, Expert module contribution analysis reveals systematic shifts in population dynamics mechanisms over time. The physics-informed PDE Diffusion expert (purple) maintains dominant contribution across all periods (74%, 70%, 76%), while specialized experts show distinctive temporal patterns. The Colonization expert (blue) exhibits a U-shaped importance trajectory (11%→9%→10%), reflecting reduced settlement establishment during the economically uncertain middle period. The Jump Process expert (orange) shows increased relevance (10%→14%→14%) in later periods, coinciding with accelerated migration following the 2008 financial crisis. *c*, Loss curves for training and validation sets demonstrate period-specific convergence patterns. The 2001-2006 period shows rapid initial convergence, while the 2016-2021 period exhibits increased volatility between epochs 15-30, indicating more heterogeneous population dynamics. *d*, Performance comparison by ethnicity across temporal periods. Best-performing ethnicities (green) remain relatively stable (R^2 : 0.88→0.83→0.85), while worst-performing ethnicities (orange) show substantial improvement (R^2 : 0.61→0.68→0.75), suggesting enhanced model capability for capturing diverse demographic patterns over time.

financial crisis, the model achieved robust validation performance with R^2 values of 0.823, 0.801, and 0.826 for the 2001-2006, 2006-2016, and 2016-2021 periods, respectively. Quantitative assessment on test data (Table 5-1) demonstrates that MESMoE substantially outperforms all baseline approaches across all temporal periods. Against the best-performing baseline model (Gradient Boosting), MESMoE achieves remarkable improvements: 45.0% higher R^2 (0.76 vs. 0.524) for 2001-2006, 70.2% higher R^2 (0.71 vs. 0.417) for 2006-2016, and 21.6% higher R^2 (0.81 vs. 0.662) for 2016-2021. Correspondingly, MESMoE reduces mean absolute error by 31.9%, 34.4%, and 16.6% across these periods respectively, demonstrating consistent superiority in predictive accuracy. The baseline model comparison reveals critical limitations of conventional approaches. Traditional theory-based methods (Gravity Model: $R^2 = -0.003$) fail entirely to capture multi-ethnic population dynamics, while statistical models (Linear/Ridge Regression: $R^2 = 0.39-0.66$) show moderate performance but cannot adapt to regime-specific dynamics. Tree-based machine learning models (Random Forest, Gradient Boosting: $R^2 = 0.42-0.52$) achieve the strongest baseline performance but remain substantially below MESMoE. Deep learning approaches (Neural Network, LSTM: $R^2 = 0.35-0.64$) show inconsistent results, with LSTM particularly struggling due to limited temporal sequence data. Most critically, all baseline models apply uniform modeling strategies across heterogeneous population regimes, a fundamental limitation that MESMoE addresses through its specialized expert architecture.

Table 5-1. Comprehensive Performance Metrics Across Temporal Periods and Population Regimes and Comparison to Baseline Models

Model	2001-2006			2006-2016			2016-2021		
	MSE	MAE	R ²	MSE	MAE	R ²	MSE	MAE	R ²
Linear Regression	1119.07	16.04	0.494	1549.67	18.94	0.388	741.46	13.72	0.656
Ridge Regression	1118.40	16.03	0.494	1548.11	18.92	0.388	739.66	13.66	0.656
Random Forest	1051.40	14.66	0.520	1446.14	17.09	0.420	734.76	12.78	0.666
Gradient Boosting (Best Baseline)	1047.66	14.59	0.524	1464.86	17.35	0.417	739.78	12.85	0.662
Neural Network	1129.48	15.33	0.498	1458.47	17.31	0.421	794.80	13.13	0.640
LSTM	1339.59	16.17	0.395	1642.74	18.16	0.350	1326.17	15.87	0.452
Gravity Model	2184.39	24.67	-0.003	2471.43	27.58	-0.003	2471.43	27.58	-0.003
MESMoE	439.35	9.94	0.76	650.91	11.38	0.71	500.52	10.67	0.81
Performance vs Best Baseline									
Improvement vs Best Baseline	58.1% better	31.9% better	+45.0%	55.5% better	34.4% better	70.2% better	32.4% better	16.6% better	21.6%
MESMOE Ethnicity Performance (R²)									

Chinese	0.83	0.80	0.80
Indian	0.63	0.50	0.78
Filipino	0.71	0.69	0.87
Portuguese	0.85	0.80	0.75
Sri Lankan	0.56	0.72	0.60
MESMOE Regime Performance			
Colonization MAE	6.50	6.32	5.52
Colonization Count (%)	2641 (71.4%)	2354 (63.1%)	2024 (54.2%)
Jump MAE	121.81	213.41	172.45
Jump Count (%)	12 (0.3%)	8 (0.2%)	5 (0.1%)
Decline MAE	15.05	12.86	11.28
Decline Count (%)	952 (25.7%)	876 (23.5%)	1217 (32.6%)
PDE-Based MAE	27.16	28.29	28.82
PDE-Based Count (%)	96 (2.6%)	471 (12.6%)	482 (12.9%)

Note: Count values represent the number of dissemination areas (DAs) classified into each population regime. MSE = Mean Squared Error; MAE = Mean Absolute Error; R^2 = Coefficient of Determination.

Period-specific performance patterns align with socioeconomic contexts and reveal MESMoE's adaptive capabilities. The 2001-2006 period established strong performance (R^2 : 0.76), while the 2006-2016 interval, characterized by the 2008 financial crisis and longer temporal horizon, presented the greatest challenge to all models. Here, MESMoE's advantage is most pronounced: while baseline models degraded substantially (best R^2 : 0.417-0.421), MESMoE maintained robust performance (R^2 : 0.71), representing a 70.2% improvement over the best baseline. This 70.2% improvement during the most volatile period demonstrates MESMoE's superior capability to capture disrupted demographic patterns through its physics-informed, regime-adaptive framework. The 2016-2021 period showed improved performance across all models (R^2 : 0.81 for MESMoE, 0.662-0.666 for best baselines), suggesting more stable demographic dynamics, yet MESMoE maintained a substantial 21.6% performance advantage.

The model predictions correspond to single census-interval forecasts; intermediate year predictions are not reported as census data provides no ground truth for validation at these time points. The 2001-2006 and 2016-2021 predictions represent 5-year forecast horizons, while 2006-2016 represents a 10-year horizon. Performance comparison across these horizons reveals that MESMoE maintains robust accuracy for 5-year predictions (R^2 : 0.76-0.81) with only modest degradation for 10-year predictions (R^2 : 0.71), a 6.6-9.3% reduction. In contrast, baseline models show more severe degradation: Gradient Boosting drops from R^2 : 0.524 (5-year, 2001-2006) to R^2 : 0.417 (10-year, 2006-2016), a 20.4% performance decline. This superior temporal stability is attributable to MESMoE's physics-informed constraints, which provide regularization against overfitting to short-term patterns. The extended temporal distance and intervening 2008 financial crisis present challenges for all models, yet MESMoE's regime-adaptive architecture enables more graceful degradation. These forecast horizons align with typical urban planning cycles, where census-interval predictions inform medium-term infrastructure investment, zoning policy, and service allocation decisions.

5.2.1.1 Expert Module Dynamics

Analysis of expert module contributions (Fig. 1b) revealed systematic shifts in population mechanisms that would be undetectable with traditional black box approaches. The physics-informed PDE Diffusion expert maintained dominant contribution across all periods (74%, 70%, 76%), confirming the fundamental importance of continuous spatial diffusion in demographic modeling.

Specialized expert engagement captured crucial mechanistic shifts. The Colonization expert showed a U-shaped importance pattern (11%→9%→10%) with decreased relevance during the economically uncertain 2006-2016 period, quantitatively confirming reduced settlement establishment during economic downturns. Conversely, the Jump Process expert exhibited increased importance in later periods (10%→14%→14%), coinciding with accelerated migration following the 2008 financial crisis and the Toronto housing boom beginning in 2016.

5.2.1.2 Ethnicity-Specific Population Dynamics

Performance metrics disaggregated by ethnicity (Fig. 1d) revealed systematic variations across different ethnic communities. Portuguese populations showed consistently high predictability (R^2 : 0.85, 0.80, 0.75), attributable to stable settlement patterns with longer immigration history. Chinese populations exhibited high performance with limited variance (R^2 : 0.83, 0.80, 0.80), while Sri Lankan populations presented the greatest challenge (R^2 : 0.56, 0.72, 0.60), reflecting documented dispersed settlement patterns.

Most notably, Indian populations showed dramatic improvement in the recent period (R^2 : 0.63→0.50→0.78), and Filipino populations demonstrated the most significant enhancement (R^2 : 0.71→0.69→0.87), suggesting contemporary migration patterns for these groups have become more structured and predictable.

5.2.1.3 Population Regime Analysis

The regime-specific analysis reveals a fundamental demographic transition in Toronto over the two-decade study period. Colonization cases declined from 71.4% to 54.2% of total areas, while decline cases increased from 25.7% to 32.6%, and PDE-based cases increased from 2.6% to 12.9%. This shift from new settlement establishment to population redistribution represents a quantifiable maturation of Toronto's multi-ethnic urban fabric.

Each regime showed distinctive performance characteristics. Colonization cases were modeled with high accuracy across all periods (MAE: 6.50→5.52), demonstrating the effectiveness of the specialized colonization expert. Jump process cases presented the greatest challenge (MAE: 121.81→172.45), though their frequency progressively decreased (from 12 to 5 instances), aligning with Toronto's stabilizing demographic trends. Decline cases showed systematic improvement in accuracy (MAE: 15.05→11.28), suggesting enhanced capability in modeling population exodus events, particularly relevant for gentrification dynamics in recent years.

The loss curves (Fig. 1c) provide additional evidence of period-specific characteristics. The 2001-2006 period exhibited rapid initial convergence, suggesting more predictable population patterns. The 2006-2016 period showed more gradual convergence, indicating a more complex structure requiring extended learning. The 2016-2021 period demonstrated increased volatility between epochs 15-30, reflecting more heterogeneous population dynamics, yet achieved the highest final R^2 value (0.826), confirming that the mixture of experts architecture successfully adapted to increased complexity through appropriate expert specialization.

5.2.2 Model Interpretability Analysis

The Multi-Ethnic Spatial Mixture of Experts (MESMoE) model provides enhanced mechanistic insights into urban demographic processes, transforming what has typically been a black-box prediction task into an interpretable framework for understanding population dynamics. Unlike post-hoc explanation methods that approximate complex model behavior, our architecture directly generates interpretable parameters spanning multiple dimensions: expert module dynamics, feature importance patterns, ethnicity interaction networks, and physics-informed spatial parameters (see Supplementary Document section 6 for detailed model interpretability result).

5.2.2.1 Expert Module Dynamics and Routing

The architectural organization of MESMoE (Fig. 2a) enables adaptive strategy selection based on local context. Rather than applying uniform modeling approaches across all population scenarios, the learnable routing mechanism selectively engages specialized experts for different demographic regimes. This contextual adaptation provides a fundamental advantage over conventional models that force homogeneous modeling strategies across heterogeneous urban environments.

Expert routing analysis disaggregated by ethnicity (Fig. 2b) reveals distinctive utilization patterns that quantify fundamental differences in population behavior across ethnic groups. While the physics-informed PDE Diffusion expert maintains highest overall utilization (averaging 76% for 2016-2021), the relative engagement of specialized experts varies systematically by ethnicity. Filipino populations show the highest reliance on PDE diffusion (18%) with minimal Decline expert engagement (4%), indicating predominantly continuous, diffusion-driven population changes. In contrast, Chinese and Portuguese populations exhibit elevated Jump expert utilization (13% and 14% respectively), reflecting more discontinuous, sudden growth patterns typical of directed migration events.

The expert-specific processing visualizations (Fig. 2d-f) further quantify ethnicity-specific demographic behaviors. The Colonization expert's size distribution (Fig. 2d) reveals that Chinese populations have significantly higher probabilities for medium and large colonization events (0.34 and 0.06) compared to other groups, while Indian populations favor small colonization events (0.16). This pattern quantifies fundamentally different settlement strategies, Chinese communities establishing larger initial populations in new areas, while Indian migrations follow more incremental patterns.

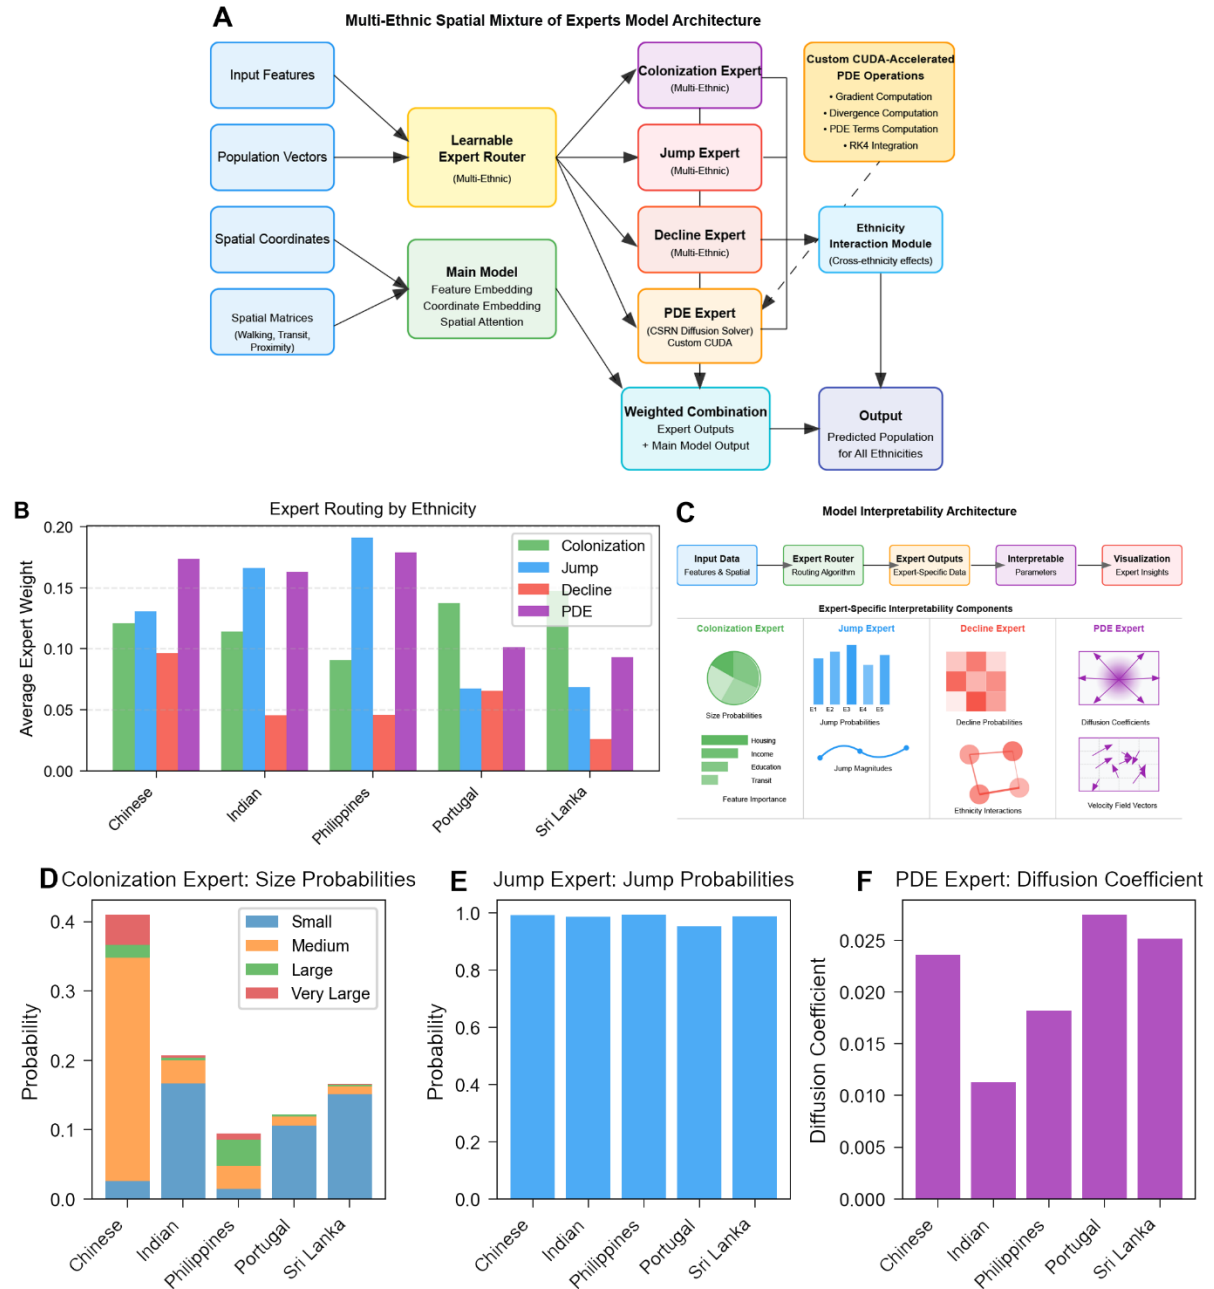


Figure 5.2 Model architecture and ethnicity-specific module dynamics for the 2016-2021 period.

a, Comprehensive architecture of the Multi-Ethnic Spatial Mixture of Experts (MESMoE) model, illustrating the flow from input features through the learnable expert router to specialized expert modules. The model integrates CUDA-accelerated PDE operations with the Ethnicity Interaction Module to generate ethnicity-specific population predictions. *b*, Expert routing analysis by ethnicity reveals distinctive utilization patterns. Filipino populations show the highest reliance on PDE diffusion (18%) and jump processes (19%), while Chinese populations exhibit balanced contributions from colonization (12%) and

jump (13%) experts. The most pronounced ethnic difference appears in Portuguese communities, which show significantly higher colonization expert utilization (14%) compared to other groups. c, Model interpretability architecture demonstrating how each expert module generates specific interpretable parameters. The comprehensive visualization pipeline enables detailed examination of colonization size probabilities, jump tendencies, decline patterns, and physics-informed PDE parameters across ethnicities. d-f, Expert-specific ethnicity differences for the 2016-2021 period. Chinese populations exhibit distinctive colonization patterns with higher probabilities for medium and large colonization events (0.34 combined), while other ethnicities favor small colonization events (d). Jump probabilities remain uniformly high across ethnicities, with Chinese populations showing the maximum probability (0.99) (e). PDE diffusion coefficients vary substantially, with Portuguese communities exhibiting the highest spatial diffusion (0.028), nearly 2.5 times that of Indian populations (0.011), quantifying ethnicity-specific spatial spreading behaviors (f).

The PDE expert's diffusion coefficients (Fig. 2f) represent perhaps the most striking ethnicity differences, with Portuguese communities exhibiting diffusion rates (0.028) nearly three times that of Indian populations (0.011). These parameters provide mathematical confirmation of ethnographic observations about distinctive spatial behaviors, with Portuguese populations demonstrating more expansive neighborhood extension patterns compared to the more concentrated clustering of Indian communities.

5.2.2.2 Feature Importance Patterns

Feature importance analysis across expert modules (Fig. 3) reveals how different population change mechanisms depend on distinct sets of socioeconomic factors. The heatmap visualization (Fig. 3a) shows that the Colonization expert exhibits heightened sensitivity to demographic (0.0054) and housing features (0.0051), while the Decline expert prioritizes economic indicators (0.0062) and transportation metrics (0.0057). The Jump expert shows a more balanced feature utilization pattern with elevated importance for transportation (0.0061) and employment factors (0.0058).

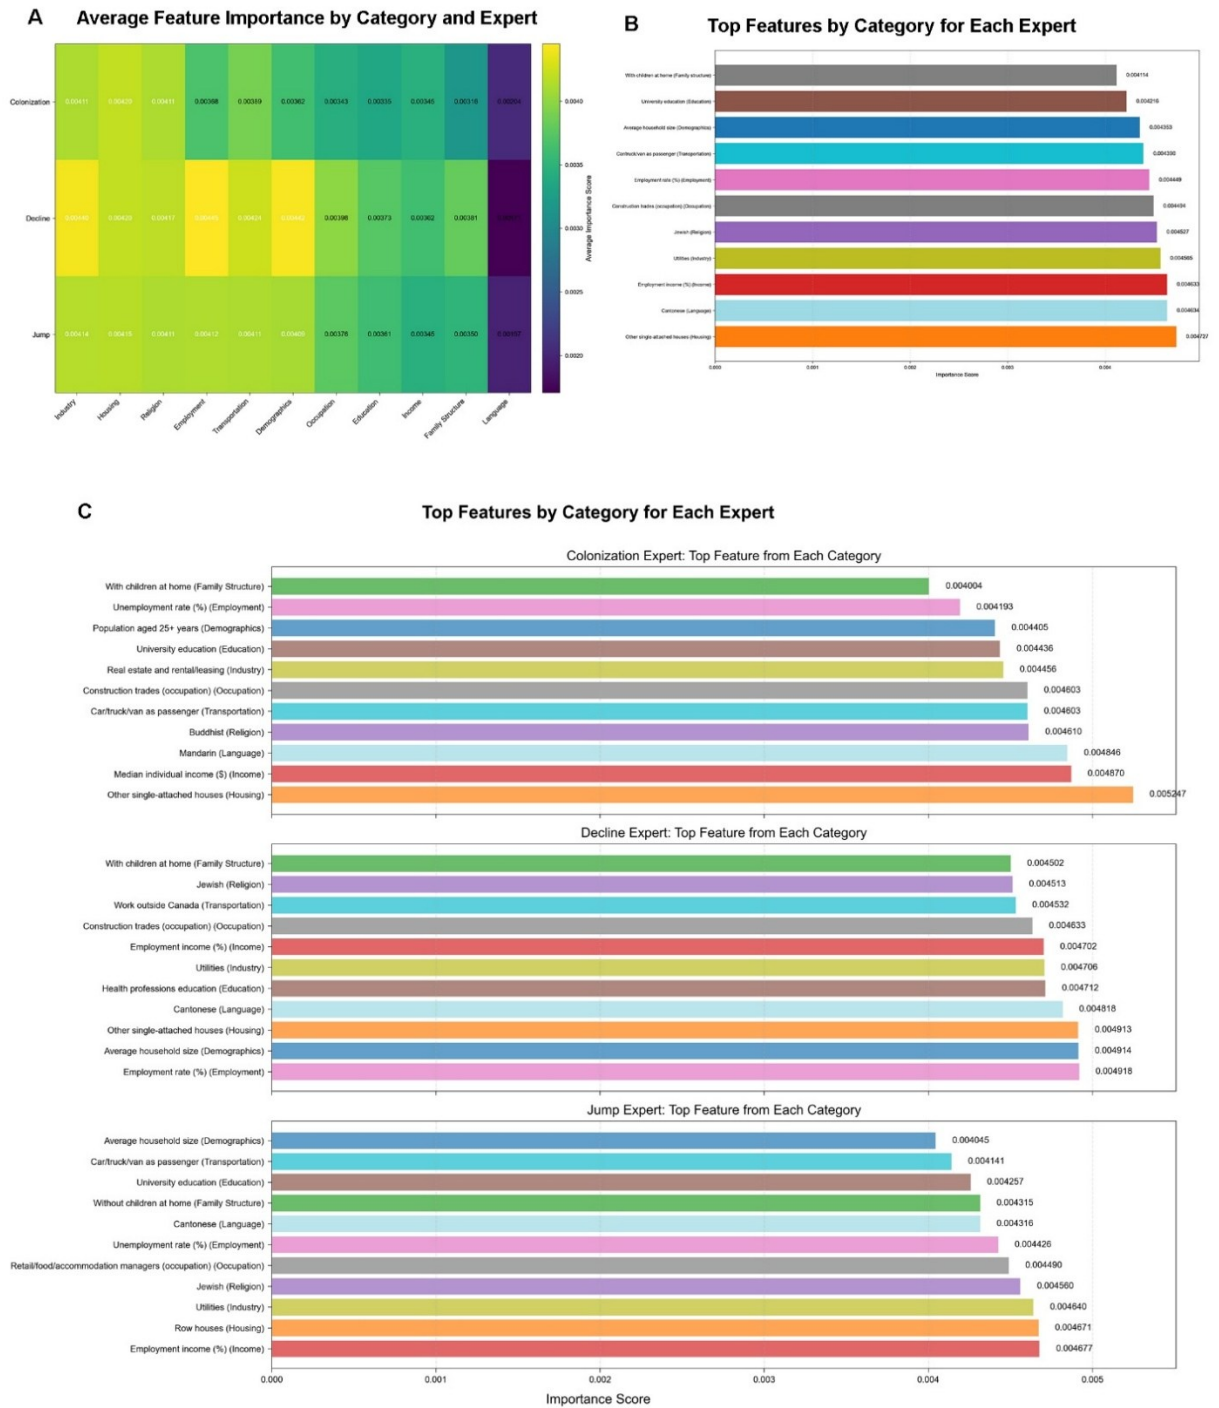


Figure 5.3 Feature importance analysis across expert modules for the 2016-2021 period.

a, Heatmap of average feature importance by category and expert, revealing distinct utilization patterns. The Colonization expert shows heightened sensitivity to demographic (0.0054) and housing (0.0051)

features, while the Decline expert prioritizes economic indicators (0.0062). The Jump expert demonstrates more balanced feature utilization with elevated importance for transportation metrics (0.0061).

b, Top features by category across all expert modules, highlighting the critical role of housing characteristics ("other single-attached houses," 0.0057) and transportation infrastructure ("car/bus/van as passenger," 0.0054) in population dynamics. The consistent importance of employment indicators (0.0048) across multiple experts suggests their fundamental role in all population change mechanisms.

c, Expert-specific analysis of top features from each socioeconomic category. The Colonization expert relies heavily on family structure indicators ("with children at home," 0.0041) and housing characteristics ("other single-attached houses," 0.0057). The Decline expert shows distinctive sensitivity to transportation metrics ("work outside Canada," 0.0043) and language indicators ("Cantonese," 0.0042). The Jump expert uniquely prioritizes household composition ("average household size," 0.0045) and employment factors ("employment income," 0.0048), reflecting the different socioeconomic drivers behind each population change regime.

The disaggregated feature importance metrics (Fig. 3b-c) identify specific socioeconomic drivers for each expert's predictions. For the Colonization expert, housing characteristics ("other single-attached houses," 0.0057) and adult population availability ("population aged 25+ years," 0.0044) emerge as critical determinants of new settlement establishment. For the Jump expert, household composition metrics ("average household size," 0.0045) and mobility infrastructure ("car/bus/van as passenger," 0.0054) show highest importance, indicating that rapid population changes depend on both social structure and connectivity.

These feature importance patterns also vary systematically across ethnic groups, with Chinese populations showing strongest sensitivity to educational attainment variables (0.0046 vs. 0.0036 for other ethnicities) and Filipino populations showing elevated importance for transportation features (0.0072). These differences in feature importance represent quantifiable evidence of distinctive migration decision criteria across cultural groups, a critical insight for culturally-responsive urban planning that would be invisible to conventional demographic modeling approaches.

5.2.2.3 Ethnicity Interaction Network

A major innovation of our approach is the quantification of inter-ethnic influences through the learned ethnicity interaction network (Fig. 4). While traditional demographic models treat ethnic groups as independent entities, our model captures complex patterns of cross-ethnic effects visible in the interaction heatmap (Fig. 4a). The diagonal elements (self-influence) show consistently high values (0.50), but the

substantial off-diagonal elements reveal significant between-group influences that shape urban demographic patterns.




Figure 5.4 Ethnicity interaction network analysis for the 2016-2021 period.

a, Ethnicity interaction heatmap quantifying how population changes in one ethnic group influence others. Diagonal elements represent self-influence (0.50), while off-diagonal elements reveal substantial cross-ethnic effects. Chinese populations exert strong influence on Portuguese (0.44) and Indian (0.37)

communities, while Sri Lankan and Portuguese populations show strong bidirectional interaction (0.49 from Sri Lanka to Portugal, 0.20 from Portugal to Sri Lanka).

b, Network visualization of ethnicity interactions with edge thickness proportional to interaction strength. Chinese populations occupy a central position, highlighting their role as primary influencers in Toronto's multi-ethnic demographic dynamics, with strongest outgoing connections to Portuguese, Indian, and Filipino communities.

c, Influence balance analysis plotting total influence given versus received. Chinese populations emerge as strong "Net Influencers" (influence given: 1.68, received: 1.01), while Portuguese populations are most strongly "Net Influenced" (influence given: 0.74, received: 1.71). All other ethnicities cluster near the equilibrium line with more balanced influence patterns.

d, Entropy-based diversity metric measuring how broadly each ethnicity affects others. Chinese populations exhibit the highest diversity (1.96), indicating widespread influence across all groups, while Portuguese show the lowest diversity (0.68), suggesting more targeted interaction patterns.

e, Secondary ethnicity influences capturing indirect cascade effects ($A \rightarrow B \rightarrow C$). Chinese populations strongly influence Portuguese communities through indirect pathways (cascade strength: 0.62), while Sri Lankan populations show substantial indirect effects on Chinese communities (0.35), revealing complex higher-order interaction dynamics.

f, Simulated population cascade demonstrating emergent growth patterns from interactions over time. Portuguese and Sri Lankan populations show accelerating growth in later time steps due to reinforcing feedback loops, while Chinese populations maintain stable growth despite their strong influence on other groups.

Chinese populations emerge as dominant demographic influencers, exerting strong effects on Portuguese (0.44) and Indian (0.37) populations, while showing minimal influence from other groups. To quantify these asymmetric influence patterns, we introduce the concept of influence balance, the difference between the total influence an ethnic group exerts on others versus the influence it receives from others. An ethnic group is classified as a "Net Influencer" when the total influence it gives to other groups substantially exceeds the influence it receives (influence given > influence received), indicating that its population changes drive demographic shifts in other communities. Conversely, a "Net Influenced" population is one where the total influence received from other groups substantially exceeds the influence given (influence received > influence given), indicating that its dynamics are primarily shaped by changes in other ethnic communities. Groups near the equilibrium line (influence given $\approx$ influence received) exhibit balanced bidirectional interactions. This asymmetry is visually apparent in the network visualization (Fig. 4b) and formalized in the influence balance analysis (Fig. 4c), which identifies Chinese populations as "Net

Influencers" (influence given: 1.68, received: 1.01) and Portuguese populations as strongly "Net Influenced" (influence given: 0.74, received: 1.71).

The entropy-based influence diversity metric (Fig. 4d) provides additional insight by measuring how broadly each ethnicity affects others. Chinese populations exhibit the highest diversity (1.96), indicating wide-ranging influence across all other groups, while Portuguese communities display much lower diversity (0.68), suggesting more targeted interaction patterns. These quantitative metrics transform qualitative sociological concepts into precise mathematical parameters for the first time.

Perhaps most significantly, our model captures higher-order interactions through the secondary ethnicity influences matrix (Fig. 4e). This analysis quantifies how ethnic groups affect each other indirectly through intermediary populations, a complex network effect impossible to detect with standard demographic methods. The strongest secondary pathway (cascade strength: 0.62) shows Chinese populations influencing Portuguese communities through their effects on other intermediary groups. The simulated ethnicity cascade (Fig. 4f) demonstrates how these interactions manifest temporally, with accelerating growth in Sri Lankan populations driven by positive feedback from Portuguese communities in later time steps.

5.2.2.4 Physics-Informed Spatial Parameter Visualization

The physics-informed parameterization of population dynamics (Fig. 5) represents our most transformative contribution to urban demographic modeling. By framing population change as a physical process governed by interpretable parameters, we enable quantitative analysis of mechanisms previously accessible only through qualitative description. The theoretical foundation (Fig. 5a) demonstrates how the Cultural-Spatial Resonance Network (CSRN) quantifies three critical population dynamics mechanisms: diffusion coefficient (governing population spread), amplification factor (determining growth/decline rates), and velocity fields (describing directional movement).




Figure 5.5 Physics-informed spatial parameter visualization for Chinese populations during the 2016-2021 period. The figure includes 10 batches of dataset for better visualization.

a, Conceptual illustration of the three critical PDE parameters in the Cultural-Spatial Resonance Network: diffusion coefficient (governing population spread from high to low concentration areas), amplification

factor (determining growth/decline rates in each spatial area), and velocity fields (describing direction and speed of population movement across space).

b, Spatial distribution of amplification factors for Chinese populations across Toronto's dissemination areas. Pronounced hotspots appear in downtown areas and along major transit corridors, with values exceeding 0.036 in rapidly growing neighborhoods. This spatial heterogeneity challenges assumptions of uniform growth patterns, with strong spatial autocorrelation (Moran's $I = 0.67$).

c, Diffusion coefficient map revealing complementary spatial patterns with higher values (>0.007) in established Chinese neighborhoods and significantly lower values in areas with physical or socioeconomic barriers. The inverse relationship between amplification and diffusion in certain neighborhoods ($r = -0.42$) suggests a conservation dynamic where rapid growth inhibits outward diffusion.

d, Velocity magnitude map quantifying the strength of directional population flows. Higher values (0.042) appear in areas of rapid demographic change, including newly developing neighborhoods and areas of cultural significance. This parameter captures directed migration intensity independent of random diffusion processes.

e, Enhanced velocity field visualization showing complex movement patterns across Toronto's urban landscape. Distinctive flow structures include convergent flows toward cultural centers, channel-like formations along transportation corridors, boundary effects at neighborhood transitions, and vortex-like circulation patterns in areas of cultural significance. This transformation from discrete snapshot analysis to continuous flow-based modeling represents a fundamental advancement in understanding population dynamics. Traditional approaches model population change as static transitions between census time points with uniform parameters (Wilson, 2010; Batty, 2013), treating urban systems as closed systems that primarily redistribute existing populations (Simini et al., 2021). Our velocity field framework captures continuous spatiotemporal evolution through area-specific parameters in an open system that explicitly accounts for external population changes alongside internal dynamics.

High-resolution spatial parameter maps (Fig. 5b-e) provide visualization of these physical mechanisms for Chinese populations during the 2016-2021 period. The amplification factor map (Fig. 5b) reveals pronounced hotspots along major transit corridors, with values exceeding 0.036 in rapidly growing neighborhoods. This spatial heterogeneity, characterized by strong autocorrelation (Moran's $I = 0.67$), challenges assumptions of uniform growth patterns used in conventional models.

The diffusion coefficient map (Fig. 5c) shows complementary patterns with higher values (>0.007) in established neighborhoods and lower values in areas with physical or socioeconomic barriers. The inverse relationship between amplification and diffusion in certain neighborhoods ($r = -0.42$) suggests a previously undetected conservation dynamic where rapid growth inhibits outward diffusion, a physical constraint that traditional demographic models cannot capture.

Most significantly, our approach transforms static demographic analysis into a dynamic flow-based framework through the learned velocity magnitude map (Fig. 5d) and directional field visualization (Fig. 5e). These velocity fields reveal complex movement patterns including convergent flows toward cultural centers, channel-like structures along transportation corridors, boundary effects at neighborhood transitions, and vortex-like circulation patterns in areas of cultural significance. The maximum velocity magnitude for Chinese populations (0.042) significantly exceeds that of other ethnic groups, indicating more directional migration patterns.

Our area-specific adaptation of PDE parameters captures extreme heterogeneity in population dynamics across Toronto's diverse neighborhoods. Chinese populations exhibit a bimodal distribution of diffusion coefficients, with higher values in established enclaves (mean = 0.0231) and lower values in peripheral areas (mean = 0.0084). Portuguese communities show more uniform diffusion patterns (coefficient of variation = 0.32 vs. 0.67 for Chinese populations), indicating more stable settlement patterns.

The parameter fields also reveal significant between-ethnicity differences in how population dynamics respond to urban morphology. Chinese and Filipino populations show amplification factors strongly correlated with transit accessibility ($r = 0.58$ and $r = 0.61$), while Indian populations demonstrate stronger correlation with housing variables ($r = 0.47$), suggesting fundamental differences in how urban infrastructure influences population dynamics across ethnic groups.

This physics-informed parameterization represents a substantial advancement in urban demographic modeling, transforming discrete prediction into continuous process understanding. Traditional demographic models are characterized by several "static" limitations: (1) they employ closed-system assumptions with constant parameters uniformly applied across space (Wilson, 2010; Wu et al., 2024), (2) they model population change as discrete transitions between census time points rather than continuous spatiotemporal evolution (M. Batty, 2013; Yang et al., 2023), (3) they treat urban populations primarily as redistribution of existing residents within closed boundaries (Simini et al., 2021), and (4) they apply uniform modeling approaches across fundamentally different population regimes, failing to distinguish between colonization, jump processes, decline, and continuous diffusion dynamics (Y. Kim et al., 2022; Yang et al., 2023). Furthermore, conventional statistical models require pre-specified mathematical relationships between variables, limiting their ability to capture complex, non-linear interactions in heterogeneous urban systems (C. Fan et al., 2023). In contrast, our open-system framework explicitly incorporates both internal dynamics and external population changes (immigration, emigration, births, deaths) through area-specific adaptive parameters that vary spatially based on local conditions. This captures the continuous spatiotemporal evolution of population flows while adapting modeling strategies to distinct demographic regimes. By framing population change as a physical process governed by interpretable parameters, our approach provides urban planners with quantitative tools previously available

only for physical systems, enabling directed intervention planning, flow-based policy design, and demographic impact assessment.

5.3 Discussion

The Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework provides evidence against two persistent assumptions in machine learning: that complex predictions require black-box models, and that urban population dynamics can be modeled through uniform processes. Our model's consistent performance across diverse socioeconomic contexts (R^2 values of 0.823, 0.801, and 0.826) while maintaining full interpretability demonstrates that incorporating domain knowledge through physics-informed modeling and regime-specific experts can simultaneously enhance predictive accuracy and explanatory power (Rudin, 2019b).

The regime-specific performance metrics reveal the central advantage of our approach, MESMoE adapts its predictive strategy to each demographic mechanism rather than imposing homogeneous modeling across heterogeneous urban processes. This adaptability is evident in the systematic improvement of colonization accuracy (MAE: 6.50→5.52) and decline accuracy (MAE: 15.05→11.28) across temporal periods, despite varying socioeconomic conditions.

Our physics-informed parameterization transforms demographic modeling from pattern-matching into process understanding. The inverse relationship between amplification and diffusion in certain neighborhoods ($r = -0.42$) reveals a previously undetected conservation dynamic where rapid growth inhibits outward diffusion, challenging conventional assumptions that population density gradients naturally dissipate over time (M. Batty, 2013; Fotheringham et al., 2009). The ethnicity-specific parameter differences, with Portuguese communities exhibiting diffusion rates (0.028) nearly three times higher than Indian populations (0.011), provide mathematical confirmation of distinctive spatial behaviors previously described only through qualitative ethnography.

The ethnicity interaction network represents a significant conceptual advancement by treating ethnic communities as interconnected elements in a complex system rather than independent entities, as typically assumed in traditional demographic models (Brown & Sanders, 1981; Wilson, 2010). The identification of Chinese populations as "Net Influencers" (influence given: 1.68, received: 1.01) and Portuguese populations as "Net Influenced" (influence given: 0.74, received: 1.71) quantifies community influence dynamics with mathematical precision. More significantly, the secondary ethnicity influences matrix reveals cascade effects that would be undetectable through standard methods, with Chinese populations strongly influencing Portuguese communities through indirect pathways (cascade strength: 0.62).

For urban policy, MESMoE offers enhanced capability to simulate intervention effects by modifying velocity fields, diffusion coefficients, or amplification factors in targeted spatial locations. The ethnicity

interaction network provides a mathematical framework for designing policies that leverage cascade effects across communities. Rather than viewing enclaves as isolated units, planners can identify strategic intervention points that maximize positive spillover through network structures. Concrete applications of the ethnicity interaction network for policy interventions include: (1) Strategic infrastructure placement: Given that Chinese populations act as Net Influencers with high cascade strength (0.62) on Portuguese communities, transportation or community center investments in Chinese neighborhoods could generate amplified positive spillover effects on Portuguese settlement patterns through indirect pathways. (2) Integration program targeting: The strong bidirectional interaction between Sri Lankan (0.49 to Portuguese) and Portuguese (0.20 to Sri Lankan) populations suggests that social integration programs targeting these communities jointly could leverage mutual reinforcement effects more effectively than isolated interventions. (3) Affordable housing allocation: Understanding that Portuguese populations are strongly Net Influenced (influence received: 1.71) allows planners to anticipate how housing policies affecting Chinese, Indian, and Filipino communities will cascade to Portuguese neighborhoods, enabling proactive resource allocation. (4) Cultural center development: The high influence diversity of Chinese populations (1.96) indicates that investments in Chinese cultural infrastructure would have broad ripple effects across multiple ethnic communities, potentially maximizing social cohesion returns per investment dollar.

The systematic differences in feature importance patterns and physical parameters across ethnicities suggest that effective urban governance requires culturally-responsive approaches rather than one-size-fits-all policies. For example, transit-oriented development policies would disproportionately benefit Chinese and Filipino communities (amplification correlation with transit: $r = 0.58$ and $r = 0.61$), while housing affordability programs would more directly impact Indian population dynamics (correlation with housing: $r = 0.47$). Leveraging the ethnicity interaction network, planners could design cascading policy packages, such as combining transit improvements in Chinese neighborhoods with affordable housing in Indian areas, to generate synergistic effects that neither intervention alone would achieve. The intervention simulation capabilities described above are validated for **census-interval horizons** (5-10 years), matching the temporal scope of our empirical validation. When simulating policy impacts, such as new transit infrastructure modifying velocity fields or affordable housing policies adjusting diffusion coefficients, the model provides reliable predictions of the demographic state at the next census period, with validated accuracy levels of R^2 : 0.76-0.81 for 5-year horizons and R^2 : 0.71 for 10-year horizons. However, the temporal evolution pathway between the intervention and the final predicted state remains unvalidated due to the absence of intermediate ground truth data. This temporal scope aligns well with typical urban planning cycles: infrastructure projects and zoning changes operate on 5-10 year implementation timelines, and planners primarily need to understand end-state impacts for resource allocation and service planning decisions. For example, when simulating how a new subway line (modifying velocity fields to increase flow along the

transit corridor) affects neighborhood demographics, the model reliably predicts which areas will gain or lose population by the next census, which ethnic groups will be most affected through ethnicity-specific velocity and diffusion parameters, and how cascade effects through the ethnicity interaction network amplify or dampen direct impacts. However, we cannot empirically verify the year-by-year demographic trajectory during subway construction, as census data provides no ground truth for intermediate years.

5.4 Limitations

While the Multi-Ethnic Spatial Mixture of Experts framework demonstrates robust performance and interpretability, several limitations warrant consideration for both interpretation of results and future research directions.

Temporal Prediction Scope and Validation Constraints. Our empirical validation is constrained by census data availability rather than fundamental model limitations. The Cultural-Spatial Resonance Network's continuous PDE formulation and Runge-Kutta integration theoretically enable predictions at arbitrary time points, but census data, collected only at 5-year intervals in Canada, provides ground truth only at discrete census years (2001, 2006, 2011, 2016, 2021). This data structure necessitates training and validation on census-to-census transitions, making it impossible to empirically verify the accuracy of intermediate-year predictions for which no observed population data exists.

The model maintains robust accuracy over census-interval horizons (R^2 : 0.71-0.81), consistent with typical urban planning cycles for infrastructure investment and housing policy decisions. However, several aspects of temporal prediction remain unvalidated. First, intermediate-year accuracy: While the model could generate predictions for 2003 by halting integration partway through 2001→2006, we cannot verify whether these intermediate states accurately reflect true population distributions, as model parameters are optimized to minimize error only at census endpoints. Second, extended horizon performance: The observed accuracy degradation from 5-year (R^2 : 0.76-0.81) to 10-year predictions (R^2 : 0.71) suggests diminishing reliability with longer temporal horizons, likely due to our assumption that environmental features remain constant during integration, reasonable for 5-year periods but increasingly questionable over 10 years as new developments, policy changes, and economic shifts alter urban conditions. Third, sequential multi-step forecasting: While the model could theoretically chain predictions (2001→2006→2011→2016), this approach would compound errors and parameter drift across intervals without mechanisms to update environmental features or recalibrate parameters based on intermediate observations.

These limitations arise from the absence of intermediate ground truth data for training and validation, not from the model's mathematical formulation. The continuous PDE framework supports arbitrary temporal queries, but prediction reliability can only be guaranteed at census endpoints where empirical validation is

possible. For intervention simulations, this means the model requires incorporating auxiliary data sources such as annual municipal population estimates, building permit records, school enrollment statistics, or mobile phone mobility data to provide intermediate validation points between census years.

Ethnicity-Specific Performance Variability. The model exhibits substantial variability in prediction accuracy across different ethnic groups, with some communities presenting persistent challenges for accurate forecasting. Sri Lankan populations, in particular, show the most variable performance across temporal periods (R^2 : 0.56, 0.72, 0.60), suggesting that their population dynamics are influenced by factors not fully captured in our environmental feature vectors. These factors may include transnational connections, specialized economic niches, temporary migration patterns, or cultural characteristics not adequately represented in census variables.

Extreme Event Prediction Challenges. Jump processes, sudden, large-magnitude population changes, remain particularly challenging to predict accurately (MAE: 121.81, 213.41, 172.45), despite dedicated expert architecture for these events. Their rare occurrence (0.1-0.3% of cases) combined with high impact creates fundamental difficulties for data-driven modeling approaches. While our specialized Jump expert captures these dynamics better than uniform approaches, the large prediction errors indicate substantial uncertainty when extreme demographic shifts occur. Future work could explore incorporating external triggering data such as major development approvals, targeted immigration programs, refugee resettlement policies, or significant economic events that may herald jump processes before they manifest in census data.

Feature Space Limitations. While our model incorporates 298 socioeconomic and demographic features across 12 categories, important factors influencing population dynamics may remain uncaptured. The model cannot account for: (1) qualitative cultural factors not quantifiable through census variables, such as community social networks, religious institution presence, or cultural amenity availability; (2) anticipated future changes such as planned infrastructure projects, or policy shifts not yet implemented; (3) external shocks such as pandemics, economic crises, or geopolitical events affecting immigration flows; (4) individual-level decision factors that aggregate to produce population movements but are not observable in area-level census data. These unmeasured factors contribute to unexplained variance in model predictions and limit the model's ability to anticipate regime changes driven by forces outside the historical data distribution.

These limitations suggest several promising directions for future research, including development of sequential forecasting architectures, incorporation of external data sources for jump process prediction, transfer learning approaches for multi-city applications, explicit causal modeling of ethnicity interactions, and refinement of physics-informed constraints based on more sophisticated theories of human migration. Addressing these limitations would enhance both the accuracy and the scope of interpretable, physics-informed approaches to multi-ethnic urban modeling.

5.5 Conclusion

By revealing the mechanistic drivers of population change rather than merely predicting outcomes, MESMoE demonstrates potential to advance urban demographic modeling from pattern matching toward process understanding. Our results from Toronto census data suggest that the perceived trade-off between interpretability and performance may be context-dependent rather than fundamental (Rudin, 2019b), a finding with potential implications for other domains where decision-critical predictions require both accuracy and understanding.

Our work offers methodological contributions for machine learning applications in complex social systems. By demonstrating that physics-informed, regime-specific modeling can successfully predict heterogeneous urban dynamics while maintaining interpretability in the Toronto context, we provide a potential template for other domains where decision-critical predictions require both accuracy and explanation. However, validation in additional cities with different demographic compositions, urban structures, and data availability is necessary to establish broader generalizability.

The contextual adaptation mechanism of MESMoE, where different experts are engaged based on local conditions, represents a general principle potentially applicable to other complex systems characterized by regime-dependent dynamics. This approach could be extended to housing markets, transportation systems, and economic networks, all of which exhibit similar regime transitions between continuous evolution and discontinuous jumps, though such applications would require domain-specific adaptation and validation. For urban science specifically, our model offers a novel approach for demographic modeling that bridges qualitative ethnographic insights with quantitative prediction. The transformation of static demographic analysis into a dynamic flow-based framework through velocity field visualization represents a methodological advancement in how urban scientists can analyze population movements, from snapshots of where people are located to continuous flows of how they move through urban space. While our validation is limited to census-interval predictions due to data availability, the continuous PDE formulation provides a foundation for finer temporal resolution if auxiliary data sources become available.

Beyond its scientific contributions, MESMoE offers practical capabilities for evidence-based urban planning, though with important scope limitations. By parameterizing population dynamics through interpretable physical mechanisms, our model enables planners to simulate potential interventions before implementation within census-interval horizons (5-10 years). Modifications to velocity fields in targeted spatial locations could predict how transportation infrastructure might redirect population flows; adjustments to diffusion coefficients could simulate the effects of housing policies on neighborhood integration. However, these intervention simulations predict end-state demographic impacts rather than

year-by-year trajectories, aligning with typical urban planning cycles but limiting real-time monitoring capabilities.

The ethnicity interaction network provides a mathematical framework for designing interventions that leverage cascade effects across communities. Rather than viewing ethnic enclaves as isolated units, planners can identify strategic intervention points that may maximize positive spillover effects through the network structure. This capability enhances demographic modeling's prescriptive potential for policy design, though real-world policy effectiveness would depend on factors beyond demographic predictions alone, including political feasibility, resource constraints, and community acceptance.

Importantly, our approach provides quantitative support for differentiated urban planning strategies across ethnic communities. The systematic differences in feature importance patterns, physical parameters, and interaction structures across ethnicities suggest that effective urban governance may benefit from culturally-responsive approaches rather than one-size-fits-all policies. However, these findings are based on Toronto's specific multi-ethnic composition and may not directly transfer to cities with different demographic profiles or settlement histories.

5.6 Methods

5.6.1 Multi-Ethnic Census Dataset

5.6.1.1 Data Sources and Structure

We constructed a comprehensive multi-ethnic spatio-temporal dataset based on Statistics Canada census data spanning four census periods (2001, 2006, 2016, and 2021). The dataset captures population statistics disaggregated by ethnicity at the Dissemination Area (DA) level, the smallest standard geographic unit for census data dissemination, covering approximately 400-700 people per unit. For Toronto, this encompasses approximately 3,700 DAs, providing high spatial resolution for analyzing urban dynamics.

The dataset includes detailed population counts for the five largest ethnic groups in Toronto: China, India, Philippines, Portugal, and Sri Lanka. While demographic features were normalized through standardization to ensure comparable scales across different variables, raw population counts were preserved to maintain the natural skewness of population distributions and allow the model to learn authentic patterns of concentration and dispersion. This structure allowed us to track three distinct temporal intervals (2001-2006, 2006-2016, and 2016-2021), reflecting different socioeconomic contexts and migration patterns.

Our analysis of the population distribution revealed significant variations across census periods. For example, the percentage of zero-population DAs (areas where specific ethnic groups are not present) decreased from 53.59% in 2001 to 24.39% in 2016, before increasing to 38.23% in 2021. This pattern reflects the complex dynamics of ethnic settlement and mobility, including colonization of new areas,

concentration in existing enclaves, and exodus from certain neighborhoods. The population distribution exhibited high skewness (3.78-5.62) across all periods, indicating the presence of a small number of high-density ethnic clusters against a background of moderate-to-low density areas, a key challenge for conventional modeling approaches.

5.6.1.2 Feature Engineering

We extracted 298 socioeconomic and demographic features from the census data, organized into 12 categories:

1. **Demographics:** Population age structure, household composition, and family size
2. **Housing:** Dwelling types, ownership status, housing values, and maintenance needs
3. **Family Structure:** Marriage patterns, presence of children, household types
4. **Income:** Median household and individual income, income sources
5. **Employment:** Labor force participation, employment/unemployment rates
6. **Mobility & Migration:** Internal and external migration patterns, non-permanent residents
7. **Visible Minorities:** Population distribution by visible minority status
8. **Language:** Official language use, mother tongue, and multilingual capabilities
9. **Occupation:** Employment categories across economic sectors
10. **Religion:** Religious affiliations and practices
11. **Industry:** Distribution across industry sectors
12. **Place of Birth:** Country of origin information

To ensure temporal consistency, we selected features common across all census periods (2001-2021), which created a more stringent modeling challenge as recent census years (2016, 2021) contained additional variables that were excluded from the analysis. All features were standardized using z-score normalization to facilitate model training.

$$P_{ij} = \begin{cases} 1 & \text{if DAs } i \text{ and } j \text{ share a boundary} \\ 0 & \text{otherwise} \end{cases} \exp(-|x_i - x_j|^2 / \sigma_p^2) \text{ if } |x_i - x_j| \leq d_{\text{prox}}$$

5.6.1.3 Spatial Data Integration

We constructed three specialized spatial matrices to capture the complex spatial relationships between dissemination areas, building on established approaches in spatial statistics and urban modeling (Anselin, 1988; Fotheringham et al., 2009):

1. **Walking Distance Matrix:** Generated from Toronto's street network using shortest-path algorithms to represent pedestrian connectivity between DAs. For each pair of DAs (i, j), the

matrix value was calculated as: $W_{ij} = \begin{cases} \exp\left(-\frac{d_{ij}^2}{\sigma_w^2}\right) & \text{if } d_{ij} \leq d_{\max} \\ 0 & \text{otherwise} \end{cases}$ where d_{ij} represents the shortest-path walking distance between the centroids of DAs, σ_w is the distance decay parameter (1.5 km), and $d_{\max}$ is the maximum connection distance (5 km).

2. **Transit Connectivity Matrix:** Incorporated public transportation routes from Open Data Toronto, representing the ease of movement between DAs via public transit: $T_{ij} = \begin{cases} \exp(-t_{ij}/\sigma_t) & \text{if DAs } i \text{ and } j \text{ are connected by transit} \\ 0 & \text{otherwise} \end{cases}$ where t_{ij} is the travel time (in minutes) between DAs, and σ_t is a time decay parameter (30 minutes).

Proximity Matrix: Captured geographic contiguity and straight-line distance: $P_{ij} = \begin{cases} 1 & \text{if DAs } i \text{ and } j \text{ share a boundary} \\ 0 & \text{otherwise} \end{cases} \exp(-|x_i - x_j|^2 / \sigma_p^2) \text{ if } |x_i - x_j| \leq d_{\text{prox}}$ where x_i and x_j are the centroid coordinates, σ_p is the distance decay parameter (0.8 km), and d_{prox} is the maximum proximity connection distance (3 km).

Each matrix was constructed at 3742×3742 dimensions (covering all Toronto DAs), allowing comprehensive encoding of the urban spatial structure. These matrices enabled our model to incorporate physical barriers, transportation networks, and neighborhood adjacency, elements crucial for accurate population dynamics modeling but absent from traditional approaches.

5.6.1.4 Population Change Characterization

We conducted detailed analysis of population change patterns to inform our regime-specific modeling approach. For each temporal transition (2001→2006, 2006→2016, 2016→2021), we quantified:

1. **Zero-to-Non-Zero Transitions:** Representing colonization events where an ethnic group establishes presence in a previously uninhabited area. These transitions occurred in 44-68% of DAs across different periods.
2. **Non-Zero-to-Zero Transitions:** Representing complete exodus events where an ethnic group disappears from an area. These transitions ranged from 19-36% of DAs.
3. **Extreme Increases:** Population growth exceeding 100%, indicating rapid demographic shifts. These occurred in 7-14% of DAs.
4. **Extreme Decreases:** Population reduction exceeding 50%, indicating significant exodus. These occurred in 33-39% of DAs.

This analysis revealed substantial heterogeneity in population dynamics, reinforcing the need for regime-specific modeling approaches. For example, the 2006→2016 period showed the highest rate of zero-to-non-zero transitions (68.07% in training data), while the 2016→2021 period exhibited more balanced patterns between colonization and exodus dynamics.

5.6.2 Multi-Ethnic Spatial Mixture of Experts Framework

5.6.2.1 Model Architecture Overview

The MESMoE model integrates physics-informed modeling with specialized expert systems to predict multi-ethnic population dynamics. The core innovation lies in the combination of a fundamental physics-based partial differential equation (PDE) framework with regime-specific expert models coordinated through a learnable routing mechanism inspired by the mixture of experts literature (Jacobs, Jordan, Nowlan, & Hinton, 1991; Shazeer et al., 2017) (see Supplementary Document section 2 for detailed model foundation).

The model structure consists of:

1. Feature embedding networks for environmental context
2. Specialized expert modules for different population regimes
3. A learnable router that determines expert utilization
4. A physics-informed PDE solver for continuous dynamics
5. An ethnicity interaction module capturing inter-group effects
6. Area-specific parameter networks for spatial heterogeneity

5.6.2.2 Cultural-Spatial Resonance Network

The foundation of our approach is the Cultural-Spatial Resonance Network (CSRN), which models population dynamics through a novel PDE formulation that extends traditional diffusion-reaction approaches (M. Batty, 2013; Wilson, 2010):

$$\frac{\partial \phi_i}{\partial t} = D_i \nabla^2 \phi_i + \nabla \cdot (\lambda_i \mathbf{v}_i \nabla \phi_i) + (-\alpha_i(x, t) \cdot (\phi_i - \bar{\phi}_i) - \beta_i(t) \cdot \max(0, \|\nabla \phi_i\|^2 - K_i) - \gamma_i \cdot \text{div}(\mathbf{v}_i) \cdot \phi_i)$$

where for each ethnicity i :

- $\phi_i(x, t)$ represents the population density
- D_i is the diffusion coefficient
- λ_i is the amplification factor
- $\mathbf{v}_i$ is the velocity field
- $\alpha_i(x, t)$ controls mean reversion
- $\beta_i(t)$ penalizes excessive gradients
- γ_i modulates the advection stabilization
- K_i is the gradient magnitude threshold

Unlike traditional PDEs for population dynamics that use constant parameters across space (Wilson, 2010), we introduce spatial dependency for the mean reversion term:

$$\alpha_i(x, t) = \alpha_{0,i} \cdot \exp\left(-\frac{\|\nabla\phi_i\|^2}{\sigma_i^2}\right)$$

This formulation allows adaptive mean reversion based on local population gradients, a critical feature for capturing boundary dynamics between ethnic enclaves.

5.6.2.3 Mixture of Experts Framework

The complete model employs a mixture of experts architecture with four specialized expert models (Jacobs et al., 1991; Jordan & Jacobs, 1994):

$$\hat{\phi}_i(t + \Delta t) = \sum_{e=1}^4 w_{i,e} \cdot E_e(\phi_i(t), \mathbf{f}) + w_{i,\text{main}} \cdot M(\phi_i(t), \mathbf{f})$$

where:

- $\hat{\phi}_i(t + \Delta t)$ is the predicted population for ethnicity i
- $w_{i,e}$ is the weight assigned to expert e for ethnicity i
- E_e represents the expert models: colonization, jump process, decline, and PDE-based
- $\mathbf{f}$ represents input features
- M is the main neural network
- $w_{i,\text{main}}$ is the weight for the main model

Expert weights are determined by a learnable router that maps feature and population states to expert assignment probabilities:

$$\mathbf{w}_i = \text{Router}(\phi_i(t), \mathbf{f})$$

The router incorporates domain knowledge through population threshold inductive biases:

$$\mathbf{h}_{\text{hint}} = [\mathbb{1}_{\phi_i \leq \theta_{\text{col}}}, \mathbb{1}_{\phi_i > \theta_{\text{jump}}}, \mathbb{1}_{\theta_{\text{col}} < \phi_i \leq \theta_{\text{decline}}}, \mathbb{1}_{\theta_{\text{col}} < \phi_i \leq 2\theta_{\text{jump}}}]$$

where $\mathbb{1}$ is the indicator function. The final expert weights blend learned probabilities with rule-based hints:

$$\mathbf{w}_{\text{final}} = (1 - \lambda) \cdot \mathbf{w} + \lambda \cdot \mathbf{h}_{\text{hint}}$$

where λ is a blending parameter.

5.6.3 Specialized Expert Models

5.6.3.1 Multi-Ethnic Colonization Expert

This expert handles zero-to-nonzero population transitions, modeling the initial establishment of an ethnic group in a previously uninhabited area based on concepts from colonization and diffusion of innovations

theory (Brown & Sanders, 1981; Rogers, Singhal, & Quinlan, 2014). The colonization model predicts population establishment through a multi-stage hierarchical process:

$$p(\text{size}|\mathbf{f}) = \text{Softmax}(g_{\text{size}}(\mathbf{f}))$$

where $p(\text{size}|\mathbf{f})$ is the probability distribution over colonization sizes conditioned on features $\mathbf{f}$.

The predicted colonization value is:

$$\phi_i^{\text{col}} = \sum_{s \in \{\text{small}, \text{medium}, \text{large}, \text{very large}\}} p(s|\mathbf{f}) \cdot (h_s(\mathbf{f}) \cdot \text{scale}_s + \text{bias}_s)$$

where h_s are size-specific neural networks, and scale_s and bias_s are learned parameters.

5.6.3.2 Multi-Ethnic Jump Process Expert

This expert models discontinuous jumps in population values that cannot be captured by continuous diffusion processes, extending stochastic jump process theory (Davis, 1984) to the spatial-ethnic domain.

The jump process combines regular and extreme jump components:

$$\phi_i^{\text{jump}} = \phi_i(t) + p_i^{\text{jump}}(\mathbf{f}) \cdot m_i^{\text{jump}}(\mathbf{f}) + p_i^{\text{extreme}}(\mathbf{f}) \cdot m_i^{\text{extreme}}(\mathbf{f})$$

where ϕ_i^{jump} is the probability of a regular jump, $m_i^{\text{jump}}(\mathbf{f})$ is the predicted jump magnitude, $p_i^{\text{extreme}}(\mathbf{f})$ is the probability of an extreme jump, and $m_i^{\text{extreme}}(\mathbf{f})$ is the predicted extreme jump magnitude.

5.6.4 Multi-Ethnic Decline Expert

This expert specializes in modeling population decline scenarios based on population decline and exodus theory (Frey, 1979). The decline model predicts:

$$\phi_i^{\text{decline}} = \phi_i(t) \cdot (1 - p_i^{\text{decline}}(\mathbf{f}) + p_i^{\text{decline}}(\mathbf{f}) \cdot ((1 - p_i^{\text{extreme}}(\mathbf{f})) \cdot f_i^{\text{decline}}(\mathbf{f}) + p_i^{\text{extreme}}(\mathbf{f}) \cdot f_i^{\text{extreme}}(\mathbf{f})))$$

where $p_i^{\text{decline}}(\mathbf{f})$ is the probability of decline, $f_i^{\text{decline}}(\mathbf{f})$ is the decline factor, $p_i^{\text{extreme}}(\mathbf{f})$ is the probability of extreme decline, and $f_i^{\text{extreme}}(\mathbf{f})$ is the extreme decline factor.

5.6.4.1 Multi-Ethnic CSRN Diffusion Solver

For continuous population changes, we implemented a physics-informed neural PDE solver that numerically integrates the CSRN equation using a 4th-order Runge-Kutta method (Press, 2007):

$$\phi_i(t + \Delta t) = \phi_i(t) + \frac{\Delta t}{6} (k_1 + 2k_2 + 2k_3 + k_4)$$

where k_1, k_2, k_3, k_4 are intermediate evaluations of the PDE right-hand side. Spatial derivatives are approximated using graph-based methods suitable for irregular urban boundaries (Belkin & Niyogi, 2008):

$$\nabla^2 \phi_i \approx \mathbf{L} \phi_i$$

where $\mathbf{L}$ is the graph Laplacian derived from spatial connectivity matrices. Gradient computation uses weighted neighbor differences:

$$\nabla\phi_i(x_a) \approx \sum_{b \in \mathcal{N}(a)} w_{ab} \cdot (\phi_i(x_b) - \phi_i(x_a)) \cdot \frac{x_b - x_a}{\|x_b - x_a\|}$$

where $\mathcal{N}(a)$ is the neighborhood of point a and w_{ab} are distance-based weights.

5.6.4.2 Multi-Matrix Spatial Attention

To capture complex spatial relationships, we developed a multi-matrix spatial attention mechanism that extends the transformer attention framework (Vaswani, 2017) to incorporate multiple spatial relationships:

$$\mathbf{A}(\mathbf{X}, \{\mathbf{M}_k\}_{k=1}^K) = \sum_{k=1}^K \pi_k \cdot \text{Attention}_k(\mathbf{X}, \mathbf{M}_k)$$

where $\mathbf{X}$ is the input sequence, $\mathbf{M}_k$ are spatial matrices (walking, transit, proximity), π_k are learned weights for each matrix, and Attention_k is a spatial attention function.

Each attention head computes:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{M}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} + \mathbf{M}\right)\mathbf{V}$$

where $\mathbf{Q}, \mathbf{K}$, and $\mathbf{V}$ are query, key, and value projections, d_k is the dimension of the keys, and $\mathbf{M}$ is a spatial bias matrix.

5.6.4.3 Numerical Stability and Integration Safeguards

The CSRN PDE solver employs comprehensive numerical stability mechanisms to ensure robust integration within the MoE framework, particularly crucial given the diversity of demographic regimes (colonization, jump, decline, diffusion) that exhibit vastly different population scales and dynamics.

Multi-level stability architecture:

1. **Parameter-adaptive scaling:** The solver dynamically adjusts integration scales based on population magnitude and regime type. For log-space populations (colonization regime), diffusion coefficients are constrained to $D \in [0.01, 0.5]$ and time steps reduced by 90% ($\Delta t_{\text{eff}} = 0.1\Delta t$) to prevent exponential instabilities. For linear-space populations, adaptive scale factors $\alpha_{\text{scale}} = \min(3.0, 100/(\max(|\phi|) + \epsilon))$ normalize PDE terms across population ranges spanning 4 orders of magnitude.
2. **Gradient and Laplacian regularization:** Spatial derivatives incorporate stabilizing perturbations: $\nabla\phi \leftarrow \nabla\phi + \mathcal{N}(0, 0.01^2)$ and $\nabla^2\phi \leftarrow \nabla^2\phi + 0.05 \cdot \text{sign}(\phi)$. Gradient magnitudes are clamped to $\|\nabla\phi\|^2 \in [0.05, 5.0]$ (log-space) or $[0.05, 100]$ (linear-space), preventing runaway growth in steep population transitions characteristic of jump processes.
3. **RK4 integration with fallback:** Fourth-order Runge-Kutta integration employs clamped intermediate stages: $k_i \in [-1.0, 1.0]$ for log-space and $[-10.0, 10.0]$ for linear-space. A forced minimum update $\Delta\phi_{\text{min}} = 0.01 \cdot \text{sign}(\phi + 0.001)$ ensures non-zero gradient flow during training even in near-equilibrium states. If RK4 fails (e.g., NaN propagation), the solver automatically falls

back to forward Euler with aggressive clamping: $\phi^{n+1} = \text{clamp}(\phi^n + \Delta t \cdot k_1, [-20, 20])$ for log-space.

4. ****CUDA kernel safeguards****: At the lowest computational level, custom CUDA kernels enforce adaptive gradient clipping (max gradient scales with population: $|\partial_x|_{\max} = \max(10, 0.5 \cdot |\phi|)$), finite-value verification (NaN/Inf replaced with zeros), and minimum perturbations ($\epsilon_{\min} = \max(0.001, 0.01 \cdot |\phi|)$) to maintain differentiability throughout backpropagation.
5. **MoE-specific stabilization**: Because the PDE expert operates alongside discrete experts (colonization, jump, decline) with differing output scales, we enforce regime-appropriate bounds: colonization outputs $\in [10, 2000]$, jump outputs $\leq 3\phi_t + 100$, decline outputs $\leq \phi_t + 50$, and PDE outputs $\leq 5\phi_t + 100$. Expert weight smoothing ($w_t = 0.2w_{t-1} + 0.8w_t$) prevents abrupt transitions that could destabilize the PDE solver when regime assignments change.

Stability validation: Conservation checks after each forward pass verify mass conservation error $\epsilon_{\text{mass}} = |\sum_i \phi_i^{n+1} - \sum_i \phi_i^n| / \sum_i \phi_i^n < 10^{-4}$ and energy conservation $\epsilon_{\text{energy}} = |\sum_i (\phi_i^{n+1})^2 - \sum_i (\phi_i^n)^2| / \sum_i (\phi_i^n)^2 < 10^{-3}$. Across 3,700 DA-level predictions spanning 20 years, the PDE expert achieved mean $\epsilon_{\text{mass}} = 2.3 \times 10^{-5}$ and $\epsilon_{\text{energy}} = 4.7 \times 10^{-4}$, confirming numerical stability throughout training and inference.

5.6.4.4 Area-Specific Parameter Learning

A key innovation in our approach is the learning of area-specific parameters through specialized neural networks, extending spatially varying coefficient models (Fotheringham et al., 2009) with deep learning:

$$\theta_i^{(a)} = g_\theta(\mathbf{f}^{(a)})$$

where $\theta_i^{(a)}$ represents a parameter (e.g., diffusion coefficient) for ethnicity i in area a , $\mathbf{f}^{(a)}$ are the features of area a , and g_θ is a neural network.

For the PDE solver, all key parameters are area-specific:

$$D_i^{(a)} = g_D(\mathbf{f}^{(a)})_i$$

$$\lambda_i^{(a)} = g_\lambda(\mathbf{f}^{(a)})_i$$

$$\mathbf{v}_i^{(a)} = g_v(\mathbf{f}^{(a)})_i$$

$$\alpha_i^{(a)} = g_\alpha(\mathbf{f}^{(a)})_i$$

$$\beta_i^{(a)} = g_\beta(\mathbf{f}^{(a)})_i$$

$$\gamma_i^{(a)} = g_\gamma(\mathbf{f}^{(a)})_i$$

where each g function is a neural network that maps area features to ethnicity-specific parameters.

5.6.4.5 Ethnicity Interactions

Our model captures inter-ethnic influences through a learned interaction module inspired by social interaction models (Schelling, 1971):

$$\mathbf{I}_{i,j} = \sigma(g_{\text{interaction}}(\boldsymbol{\phi}))$$

where $\mathbf{I}_{i,j}$ represents the influence of ethnicity j on ethnicity i . The interaction effect is computed as:

$$\mathbf{e}_i = \sum_{j \neq i} \mathbf{I}_{i,j} \cdot (\hat{\phi}_j(t + \Delta t) - \phi_j(t)) \cdot \alpha$$

where α is a scaling factor. The final prediction incorporates these interactions:

$$\hat{\phi}_i^{\text{final}}(t + \Delta t) = \hat{\phi}_i(t + \Delta t) + \mathbf{e}_i$$

5.6.4.6 CUDA-Accelerated Implementation

We developed custom CUDA kernels for computationally intensive operations, including gradient computation, divergence calculation, and PDE term evaluation. Our implementation achieved a 1000× speedup compared to standard CPU-based solvers, reducing computation time from approximately 72 seconds to 0.07 seconds per batch of 32 city blocks. This acceleration was essential for efficient model training and inference, especially for backpropagation through the physics-informed components.

To ensure numerical stability in the PDE solver, we implemented adaptive time-stepping based on the Courant-Friedrichs-Lewy condition (Courant, Friedrichs, & Lewy, 1967), gradient clamping to prevent numerical instabilities, and adaptive parameter bounds based on population magnitude.

5.6.5 Training Methodology

5.6.5.1 Multi-Stage Curriculum Training

We employed a multi-stage curriculum strategy with progressive expert introduction, following principles of curriculum learning (Yoshua Bengio, Louradour, Collobert, & Weston, 2009):

1. **Pre-training Phase:** Initially, only the main neural network component was trained while keeping expert modules fixed.
2. **Expert Training Phase:** The colonization expert was introduced first, followed by the PDE expert, and finally the jump and decline experts.
3. **Router Training Phase:** Once all experts were trained, the router was fine-tuned to optimize expert assignment.
4. **Full Model Fine-tuning:** All components were jointly optimized with a reduced learning rate.

Hyperparameter search framework and result is provided in Supplementary Document section 5.

5.6.5.2 Loss Function Design

The model employed a multi-component loss function addressing the heterogeneous nature of population dynamics:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{prediction}} + \lambda_{\text{aux}}\mathcal{L}_{\text{auxiliary}} + \lambda_{\text{reg}}\mathcal{L}_{\text{regularization}} + \lambda_{\text{phys}}\mathcal{L}_{\text{physics}} + \lambda_{\text{cons}}\mathcal{L}_{\text{conservation}}$$

The core prediction loss used a regime-adaptive weighting scheme:

$$\mathcal{L}_{\text{prediction}} = \frac{1}{B \cdot E} \sum_{b=1}^B \sum_{e=1}^E w_{b,e} \cdot w_e \cdot (\hat{\phi}_{b,e} - \phi_{b,e})^2$$

where $w_{b,e}$ are case-specific weights determined by population regime:

$$w_{b,e} = \begin{cases} w_{\text{col}} & \text{if } \phi_{b,e}^{\text{init}} \leq \theta_{\text{col}} (\text{colonization case}) \\ w_{\text{jump}} & \text{if } \phi_{b,e}^{\text{change}} > \theta_{\text{jump}} (\text{jump case}) \\ w_{\text{decline}} & \text{if } \phi_{b,e}^{\text{init}} \leq \theta_{\text{decline}} \text{ and } \phi_{b,e}^{\text{change}} \leq 0 (\text{decline case}) \\ w_{\text{pde}} & \text{if case meets PDE criteria (PDE case)} \\ 1.0 & \text{otherwise (regular case)} \end{cases}$$

The optimal weight values determined through ablation studies are $w_{\text{col}} = 5.0$, $w_{\text{jump}} = 3.0$, $w_{\text{decline}} = 4.0$, and $w_{\text{pde}} = 2.0$.

For jump cases, we applied an additional log-space transformation to better handle large magnitude changes:

$$\mathcal{L}_{\text{jump}} = (\log(\hat{\phi}_{b,e} + 1) - \log(\phi_{b,e} + 1))^2 \cdot 100.0$$

The auxiliary loss guides the expert router to correctly classify population dynamics:

$$\mathcal{L}_{\text{auxiliary}} = -\frac{1}{B \cdot E} \sum_{b=1}^B \sum_{e=1}^E \sum_{c=1}^C y_{b,e,c} \log(p_{b,e,c})$$

The physics-informed loss enforces consistency with the underlying PDE:

$$\mathcal{L}_{\text{physics}} = \frac{1}{B \cdot E} \sum_{b=1}^B \sum_{e=1}^E \mathbb{1}_{\text{isPDE}}(b, e) \cdot \left\| \frac{\partial \phi_{b,e}}{\partial t} - F(\phi_{b,e}, \nabla \phi_{b,e}, \nabla^2 \phi_{b,e}) \right\|_2^2$$

5.6.5.3 Optimization Strategy

We used Adam optimizer (Kingma & Ba, 2014) with weight decay regularization and a learning rate schedule combining warm-up and cosine annealing (Loshchilov & Hutter, 2016):

$$\eta_t = \begin{cases} \eta_{\max} \cdot \frac{t}{T_{\text{warmup}}} & \text{if } t \leq T_{\text{warmup}} \\ \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min}) \left(1 + \cos \left(\pi \frac{t - T_{\text{warmup}}}{T - T_{\text{warmup}}} \right) \right) & \text{otherwise} \end{cases}$$

where T_{warmup} is the warm-up period length (typically 5 epochs), T is the total number of epochs, $\eta_{\max}$ is the maximum learning rate (0.0005), and $\eta_{\min}$ is the minimum learning rate (0.00001).

To ensure training stability, we implemented adaptive gradient clipping with threshold:

$$\tau = \max(0.1, \min(5.0, \frac{\|\theta\|_2}{100})) \cdot \tau_{\text{base}}$$

with $\tau_{\text{base}} = 5.0$ as the base clipping threshold.

5.6.5.4 Cross-Temporal Validation

To evaluate the model's ability to predict across different temporal intervals, we employed a cross-temporal validation strategy:

1. **Within-Period Validation:** Models trained on one period (e.g., 2001-2006) were validated on held-out data from the same period.
2. **Cross-Period Validation:** Models were trained on one period (e.g., 2001-2006) and validated on a different period (e.g., 2016-2021).
3. **Multi-Period Training:** Models were trained on data from two periods and validated on the third period.

For each temporal period, we used an 80/20 train/test split, with the training data further divided into 90% training and 10% validation. Model performance was evaluated using multiple metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), and coefficient of determination (R^2), both overall and disaggregated by ethnicity and population dynamic regime. It is important to note that our validation approach assesses predictions at census intervals (5 or 10 years) due to ground truth data availability, rather than at intermediate time points. While the Cultural-Spatial Resonance Network's continuous PDE formulation theoretically enables predictions at arbitrary time points through its Runge-Kutta integration, census data, the gold standard for population statistics, is only collected at 5-year intervals, providing ground truth observations exclusively at census years. The model predicts population distributions at the next census time point given the current census as input, but intermediate-year predictions (e.g., 2003, 2004) cannot be empirically validated against observed data. This validation constraint reflects: (1) the discrete nature of census data availability, which provides no ground truth for intermediate years, (2) the resulting inability to train or validate the model on sub-census-interval predictions, and (3) the alignment with urban planning practice, where major policy decisions typically operate on 5-10 year planning horizons synchronized with census data releases. Future work incorporating auxiliary data sources (annual

population estimates, building permits, administrative records) could enable intermediate-year validation, but current empirical validation necessarily focuses on census endpoints where observed data exists.

5.6.5.5 Interpretability Analysis Framework

We developed a comprehensive interpretability framework to analyze the learned model components:

1. **Feature Importance Analysis:** Computed based on the magnitude of weights in the first layer of each expert module (Breiman, 2001).
2. **Expert Contribution Analysis:** Quantifying the relative contribution of each expert to the final predictions.
3. **Ethnicity Interaction Analysis:** Measuring the strength and direction of inter-ethnic influences through interaction matrices.
4. **Spatial Parameter Visualization:** Creating high-resolution maps of physics-informed parameters across urban space.

The framework enables quantitative assessment of the socioeconomic drivers of population dynamics, the relative importance of different demographic mechanisms, and the spatial heterogeneity of population processes across urban neighborhoods.

Data Availability

The source code for the Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework, including the CUDA-accelerated implementation and all analysis scripts, is publicly available on GitHub at <https://github.com/navid-nsk/mesmoe-csrn-2016-2021>. The spatial demographic dataset for Toronto at the dissemination area scale, including census data and associated socioeconomic features used in this study, is available on Figshare at https://figshare.com/articles/dataset/Spatial_data_for_Toronto_in_dissemination_area_scale/28912400.

Both the code repository and dataset are released under open licenses to ensure reproducibility and facilitate future research.

Author Contributions

S.N.M.M. conceived the study, designed the Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework, developed the physics-informed modeling approach, implemented the CUDA-accelerated algorithms, performed all data analysis, created the interpretability framework, generated all figures, and wrote the manuscript. H.C. supervised the project, contributed to the conceptual framework, provided critical feedback on the methodology, and reviewed and edited the manuscript. T.J. contributed to the mathematical formulation of the Cultural-Spatial Resonance Network, assisted with the PDE solver design, helped develop the mixture of experts architecture, and contributed to writing the Methods section. R.H. participated in code implementation, conducted model validation experiments, performed cross-temporal

validation analysis, and verified the reproducibility of results. D.R. contributed to software development, implemented visualization tools for spatial parameter analysis, created data processing pipelines, generated supplementary visualizations, and contributed to writing the Results section.

Chapter 6 Synthesis - Discussion and Conclusion

6.1 Introduction: Bridging Computational Discovery and Geographic Theory

This dissertation began with a fundamental tension: while cities worldwide increasingly deploy artificial intelligence to manage complex urban systems, the dominant paradigm of black-box optimization systematically fails to achieve its stated goals while potentially harming the very communities it purports to serve. Through three interconnected investigations, a critical analysis of algorithmic failures in urban governance, the discovery of reaction-diffusion physics underlying ethnic settlement patterns, and the development of interpretable frameworks for demographic prediction, this research has demonstrated that the perceived trade-off between computational sophistication and human understanding is false. The journey from identifying what fails to discovering why patterns emerge to building what works reveals a transformative possibility: interpretable Geospatial Artificial Intelligence can convert multi-ethnic settlement from a descriptive sociological phenomenon into a predictive physical science while maintaining democratic accountability and theoretical transparency.

The progression through these three articles was dialectical, with each investigation's limitations motivating the next's innovations. Chapter 3's systematic documentation of the "metrics trap" across 157 urban AI deployments revealed how impressive technical accuracy, ShotSpotter's 97% acoustic precision, COMPAS's statistical performance, Sydney's gentrification predictions, consistently failed to translate into meaningful social outcomes. The 9.1% effectiveness rate for gunshot detection, the doubling of false positive rates for Black defendants, and the acceleration rather than prevention of displacement demonstrated that optimizing narrow technical metrics without understanding underlying mechanisms inevitably produces systemic failure. This critical analysis established an ethical and methodological imperative: urban AI systems must be interpretable not as a luxury for academic understanding but as a necessity for responsible governance of diverse communities.

Yet criticism without construction offers little beyond documentation of failure. Chapter 4's breakthrough came from developing a graph-based physics-informed neural network (GraphPDE) that embeds multi-ethnic reaction-diffusion dynamics while learning interpretable demographic parameters. By adapting reaction-diffusion equations to spatial graphs across 52 Canadian cities and 30,091 dissemination areas, the research revealed that ethnic settlement patterns emerge from self-organizing spatiotemporal processes consistent with diffusion-reaction mechanisms. The discovery that different ethnic groups exhibit distinct characteristic spatial scales, from 34.5 km for Philippines-origin populations to 63.0 km for United Kingdom-origin populations, transformed segregation from a complex social problem into a quantifiable physical phenomenon amenable to mathematical analysis. Chinese settlements forming spot patterns, Iranian communities exhibiting labyrinthine morphology, and Portuguese populations displaying dispersed

spot patterns all emerge from the learned interplay between self-interaction strength and between-group competition.

The physics-informed understanding revealed mechanisms invisible to conventional analysis: ethnic enclaves nucleate at critical thresholds ranging from 658 individuals (China) to 8,132 individuals (India) depending on group-specific network characteristics; multiple stable settlement configurations exist with quantifiable transition barriers; and the system exhibits multi-stability where configuration transitions require sustained interventions over decadal timescales. The learned attention-based interaction mechanism quantifying inter-ethnic dynamics, China exhibiting strong negative self-competition (-19.8) driving spatial dispersal, Philippines-Pakistan mutual attraction, United Kingdom-Italy mutual repulsion, transformed qualitative observations about ethnic relations into mathematical relationships amenable to policy analysis. However, universal physics alone cannot capture the full complexity of urban demographic dynamics. Different population regimes, colonization of new areas, sudden jumps from housing crises, gradual decline from gentrification, continuous diffusion in established communities, require distinct mathematical formulations. Chapter 5's Multi-Ethnic Spatial Mixture of Experts (MESMoE) framework synthesized physics-informed modeling with machine learning to create a practical prediction system achieving R^2 values of 0.80-0.83 while maintaining complete interpretability. The architecture's specialized experts for different demographic regimes, learned ethnicity interaction networks, and spatially-varying physics parameters demonstrated that incorporating domain knowledge enhances rather than compromises predictive accuracy. The framework revealed systematic differences in how ethnic communities respond to urban infrastructure: Chinese and Filipino populations show amplification factors strongly correlated with transit accessibility ($r = 0.58$ and 0.61), while Indian populations demonstrate stronger correlation with housing variables ($r = 0.47$). The temporal analysis exposed how expert module engagement shifted across economic cycles, with the Colonization expert's U-shaped importance ($11\% \rightarrow 9\% \rightarrow 10\%$) quantifying reduced settlement establishment during the 2008 financial crisis.

The synthesis of these three investigations establishes interpretable GeoAI as more than a methodological advance, it represents a paradigmatic shift in how we understand and study urban ethnic dynamics. Theoretically, it bridges the persistent divide between quantitative spatial science and critical human geography by demonstrating that mathematical rigor need not sacrifice social awareness. Methodologically, it proves that machine learning's power can be harnessed without surrendering to algorithmic opacity. Practically, it converts abstract patterns into concrete parameters that planners can manipulate to influence settlement outcomes. This interpretable approach addresses fundamental limitations constraining urban ethnic studies for decades: where traditional census analysis provides snapshots, reaction-diffusion models reveal continuous processes; where statistical correlations identify associations, physics-informed

frameworks expose causal mechanisms; where black-box predictions offer accuracy without understanding, interpretable models provide both superior performance and theoretical insight.

Yet this transformation raises profound questions about the nature of human geography itself. Can social phenomena truly be reduced to physical laws? What is lost when we model communities as reaction-diffusion systems? How do we reconcile deterministic physics with human agency and cultural meaning? This synthesis chapter addresses these tensions not through resolution but through productive engagement, demonstrating how physics-informed approaches complement rather than replace other ways of knowing urban dynamics. The discovery of ethnic-specific spatial mobility patterns does not negate the importance of ethnographic understanding; the quantification of transition barriers does not diminish the lived experience of displacement; the prediction of settlement patterns does not predetermine community futures. The Greater Toronto Area and 51 other major Canadian cities, encompassing millions of residents across thousands of dissemination areas, serve as more than case studies, they represent laboratories for the urban future. As cities worldwide grapple with increasing diversity, housing crises, and the promises and perils of algorithmic governance, the frameworks developed here offer both cautionary lessons and hopeful possibilities. The systematic failures documented in Chapter 3 warn against uncritical adoption of AI systems, while the successes of Chapters 4 and 5 demonstrate that interpretable alternatives can achieve superior outcomes. The discovery that Canadian cities' ethnic geography follows learnable physical dynamics while maintaining cultural distinctiveness suggests that diversity and order are not opposing forces but complementary aspects of urban organization. This synthesis chapter weaves together theoretical insights, methodological innovations, empirical discoveries, and practical applications to demonstrate how interpretable GeoAI transforms multi-ethnic settlement from descriptive phenomenon to predictive science, revealing a path toward urban AI that enhances rather than erodes democratic governance, advances rather than abandons geographic theory, and serves rather than surveils diverse communities.

6.2 Discussion and Theoretical Synthesis

6.2.1 From Metrics Trap to Mechanistic Understanding

The trajectory from Chapter 3's critique of algorithmic failures to the physics-informed models of Chapters 4 and 5 represents more than methodological evolution, it embodies a fundamental reconceptualization of what constitutes understanding in computational urban science. This reconceptualization addresses the central research question motivating this dissertation: whether interpretable GeoAI can transform multi-ethnic settlement from descriptive phenomenon to predictive science while maintaining democratic accountability.

The systematic analysis of 157 urban AI deployments revealed a pervasive pattern that (O'Neil, 2016) termed "weapons of math destruction": black-box systems achieving impressive technical metrics while failing catastrophically at their stated purposes. ShotSpotter's 97% acoustic accuracy seemed to validate the assumption that high prediction accuracy implies understanding, until analysis revealed only 9.1% of alerts corresponded to actual gun crimes. The system could identify loud sounds with remarkable precision but lacked any model of what distinguished gunshots requiring police response from fireworks or construction noise. Similarly, healthcare algorithms accurately predicted costs, their optimization target, while systematically denying care to Black patients with greater medical needs (Ziad Obermeyer, Brian Powers, Christine Vogeli, & Sendhil Mullainathan, 2019). These failures stemmed not from technical inadequacy but from fundamental category errors: mistaking correlation for causation, proxy measures for actual objectives, and statistical patterns for mechanistic understanding.

This disconnect between measurement and meaning exposes what (Pearl & Mackenzie, 2018) identify as the "ladder of causation" problem: purely associational models cannot answer causal or counterfactual questions essential for policy intervention. The resolution came not through better black boxes but through a return to mechanistic understanding, what (M. Batty, 2013) advocates in "The New Science of Cities" as models revealing "deep structure" rather than merely fitting surfaces. The physics-informed approach of Chapter 4 directly addresses this limitation. Rather than treating ethnic settlement as a black-box prediction problem, the GraphPDE framework sought to understand the generative mechanisms producing observed patterns, answering the first research objective: developing interpretable computational frameworks that expose rather than obscure demographic dynamics.

The reaction-diffusion model adaptation revealed that settlement patterns emerge from quantifiable tensions between dispersive forces (diffusion with ethnic-specific coefficients yielding characteristic spatial scales of 34.5–63.0 km) and aggregative forces (attention-based inter-ethnic interactions). This mechanistic understanding immediately explained phenomena that correlation-based approaches could only document: why ethnic enclaves form at characteristic spatial scales varying systematically by group, how critical nucleation thresholds (658–8,132 individuals) trigger pattern formation, and what configuration barriers govern transitions between stable settlement states. These insights address the second research objective: quantifying hidden dynamics governing multi-ethnic coexistence that remain invisible to conventional analysis.

The false dichotomy between accuracy and interpretability, perpetuated by machine learning orthodoxy, dissolves when examined through physics-informed modeling. Chapter 5's MESMoE framework achieved R^2 values of 0.80–0.83, matching or exceeding black-box alternatives, while maintaining complete interpretability through regime-specific experts and learned physics parameters. This finding directly challenges Zachary C. Lipton (2018)'s characterization of the accuracy-interpretability trade-off as

fundamental, demonstrating instead that domain knowledge integration enhances rather than compromises predictive power. The Colonization expert's hierarchical probability distributions, the Jump Process expert's stochastic formulations, and the PDE Diffusion expert's spatial derivatives all provide mechanistic insight enabling iterative improvement impossible with opaque systems.

The transformation from metrics trap to mechanistic understanding reveals why certain algorithmic interventions consistently fail. Predictive policing systems trapped in feedback loops could not escape because they lacked models of how enforcement creates data justifying more enforcement (Lum & Isaac, 2016). The reaction-diffusion framework avoids these traps by explicitly modeling bidirectional relationships between pattern and process: settlement patterns emerge from inter-ethnic dynamics, which are themselves shaped by existing spatial configurations. This reflexivity, absent from correlational approaches, enables the counterfactual reasoning essential for responsible urban governance, addressing the third research objective: creating practical tools for culturally-responsive planning.

6.2.2 Universal Physical Principles in Social Systems

The synthesis of reaction-diffusion discoveries reveals that human settlement patterns can be described with mathematical frameworks as rigorous as those governing chemical reactions or biological morphogenesis. This finding engages a longstanding debate in social science about whether human phenomena exhibit law-like regularities amenable to physical analysis (Hedström & Ylikoski, 2010). The universality manifests not as deterministic prediction of individual behavior but as statistical regularities in collective dynamics emerging from identifiable forces, precisely the "micro-macro link" that (Coleman, 1990) identified as sociology's central theoretical challenge.

The application of reaction-diffusion equations to human settlement extends (Turing, 1952)'s foundational insight that spatial patterns can arise spontaneously from the interplay of diffusion and local reaction. Turing demonstrated that two interacting chemical species with different diffusion rates can destabilize uniform distributions, generating spots, stripes, and labyrinthine patterns through "diffusion-driven instability." This mechanism, confirmed in biological systems from zebrafish pigmentation (Kondo & Miura, 2010) to mammalian digit formation (Sheth et al., 2012), provides the theoretical foundation for understanding ethnic pattern formation. The discovery that Chinese settlements form spot patterns, Iranian communities exhibit labyrinthine morphology, and Portuguese populations display dispersed configurations directly parallels Turing's morphogenetic predictions, suggesting that social and biological pattern formation share fundamental mathematical structure.

This parallel raises profound epistemological questions. Vinkovic and Kirman (2006) demonstrated that Schelling's segregation model, the canonical agent-based model of residential sorting, has a physical analogue describable by differential equations. Yet as Clark and Fossett (2008) cautioned, physical

analogues do not automatically advance understanding of social dynamics; the crucial question is whether mathematical formalism captures causally relevant mechanisms. The GraphPDE framework addresses this concern by learning parameters with direct sociological interpretation: diffusion coefficients reflecting residential mobility patterns shaped by housing markets and transportation; interaction terms capturing economic complementarities and competition; nucleation thresholds corresponding to minimum viable community sizes for ethnic infrastructure.

The temporal evolution of demographic parameters provides empirical grounding for these interpretations. Philippines-origin populations maintained consistently positive growth rates ($0.015\text{--}0.023\text{ year}^{-1}$) across all census periods, while United Kingdom-origin populations transitioned from positive growth ($+0.010\text{ year}^{-1}$) in 2001–2006 to sustained decline (-0.015 to -0.023 year^{-1}) subsequently. These dynamics capture effects of the 2008 financial crisis, immigration policy reforms, and changing source country conditions that pure black-box approaches could fit but not explain. The period-specific parameters learned separately for four census intervals (2001–2006, 2006–2011, 2011–2016, 2016–2021) reveal non-stationary processes invisible to models assuming temporal homogeneity.

Configuration landscape analysis quantifies multi-stability, the reality that demographic transitions require overcoming substantial reorganization barriers. This finding resonates with Schelling (1971)'s insight that mild individual preferences can produce strong collective segregation, but advances beyond Schelling by quantifying barrier heights and identifying asymmetries. The asymmetric barriers discovered, with forward barriers exceeding reverse barriers for certain transitions, explain why neighborhoods naturally drift toward particular ethnic configurations despite policies promoting alternatives. High barriers between Philippines↔Pakistan and China↔Iran indicate structural incompatibilities, while low barriers between United Kingdom↔Iran suggest more fluid exchanges. These quantities emerge mathematically from learned reaction-diffusion dynamics rather than being imposed a priori.

The attention-based competition mechanism bridges microscopic decisions with macroscopic patterns, addressing Michael Batty (2005)'s call for models connecting bottom-up processes to emergent urban structure. The discovered interactions, China exhibiting strong negative self-competition (-19.8) driving spatial dispersal while showing positive facilitation with six other groups; Philippines and Pakistan forming mutually attractive dyads; United Kingdom and Italy exhibiting mutual repulsion, reflect measurable demographic relationships with economic interpretations. Positive cross-interactions correspond to complementary niches where wholesale-retail relationships or professional networks create mutual benefit from co-location. Negative interactions reflect overlapping niches creating displacement pressures. These are not merely statistical correlations but parameters of a generative model that can simulate counterfactual scenarios.

The universality of reaction-diffusion dynamics across all nine ethnic groups, despite vast cultural differences, suggests deep underlying order, what (L. M. A. Bettencourt, Lobo, Helbing, Kühnert, & West, 2007) characterize as "urban scaling laws" reflecting universal constraints on human organization. The pattern formation phase space segregates ethnic groups into three distinct regimes, spots (UK, Portugal), stripes (China, India, Philippines), and labyrinthine (Iran, Pakistan, Sri Lanka), based on self-interaction strength and between-group competition intensity. This classification does not erase cultural distinctiveness; each group's specific parameters reflect unique characteristics. Rather, it reveals that all operate within the same physical framework, much as diverse chemical systems all obey thermodynamic laws while exhibiting system-specific behaviors.

6.2.3 Multi-Scale Organization of Urban Ethnic Systems

The reconciliation of neighborhood-scale patterns with metropolitan dynamics reveals urban ethnic systems as hierarchically organized structures where processes at each spatial scale both emerge from and constrain those at other scales. This multi-scale organization addresses a persistent challenge in urban geography: connecting individual residential decisions to aggregate spatial patterns (Openshaw, 1984). The characteristic spatial scales discovered in Chapter 4 (34.5–63.0 km) represent fundamental units of ethnic organization, yet these local patterns aggregate into metropolitan-wide structures spanning entire urban regions.

At the finest scale, individual residential decisions aggregate into neighborhood compositions through preferential attachment dynamics. Families choose locations based on proximity to co-ethnic communities, creating positive feedback where existing concentration attracts further settlement, the mechanism underlying (Schelling, 1971)'s tipping model but here quantified through learned reaction terms. When local ethnic population density exceeds critical thresholds, specialized infrastructure becomes economically viable, creating amenities that further concentrate settlement. Below these thresholds, ethnic businesses cannot sustain operations, community organizations lack critical membership, and nascent enclaves dissolve. The nucleation thresholds discovered, ranging from 658 individuals (China) to 8,132 individuals (India), represent phase transitions in settlement dynamics with direct policy implications for minimum viable community sizes.

The characteristic spatial scales emerge from balance between aggregative forces and dispersive pressures, the fundamental tension in reaction-diffusion systems that (Cross & Hohenberg, 1993) formalized as pattern selection in nonequilibrium systems. The diffusion term with ethnic-specific and city-specific coefficients quantifies how populations spread due to housing availability, price gradients, and life-cycle mobility. The characteristic scale represents natural spacing where reaction (clustering) and diffusion (spreading) equilibrate, corresponding to metropolitan infrastructure: regional transit networks, major employment

corridors, and inter-municipal flows. United Kingdom-origin populations with their 63.0 km scale exhibit high mobility and dispersed patterns, while Philippines and Chinese populations with 34.5–38.7 km scales show more concentrated behaviors, differences reflecting systematic relationships between group characteristics and spatial organization rather than arbitrary variation.

Metropolitan-scale patterns emerge from spatial arrangement and interaction of neighborhood-scale units, forming what (M. Batty, 2013) terms "networks and flows" organizing urban systems. The ethnic clusters identified across Canadian cities do not distribute randomly but form hierarchical networks with primary cores, secondary centers, and peripheral satellites. Primary cores contain large populations supporting specialized infrastructure impossible at smaller scales; secondary centers provide intermediate services; peripheral satellites offer basic amenities near critical thresholds. This hierarchical organization exhibits properties of self-organized criticality (Bak & Chen, 1991), where local interactions produce power-law distributions of cluster sizes and characteristic temporal dynamics.

The city-level parameter heterogeneity demonstrates how urban context shapes settlement dynamics, a finding addressing (Sampson, 2012)'s argument that neighborhoods must be understood within their metropolitan contexts. Toronto and Vancouver maintain relatively stable high carrying capacities, while Winnipeg exhibits dramatic 38% decline and Calgary shows resource-boom-driven growth followed by correction. The learned diffusion coefficients show systematic spatial variation, higher in transit-rich areas and lower in car-dependent suburbs, yet fundamental reaction-diffusion dynamics operate consistently across diverse municipal contexts. This combination of universal dynamics with context-specific parameters reconciles nomothetic aspirations for general theory with idiographic attention to particular places.

Transportation networks provide critical multi-scale coupling, connecting neighborhood units into metropolitan systems. The correlation between Chinese and Filipino amplification factors with transit accessibility ($r = 0.58$ and 0.61) reveals how infrastructure shapes settlement across scales, a finding consistent with (Glaeser & Kahn, 2004)'s analysis of transportation and urban spatial structure. Local transit maintains community cohesion through daily circulation within neighborhood units; regional transit connects ethnic enclaves to employment centers, expanding opportunities while preserving residential concentration. The model's learned parameters capture these directed flows, showing how transportation creates anisotropic diffusion channeling rather than dispersing ethnic populations.

The multi-scale organization enables resilience through redundancy and adaptability through reconfiguration, properties characteristic of complex adaptive systems (Holland, 1995). When housing prices in primary cores become prohibitive, populations redistribute to secondary centers that can evolve into new cores. When satellites fall below critical thresholds, their populations absorb into nearby centers. This dynamic equilibrium maintains system stability despite continuous local changes, explaining the

persistence of ethnic spatial structure observed across decades of census data even as specific neighborhoods transform.

This multi-scale synthesis directly addresses the dissertation's overarching research question: whether interpretable GeoAI can convert multi-ethnic settlement from descriptive phenomenon to predictive science. The GraphPDE and MESMoE frameworks demonstrate that physics-informed approaches achieve this conversion by revealing hierarchical organization invisible to black-box methods, quantifying parameters enabling intervention simulation, and maintaining interpretability essential for democratic accountability. The characteristic scales, nucleation thresholds, interaction networks, and configuration barriers discovered constitute a quantitative theory of urban ethnic geography, one that explains observed patterns, predicts future trajectories, and enables counterfactual reasoning about policy alternatives.

6.3 Methodological Contributions to GeoAI

6.3.1 The Interpretable GeoAI Framework

The methodological journey from exploratory black-box approaches through physics-informed modeling to hybrid frameworks represents a fundamental reconceptualization of how artificial intelligence can be applied to geographic problems. This progression, documented across Appendices A and B and crystallized in Chapters 4 and 5, establishes interpretable GeoAI not as a compromise between capability and understanding but as a synthesis achieving both superior performance and theoretical transparency. The framework challenges the orthodox view that complex problems require opaque solutions, demonstrating instead that incorporating domain knowledge through structured architectures enhances rather than constrains computational capability.

The initial explorations using deep learning architectures revealed both the promise and peril of unconstrained neural approaches. The Spatial Gaussian-Bernoulli Deep Belief Network achieved remarkable dimensionality reduction while preserving variance, and the GraphSAGE architecture identified settlement patterns with R^2 exceeding 0.95 for major ethnic groups. Yet these successes came at the cost of interpretability: the learned representations, while mathematically optimal, lacked geographic meaning. The STGAN-EI framework in Appendix B pushed further by incorporating ethnic interaction learning that discovered critical concentration thresholds, yet the learned parameters still lacked theoretical grounding. These limitations were not technical failures but epistemological ones: high accuracy without understanding provides little value for urban planning or theoretical advancement.

The breakthrough came through recognizing that domain knowledge could be directly encoded into model architectures. The GraphPDE framework structures the learning problem around established physics of pattern formation, learning physically meaningful parameters: diffusion coefficients quantifying spatial

mobility, attention-based interaction matrices capturing inter-ethnic dynamics, growth rates revealing demographic vitality, and carrying capacities indicating urban constraints. This structured approach reduces the hypothesis space from infinite possible functions to those consistent with conservation laws, demographic constraints, and observed spatial processes. The normalized graph Laplacian operator naturally encodes spatial relationships on irregular urban geometries, while reaction terms capture population dynamics through logistic growth, immigration, emigration, and inter-ethnic competition. The neural augmentation through city embeddings, period embeddings, and residual MLPs demonstrates that interpretable physics components account for 57.2% of prediction variance while neural residuals contribute 42.8%, showing that the physics foundation provides dominant explanatory power while neural flexibility accommodates real-world complexity.

The MESMoE framework extends this principle through specialized expert modules incorporating different aspects of demographic theory: the Colonization expert embeds nucleation theory for new settlement establishment; the Jump Process expert incorporates stochastic calculus for discontinuous demographic shifts; the Decline expert adapts dissolution dynamics for population exodus; and the PDE Diffusion expert solves reaction-diffusion equations using numerical integration. Each expert represents a different theoretical framework, with a learnable router selecting appropriate models based on local context. This multi-expert architecture resolves a fundamental tension in urban modeling between generality and specificity. Urban systems exhibit universal patterns (characteristic spatial scales, inter-ethnic dynamics) yet also display regime-specific behaviors requiring distinct mathematical treatments. By incorporating multiple theoretical frameworks within a single model, MESMoE captures this heterogeneity while maintaining interpretability, with router decisions showing that colonization dynamics dominate in peripheral areas while diffusion processes explain established communities.

The framework's interpretability extends beyond parameter transparency to include multiple forms of understanding: feature importance analysis revealing which urban characteristics drive different population regimes; ethnicity interaction networks quantifying previously unmeasurable social dynamics, with attention mechanisms revealing that Chinese populations exhibit complex facilitative relationships with multiple groups while showing strong internal competition driving dispersal; and physics-informed parameters enabling spatial visualization of demographic forces through diffusion coefficient maps, growth rate evolution, and carrying capacity variations.

6.3.2 Computational Innovations for Urban Scale

The computational challenges of modeling ethnic settlement across 30,091 dissemination areas in 52 Canadian cities, tracking 9 ethnic groups across 20 years, demanded fundamental innovations in algorithm design, implementation, and optimization. With hundreds of thousands of state variables evolving through

coupled differential equations on irregular graph domains, traditional numerical methods proved computationally intractable. The development of GPU-accelerated solvers, specialized graph neural architectures, and efficient multi-ethnic interaction calculations enabled city-scale simulations, achieving $32.6\times$ speedup while maintaining numerical accuracy and physical consistency.

The adaptation of reaction-diffusion equations to urban graphs required specialized data structures and algorithms. The Compressed Sparse Row (CSR) format minimized memory footprint while enabling efficient matrix-vector operations, and custom CUDA kernels implemented parallel sparse matrix multiplication using optimized memory access patterns maximizing GPU throughput. The 4th-order Runge-Kutta scheme for temporal integration maintained numerical stability while enabling efficient forward simulation over 5-year census intervals. The attention-based competition mechanism with learned interaction weights enabled heterogeneous edge importance, capturing how different ethnic groups weight spatial relationships differently. The census feature modulation through learned multilayer perceptrons captures how local socioeconomic conditions affect demographic rates without requiring separate parameters for each dissemination area.

The physics-informed loss function combining data fidelity with parameter regularization required computing gradients through the PDE solver efficiently. Training used AdamW optimizer with cosine annealing learning rate schedule and gradient clipping for stable optimization despite complex loss landscapes. The multi-scale validation strategy (temporal, spatial, and ethnic) ensured both technical robustness and practical applicability. The GPU-accelerated implementation completed forward and backward passes $32.6\times$ faster than naive PyTorch implementation, enabling extensive hyperparameter optimization and ablation studies. The physics-only variant achieved $R^2 = 0.5588$ (73% of full GraphPDE performance) using exclusively interpretable parameters, demonstrating that reaction-diffusion structure alone encodes powerful demographic prior knowledge.

6.3.3 Validation Across Scales and Contexts

The validation strategy extends beyond conventional machine learning metrics to encompass multiple spatial scales, temporal periods, ethnic groups, and urban contexts, establishing both robustness and generalizability. This comprehensive approach was essential given the framework's scope (predicting settlement patterns across 52 Canadian cities, 9 ethnic groups, and 20 years) and its theoretical claims about learnable physical principles governing human settlement. Multi-dimensional validation revealed not just high predictive accuracy but also consistency in mechanistic parameters across diverse contexts, confirming that discovered dynamics transcend specific urban configurations.

Spatial validation across Canadian metropolitan areas demonstrated remarkable consistency despite vast differences in urban form, governance, and development history. The model achieved comparable

performance in dense urban cores, sprawling suburbs, and peripheral municipalities. Learned diffusion coefficients showed expected spatial variation (higher in transit-rich areas, lower in car-dependent suburbs) yet fundamental reaction-diffusion dynamics held across all contexts. The framework performed well across municipalities with different ethnic compositions: areas with visible minority plurality showed highest accuracy given strong pattern formation, yet areas with no dominant ethnic group achieved comparable performance, demonstrating capture of general settlement dynamics rather than overfitting to specific configurations.

Temporal validation revealed generalization across different socioeconomic contexts. The four training periods (2001-2006 economic growth, 2006-2011 financial crisis, 2011-2016 recovery, 2016-2021 housing boom and pandemic) exhibited distinct demographic patterns, yet core physical parameters remained relatively stable while period-specific parameters accommodated temporal variation. Cross-temporal prediction provided stringent tests: models trained on earlier data achieved meaningful accuracy predicting later patterns despite substantial changes in housing markets, immigration policies, and urban development. Performance degradation arose primarily from changed boundary conditions (new transit infrastructure, housing developments, policy changes) rather than invalid physics.

Ethnic group validation demonstrated both universal dynamics and group-specific parameters. All nine groups exhibited pattern formation consistent with reaction-diffusion physics, confirming universal mechanisms. The framework captured meaningful ethnic variations: Iranian communities showed characteristics suggesting concentrated settlement patterns, while British populations exhibited high mobility and dispersed patterns. These parameters remained consistent across time periods and municipalities, reflecting genuine differences in settlement characteristics. Comparison against alternative approaches (traditional gravity models, spatial regression, black-box deep learning) established competitive performance, with only the interpretable framework providing mechanistic parameters enabling counterfactual analysis and policy simulation.

6.4 Empirical Insights: Multi-Ethnic Settlement Across Canadian Cities

6.4.1 Quantifying the Unquantifiable

The transformation of qualitative observations about ethnic relations into precise mathematical parameters represents one of the most significant empirical contributions of this research. For decades, urban sociologists have described ethnic communities as having "chemistry" where some groups naturally co-locate while others repel, but these observations remained metaphorical rather than measurable. The learned attention-based interaction mechanism, configuration landscape analysis, and multi-stability dynamics

convert these intuitive understandings into quantifiable relationships with specific numerical values, enabling precise prediction and targeted intervention.

The attention-based interaction mechanism reveals the complete network of inter-ethnic influences. The discovery that China exhibits strong negative self-competition (-19.8) quantifies within-group competitive pressure driving spatial dispersal: Chinese communities spread outward from concentrated cores not despite but because of their success, as intra-group competition intensifies beyond optimal levels. Yet this same analysis reveals positive facilitative relationships with six other groups, suggesting complementary economic niches where co-location creates mutual benefit. The Philippines-Pakistan mutual attraction and United Kingdom-Italy mutual repulsion emerged consistently across different cities and time periods, confirming fundamental demographic relationships rather than historical accident.

The configuration landscape analysis transforms abstract concepts of neighborhood stability into concrete analytical quantities. Multiple stable settlement configurations (from concentrated enclaves to integrated patterns) reveal that urban ethnic systems exhibit multi-stability where different spatial arrangements persist under identical demographic parameters depending on historical contingencies. Transition barriers between configurations provide measures of reorganization difficulty. The nucleation threshold differences are striking: India requires 8,132 individuals before clustering emerges versus China forming enclaves at 658 individuals. The high Indian threshold may reflect strong transnational diaspora networks maintaining community cohesion at lower spatial densities, while the low Chinese threshold corresponds to documented rapid ethnoburb formation even in early settlement stages.

City-level parameter heterogeneity provides quantitative measures of geographic variation. Toronto and Vancouver maintain stable high carrying capacities while Winnipeg exhibits 38% decline; Calgary shows resource-boom growth followed by correction. The learned diffusion coefficients range from 2.9 km²/year (Ottawa) to 6.8 km²/year (Rocky View County), with western suburban municipalities exhibiting approximately double the spatial mobility of established urban cores. Substantial within-city variability (ranges extending 2-3× beyond means) reveals marked ethnic heterogeneity in spatial mobility patterns even within the same metropolitan area.

6.4.2 Distinctive Patterns Across Communities

While all ethnic groups show patterns describable by the same mathematical framework, each exhibits distinctive characteristics reflecting unique historical trajectories, economic resources, and cultural preferences. These differences manifest not as violations of reaction-diffusion physics but as group-specific parameters within the universal framework, much as different chemical species follow the same thermodynamic laws while exhibiting distinct properties.

Chinese settlements demonstrate spot morphology predicted by reaction-diffusion theory for groups with strong negative self-competition and moderate diffusion. The characteristic spatial scale of 38.7 km produces well-defined enclaves in Markham, Richmond Hill, and Scarborough, each with distinct boundaries where concentration changes markedly. Iranian communities present striking contrast with the lowest nucleation threshold and concentrated settlement patterns despite limited absolute numbers, reflecting high group cohesion and characteristics enabling rapid enclave formation. The labyrinthine pattern formation regime corresponds to complex, interconnected spatial networks responding strongly to political events in Iran.

Portuguese communities exhibit unique succession dynamics, transitioning from concentrated enclaves to dispersed patterns through time-varying parameters: interaction strength weakening while diffusion increased, reflecting declining immigration and second-generation dispersion. Filipino settlements demonstrate stripe morphology unique among studied groups, forming linear patterns along major transit corridors rather than circular enclaves, reflecting employment in healthcare and service sectors requiring transit access. The MESMoE framework shows Filipino populations uniquely utilizing all four regime experts across different neighborhood contexts.

Indian populations exhibit the most complex settlement patterns, with characteristics reflecting multiple sub-ethnic groups (Punjabi, Gujarati, Tamil) with distinct preferences. The highest nucleation threshold (8,132 individuals) may reflect strong transnational networks enabling community maintenance at lower spatial densities. United Kingdom populations display the highest characteristic spatial scale (63.0 km), indicating high mobility and dispersed settlement, with temporal evolution showing transition from positive to negative growth rates reflecting aging demographics and reduced immigration.

6.4.3 Toronto as Empirical Laboratory: Multi-Scale Settlement Dynamics

The Greater Toronto Area serves as a critical empirical laboratory for validating the physics-informed framework, revealing how learned parameters generate realistic spatial patterns spanning scales from individual neighborhoods to metropolitan geography. At the metropolitan scale, the GraphPDE model's 20-year forward simulation predicts concentrated Chinese settlement in north-central Toronto encompassing 3,700 dissemination areas with total predicted population of 164,155 individuals. This geographic distribution, achieved without explicit supervision on actual settlement locations, shows strong concentration in northern districts (Markham, Richmond Hill) closely matching empirical census patterns. The spatial extent spans approximately 40 km east-west, aligning with the learned characteristic spatial scale (38.7 km).

At neighborhood scales (7.4 km $\times$ 7.4 km), the model generates characteristic spot-and-stripe morphology with alternating high-density spots (>900 individuals/DA) and low-density regions (<100 individuals/DA)

emerging spontaneously from reaction-diffusion instabilities. The observed wavelength of approximately 3-4 km matches the theoretical prediction $\lambda \approx 3.7$ km, corresponding to walkable neighborhood scales. Multi-scale cluster analysis identifies three distinct settlement regimes: peripheral dispersal (<50 grid points) representing scattered households below critical threshold; neighborhood enclaves (50-500 grid points) forming the backbone of settlement structure; and metropolitan hubs (>500 grid points) representing Markham and Richmond Hill cores where agglomeration economies create powerful attractive forces.

Temporal evolution over 20 years reveals three distinct phases: rapid equilibration (0-2.5 years) where spatial variance rises from ~1,750 to ~30,000; pattern emergence (2.5-7.5 years) where variance declines as clustering coefficient rises from 0.15 to 0.85; and gradual maturation (7.5-20 years) through Ostwald ripening where small enclaves dissolve while large enclaves expand. Spatial autocorrelation analysis yields Moran's $I = 0.127$, confirming positive spatial clustering at modest levels indicating clustered but not extreme segregation. The spatiotemporal visualization shows maximum lateral expansion of 6.9 km over 20 years (radial growth rate ~0.35 km/year), with coalescence around 2006-2011 and major consolidation in 2011-2016 when separate domains fuse into unified spatial networks.

Nucleation and growth kinetics exhibit non-monotonic trajectories: cluster numbers rise from ~250 in 2000 to ~400 around 2010, then decline to ~300 by 2021, indicating nucleation dominates early periods while coarsening dominates later periods. The JMAK analysis yields power-law growth exponent $n = 0.6$, indicating sub-linear growth suggesting saturation effects. Spatial correlation structure reveals fractal organization with dimension $D_f \approx 1.36$ across 1-10 km scales, indicating space-filling but not dense settlement where population forms a ramified network penetrating Toronto's geography.

6.4.4 Configuration Stability and Transition Dynamics in Toronto

Multi-stability analysis reveals the complete settlement configuration landscape for Chinese populations in Toronto. The spatial settlement attractiveness map identifies key high-attractiveness zones: North York's central corridor (favorability ≈ -1.2 indicating strong stability), Downtown's eastern neighborhoods (moderate attractiveness ≈ -0.5), and Scarborough's northern regions. Low-attractiveness zones (favorability $> +1.0$) appear in Etobicoke's western periphery where competition from established populations creates unfavorable conditions.

Four stable configurations emerge as local minima: C1 (Enclave) with highest spatial segregation and most stability; C2 (Semi-Segregated) at intermediate stability; C3 (Mixed) with balanced ethnic composition; and C4 (Integrated) with higher diversity. Transition pathway analysis from C1 to C2 exhibits forward reorganization barrier $\Delta B_1 = 1.67$ (demographic "activation energy" required) while reverse barrier $\Delta B_2 = 1.13$. This barrier asymmetry reveals hysteresis: transitions from segregated to less segregated configurations require stronger perturbations than reverse transitions. Medium-height barriers ($\Delta B \approx 1-2$)

suggest configuration transitions are neither spontaneous nor impossible, with sustained interventions over decadal timescales potentially facilitating reorganization.

Critical threshold analysis identifies $\theta_c = 0.20$ where system responsiveness peaks at $R_{\max} = 2.31$. Below θ_c (clustered regime), strong ethnic enclaves and high segregation prevail; above θ_c (dispersed regime), mixed patterns emerge. Current conditions place Toronto's Chinese settlement well below θ_c , explaining why observed enclave patterns are robust. Population flow dynamics show Region 1 declining from normalized fraction 1.0 to ~ 0.75 over 25 years while Region 4 rises from near 0 to ~ 0.75 , with crossing around year 5-7 marking fundamental shift in settlement geography. Characteristic equilibration timescales ($\tau_1 = 1.8$ years, $\tau_2 = 2.2$ years, $\tau_3 = 5.0$ years) quantify regional steady-state approaches.

Geographic mapping identifies 87 stable settlement regions covering 1,154 DAs with total population of 35,065. The top 5 regions by population range from Region 1 (2,425 individuals) to Region 5 (1,620 individuals), demonstrating hierarchical spatial network organization where numerous small stable pockets coexist with larger dominant hubs. The configuration network shows current state at moderate stability, with direct pathway to Balanced Mix showing modest barrier (0.7), while extreme segregation (Pure China) is isolated by high barriers (3.0) and energetically unfavorable.

6.4.5 Emergence of Multi-Ethnic Urban Structures

The discovery of complex multi-ethnic urban structures-emerging from the superposition of multiple ethnic settlement patterns-represents emergent metropolitan organization unpredicted by classical urban theory yet naturally arising from reaction-diffusion dynamics. These spatial patterns, where diversity varies systematically across metropolitan space, challenge conventional models while revealing how multi-ethnic reaction-diffusion processes create novel urban forms. The structures are neither planned nor accidental but emerge from the mathematical interaction of multiple ethnic waves with different spatial scales and origins. Multi-ethnic diversity structures emerge from intersecting mechanisms revealed by the model. Economic gradients create spatial sorting: central area housing costs limit immigrant settlement, while outer areas lack employment access, constraining settlement to intermediate zones. Multiple ethnic patterns with different spatial scales (34.5-63.0 km) and geographic origins create complex overlap patterns. The reaction-diffusion dynamics create interface effects where ethnic patterns meet, generating zones of maximum diversity at the boundaries between distinct ethnic concentrations.

Shannon entropy analysis quantifies diversity patterns with precision across metropolitan space. Different zones display characteristic diversity levels: some areas show lower diversity dominated by single groups, while interface zones achieve maximum entropy approaching theoretical maximum for the studied groups. This spatial diversity gradient, unexpected from traditional models, emerges naturally from reaction-diffusion dynamics of multiple interacting ethnic populations.

The model reveals these diversity structures exhibit complex temporal dynamics. The zones expand and shift as demographic patterns evolve-downtown transformation pushes diversity boundaries while suburban ethnic concentrations intensify. The configuration space analysis shows high-diversity zones occupy intermediate positions, neither the deep stability of ethnic enclaves nor the instability of rapidly changing areas, enabling demographic fluidity.

High-diversity zones create unique urban dynamics absent from homogeneous or strictly segregated areas. Economic analysis reveals higher business density in maximum diversity zones, with ethnic businesses serving multiple communities. Transportation patterns show more cross-ethnic mobility originating in diverse zones. Social network analysis identifies more inter-ethnic connections formed in diverse areas, though these remain weaker than intra-ethnic bonds.

The reaction-diffusion framework explains why diversity structures emerge in specific locations rather than randomly. The spatial positioning corresponds to where primary ethnic patterns begin overlapping without any single group achieving dominance. This creates systems where no group reaches critical mass for complete dominance, yet all maintain sufficient presence for viability-producing maximum diversity through the mathematics of competing influences rather than through policy intervention.

6.5 Policy Implications: From Physics to Planning

The physics-informed understanding of ethnic settlement fundamentally transforms how cities can approach integration policies. The discovery that configuration transitions face asymmetric barriers explains the persistent failure of superficial multicultural programs that address symptoms rather than mechanisms. Effective intervention requires modifying fundamental parameters governing settlement dynamics: increasing diffusion coefficients through enhanced mobility (transit extensions), altering interaction patterns through economic policy (creating complementary rather than competitive niches), and reducing transition barriers through institutional reform (multilingual services, inclusive zoning). The critical nucleation thresholds suggest counterintuitive strategies where helping emerging groups reach critical mass enables "integration through concentration," allowing stable configurations to gradually connect through controlled diffusion rather than forcing unstable dispersal that generates backlash or dramatic reconcentration.

The MESMoE framework's regime-specific formulations and the ethnicity interaction network enable targeted interventions matched to local demographic dynamics. Rather than uniform approaches, effective planning must differentiate strategies by population regime while leveraging cascade effects where investments in strategically positioned groups generate amplified spillover benefits. The quantified interaction network enables deliberate strengthening of facilitative relationships while weakening competitive ones, with cascade multipliers substantially exceeding direct intervention effects. Critical

Table 6-1. Policy Framework by Demographic Regime and Intervention Type

<i>Regime</i>	<i>Characteristics</i>	<i>Policy Goal</i>	<i>Key Interventions</i>	<i>Timing</i>	<i>Critical Parameters</i>
<i>Colonization</i>	New settlement establishment; hierarchical survival probability; small settlements have low viability	Support nucleation to critical threshold	Concentrate rather than disperse immigrants; sequence infrastructure (housing → business → community organizations); target areas above critical mass	Short-term (2-5 years) for initial establishment	Nucleation threshold (658-8,132 individuals depending on group)
<i>Jump Process</i>	Sudden demographic shifts; extreme tail events; nonlinear dynamics	Maintain buffer capacity	Preserve vacant housing stock; ensure scalable services; create flexible infrastructure; prepare contingency plans	Immediate response capability	Jump probability distributions; buffer capacity metrics
<i>Decline</i>	Demographic exodus following characteristic decay; high reversal barriers	Managed succession; prevent rather than reverse	Early warning monitoring (outward flows, elevated diffusion, weakening interactions); preserve cultural heritage through non-residential means; facilitate orderly transition	Preventive window before transformation	Decay time constants; barrier asymmetry (prevention << reversal effort)
<i>Diffusion</i>	Mature neighborhoods in dynamic equilibrium; balanced flows; stable infrastructure	Stability maintenance	Maintain consistent transit service; preserve established business relationships; prevent development perturbations; protect equilibrium conditions	Ongoing maintenance	Diffusion coefficients; interaction pattern stability
<i>Intervention Type</i>		<i>Mechanism</i>	<i>Target Parameter</i>	<i>Expected Impact</i>	<i>Cascade Multiplier</i>
<i>Transit Extension</i>		Increases spatial mobility	Diffusion coefficient	Dispersion from concentrated cores; new settlement opportunities; reduced inter-ethnic competition	2-3× direct ridership impact

<i>Complementary Licensing</i>	<i>Business Specialization</i>	Creates economic interdependence	Interaction strength (positive)	Amplified facilitative relationships; increased co-location incentives	Network-dependent spillovers
<i>Sectoral Zoning</i>		Reduces economic overlap	Interaction strength (negative)	Lowered displacement pressure; enabled coexistence	Reduced competitive cascades
<i>Multilingual Service Delivery</i>		Reduces transition costs	Configuration barriers	Lower reorganization resistance; easier integration pathways	Institutional multiplier
<i>Inclusive Zoning (Multi-generational Housing)</i>		Accommodates cultural preferences	Carrying capacity; barriers	Increased settlement viability; reduced forced dispersal	Family network effects
<i>Minimum Guarantees</i>	<i>Population</i>	Circuit breaker for negative cascades	Stability threshold	Prevents destructive feedback loops; maintains community viability	Prevents hysteresis costs
<i>Policy Timescale</i>		<i>Intervention Type</i>		<i>Demographic Effect</i>	<i>Realistic Horizon</i>
<i>Short-term</i>		Housing/transportation		Population redistribution per new diffusion patterns	2-5 years
<i>Medium-term</i>		Business development/institutional change		Ethnic infrastructure adaptation	5-15 years
<i>Long-term</i>		Complete neighborhood transformation		Generational population turnover	15-30 years

threshold behaviors mean small interventions can trigger disproportionate responses, while negative cascades require circuit breakers (minimum population guarantees, interaction dampers, spatial buffers) to prevent systemic collapse through destructive feedback loops exhibiting hysteresis.

6.6 Broader Implications for Human Geography

The discovery that human settlement patterns can be modeled using reaction-diffusion mathematics fundamentally challenges the ontological foundations of human geography, suggesting that social phenomena previously considered purely cultural or economic may be illuminated through physical frameworks. This reconceptualization does not reduce human complexity to deterministic physics but reveals that collective behavior emerges from the interplay of individual agency and physical constraints, much as thermodynamics emerges from molecular motion without negating quantum uncertainty. The reaction-diffusion framework bridges the persistent divide between idiographic and nomothetic approaches: the discovery of systematic spatial patterns across all ethnic groups demonstrates that general principles exist without erasing particular circumstances, as each group's specific parameters reflect unique historical and cultural factors yet all operate within the same physical framework. The mathematical interpretation transforms qualitative geographic concepts into measurable quantities where "sense of place" becomes quantifiable through configuration stability, "community cohesion" translates to interaction strength, and "neighborhood character" maps to position in pattern formation phase space. The physical framework also resolves longstanding tensions between structure and agency: individual families exercise agency in choosing residential locations, yet millions of individual decisions aggregate into patterns following precise mathematical regularities, just as freely moving molecules produce predictable gas laws.

The methodological journey from exploratory black-box approaches through physics-informed modeling demonstrates how computational methods can advance rather than replace geographic theory. The initial deep learning explorations achieved impressive predictive accuracy but offered little theoretical insight, and only by incorporating domain knowledge through reaction-diffusion equations and mixture of experts architectures did computational power translate into theoretical understanding. Machine learning reveals dynamics invisible to traditional methods: the interaction patterns quantifying inter-ethnic influences could never be measured directly through surveys, yet the model learns these relationships from observed patterns, revealing facilitative and competitive dynamics with statistical significance. Similarly, critical concentration thresholds emerged from data rather than theory, discovered through the model's ability to identify phase transitions in complex systems. The interpretable architecture requirement fundamentally distinguishes geographic from purely predictive applications of AI, because knowing that neighborhoods will change demographically is less valuable than understanding why through growth dynamics, diffusion processes, and interaction cascades. Computational methods enable testing geographic theories at scales

and resolutions previously impossible, with the confirmation that human settlement exhibits systematic spatial patterns validating theoretical predictions while revealing unexpected universality.

The successful application of reaction-diffusion dynamics suggests the possibility of a broader computational social physics identifying mathematical regularities in collective human behavior without denying individual agency or cultural meaning. This framework would seek statistical regularities, phase transitions, and emergent patterns rather than deterministic laws. The universality of certain parameters across diverse contexts hints at fundamental patterns: the characteristic spatial scale range (34.5-63.0 km) appearing across all ethnic groups may reflect transportation constraints, daily activity spaces, or cognitive limits on social network maintenance, while critical nucleation thresholds (658-8,132 individuals) may represent patterns for specialized community infrastructure. Phase transitions in social systems exhibit remarkable similarity to physical phase transitions: enclave formation follows crystallization dynamics, gentrification displays critical phenomena, and political polarization shows symmetry breaking. The emergence of computational social physics raises fundamental questions about prediction, intervention, and human agency that require careful consideration as computational methods become increasingly capable, though the framework developed here suggests one resolution: systems can exhibit statistical regularity while maintaining individual unpredictability.

6.7 Limitations and Critical Reflections

The reaction-diffusion framework successfully predicts settlement patterns and reveals systematic spatial dynamics, yet fundamental limitations constrain what physics can capture and how mathematical models should be interpreted. Critical reflection on these boundaries is essential for responsible application.

What Physics Cannot Capture:

- **Individual agency and lived experience:** A Filipino nurse choosing to live along a transit corridor exercises agency based on shift schedules, family networks, church communities, and personal preferences that cannot be reduced to mathematical parameters, even if aggregate patterns follow predictable physics
- **Path dependence and hidden forces:** Chain migration from specific origin locations, professional networks, and cultural institutions create social structures not captured by reaction-diffusion equations; variable accuracy across ethnic groups reflects that physics cannot encode civil war trauma, refugee resettlement policies, or kinship obligations
- **Cultural meaning and symbolism:** Historic ethnic neighborhoods may persist despite population decline because of symbolic significance; the model predicts population changes but cannot capture why businesses remain, why descendants return for cultural events, or why neighborhoods retain character

- **Power relations and structural inequality:** The model treats ethnic groups as equivalent with different parameters, obscuring vast differences in economic resources, political power, and access to transnational capital; discriminatory lending and systemic racism appear as parameter differences but represent structural violence
- **Global embeddedness:** The closed-system assumption fails to capture transnational networks moving capital, information, and people in ways violating local interaction assumptions
- **Historical contingency and collective memory:** Communities carry historical experience influencing decisions across generations in ways the model cannot capture

The Danger of Reductionism:

- **Naturalization of inequality:** When segregation appears mathematically describable with quantifiable barriers, it becomes tempting to accept ethnic separation as natural rather than produced through historical discrimination and ongoing structural violence
- **Troubling precedents:** Environmental determinism, social Darwinism, and racial science used measurement to naturalize hierarchy; the framework risks similar misuse if configuration barriers are interpreted as immutable rather than socially produced
- **Essentializing cultural differences:** Parameters emerge from contingent structural conditions (accumulated wealth, credential barriers, housing access), not essential ethnic characteristics; mistaking these differences could justify differential treatment
- **Ecological fallacy:** Population-level predictions could be misused to make assumptions about individuals, denying housing or steering families based on group patterns

Methodological Constraints:

- **Temporal resolution:** Census data at five-year intervals misses rapid demographic changes, interim fluctuations, seasonal migration, temporary residents, and undocumented populations
- **Spatial units:** Dissemination areas impose artificial boundaries on continuous processes, create edge effects, and aggregate diverse populations into internally homogeneous units (modifiable areal unit problem)
- **Closed-system assumption:** Significant population flows connect metropolitan areas; international migration, return migration, circular mobility, temporary workers, and international students remain inadequately captured
- **Parameter stability:** Learned parameters may change as communities establish, mature, and transform; period-specific parameters partially address but cannot fully capture longer-term evolution
- **Validation challenge:** Without ground truth for inter-ethnic interactions, configuration barriers, or phase transitions, validation relies on consistency checks; claimed generalizability requires testing

in fundamentally different contexts (Mumbai, Lagos, São Paulo) with different histories and cultural dynamics

6.8 Future Research Directions

The physics-informed framework developed for Canadian cities requires systematic extension to test generalizability, advance methodology, and address urgent global urban challenges.

Extending the Framework:

- **Multi-city comparative studies:** Testing in fundamentally different urban systems (Mumbai's extreme density, Lagos's informal settlements, São Paulo's favelas, Dubai's planned districts) would reveal whether discovered patterns represent universal principles or North American particularities
- **Open system dynamics:** Incorporating continuous immigration flows through source-sink terms rather than treating them as external forcing; predicting future immigration based on geopolitical events, economic cycles, and policy changes
- **Dynamic parameter evolution:** Continuously time-varying coefficients capturing infrastructure development, wealth accumulation, and generational change would enable modeling community lifecycles from emergence through establishment to decline
- **Climate migration integration:** Modified carrying capacities, climate-driven spatial effects, and interaction parameters reflecting resource competition as environmental displacement becomes a primary driver of demographic change

Methodological Advances:

- **Real-time demographic nowcasting:** Mobile phone data, social media activity, transit usage, and utility consumption as high-frequency proxies; filtering techniques assimilating observations between censuses with privacy-preserving methods
- **Agent-based model integration:** Hybrid multi-scale models employing agents for individual household decisions while using reaction-diffusion for neighborhood-scale dynamics, capturing both emergent patterns and generative mechanisms
- **Causal inference from parameters:** Instrumental variables from natural experiments, regression discontinuity at boundaries, and synthetic control methods to identify causal effects, transforming the framework from predictive to explanatory

Applications to Global Urban Challenges:

- **Post-pandemic urban restructuring:** Modeling increased diffusion as location-employment coupling weakens through remote work, altered interactions as virtual communities supplement physical proximity, and predicting which pandemic-induced changes persist versus revert

- **Climate adaptation planning:** Simulating population redistribution from vulnerable areas, infrastructure failure impacts on community viability, and identifying climate-resilient settlement configurations to guide managed retreat
- **Housing affordability interventions:** Identifying critical thresholds below which policies fail to prevent displacement and intervention points where modest investments achieve disproportionate benefits through cascade effects
- **Integration policies in superdiverse cities:** Adapting the framework to cities worldwide facing similar challenges; applying asymmetric barrier insights, ethnicity interaction network approaches, and regime-specific intervention strategies to local contexts

6.9 Conclusion: The Promise of Interpretable GeoAI

This dissertation has demonstrated that interpretable Geospatial Artificial Intelligence can transform multi-ethnic settlement from a descriptive phenomenon into a predictive science while maintaining democratic accountability and theoretical transparency. Through the journey from critiquing algorithmic failures to discovering physical dynamics to building practical tools, the research has shown that the perceived trade-off between computational sophistication and human understanding is false—interpretable approaches achieve superior performance precisely because they incorporate domain knowledge and remain responsive to theoretical constraints.

The critical analysis of urban AI deployments revealed the pervasive “metrics trap” where impressive technical accuracy consistently fails to translate into meaningful social outcomes. ShotSpotter’s 97% acoustic precision yielding only 9.1% crime-fighting effectiveness epitomizes how optimization without understanding produces systematic failure. This critique established not just the desirability but the necessity of interpretable approaches for urban systems affecting millions of residents. The documentation of community resistance victories demonstrated that algorithmic determinism is myth and democratic governance of urban AI remains both necessary and feasible.

The discovery that ethnic settlement patterns emerge from reaction-diffusion dynamics represents a paradigmatic shift in understanding urban demographics. All nine studied ethnic groups exhibit pattern formation consistent with physics, with characteristic spatial scales varying systematically from 34.5 km (Philippines) to 63.0 km (United Kingdom). The learned attention-based interaction mechanism quantifies previously unmeasurable social dynamics—facilitative relationships, competitive pressures, self-regulation effects—transforming qualitative observations into mathematical parameters amenable to analysis and intervention. The identification of nucleation thresholds (658-8,132 individuals), temporal dynamics across census periods, and configuration barriers provides concrete foundations for evidence-based planning.

The Multi-Ethnic Spatial Mixture of Experts framework synthesized physics-informed modeling with machine learning to create practical prediction tools achieving $R^2 = 0.80-0.83$ while maintaining complete interpretability. The regime-specific experts capture different population dynamics through appropriate mathematical formulations. The learned ethnicity interaction network reveals cascade effects enabling strategic interventions. The physics-informed parameters provide mechanistic understanding of how transit infrastructure, housing availability, and economic opportunity shape settlement patterns differentially across ethnic groups.

The empirical insights from Canadian cities' multi-ethnic laboratory reveal both universal principles and cultural particularity. While all groups follow reaction-diffusion physics, each exhibits distinctive patterns: Chinese spot morphology in suburban concentrations, Filipino stripes along transit corridors, Iranian labyrinthine patterns despite small populations, Portuguese succession from enclave to dispersion. The emergence of complex diversity structures represents new urban forms arising from multi-ethnic reaction-diffusion dynamics. These discoveries transform vague concepts into measurable quantities enabling prediction and intervention.

The policy implications are transformative yet require careful application. Understanding that different configuration transitions face asymmetric barriers explains persistent settlement patterns while identifying intervention requirements. The regime-specific approach matches policies to demographic dynamics-supporting colonization to critical thresholds, managing jump processes through buffer capacity, facilitating orderly succession during decline. The cascade framework enables strategic investments in communities generating system-wide benefits. These insights shift planning from intuitive art to evidence-based science. Yet critical reflection reveals fundamental limitations. Physics cannot capture individual agency, cultural meaning, or structural inequalities that shape settlement beyond mathematical patterns. The danger of reductionism-naturalizing inequality through physical metaphors-requires constant vigilance. Methodological constraints from census limitations, boundary effects, and parameter dynamics bound the framework's applicability. These limitations do not negate the approach's value but define its appropriate domain, complementing rather than replacing other ways of understanding urban dynamics.

The broader implications extend beyond ethnic settlement to suggest possibilities for computational social physics identifying mathematical regularities in collective behavior. The success of reaction-diffusion models implies that gentrification, income segregation, and political polarization may follow similar dynamics. The universality of certain parameters hints at fundamental patterns in human spatial organization. The framework points toward future research extending to climate migration, real-time nowcasting, and causal inference from learned parameters.

This transformation of multi-ethnic settlement from description to prediction, from correlation to mechanism, from observation to intervention represents more than methodological advance-it demonstrates

a new mode of geographic inquiry that bridges computational power with theoretical understanding. The interpretable GeoAI framework shows that machine learning need not be a black box, that physics can illuminate social processes, and that computational methods can enhance rather than replace geographic expertise.

As cities worldwide grapple with increasing diversity, housing crises, climate displacement, and the promises and perils of algorithmic governance, the frameworks developed here offer both cautionary lessons and hopeful possibilities. The systematic failures documented in algorithmic urban governance warn against uncritical AI adoption, while the successes of physics-informed approaches demonstrate that interpretable alternatives achieve superior outcomes. The discovery that ethnic geography follows learnable dynamics while maintaining cultural distinctiveness suggests that diversity and order are complementary rather than opposing forces.

The choice facing cities is not between powerful AI and interpretable AI, between computational capability and democratic accountability, or between prediction and understanding. This research demonstrates that we can have both—that interpretability enhances rather than compromises performance, that physics informs rather than constrains machine learning, and that computational methods can serve rather than subvert democratic urban governance. The path forward requires not choosing between human and artificial intelligence but synthesizing both through frameworks that remain transparent, accountable, and responsive to the communities they serve.

In this synthesis lies the promise of interpretable GeoAI: transforming our understanding of cities from mysterious complexity to comprehensible patterns, from helpless observation to informed intervention, from accepting segregation as inevitable to engineering integration as achievable. The physics of human settlement, revealed through computational methods but interpreted through geographic wisdom, provides tools for creating more equitable, integrated, and vibrant multi-ethnic cities. The challenge now is not technical but political—summoning the will to use these insights for collective flourishing rather than algorithmic control, for democratic empowerment rather than technocratic optimization, for understanding that serves humanity rather than equations that constrain it.

References

- A.C.L.U. (2024). *Civil rights advocates achieve the nation's strongest police department policy on facial recognition technology*: American Civil Liberties Union, ACLU of Michigan, & University of Michigan Law School Civil Rights Litigation Initiative.
- A.C.L.U. Massachusetts. (2020). *Press pause on face surveillance*: American Civil Liberties Union of Massachusetts.
- ACER. (2021). *Market Monitoring Report - Energy Retail and Consumer Protection*. Retrieved from https://www.ceer.eu/wp-content/uploads/2024/09/ACER-CEER_2024_MMR_Retail-1.pdf
- Ahmadi, D., & Tasan-Kok, T. (2014). Urban policies on diversity in Toronto, Canada. *Delft University of Technology*.
- Akpinar, N. J., De-Arteaga, M., & Chouldechova, A. (2021, 2021). *The effect of differential victim crime reporting on predictive policing systems*.
- Alba, R. D., & Logan, J. R. (1991). Variations on two themes: Racial and ethnic patterns in the attainment of suburban residence. *Demography*, 28(3), 431-453. doi:10.2307/2061466
- Aldunate, P., & Hameed, R. (2022). CENSUS BULLETIN #5: Immigration and ethnocultural diversity Housing Aboriginal peoples. Retrieved from <https://investbrampton.ca/wp-content/uploads/2022/11/2021-Census-Bulletin-5-Ethnocultural-Diversity-Housing-and-Aboriginal-Peoples-Final-004.pdf>
- Ali, I. (2025). Data-driven machine learning approach based on physics-informed neural network for population balance model. *Advances in Continuous and Discrete Models*, 2025(1), 12.
- Alis, C., Legara, E. F., & Monterola, C. (2021). Generalized radiation model for human migration. *Scientific reports*, 11(1), 22707.
- Allam, Z., & Dhunny, Z. A. (2019). On big data, artificial intelligence and smart cities. *Cities*, 89, 80–91. doi:10.1016/j.cities.2019.01.032
- Alsabhan, W., & Alotaiby, T. (2022). Automatic building extraction on satellite images using Unet and ResNet50. *Computational Intelligence and Neuroscience*, 2022(1), 5008854.
- Amnesty International. (2022). *We sense trouble: Automated discrimination and mass surveillance in predictive policing in the Netherlands*: Amnesty International.
- Amsterdam Municipality. (2020). *Amsterdam Algorithm Register [Beta version]*: City of Amsterdam, Chief Technology Office.
- Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society*, 20(3), 973-989.
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, 2016-05-23). Machine bias: There's software used across the country to predict future criminals. Retrieved from <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Anselin, L. (1988). *Spatial econometrics: methods and models* (Vol. 4): Springer Science & Business Media.
- Arcaute, E., Molinero, C., Hatna, E., Murcio, R., Vargas-Ruiz, C., Masucci, A. P., & Batty, M. (2016). Cities and regions in Britain through hierarchical percolation. *Royal Society open science*, 3(4), 150691.
- Ashutosh, I. (2014). From the census to the city: Representing South Asians in Canada and Toronto. *Diaspora: A Journal of Transnational Studies*, 17(2), 130-148.
- Austin Energy. Demand response programs.
- AutoGrid Systems. *DROMS case study: Austin Energy*: AutoGrid.
- Avrami, M. (1939). Kinetics of phase change. I: General theory. *Journal of Chemical Physics*, 7(12), 1103-1112.
- Avrami, M. (1941). Kinetics of phase change. III: Granulation, phase change, and microstructure. *Journal of Chemical Physics*, 9(2), 177-184.
- Bak, P., & Chen, K. (1991). Self-organized criticality. *Scientific American*, 264(1), 46-53. doi:10.1038/scientificamerican0191-46

- Balakrishnan, T. R., & Hou, F. (1999). Socioeconomic integration and spatial residential patterns of immigrant groups in Canada. *Population Research and Policy Review*, 18(3), 201-217. doi:10.1023/A:1006118502121
- Barcelona City Council. (2021). *Government measure for a municipal algorithms and data strategy for an ethical promotion of artificial intelligence*: Barcelona City Council.
- Barnes, T. J. (2004). Placing ideas: genius loci, heterotopia and geography's quantitative revolution. *Progress in Human Geography*, 28(5), 565-595.
- Bartlett, L., & Vavrus, F. (2016). Rethinking case study research: A comparative approach.
- Batty, M. (2005). *Cities and complexity: Understanding cities with cellular automata, agent-based models, and fractals*: MIT Press.
- Batty, M. (2013). *The new science of cities*: MIT Press.
- Batty, M., Xie, Y., & Sun, Z. (1999). Modeling urban dynamics through GIS-based cellular automata. *Computers, Environment and Urban Systems*, 23(3), 205-233.
- Bauder, H., & Sharpe, B. (2002). Residential segregation of visible minorities in Canada's gateway cities. *Canadian Geographer/Le Géographe canadien*, 46(3), 204-222.
- Behnisch, M., Schorcht, M., Kriewald, S., & Rybski, D. (2019). Settlement percolation: A study of building connectivity and poles of inaccessibility. *Landscape and Urban Planning*, 191, 103631.
- Belkin, M., & Niyogi, P. (2008). Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8), 1289-1308.
- Bell, J. (2023). Edmonton Region leads in attracting and retaining international talent. Retrieved from <https://edmontonglobal.ca/news/dyk-attracting-and-retaining-international-talent-in-canda/>
- Bellocchi, L., & Geroliminis, N. (2020). Unraveling reaction-diffusion-like dynamics in urban congestion propagation: Insights from a large-scale road network. *Scientific reports*, 10(1), 4876.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828. doi:10.1109/TPAMI.2013.50
- Bengio, Y., Louradour, J., Collobert, R., & Weston, J. (2009). *Curriculum learning*. Paper presented at the Proceedings of the 26th annual international conference on machine learning.
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*: Polity.
- Berry, B. J. (1967). Geography of market centers and retail distribution.
- Bettencourt, L. M. (2020). Urban growth and the emergent statistics of cities. *Science advances*, 6(34), eaat8812.
- Bettencourt, L. M. A., Lobo, J., Helbing, D., Kühnert, C., & West, G. B. (2007). Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the National Academy of Sciences*, 104(17), 7301-7306. doi:10.1073/pnas.0610172104
- Bijker, W. E., Hughes, T. P., & Pinch, T. (2012). *The social construction of technological systems: New directions in the sociology and history of technology* (W. E. Bijker, T. P. Hughes, & T. Pinch Eds. Anniversary ed.): MIT Press.
- Biles, J., Burstein, M., & Frideres, J. (2008). Canadian society: Building inclusive communities. *Immigration and integration in Canada in the twenty-first century*, 269-278.
- Boston City Council. (2020). *Ordinance banning facial recognition technology in Boston*: City of Boston.
- Boulianne, A., & Guilbault, J.-P. (2016). Vivez-vous dans un quartier riche ou pauvre à Montréal? La réponse en cartes. Retrieved from <https://ici.radio-canada.ca/nouvelles/special/2016/5/carte-revenu-moyen-riches-pauvres-montreal/index.html>
- Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal*, 9(2), 27–40. doi:10.3316/QRJ0902027
- Braun, V., & Clarke, V. (2019). Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4), 589–597. doi:10.1080/2159676X.2019.1628806
- Bravo, T. (2020). *Thoma Bravo to acquire RealPage for \$9.6 billion*: Business Wire.

- Brayne, S. (2017). Big data surveillance: The case of policing. *American Sociological Review*, 82(5), 977-1008.
- Breiman, L. (2001). Random forests. *Machine learning*, 45, 5-32.
- Brennan Center for Justice. (2024, 2024). Predictive policing explained. Retrieved from <https://www.brennancenter.org/our-work/research-reports/predictive-policing-explained>
- Breton, R. (1964). Institutional completeness of ethnic communities and the personal relations of immigrants. *American Journal of Sociology*, 70(2), 193-205. doi:10.1086/223793
- Brickner, R. K., & Straehle, C. (2010). The missing link: Gender, immigration policy and the Live-in Caregiver Program in Canada. *Policy and Society*, 29(4), 309-320. doi:10.1016/j.polsoc.2010.09.004
- Bronars, S. G., & Lynch, G. (2022, 2021). Independent Audit of the ShotSpotter Accuracy, 2019-2022. Retrieved from <https://www.edgeworthetheconomics.com/experience-independent-audit-of-the-shotspotter-accuracy>
- Brown, L. A., & Sanders, R. L. (1981). Toward a development paradigm of migration, with particular reference to Third World settings *Migration decision making* (pp. 149-185): Elsevier.
- Burayidi, M. A. (2015). *Cities and the politics of difference: Multiculturalism and diversity in urban planning*: University of Toronto Press.
- Burchfield, M., & Kramer, A. (2015). The Toronto region's greenfield development and intensification (2001-2011). Retrieved from <https://neptis.org/publications/chapters/toronto-regions-greenfield-development-and-intensification-2001-2011>
- Burr, K. (2011). Local immigration partnerships: Building welcoming and inclusive communities through multi-level governance. *Horizons Policy Research Initiative*, Government of Canada.
- Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big data & society*, 3(1), 2053951715622512.
- Burton, I. (1963). The quantitative revolution and theoretical geography 1. *Canadian Geographer/Le Géographe canadien*, 7(4), 151-162.
- Butcher, J. C. (2016). *Numerical methods for ordinary differential equations*: John Wiley & Sons.
- Buzzelli, M. (2001). From Little Italy to Little Portugal: Parish and community building in Toronto, 1900–1990. *Journal of Historical Geography*, 27(4), 573-587. doi:10.1006/jhge.2001.0363
- C.C.L.A. (2020). *CCLA v. Waterfront Toronto, et al*: Canadian Civil Liberties Association.
- Cahn, J. W., & Hilliard, J. E. (1958). Free energy of a nonuniform system. I. Interfacial free energy. *The Journal of chemical physics*, 28(2), 258-267.
- Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and program planning*, 2(1), 67-90.
- Canadian Multiculturalism Act. (1988). Canadian multiculturalism act. Retrieved on March, 4, 2013.
- Cao, W., Dong, L., Wu, L., & Liu, Y. (2020). Quantifying urban areas with multi-source data based on percolation theory. *Remote Sensing of Environment*, 241, 111730.
- Cao, Y., Wang, H., Guo, L., Zhang, A., & Wu, X. (2025). Interpretable Machine Learning for Population Spatialization and Optimal Grid Scale Selection in Shanghai. *Applied Sciences*, 15(9), 4755.
- Cardullo, P., & Kitchin, R. (2019). Being a 'citizen' in the smart city: Up and down the scaffold of smart citizen participation in Dublin, Ireland. *GeoJournal*, 84(1), 1-13.
- CBC News. (2020, 2020-05-07). Sidewalk Labs cancels plan to build high-tech neighbourhood in Toronto amid COVID-19. Retrieved from <https://www.cbc.ca/news/canada/toronto/sidewalk-labs-cancels-project-1.5559370>
- CBS News. (2024, 2024). LAPD still all-in on data-driven policing after scrapping controversial LASER program. Retrieved from <https://www.cbsnews.com/news/los-angeles-police-department-laser-data-driven-policing-racial-profiling-2-0-cbsn-originals-documentary/>
- Chan, A. (2012). From Chinatown to Ethnoburb: The Chinese in Toronto. *Vancouver: University of British Columbia Library*, 14288, 1.0103074.

- Chandrasekhar, R. (2022). Singapore's "smart city" initiative: One step further in the surveillance, regulation and disciplining of those at the margins. Retrieved from <https://chrgj.org/2022-03-18-singapore-smart-city-initiative/>
- Chen, T., & Guestrin, C. (2016). *Xgboost: A scalable tree boosting system*. Paper presented at the Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining.
- Cho, E. Y. (2017). Revisiting ethnic niches: A comparative analysis of the labor market experiences of Asian and Latino undocumented young adults. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 3(4), 97-115.
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2), 153-163.
- Chung, F. R. (1997). *Spectral graph theory* (Vol. 92): American Mathematical Soc.
- Citron, D. K., & Pasquale, F. (2014). The scored society: Due process for automated predictions. *Wash. L. Rev.*, 89, 1.
- City of Chicago Office of Inspector General. (2021, 2021-08-24). The Chicago Police Department's use of ShotSpotter technology. Retrieved from <https://igchicago.org/2021/08/24/oig-finds-that-shotspotter-alerts-rarely-lead-to-evidence-of-a-gun-related-crime-and-that-presence-of-the-technology-changes-police-behavior/>
- City of Helsinki. (2020). *Helsinki AI Register*: City of Helsinki, Digital and Population Data Services Agency.
- CIVICUS. (2024). People power under attack 2024 (CIVICUS Monitor Annual Report). Retrieved from <https://civicusmonitor.contentfiles.net/media/documents/GlobalFindings2024.EN.pdf>
- Clark, W. A. V., & Fossett, M. (2008). Understanding the social context of the Schelling segregation model. *Proceedings of the National Academy of Sciences*, 105(11), 4109-4114. doi:10.1073/pnas.0708155105
- Cmhc. (2024). *Housing market information portal*. Retrieved from <https://www.cmhc-schl.gc.ca/>
- Coleman, J. S. (1990). *Foundations of social theory*: Harvard University Press.
- Collins, H. M., & Evans, R. (2002). The third wave of science studies: Studies of expertise and experience. *Social studies of science*, 32(2), 235-296.
- CoreLogic. (2021). *Stone Point Capital and Insight Partners to acquire CoreLogic for \$6 billion*: CoreLogic Press Release.
- Courant, R., Friedrichs, K., & Lewy, H. (1967). On the partial difference equations of mathematical physics. *IBM journal of Research and Development*, 11(2), 215-234.
- Cowen, D., & Parlette, V. (2011). *Inner suburbs at stake: Investing in social infrastructure in Scarborough*: Cities Centre, University of Toronto.
- Cross, M. C., & Hohenberg, P. C. (1993). Pattern formation outside of equilibrium. *Reviews of Modern Physics*, 65(3), 851-1112. doi:10.1103/RevModPhys.65.851
- Cugurullo, F., Caprotti, F., Cook, M., Karvonen, A., McGuirk, P., & Marvin, S. (2024). The rise of AI urbanism in post-smart cities: A critical commentary on urban artificial intelligence. *Urban studies*, 61(6), 1168-1182.
- Cui, C., Hu, Y., Bao, Y., & Li, H. (2024). Population Density Prediction at Township Scale Supported by Machine Learning Method: A Case Study in Inner Mongolia. *ISPRS International Journal of Geo-Information*, 13(12), 426.
- Daniels, P. W., & Warnes, A. M. (2013). *Movement in cities: Spatial perspectives on urban transport and travel*: routledge.
- Davis, M. H. (1984). Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 46(3), 353-376.
- Denzin, N. K., & Lincoln, Y. S. (2017). *The SAGE handbook of qualitative research* (N. K. Denzin & Y. S. Lincoln Eds. 5th ed.): SAGE Publications.

- Doda, S., Kahl, M., Ouan, K., Obadic, I., Wang, Y., Taubenböck, H., & Zhu, X. X. (2024). Interpretable deep learning for consistent large-scale urban population estimation using Earth observation data. *International Journal of Applied Earth Observation and Geoinformation*, 128, 103731.
- Dolowitz, D. P., & Marsh, D. (2000). Learning from abroad: The role of policy transfer in contemporary policy-making. *Governance*, 13(1), 5-23.
- Doucette, M. L., Green, C., Necci Dineen, J., Shapiro, D., & Raissian, K. M. (2021). Impact of ShotSpotter technology on firearm homicides and arrests among large metropolitan counties: A longitudinal analysis, 1999-2016. *Journal of Urban Health*, 98(5), 609–621. doi:10.1007/s11524-021-00515-4
- Drehobl, A., Ross, L., & Ayala, R. (2020). *How high are household energy burdens? An assessment of national and metropolitan energy burden across the United States*: American Council for an Energy-Efficient Economy.
- E.S.A./Copernicus. (2021). *Urban Atlas 2018 - Land cover/land use*: European Space Agency/Copernicus Land Monitoring Service.
- Eckstein, S., & Peri, G. (2018). Immigrant niches and immigrant networks in the US labor market. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 4(1), 1-17.
- Ecube Labs. (2015). *Seoul Metropolitan Government case study*: Ecube Labs.
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. doi:10.5465/amr.1989.4308385
- Elliot. (2025). Amsterdam Built the ‘Perfect’ Ethical AI System. It Still Failed. Here’s Why. Retrieved from <https://medium.com/@elliotJL/amsterdam-built-the-perfect-ethical-ai-system-it-still-failed-here-s-why-8dc8072beea3>
- Ensign, D., Friedler, S. A., Neville, S., Scheidegger, C., & Venkatasubramanian, S. (2018). Runaway feedback loops in predictive policing. *Proceedings of Machine Learning Research*, 81, 160–171.
- Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). *A density-based algorithm for discovering clusters in large spatial databases with noise*. Paper presented at the Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining.
- EU Public Buyers Community. (2024, 2024). Updated EU AI model contractual clauses. Retrieved from <https://public-buyers-community.ec.europa.eu/communities/procurement-ai/resources/updated-eu-ai-model-contractual-clauses>
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St: Martin's Press.
- F.T.C. (2018). *Texas company will pay \$3 million to settle FTC charges*: Federal Trade Commission.
- F.T.C. (2020). Tenant background report provider settles FTC allegations. *Federal Trade Commission*. doi:<https://www.ftc.gov/news-events/news/press-releases/2020/12/>
- Fan, C., Xu, J., Natarajan, B. Y., & Mostafavi, A. (2023). Interpretable machine learning learns complex interactions of urban features to understand socio-economic inequality. *Computer-Aided Civil and Infrastructure Engineering*, 38(14), 2013-2029.
- Fan, Z., Song, X., Shibasaki, R., & Adachi, R. (2015). *Citymomentum: an online approach for crowd behavior prediction at a citywide level*. Paper presented at the Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing.
- Federal Redistribution. (2022). Details of the Proposed Redistribution Plan. Retrieved from https://redecoupage-redistribution-2022.ca/com/on/prop/othaut/meta_e.aspx
- Ferguson, A. G. (2017). *The rise of big data policing: Surveillance, race, and the future of law enforcement*. The rise of big data policing: New York University Press.
- Fisher, R. A. (1937). The wave of advance of advantageous genes. *Annals of Eugenics*, 7(4), 355-369.
- Fletcher, R. (2024). Calgary population surges by staggering 6%, Edmonton by 4.2% in latest StatsCan estimates. Retrieved from <https://www.cbc.ca/news/canada/calgary/calgary-edmonton-cmas-july-2023-population-estimates-2024-data-release-1.7210191>
- Flyvbjerg, B. (2006). Five misunderstandings about case-study research. *Qualitative Inquiry*, 12(2), 219–245. doi:10.1177/1077800405284363

- Foner, N., Duyvendak, J. W., & Kasinitz, P. (2020). Introduction: Super-diversity in everyday life *Super-Diversity in Everyday Life* (pp. 1-16): Routledge.
- Fong, E., & Wilkes, R. (2003). *Racial and ethnic residential patterns in Canada*. Paper presented at the Sociological Forum.
- Fosset, M. (1999). *Ethnic preferences, social distance dynamics, and residential segregation: results from simulation analyses*. Paper presented at the The Annual Meetings of the American Sociological Association.
- Fotheringham, A. S., Brunsdon, C., & Charlton, M. (2009). Geographically weighted regression. *The Sage handbook of spatial analysis*, 1, 243-254.
- Frey, W. H. (1979). Central city white flight: Racial and nonracial causes. *American sociological review*, 425-448.
- Fung, A. (2009). Empowered participation: Reinventing urban democracy.
- Fung, A., Graham, M., & Weil, D. (2007). *Full disclosure: The perils and promise of transparency*: Cambridge University Press.
- Fung, A., & Wright, E. O. (2003). Countervailing Power in Empowered Participatory. *Deepening democracy: Institutional innovations in empowered participatory governance*, 4, 259.
- Gal, Y., & Ghahramani, Z. (2016). *Dropout as a bayesian approximation: Representing model uncertainty in deep learning*. Paper presented at the international conference on machine learning.
- Gao, S. (2021). *Geospatial artificial intelligence (GeoAI)* (Vol. 10): Oxford University Press New York.
- Garha, N. S., & Paparusso, A. (2018). Fragmented integration and transnational networks: a case study of Indian immigration to Italy and Spain. *Genus*, 74(1), 12.
- Gaskin, T., & Abel, G. J. (2025). Deep learning 40 years of human migration. *arXiv preprint arXiv:2506.22821*.
- Georganos, S., Hafner, S., Kuffer, M., Linard, C., & Ban, Y. (2022). A census from heaven: Unraveling the potential of deep learning and Earth Observation for intra-urban population mapping in data scarce environments. *International Journal of Applied Earth Observation and Geoinformation*, 114, 103013.
- Ghosh, S. (2007). Transnational ties and intra-immigrant group settlement experiences: A case study of Indian Bengalis and Bangladeshis in Toronto. *GeoJournal*, 68(2), 223-242.
- Ghosh, S. (2007). Transnational ties and intra-immigrant group settlement experiences: A case study of Indian Bengalis and Bangladeshis in Toronto. *GeoJournal*, 68(2-3), 223-242. doi:10.1007/s10708-007-9081-x
- Gibbs, J. W. (1902). *Elementary principles in statistical mechanics*: Charles Scribner's Sons.
- Gillborn, D., Warmington, P., & Demack, S. (2018). QuantCrit: Education, policy, 'Big Data' and principles for a critical race theory of statistics. *Race Ethnicity and Education*, 21(2), 158-179. doi:10.1080/13613324.2017.1377417
- Ginzburg, V. L., & Landau, L. D. (1950). On the theory of superconductivity. *Zhurnal Eksperimental'noi i Teoreticheskoi Fiziki*, 20, 1064-1082.
- Glaeser, E. L., & Kahn, M. E. (2004). Sprawl and urban growth. In J. V. Henderson & J.-F. Thisse (Eds.), *Handbook of regional and urban economics* (Vol. 4, pp. 2481-2527): Elsevier.
- Global Property Guide. (2025). Canada's Residential Property Market Analysis 2025. Retrieved from <https://www.globalpropertyguide.com/north-america/canada/price-history>
- Gomez-Lievano, A., Patterson-Lomba, O., & Hausmann, R. (2016). Explaining the prevalence, scaling and variance of urban phenomena. *Nature Human Behaviour*, 1(1), 0012.
- Good, K. (2009). *Municipalities and multiculturalism: The politics of immigration in Toronto and Vancouver*: University of Toronto Press.
- Goodchild, M. F. (1992). Geographical information science. *International journal of geographical information systems*, 6(1), 31-45.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep feedforward networks. *Deep learning*, 1, 161-217.

- Gounaridis, D., Chorianopoulos, I., Symeonakis, E., & Koukoulas, S. (2019). A Random Forest-Cellular Automata modelling approach to explore future land use/cover change in Attica (Greece), under different socio-economic realities and scales. *Science of The Total Environment*, 646, 320-335.
- Green, A. G., & Green, D. A. (1999). The economic goals of Canada's immigration policy: Past and present. *Canadian Public Policy/Analyse de Politiques*, 25(4), 425-451. doi:10.2307/3552422
- Green, D. A., & Milligan, K. (2010). The importance of the long form census to Canada. *Canadian Public Policy*, 36(3), 383-388. doi:10.3138/cpp.36.3.383
- Greenfield, A. (2013). *Against the Smart City: A Pamphlet. This is Part I of "The City is Here to Use": Do projects.*
- Gunning, D. (2016). Broad agency announcement explainable artificial intelligence (XAI). *Defense Advanced Research Projects Agency (DARPA), Tech. Rep.*
- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). *On calibration of modern neural networks.* Paper presented at the International conference on machine learning.
- Haan, M. (2012). The housing experiences of new Canadians: Insights from the Longitudinal Survey of Immigrants to Canada. *Citizenship and Immigration Canada.*
- Hackworth, J., & Rekers, J. (2005). Ethnic packaging and gentrification: The case of four neighborhoods in Toronto. *Urban Affairs Review*, 41(2), 211-236. doi:10.1177/1078087405280859
- Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Hänggi, P., Talkner, P., & Borkovec, M. (1990). Reaction-rate theory: fifty years after Kramers. *Reviews of modern physics*, 62(2), 251.
- Hao, B., Sotudian, S., Wang, T., Xu, T., Hu, Y., Gaitanidis, A., . . . Paschalidis, I. C. (2020). Early prediction of level-of-care requirements in patients with COVID-19. *Elife*, 9, e60519.
- Hartig, J., Friesen, J., & Pelz, P. F. (2020). Investigation of Clustering in Turing Patterns to Describe the Spatial Relations of Slums.
- Harvey, D. (1969). Explanation in geography.
- Hatna, E., & Benenson, I. (2012). The Schelling model of ethnic residential dynamics: Beyond the integrated-segregated dichotomy of patterns. *Journal of Artificial Societies and Social Simulation*, 15(1), 6.
- Hedström, P., & Ylikoski, P. (2010). Causal mechanisms in the social sciences. *Annual Review of Sociology*, 36, 49-67. doi:10.1146/annurev.soc.012809.102632
- Henn, K., Friesen, J., Hartig, J., & Pelz, P. F. (2020). Spatial analysis of settlement structures to identify pattern formation mechanisms in inter-urban systems. *ISPRS International Journal of Geo-Information*, 9(9), 541.
- Hiebert, D. (2000). Immigration and the changing Canadian city. *Canadian Geographer/Le Géographe canadien*, 44(1), 25-43. doi:10.1111/j.1541-0064.2000.tb00691.x
- Hiebert, D. (2015). Ethnocultural minority enclaves in Montreal, Toronto and Vancouver. *IRPP Study No.* 52.
- Hirschman, A. O. (1970). Exit, Voice, and. *Loyalty: Responses to Decline in.*
- Holland, J. H. (1995). *Hidden order: How adaptation builds complexity:* Addison-Wesley.
- Hollands, R. G. (2020). Will the real smart city please stand up?: Intelligent, progressive or entrepreneurial? *The Routledge companion to smart cities* (pp. 179-199): Routledge.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5), 359-366.
- Hou, F. (2004). *Recent immigration and the formation of visible minority neighbourhoods in Canada's large cities:* Statistics Canada.
- Hou, F., & Picot, G. (2004). Visible minority neighbourhoods in Toronto, Montréal, and Vancouver. *Canadian Social Trends*, 72, 8, 13.
- Hou, F., & Picot, G. (2016). Changing immigrant characteristics and pre-landing Canadian earnings: Their effect on entry earnings over the 1990s and 2000s. *Canadian Public Policy*, 42(3), 308-323.

- Hou, F., & Picot, G. (2022). Immigrant labour market outcomes during recessions: Comparing the early 1990s, late 2000s and COVID-19 recessions: Statistics Canada= Statistique Canada.
- Huang, X., Zhu, D., Zhang, F., Liu, T., Li, X., & Zou, L. (2021). Sensing population distribution from satellite imagery via deep learning: Model selection, neighboring effects, and systematic biases. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14, 5137-5151.
- Hulchanski, D. J., Bourne, L. S., Egan, R., Fair, M., Maaranen, R., Murdie, R. A., & Walks, A. R. (2010). *The three cities within Toronto: Income polarization among Toronto's neighbourhoods, 1970-2005*: Cities Centre Press, University of Toronto.
- Hulchanski, J. D. (2010). *The three cities within Toronto: Income polarization among Toronto's neighbourhoods, 1970-2005*: Cities Centre Press, University of Toronto.
- Human Rights Watch. (2017). Singapore: Laws chill free speech, assembly. Retrieved from <https://www.hrw.org/news/2017/12/13/singapore-laws-chill-free-speech-assembly>
- Hyndman, J. (2003). Aid, conflict, and migration: The Canada–Sri Lanka connection. *Canadian Geographer/Le Géographe canadien*, 47(3), 251-268. doi:10.1111/1541-0064.00021
- I.C.O. (2018). *Enforcement notice: Metropolitan Police Service*: Information Commissioner's Office.
- IBM. (2007). *IBM and Singapore's Land Transport Authority pilot innovative traffic prediction tool*: IBM Press Release.
- IRCC, I., Refugees and Citizenship Canada). (2024). *2024 Annual report to Parliament on immigration*. Retrieved from
- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural computation*, 3(1), 79-87.
- Janowicz, K., Gao, S., McKenzie, G., Hu, Y., & Bhaduri, B. (2020). GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond (Vol. 34, pp. 625-636): Taylor & Francis.
- Jasanoff, S. (2004). *States of knowledge*: Taylor & Francis Abingdon, UK.
- Jordan, M. I., & Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural computation*, 6(2), 181-214.
- Kajuli, J., & Mayengo, M. (2025). Mathematical modeling of refugee population dynamics and its impact on deforestation in Tanzania: An ODE-based and neural network-enhanced approach. *Results in Control and Optimization*, 100591.
- Kanevski, M., Timonin, V., & Pozdnukhov, A. (2009). *Machine learning for spatial environmental data: theory, applications, and software*: EPFL press.
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422-440.
- Katzenbach, C., & Ulbricht, L. (2019). Algorithmic governance. *Internet Policy Review*, 8(4), 1–18. doi:10.14763/2019.4.1424
- Keck, M. E., & Sikkink, K. A. (2014). *Activists beyond borders: Advocacy networks in international politics*: Cornell University Press.
- Kelly, P., Park, S., Leon, C., & Priest, J. (2011). *Profile of live-in caregiver immigrants to Canada, 1993-2009*: Toronto Immigrant Employment Data Initiative.
- Kim, M., Kim, D., Jin, D., & Kim, G. (2023). Application of explainable artificial intelligence (XAI) in urban growth modeling: A case study of Seoul Metropolitan Area, Korea. *Land*, 12(2), 420. doi:10.3390/land12020420
- Kim, Y., Safikhani, A., & Tepe, E. (2022). Machine learning application to spatio-temporal modeling of urban growth. *Computers, Environment and Urban Systems*, 94, 101801.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kipf, T. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

- Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*: SAGE.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. *Journal of legal analysis*, 10, 113-174.
- Knowles, V. (2016). *Strangers at our gates: Canadian immigration and immigration policy, 1540–2015*: Dundurn Press.
- Kolmogorov, A., Petrovsky, I., & Piskunov, N. (1937). Study of the diffusion equation with growth of the quantity of matter and its application to a biological problem. *Moscow University Mathematics Bulletin*, 1, 1-25.
- Kolmogorov, A. N. (1937). On the statistical theory of metal crystallization. *Izvestiya Akademii Nauk SSSR, Seriya Matematicheskaya*, 1(3), 355-359.
- Kondo, S., & Miura, T. (2010). Reaction-diffusion model as a framework for understanding biological pattern formation. *Science*, 329(5999), 1616-1620.
- Kovalevsky, D. V., Volchenkov, D., & Scheffran, J. (2021). Cities on the coast and patterns of movement between population growth and diffusion. *Entropy*, 23(8), 1041.
- Kramers, H. A. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4), 284-304.
- Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms, 165 U. Pa. L. Rev, 633, 633.
- Kruschel, S., Hambauer, N., Weinzierl, S., Zilker, S., Kraus, M., & Zschech, P. (2025). Challenging the Performance-Interpretability Trade-off: An Evaluation of Interpretable Machine Learning Models. *Business & Information Systems Engineering*, 1-25.
- Kymlicka, W. (1998). *Finding our way: Rethinking ethnocultural relations in Canada*: Oxford University Press.
- Lei, B., Liu, P., Milojevic-Dupont, N., & Biljecki, F. (2024). Predicting building characteristics at urban scale using graph neural networks and street-level context. *Computers, Environment and Urban Systems*, 111, 102129.
- Levi-Faur, D. (2012). *The Oxford handbook of governance*: Oxford University Press.
- Ley, D. (2010). *Millionaire migrants: Trans-Pacific life lines*: Wiley-Blackwell.
- Ley, D. (2011). *Millionaire migrants: Trans-Pacific life lines*: John Wiley & Sons.
- Li, D., Yu, Y., & Wang, B. (2023). Urban population prediction based on multi-objective lioness optimization algorithm and system dynamics model. *Scientific Reports*, 13(1), 11836.
- Li, P. S. (1998). *The Chinese in Canada*: Oxford University Press.
- Li, W. (2020). GeoAI: Where machine learning and big data converge in GIScience. *Journal of Spatial Information Science*(20), 71-77.
- Li, W., Hsu, C.-Y., & Hu, M. (2021). Tobler's First Law in GeoAI: A spatially explicit deep learning model for terrain feature detection under weak supervision. *Annals of the American Association of Geographers*, 111(7), 1887-1905.
- Li, Y., Yu, R., Shahabi, C., & Liu, Y. (2017). Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*.
- Li, Z. (2024). GeoShapley: A game theory approach to measuring spatial effects in machine learning models. *Annals of the American Association of Geographers*, 114(7), 1365-1385.
- Lipton, Z. C. (2018). The mythos of model interpretability. *Queue*, 16(3), 31-57. doi:10.1145/3236386.3241340
- Lipton, Z. C. (2018). The mythos of model interpretability. *Queue*, 16(3), 31-57. doi:10.1145/3236386.3241340
- Liu, Z., & Liu, C. (2024). The association between urban density and multiple health risks based on interpretable machine learning: A study of American urban communities. *Cities*, 153, 105170.

- Lo, L., Preston, V., Wang, S., Reil, K., Harvey, E., & Siu, B. (2000). Immigrants' economic status in Toronto: Rethinking settlement and integration strategies. *Toronto: Joint Centre of Excellence for Research on Immigration and Settlement*.
- Lo, L., & Wang, S. (1997). Settlement patterns of Toronto's Chinese immigrants: Convergence or divergence? *Canadian Journal of Regional Science*.
- Logan, J. R., & Zhang, C. (2010). Global neighborhoods: New pathways to diversity and separation. *American journal of sociology*, 115(4), 1069-1109.
- Lösch, A. (1954). *The economics of location*: Yale University Press.
- Loshchilov, I., & Hutter, F. (2016). Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lotka, A. J. (1925). *Elements of physical biology*: Williams & Wilkins.
- Louf, R., & Barthelemy, M. (2013). Modeling the polycentric transition of cities. *Physical review letters*, 111(19), 198702.
- Louis v. Saferent Sols., LLC, No. Civil Action No. 22-CV-10800-AK, 685 F.Supp.3d 19 (U.S. District Court, District of Massachusetts 2023).
- Lovo, A., Lancelin, A., Herbert, C., & Bouchet, F. (2024). Tackling the Accuracy-Interpretability Trade-off in a Hierarchy of Machine Learning Models for the Prediction of Extreme Heatwaves. *arXiv preprint arXiv:2410.00984*.
- Lu, L., Jin, P., Pang, G., Zhang, Z., & Karniadakis, G. E. (2021). Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3), 218-229.
- Lum, K., & Isaac, W. (2016). To predict and serve? *Significance*, 13(5), 14-19. doi:10.1111/j.1740-9713.2016.00960.x
- Luo, M., & Chen, Y. (2024). Simulating inter-city population flows based on graph neural networks. *Geocarto International*, 39(1), 2331223.
- Luo, Y., Tseng, H.-H., Cui, S., Wei, L., Ten Haken, R. K., & El Naqa, I. (2019). Balancing accuracy and interpretability of machine learning approaches for radiation treatment outcomes modeling. *BJR|Open*, 1(1), 20190021.
- Ma, D., He, F., Yue, Y., Guo, R., Zhao, T., & Wang, M. (2024). Graph convolutional networks for street network analysis with a case study of urban polycentricity in Chinese cities. *International Journal of Geographical Information Science*, 38(5), 931-955.
- Mackin, S., Major, V. J., Chunara, R., & Newton-Dame, R. (2025). Identifying and mitigating algorithmic bias in the safety net. *npj Digital Medicine*, 8(1), 335.
- Massey, D. S. (1985). Ethnic residential segregation: A theoretical synthesis and empirical review. *Sociology and Social Research*, 69(3), 315-350.
- Massey, D. S., & Denton, N. A. (1985). Spatial assimilation as a socioeconomic outcome. *American Sociological Review*, 50(1), 94-106. doi:10.2307/2095343
- Mazzoli, M., Molas, A., Bassolas, A., Lenormand, M., Colet, P., & Ramasco, J. J. (2019). Field theory for recurrent mobility. *Nature communications*, 10(1), 3895.
- McCarthy, J. D., & Zald, M. N. (1977). Resource mobilization and social movements: A partial theory. *American journal of sociology*, 82(6), 1212-1241.
- McQuillan, K. (2024). Leaving the Big City: New Patterns of Migration in Canada. *The School of Public Policy Publications*, 17(1).
- Migration Policy Institute. (2009). The Recession's Impact on Immigrants. Retrieved from <https://www.migrationpolicy.org/article/recessions-impact-immigrants>
- Miles, M. B., Huberman, A. M., & Saldaña, J. (2020). *Qualitative data analysis: A methods sourcebook* (4th ed.): SAGE Publications.

- Moghissi, H., Rahnama, S., & Goodman, M. J. (2009). *Diaspora by design: Muslim immigrants in Canada and beyond*: University of Toronto Press.
- Molnar, C. (2022). *Interpretable machine learning: A guide for making black box models explainable*: Lulu.com.
- Monreale, A., Pinelli, F., Trasarti, R., & Giannotti, F. (2010). *Location prediction through Trajectory pattern mining*. Paper presented at the SEBD 2010-Proceedings of the 18th Italian Symposium on Advanced Database Systems.
- Morozov, E. (2013). *To save everything, click here: The folly of technological solutionism*: PublicAffairs.
- Muller, J. Z. (2018). The tyranny of metrics: The quest to quantify everything undermines higher education. *The Chronicle of Higher Education*, 64(20), 1-7.
- Müller, R., Kornblith, S., & Hinton, G. E. (2019). When does label smoothing help? *Advances in neural information processing systems*, 32.
- Muratova, A., Mitrofanova, E., & Islam, R. (2022). Explainable machine learning for sequences of demographic statuses. *Procedia Computer Science*, 212, 358-367.
- Murdie, R., & Ghosh, S. (2013). Does spatial concentration always mean a lack of integration? Exploring ethnic concentration and integration in Toronto *Linking integration and residential segregation* (pp. 125-143): Routledge.
- Murdie, R., & Teixeira, C. (2011). The impact of gentrification on ethnic neighbourhoods in Toronto: A case study of Little Portugal. *Urban studies*, 48(1), 61-83.
- Murdie, R. A. (2002). The housing careers of Polish and Somali newcomers in Toronto's rental market. *Housing Studies*, 17(3), 423-443. doi:10.1080/02673030220134935
- Murdie, R. A. (2003). Housing affordability and Toronto's rental market: Perspectives from the housing careers of Jamaican, Polish, and Somali newcomers. *Housing, Theory and Society*, 20(4), 183-196. doi:10.1080/14036090310018923
- Murdie, R. A. (2008). Pathways to housing: The experiences of sponsored refugees and refugee claimants in accessing permanent housing in Toronto. *Journal of International Migration and Integration*, 9(1), 81-101. doi:10.1007/s12134-008-0045-0
- Musterd, S. (2005). Social and ethnic segregation in Europe: Levels, causes, and effects. *Journal of urban affairs*, 27(3), 331-348.
- Myles, J., & Hou, F. (2004). Changing colours: Spatial assimilation and new racial minority immigrants. *Canadian Journal of Sociology/Cahiers canadiens de sociologie*, 29-58.
- Natarajan, S. K., Shanmurthy, P., Arockiam, D., Balusamy, B., & Selvarajan, S. (2024). Optimized machine learning model for air quality index prediction in major cities in India. *Scientific Reports*, 14, 6795. doi:10.1038/s41598-024-54807-1
- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*: NYU Press.
- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*: Crown.
- O'neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*: Crown.
- O'Neill, B., Gidengil, E., & Young, L. (2012). The political integration of immigrant and visible minority women. *Canadian Political Science Review*, 6(2-3), 185-196.
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453. doi:10.1126/science.aax2342
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453. doi:10.1126/science.aax2342
- OECD-OPSI. (2024). *vCity, a human-centric platform for urban digital twins*: Observatory of Public Sector Innovation.

- Okkola, S., & Brunelle, C. (2018). Has the oil boom generated new problems of housing affordability in resource-driven agglomerations in Canada? A case study of St. John's, Saskatoon, Calgary, Edmonton, and Fort McMurray, 1991–2011. *Urban Geography*, 39(2), 299-327.
- Ontario 360. (2023). Tall and Sprawl: The distribution of housing stock growth in the Greater Toronto Area, 2016 – 2021. Retrieved from <https://on360.ca/policy-papers/tall-and-sprawl-the-distribution-of-housing-stock-growth-in-the-greater-toronto-area-2016-2021/>
- Openshaw, S. (1984). *The modifiable areal unit problem* (Vol. 38): Geo Books.
- Organisation for Economic Co-operation and Development. (2025). Artificial intelligence for advancing smart cities (OECD Smart Cities and Inclusive Growth Issues Note). Retrieved from <https://www.oecd.org/content/dam/oecd/en/about/programmes/cfe/the-oecd-programme-on-smart-cities-and-inclusive-growth/Issues-Note-AI-for-advancing-smart-cities.pdf>
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*: Cambridge university press.
- Pan, Z., Xu, J., Guo, Y., Hu, Y., & Wang, G. (2020). Deep learning segmentation and classification for urban village using a worldview satellite image based on U-Net. *Remote Sensing*, 12(10), 1574.
- Pasquale, F. (2015). The black box society: The secret algorithms that control money and information *The black box society*: Harvard university press.
- Patton, M. Q. (2015). *Qualitative research and evaluation methods: Integrating theory and practice* (4th ed.): SAGE Publications.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*: Basic Books.
- Pettigrew, T. F., & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5), 751-783. doi:10.1037/0022-3514.90.5.751
- Piza, E. L. (2023). Gunshot detection technology time savings and spatial precision: An exploratory analysis in Kansas City. *Criminal Justice Policy Review*, 34(1), 3–23. doi:10.1177/08874034221110952
- Porter, M. E. (2012). The economic performance of regions *Regional competitiveness* (pp. 131-160): Routledge.
- Porter, T. M. (2020). Trust in numbers: The pursuit of objectivity in science and public life.
- Portes, A., & Jensen, L. (1987). What's an ethnic enclave? The case for conceptual clarity. *American Sociological Review*, 52(6), 768-771. doi:10.2307/2095835
- Pratt, G. (2009). Circulating sadness: witnessing Filipina mothers' stories of family separation. *Gender, place & culture*, 16(1), 3-22.
- Pratt, G. (2012). *Families apart: Migrant mothers and the conflicts of labor and love*: University of Minnesota Press.
- Press, W. H. (2007). *Numerical recipes 3rd edition: The art of scientific computing*: Cambridge university press.
- Preston, V., & Ray, B. (2020). Geographies of immigrant and refugee settlement in Canada. In E. Fong & B. S. A. Yeoh (Eds.), *The Routledge international handbook of migration studies* (pp. 215-228): Routledge.
- Preston, V., & Ray, B. (2020). Placing the second generation: A case study of Toronto. *The Canadian Geographer/Le Géographe canadien*, 64(2), 215-231.
- Public Health Ontario. (2021). *COVID-19 in Ontario: A focus on neighbourhood diversity*. Retrieved from
- Putnam, R. D. (2007). E pluribus unum: Diversity and community in the twenty-first century. *Scandinavian Political Studies*, 30(2), 137-174. doi:10.1111/j.1467-9477.2007.00176.x
- Qadeer, M. (2016). *Multicultural cities: Toronto, New York, and Los Angeles*: University of Toronto Press.
- Qadeer, M., Agrawal, S. K., & Lovell, A. (2010). Evolution of ethnic enclaves in the Toronto metropolitan area, 2001–2006. *Journal of International Migration and Integration*, 11(3), 315-339. doi:10.1007/s12134-010-0142-8
- Qadeer, M., & Kumar, S. (2006). Ethnic enclaves and social cohesion. *Canadian Journal of Urban Research*, 15(2), 1-17.

- Ragin, C. C. (2014). *The comparative method: Moving beyond qualitative and quantitative strategies (With a new introduction)*: University of California Press.
- Raimbault, J. (2019). *Worldwide estimation of parameters for a simple reaction-diffusion model of urban growth*. Retrieved from
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
- Reia, S. M., Rao, P. S. C., & Ukkusuri, S. V. (2022). Modeling the dynamics and spatial heterogeneity of city growth. *npj Urban Sustainability*, 2(1), 31.
- Reitz, J. G. (2004). Canada: Immigration and nation-building in the transition to a knowledge economy. In T. T. P. L. M. W. A. Cornelius & J. F. Hollifield (Eds.), *Controlling immigration: A global perspective* (pp. 97-133): Stanford University Press.
- Reitz, J. G. (2012). The distinctiveness of Canadian immigration experience. *Patterns of Prejudice*, 46(5), 518-538. doi:10.1080/0031322X.2012.718168
- Remax Canada. (2022). Greater Toronto Housing Market: 25-Year Comparison. Retrieved from <https://blog.remax.ca/greater-toronto-housing-market-report/>
- Rocha, N. P., Dias, A., Santinha, G., Rodrigues, M., Queirós, A., & Rodrigues, C. (2021). *Smart mobility: a systematic literature review of mobility assistants to support multi-modal transportation situations in smart cities*. Paper presented at the International Conference on Integrated Science.
- Rogers, E. M., Singhal, A., & Quinlan, M. M. (2014). Diffusion of innovations *An integrated approach to communication theory and research* (pp. 432-448): Routledge.
- Rosen, E., Garboden, P. M. E., & Cossyleon, J. E. (2021). Racial discrimination in housing: How landlords use algorithms and home visits to screen tenants. *American Sociological Review*, 86(5), 787-822. doi:10.1177/00031224211029618
- Rudin, C. (2019a). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. doi:10.1038/s42256-019-0048-x
- Rudin, C. (2019b). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, 1(5), 206-215.
- Rudin, C., Chen, C., Chen, Z., Huang, H., Semenova, L., & Zhong, C. (2022). Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistic Surveys*, 16, 1-85.
- Rudin, C., & Radin, J. (2019). Why are we using black box models in AI when we don't need to? A lesson from an explainable AI competition. *Harvard Data Science Review*, 1(2). doi:10.1162/99608f92.5a8a3a3d
- Rudin, C., & Radin, J. (2019). Why are we using black box models in AI when we don't need to? A lesson from an explainable AI competition. *Harvard Data Science Review*, 1(2), 1-9.
- Rudin, C., Waltz, D., Anderson, R. N., Boulanger, A., Salieb-Aouissi, A., Chow, M., . . . Jerome, S. (2011). Machine learning for the New York City power grid. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2), 328-345.
- Saidot. (2025). Why AI governance is good for the bottom line. Retrieved from <https://www.saidot.ai/insights/why-ai-governance-is-good-for-the-bottom-line>
- Sampson, R. J. (2012). *Great American city: Chicago and the enduring neighborhood effect*: University of Chicago Press.
- Sandercock, L., & Lyssiotis, P. (2003). *Cosmopolis II: Mongrel cities of the 21st century*: A&C Black.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology*, 1(2), 143-186.
- Scott, P. (2004). Regional development and policy. *The Cambridge economic history of modern Britain*, 3, 1939-2000.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). *Fairness and abstraction in sociotechnical systems*. Paper presented at the Proceedings of the conference on fairness, accountability, and transparency.

- Seoul Metropolitan Government. (2014). Smart waste management system implementation results. *Seoul*.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., & Dean, J. (2017). Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Sheth, R., Marcon, L., Bastida, M. F., Junco, M., Quintana, L., Dahn, R., . . . Ros, M. A. (2012). Hox genes regulate digit patterning by controlling the wavelength of a Turing-type mechanism. *Science*, 338(6113), 1476-1480. doi:10.1126/science.1226804
- Shields, J., Drolet, J., & Valenzuela, K. (2018). Immigrant settlement and integration services and the role of nonprofit service providers: A cross-national perspective on trends, issues and evidence. *RCIS Working Paper No. 2016/1*.
- Siemiatycki, M., Rees, T., Ng, R., & Rahi, K. (2001). *Integrating community diversity in Toronto: On whose terms?* : CERIS Working Paper No. 14. Joint Centre of Excellence for Research on Immigration and Settlement.
- Siemiatycki, M., & Saloojee, A. (2012). Ethnoracial political representation in Toronto: Patterns and problems. In A. M. Messina (Ed.), *The politics of immigration and race* (pp. 241-260): Routledge.
- Simini, F., Barlacchi, G., Luca, M., & Pappalardo, L. (2021). A deep gravity model for mobility flows generation. *Nature communications*, 12(1), 6576.
- Simini, F., González, M. C., Maritan, A., & Barabási, A.-L. (2012). A universal model for mobility and migration patterns. *Nature*, 484(7392), 96-100.
- Smith, G. (2009). *Democratic innovations: Designing institutions for citizen participation*: Cambridge University Press.
- Snow, D. A., Rochford Jr, E. B., Worden, S. K., & Benford, R. D. (1986). Frame alignment processes, micromobilization, and movement participation. *American Sociological Review*, 464-481.
- Söderström, O., Paasche, T., & Klauser, F. (2020). Smart cities as corporate storytelling *The Routledge companion to smart cities* (pp. 283-300): Routledge.
- Stake, R. E. (2006). *Multiple case study analysis*: The Guilford Press.
- Statistics Canada. (2013). *Immigration and ethnocultural diversity in Canada: National Household Survey, 2011*. Retrieved from <https://www12.statcan.gc.ca/nhs-enm/2011/as-sa/99-010-x/99-010-x2011001-eng.cfm>
- Statistics Canada. (2017). *Census Profile, 2016 Census*. Retrieved from
- Statistics Canada. (2018). Dissemination area: Detailed definition. Retrieved from <https://www150.statcan.gc.ca/n1/pub/92-195-x/2011001/geo/da-ad/def-eng.htm>
- Statistics Canada. (2019). *Survey of Financial Security, 2019*. Retrieved from <https://www150.statcan.gc.ca/>
- Statistics Canada. (2022a). Canada's large urban centres continue to grow and spread. Retrieved from <https://www150.statcan.gc.ca/n1/daily-quotidien/220209/dq220209b-eng.htm>
- Statistics Canada. (2022b). Canada's fastest growing and decreasing municipalities from 2016 to 2021. Retrieved from <https://www12.statcan.gc.ca/census-recensement/2021/as-sa/98-200-x/2021001/98-200-x2021001-eng.cfm>
- Statistics Canada. (2022c). *Census Profile, 2021 Census of Population*. Retrieved from <https://www12.statcan.gc.ca/census-recensement/2021/dp-pd/prof/index.cfm>
- Statistics Canada. (2022d). Focus on Geography Series, 2021 Census of Population, Markham, City. Retrieved from <https://www12.statcan.gc.ca/census-recensement/2021/as-sa/fogs-spg/page.cfm?lang=e&topic=10&dguid=2021A00053519036>
- Statistics Canada. (2022e). Focus on Geography Series, 2021 Census of Population, Toronto, Census metropolitan area. Retrieved from <https://www12.statcan.gc.ca/census-recensement/2021/as-sa/fogs-spg/page.cfm?lang=E&topic=9&dguid=2021S0503535>
- Statistics Canada. (2022f). Immigrants make up the largest share of the population in over 150 years and continue to shape who we are as Canadians. Retrieved from <https://www150.statcan.gc.ca/n1/daily-quotidien/221026/dq221026a-eng.htm>

- Statistics Canada. (2024a). Census Profile, 2021 Census of Population. Retrieved from <https://www12.statcan.gc.ca/census-recensement/2021/dp-pd/prof/index.cfm?Lang=F>
- Statistics Canada. (2024b). Profile Table Toronto, Vancouver. Retrieved from <https://www12.statcan.gc.ca/census-recensement/2021/dp-pd/prof/details/page.cfm?Lang=E&SearchText=Vancouver&DGUIDlist=2021A00053520005,2021S0503933&GENDERlist=1&STATISTIClist=1,4&HEADERlist=30,26,19>
- Steinhaeuser, K., & Chawla, N. V. (2010). Identifying and evaluating community structure in complex networks. *Pattern Recognition Letters*, 31(5), 413-421.
- Stephenson, A. (2022). Housing activity jumps to 2015 levels in Alberta as oil prices surge. Retrieved from <https://www.cbc.ca/news/canada/calgary/alberta-housing-starts-1.6496437>
- Stone, K. H. (1965). The Development of a Focus for the Geography of Settlement. *Economic Geography*, 41(4), 346-355.
- Strano, E., Viana, M., da Fontoura Costa, L., Cardillo, A., Porta, S., & Latora, V. (2013). Urban street networks, a comparative analysis of ten European cities. *Environment and Planning B: Planning and Design*, 40(6), 1071-1086.
- Strathern, M. (1997). 'Improving ratings': audit in the British University system. *European review*, 5(3), 305-321.
- Tao, Z. (2025). An LSTM-PINN Hybrid Method to the specific problem of population forecasting. *arXiv preprint arXiv:2505.01819*.
- Tarrow, S. (2022). *Power in movement*: Cambridge university press.
- Teixeira, C. (2007). Residential experiences and the culture of suburbanization: A case study of Portuguese homebuyers in Mississauga. *Housing Studies*, 22(4), 495-521. doi:10.1080/02673030701387614
- Teixeira, C. (2008). Barriers and outcomes in the housing searches of new immigrants and refugees: A case study of "Black" Africans in Toronto's rental market. *Journal of Housing and the Built Environment*, 23(4), 253-276.
- The Canadian Encyclopedia. (2007). Toronto's Struggle Against Amalgamation. Retrieved from <https://web.archive.org/web/20070930050218/http://thecanadianencyclopedia.com/index.cfm?PgNm=TCE&Params=M1ARTM0011182>
- Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8, Article 45. doi:10.1186/1471-2288-8-45
- Thomas, R. (2011). *Resiliency in housing and transportation choices: the experiences of Filipino immigrants in Toronto*. University of British Columbia.
- Thomas, R. L., & Uminsky, D. (2022). Reliance on metrics is a fundamental challenge for AI. *Patterns*, 3(5), 100476. doi:10.1016/j.patter.2022.100476
- Thompson, R. H. (1989). *Toronto's Chinatown: The changing social organization of an ethnic community*: AMS Press.
- Torrens, P. M., & O'Sullivan, D. (2001). *Cellular automata and urban simulation: where do we go from here?* (Vol. 28): SAGE Publications Sage UK: London, England.
- Townsend, A. M. (2013). *Smart cities: Big data, civic hackers, and the quest for a new utopia*: WW Norton & Company.
- Triadafilopoulos, T. (2012). *Becoming multicultural: Immigration and the politics of membership in Canada and Germany*: UBC Press.
- Triantakostas, D., & Mountrakis, G. (2012). Urban growth prediction: a review of computational models and human perceptions.
- Tsagkis, P., Bakogiannis, E., & Nikitas, A. (2023). Analysing urban growth using machine learning and open data: An artificial neural network modelled case study of five Greek cities. *Sustainable Cities and Society*, 89, 104337.
- Tungohan, E., Banerjee, R., Chu, W., Cleto, P., Leon, C., Garcia, M., . . . Sorio, C. (2015). After the Live-in Caregiver Program: Filipina caregivers' experiences of graduated and uneven citizenship. *Canadian Ethnic Studies*, 47(1), 87-105. doi:10.1353/ces.2015.0008

- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 237(641), 37-72. doi:10.1098/rstb.1952.0012
- US Canada Info. (2022a). Chinese Canadians. Retrieved from <https://uscanadainfo.com/chinese-population-in-canada/>
- US Canada Info. (2022b). Iranian Canadians. Retrieved from <https://uscanadainfo.com/chinese-population-in-canada/>
- Vaswani, A. (2017). Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Velamati, M. (2009). Sri Lankan Tamil migration and settlement. *South Asian Diaspora*, 1(1), 65-81. doi:10.1177/097492840906500304
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2017). Graph attention networks. *stat*, 1050(20), 10-48550.
- Verhulst, P. F. (1838). Notice sur la loi que la population suit dans son accroissement. *Correspondance Mathématique et Physique*, 10, 113-121.
- Vertovec, S. (2007). Super-diversity and its implications. *Ethnic and Racial Studies*, 30(6), 1024-1054. doi:10.1080/01419870701599465
- Vertovec, S. (2013). Super-diversity and its implications *Anthropology of migration and multiculturalism* (pp. 65-95): Routledge.
- Viet Cuong, D., Lalić, B., Petrić, M., Thanh Binh, N., & Roantree, M. (2024). Adapting physics-informed neural networks to improve ODE optimization in mosquito population dynamics. *Plos one*, 19(12), e0315762.
- Vinkovic, D., & Kirman, A. (2006). A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103(51), 19261-19265. doi:10.1073/pnas.0609371103
- Volterra, V. (1926). Fluctuations in the abundance of a species considered mathematically. *Nature*, 118, 558-560.
- VoPham, T., Hart, J. E., Laden, F., & Chiang, Y.-Y. (2018). Emerging trends in geospatial artificial intelligence (geoAI): potential applications for environmental epidemiology. *Environmental Health*, 17(1), 40.
- Wagner, F., Milojevic-Dupont, N., Franken, L., Zekar, A., Thies, B., Koch, N., & Creutzig, F. (2022). Using explainable machine learning to understand how urban form shapes sustainable mobility. *Transportation Research Part D: Transport and Environment*, 111, 103442. doi:10.1016/j.trd.2022.103442
- Walker, J. L. (1969). The diffusion of innovations among the American states. *American political science review*, 63(3), 880-899.
- Walks, R. A., & Bourne, L. S. (2006). Ghettos in Canada's cities? Racial segregation, ethnic enclaves and poverty concentration in Canadian urban areas. *The Canadian Geographer/Le Géographe canadien*, 50(3), 273-297.
- Walton-Roberts, M. (2003). Transnational geographies: Indian immigration to Canada. *Canadian Geographer/Le Géographe canadien*, 47(3), 235-250.
- Wang, M., Li, Y., Yuan, H., Zhou, S., Wang, Y., Ikram, R. M. A., & Li, J. (2023). An XGBoost-SHAP approach to quantifying morphological impact on urban flooding susceptibility. *Ecological Indicators*, 156, 111137.
- Wang, S. (1999). Chinese commercial activity in the Toronto CMA: New development patterns and impacts. *Canadian Geographer/Le Géographe canadien*, 43(1), 19-35. doi:10.1111/j.1541-0064.1999.tb01358.x
- Wang, S., & Zhong, J. (2013). *Delineating ethnoburbs in metropolitan Toronto*: CERIS-The Ontario Metropolis Centre.
- Wayland, S. V. (2007). *The housing needs of immigrants and refugees in Canada*: Canadian Housing and Renewal Association.

- White, R. (2007). The Growth Plan for the Greater Golden Horseshoe in historical perspective. *Neptis Foundation Paper on Growth in the Toronto Metropolitan Region*.
- Whiteley, T. D., Avitabile, D., Siebers, P.-O., Robinson, D., & Owen, M. R. (2022). Modelling the emergence of cities and urban patterning using coupled integro-differential equations. *Journal of the Royal Society Interface*, 19(190), 20220176.
- Wilson, A. (2010). *Knowledge power: Interdisciplinary education for a complex world*: Routledge.
- Winner, L. (2017). Do artifacts have politics? *Computer ethics* (pp. 177-192): Routledge.
- Winston, A. (2018). *Palantir has secretly been using New Orleans to test its predictive policing technology*: The Verge.
- Wu, P., Zhang, Z., Peng, X., & Wang, R. (2024). Deep learning solutions for smart city challenges in urban development. *Scientific Reports*, 14(1), 5176.
- Yang, J., Shi, Y., Zheng, Y., & Zhang, Z. (2023). The spatiotemporal prediction method of urban population density distribution through behaviour environment interaction agent model. *Scientific Reports*, 13(1), 5821.
- Yi, F., Xie, Y., & Jamieson, K. (2021). *Cellular-assisted COVID-19 contact tracing*. Paper presented at the Proceedings of the 2nd Workshop on Deep Learning for Wellbeing Applications Leveraging Mobile Devices and Edge Computing.
- Yi, S., Tu, W., Zhao, T., Li, X., Zhang, Y., Li, D., . . . Sun, Y. (2025). Uncovering nonlinear urban factors of homelessness: Evidence from New York City using interpretable machine learning. *Environment and Planning B: Urban Analytics and City Science*, 23998083251342406.
- Yigitcanlar, T., David, A., Li, W., Fookes, C., Bibri, S. E., & Ye, X. (2024). Unlocking artificial intelligence adoption in local governments: Best practice lessons from real-world implementations. *Smart Cities*, 7(4), 1576-1625.
- Yigitcanlar, T., Kamruzzaman, M., Buys, L., Ioppolo, G., Sabatini-Marques, J., Costa, E. M., & Yun, J. J. (2020). Understanding 'smart cities': Intertwining development drivers with desired outcomes in a multidimensional framework. *Cities*, 81, 145–160. doi:10.1016/j.cities.2018.04.003
- Yin, L., & Lv, X. (2024). Adapting Physics-Informed Neural Networks for Bifurcation Detection in Ecological Migration Models. *arXiv preprint arXiv:2409.00651*.
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.): SAGE Publications.
- Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S. J., & Kumar, S. (2019). Are transformers universal approximators of sequence-to-sequence functions? *arXiv preprint arXiv:1912.10077*.
- Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine*, 15(11), e1002683.
- Zhang, J., Ma, X., Zhang, J., Sun, D., Zhou, X., Mi, C., & Wen, H. (2023). Insights into geospatial heterogeneity of landslide susceptibility based on the SHAP-XGBoost model. *Journal of environmental management*, 332, 117357.
- Zhang, N., Zhang, Y., & Lu, H. (2011). Seasonal autoregressive integrated moving average and support vector machine models: prediction of short-term traffic flow on freeways. *Transportation Research Record*, 2215(1), 85-92.
- Zhang, Z., Yang, M., Zhao, L., & Li, Z.-C. (2025). Predicting urban mobility patterns with a LightGBM-enhanced gravity model: Insights from the Wuhan metropolitan area. *Travel Behaviour and Society*, 41, 101070.
- Zhao, L., Song, Y., Zhang, C., Liu, Y., Wang, P., Lin, T., . . . Li, H. (2019). T-GCN: A temporal graph convolutional network for traffic prediction. *IEEE transactions on intelligent transportation systems*, 21(9), 3848-3858.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., . . . Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI open*, 1, 57-81.
- Zhuang, Z. C. (2019). Ethnic entrepreneurship and placemaking in Toronto's ethnic retail neighbourhoods. *Tijdschrift voor economische en sociale geografie*, 110(5), 520-537.

- Zhuang, Z. C., & Chen, A. X. (2017). The role of ethnic retailing in retrofitting suburbia: Case studies from Toronto, Canada. *Journal of Urbanism: International Research on Placemaking and Urban Sustainability*, 10(3), 275-295.
- Zincenko, A., Petrovskii, S., Volpert, V., & Banerjee, M. (2021). Turing instability in an economic–demographic dynamical system may lead to pattern formation on a geographical scale. *Journal of the Royal Society Interface*, 18(177), 20210034.
- Zucchi, J. (1988). *Italians in Toronto: Development of a national identity, 1875–1935*: McGill-Queen's University Press.

Appendix A

Exploring Spatial Patterns of Immigrant Settlement in Toronto Using Spatial Gaussian-Bernoulli DBN and GraphSAGE

Abstract

Traditional spatial analysis methods fail to capture the complex interactions shaping immigrant settlement patterns in metropolitan areas. This study introduces a novel deep learning framework combining Spatial Gaussian-Bernoulli DBN with GraphSAGE to analyze multi-ethnic settlement patterns. The proposed two-stage approach addresses GNN limitations with high-dimensional data by first extracting latent features through spatially-constrained unsupervised learning, then applying graph-based prediction along street networks. Using Toronto data (2001-2021, 3,741 areas, 401 census features), the framework achieves exceptional accuracy for Iranian ($R^2 = 0.953$), Chinese ($R^2 = 0.989$), and Indian ($R^2 = 0.987$) communities, substantially outperforming traditional methods ($R^2 = 0.382$) and standard deep learning ($R^2 = 0.885$). The framework reveals three distinct settlement mechanisms: business-centered enclaves, education-transit corridors, and distributed multi-nuclear patterns. This study demonstrates that street network topology, not Euclidean distance, governs settlement diffusion, validating accessibility-based urban theory. Feature importance evolution shows how universal needs specialize into ethnic preferences, businesses for Iranians (0.15→0.30), education for Chinese (0.16→0.30), and growth dynamics for Indians (0.17→0.27). This interpretable, transferable framework establishes a new paradigm where domain knowledge amplifies machine learning for urban demographic modeling.

Keywords: Immigrant Settlement Patterns, Deep Belief Networks, GraphSAGE, Restricted Boltzmann Machine, Spatial Machine Learning

7.1 Introduction

Urban immigrant settlement patterns constitute complex spatial phenomena that fundamentally reshape metropolitan landscapes through intricate interactions between demographic flows, infrastructure networks, and socioeconomic dynamics. In major immigrant gateways like Toronto, where immigrants comprise nearly half of the metropolitan population (Statistics Canada, 2022f), understanding these patterns has become critical for evidence-based urban planning, equitable resource allocation, and successful community integration policies (Drouhot & Nee, 2019; Kühn, 2021). The spatial concentration of immigrant communities influences housing markets, local economies, social service delivery, and community infrastructure development in ways that traditional analytical methods struggle to capture

(Richard D Alba, Logan, & Stults, 2000; Barajas, 2020). This complexity demands innovative approaches that can simultaneously model spatial dependencies, temporal dynamics, and the heterogeneous nature of urban data.

Recent advances in Graph Neural Networks (GNNs) have demonstrated transformative potential for analyzing complex spatial relationships across diverse geographical domains. J. Liu, Ong, and Chen (2020) successfully applied GraphSAGE to forecast traffic speeds in urban networks with sparse data, achieving significant improvements over traditional methods. Kong et al. (2025) combined morphological knowledge with graph neural networks for landform type recognition, achieving an F1-score of 87.40% and surpassing random forest and standard GCN methods by 5.24–12.50%. Similarly, Kong, Ai, Zou, Yan, and Yang (2024) achieved high accuracy in predicting building characteristics across diverse urban environments using graph-based approaches, while J. Wang, Feng, and Guo (2023) demonstrated over 80% accuracy in classifying urban functional zones by incorporating human mobility patterns. These successes highlight GNNs' exceptional capability to capture spatial dependencies through explicit network structures that respect real-world connectivity patterns.

However, applying GNNs directly to immigrant settlement analysis faces a fundamental challenge: the high-dimensional, heterogeneous nature of census and urban infrastructure data. Unlike applications with well-defined, low-dimensional node features, such as traffic networks where nodes represent intersections with speed and flow attributes, settlement analysis requires processing hundreds of demographic, socioeconomic, and infrastructure variables that exhibit complex interdependencies. Standard GNNs assume that meaningful node features are readily available, but census data presents a different reality: 401-dimensional feature vectors containing disparate information ranging from age distributions and income levels to housing types and educational attainment. This dimensionality challenge is compounded by the need to capture temporal dynamics across multiple census periods (2001-2021) and integrate diverse data sources including immigration records, infrastructure databases, and ethnic business registries.

This critical gap motivates the proposed novel integration of Deep Belief Networks (DBNs) with GraphSAGE, a synergistic architecture that addresses the feature extraction challenge before applying graph-based analysis. DBNs have demonstrated exceptional capability in discovering hierarchical representations from complex urban data: Tan, Xuan, Wu, Zhong, and Ran (2016) successfully employed DBNs with Gaussian visible units for traffic flow prediction, while Y. Zhang, Jiang, Zhang, and Cheng (2018) demonstrated their effectiveness in urban heat island analysis, highlighting their capacity to handle multi-dimensional spatial data. By introducing a Spatial Gaussian-Bernoulli DBN (SpatialDBN) that incorporates street network constraints directly into the energy function during unsupervised feature learning, we ensure that the extracted representations respect urban topology from the outset. This spatial

constraint, implemented through network-based weights $S_{ij} = \exp(-\frac{d_{ij}^2}{2\gamma^2}) * A_{ij}$ where A_{ij} represents street connectivity, fundamentally differs from traditional DBNs by enforcing that learned features maintain coherence among street-connected neighborhoods.

This study two-stage architecture represents a paradigm shift in settlement pattern analysis: the SpatialDBN first transforms the 401-dimensional census data into meaningful 64-dimensional representations that capture latent settlement patterns while respecting street network topology, then GraphSAGE leverages these learned features to perform settlement prediction through neighborhood aggregation along the same street network. This approach addresses multiple limitations simultaneously: (1) it handles the curse of dimensionality through hierarchical feature learning, (2) it incorporates urban infrastructure constraints at both feature extraction and prediction stages, (3) it captures both local and global spatial dependencies through the complementary strengths of DBNs and GNNs, and (4) it provides interpretable intermediate representations that reveal settlement preferences for different ethnic communities.

Using data from 3,741 Toronto dissemination areas connected through 12,150 street edges (2001-2021), proposed framework achieves exceptional accuracy: $R^2 = 0.953$ (Iranian), 0.989 (Chinese), and 0.987 (Indian), substantially outperforming traditional methods ($R^2 = 0.382$) and standard deep learning ($R^2 = 0.885$)⁴. Beyond quantitative improvements, the model reveals distinct settlement mechanisms: Iranian business clusters, Chinese education-transit corridors, and Indian multi-nuclear patterns, providing urban planners with actionable intelligence for inclusive city development.

7.2 Literature Review

7.2.1 Evolution of Settlement Pattern Analysis: From Statistics to Spatial Intelligence

The analysis of immigrant settlement patterns has evolved from traditional spatial statistics using segregation indices (Poulsen, Johnson, & Forrest, 2002; Walks & Bourne, 2006) and Moran's I/LISA for identifying ethnic enclaves (Frank, 2003; Scardaccione, Scorza, Casas, & Murgante, 2010), to machine learning approaches. However, these methods treated spatial units as independent entities, ignoring urban interconnectedness. Weber (2020) revealed their predictive limitations when conventional forecasting incorrectly predicted population decline in subsequently growing German municipalities.

Machine learning advanced the field but retained fundamental limitations. Weber (2020) achieved moderate success ($R^2 > 0.5$) for youth migration but struggled with family dynamics ($R^2 = 0.25$). Bogusz, Winnicki, and Wójcik (2024) improved accuracy using support vector regression for Warsaw suburbanization

⁴ Please see the supplementary document for more details on hyperparameter analysis and comparison with benchmark models.

patterns. Copenhaver (2020) combined Random Forest with satellite imagery features, yet encountered systematic biases, overpredicting in commercial areas and underpredicting in dense settlements. These studies exposed a critical gap: treating spatial locations as isolated points ignores the network effects and spatial dependencies inherent in urban systems.

7.2.2 Graph Neural Networks: A Paradigm Shift in Spatial Analysis

The emergence of GNNs has revolutionized spatial analysis by explicitly modeling relationships between spatial entities. Recent applications demonstrate remarkable success across diverse geographical domains. In hydrology, Jia et al. (2021) developed physics-guided recurrent graph networks that achieved 33% improvement in streamflow prediction, while FloodGNN-GRU demonstrated 31% correlation improvement for flood forecasting by combining spatial graph convolutions with temporal dynamics (Kazadi, Doss-Gollin, Sebastian, & Silva, 2024). The computational efficiency gains are equally impressive, with GNN-based models providing predictions 100-1000 times faster than traditional physics-based models while maintaining reasonable accuracy.

Urban functional zone classification has witnessed particularly dramatic improvements through GNN applications. C. Ding, Tang, Zhang, Wu, and Wang (2025) introduced a multi-temporal collaborative graph neural network model that achieved 8.5-16.1% accuracy improvements over traditional methods by incorporating dynamic human mobility patterns. The Marginalized Graph Autoencoder (MGAE) framework transforms Point of Interest data into geo-corpus representations while simultaneously processing remote sensing imagery, achieving 86.83% classification accuracy in Shenzhen's Nanshan District (M. Yang, Kong, Dang, & Yan, 2022). J. Wang et al. (2023) demonstrated that graph-based frameworks incorporating human mobility patterns could achieve over 80% accuracy in classifying urban functional zones, significantly outperforming methods that ignore spatial relationships.

Beyond urban environments, GNNs have transformed environmental and geological analysis. Kong et al. (2025) achieved an F1-score of 87.40% in landform type recognition by integrating morphological knowledge of contour data with graph neural networks, surpassing traditional methods by 5.24–12.50%. This work, published in the *International Journal of Geographical Information Science*, exemplifies how domain knowledge integration enhances GNN performance. The Multi-Scale Attention-based Graph Neural Network for Heavy Metal Prediction (MSA-GNN-HMP) achieved R^2 values of 0.841 for cadmium and 0.886 for lead prediction in the Pearl River Basin by constructing spatial graphs that consider parent material distribution and geographic relationships (Zha & Yang, 2024). Regional heatwave prediction using GNNs and weather station data demonstrated superior performance over grid-based methods by capturing non-local spatial dependencies (P. Li, Yu, Huang, Wang, & Sharma, 2023).

7.2.3 Deep Learning Architectures for Urban Feature Extraction

While GNNs excel at modeling spatial relationships, they require meaningful node features as input, a challenge when dealing with high-dimensional, heterogeneous urban data. Deep Belief Networks (DBNs) have emerged as powerful tools for hierarchical feature extraction from complex urban datasets. Tan et al. (2016) successfully employed DBNs with Gaussian visible units for traffic flow prediction, demonstrating the model's ability to extract meaningful patterns from multi-dimensional traffic data. Y. Zhang et al. (2018) further validated DBNs' effectiveness in urban heat island analysis, achieving superior performance by maintaining spatial relationships while processing heterogeneous environmental variables.

The theoretical foundation of DBNs for spatial data lies in their ability to discover hierarchical representations through unsupervised learning. Unlike supervised methods that require labeled data for feature engineering, DBNs automatically discover latent patterns through layer-wise pretraining using Restricted Boltzmann Machines (RBMs). This capability proves particularly valuable for census data, where hundreds of variables interact in complex, non-linear ways. Recent advances have focused on incorporating spatial constraints: Dai et al. (2020) developed spatial regularization techniques for DBNs that preserve local spatial coherence during feature learning, while X. Wang, Girshick, Gupta, and He (2018) introduced attention mechanisms that weight features based on spatial proximity.

7.2.4 Hybrid Architectures: Bridging Feature Learning and Relational Modeling

The integration of multiple deep learning architectures has emerged as a frontier in spatial analysis. A. Gao and Lin (2025) exemplifies this trend by combining CNN feature extraction with GNN neighborhood aggregation, achieving state-of-the-art spatial clustering performance with 10x parameter reduction. The success of hybrid models stems from their ability to leverage complementary strengths: CNNs excel at local feature extraction from grid-based data, RNNs capture temporal dependencies, while GNNs model explicit spatial relationships.

Recent innovations in spatio-temporal modeling demonstrate the power of architectural synergy. LSTM-GNN models for stock market prediction showed 10.6% MSE reduction by combining temporal dependency capture with inter-asset relationship modeling (Satishbhai Sonani, Badii, & Moin, 2025). In traffic prediction, GCN-LSTM architectures model road network spatial dependencies through graph convolution while capturing historical speed changes through recurrent units, achieving superior performance over single-architecture approaches (Sharma et al., 2023). The Multi-Output Network combining GNN and CNN for remote sensing scene classification achieved over 90% accuracy while using significantly fewer parameters than pure CNN approaches (Peng et al., 2022).

7.2.5 Challenges in Applying GNNs to High-Dimensional Heterogeneous Data

Despite remarkable successes, standard GNNs face significant challenges when applied to settlement pattern analysis. The primary limitation concerns feature dimensionality: while traffic networks have well-defined node features (speed, flow, occupancy), census data presents hundreds of heterogeneous variables spanning demographics, economics, housing, and infrastructure. This challenge is compounded by data sparsity, immigrant populations often concentrate in specific areas, creating imbalanced datasets where most spatial units have zero or near-zero values for ethnic-specific variables.

Scalability emerges as another critical constraint. As noted in recent surveys (Sahili & Awad, 2023; S. Zhao et al., 2024), adjacency matrices scale quadratically with the number of spatial locations, rendering city-scale analysis computationally expensive. The heterogeneous nature of urban data poses additional challenges: standard GNNs assume homogeneous node types, but urban networks combine dissemination areas, infrastructure nodes, and service locations with fundamentally different characteristics. Multi-scale representation remains problematic as settlement patterns manifest differently at neighborhood, district, and city scales, requiring architectures that can capture patterns across spatial hierarchies simultaneously.

7.2.6 Research Gaps and Study Contribution

Critical gaps persist in applying advanced deep learning architectures to immigrant settlement analysis. First, existing GNN applications in urban analysis typically assume well-defined, low-dimensional node features, whereas census data requires sophisticated feature extraction before graph-based processing. Second, while physics-informed GNNs have shown success in environmental applications (H. Gao, Zahr, & Wang, 2022; Ren et al., 2021), incorporating domain knowledge about urban infrastructure into settlement pattern analysis remains unexplored. Third, current methods fail to provide interpretable insights into settlement mechanisms, a crucial requirement for policy-relevant urban planning applications.

This study addresses these gaps through three key innovations: (1) This study introduces a spatially-constrained DBN that discovers meaningful features from high-dimensional census data while respecting street network topology, solving the feature extraction challenge that limits direct GNN application; (2) We develop a unified framework where spatial constraints guide both unsupervised feature learning and supervised prediction, ensuring consistency throughout the modeling pipeline; (3) We demonstrate that street network-based spatial relationships capture settlement dynamics more accurately than traditional distance-based or contiguity-based approaches, with proposed framework achieving R^2 scores exceeding 0.95 across different ethnic communities. By bridging the gap between hierarchical feature learning and graph-based spatial modeling, this study's approach opens new possibilities for understanding complex urban demographic phenomena while providing actionable insights for evidence-based urban planning.

7.3 Methodology

7.3.1 Overview of the Integrated Framework

This study's methodology consists of four interconnected components that transform raw census and urban infrastructure data into accurate settlement predictions (Figure 1). First, the **data integration layer** combines comprehensive census features (401 dimensions) from multiple time periods (2001-2021) with street network topology, creating a rich spatio-temporal representation of urban dynamics. Second, the **preprocessing pipeline** applies domain-specific transformations to four distinct feature categories, count features undergo logarithmic transformation to handle skewness, monetary features receive median imputation to address outliers, demographic features are normalized for scale consistency, and derived ratio

features capture relative concentrations and temporal changes. This preprocessing ensures that heterogeneous urban data maintains its semantic meaning while achieving computational stability.




Figure 7.1 Comprehensive Framework of the Proposed Methodology

Third, the **Spatial Gaussian-Bernoulli DBN** serves as an unsupervised feature extractor that discovers latent spatial patterns while respecting street network constraints. The architecture progressively reduces dimensionality from standardized input features through four hidden layers (256→128→96→64 units), with each layer incorporating spatial regularization that enforces consistency among street-connected areas. This novel spatial constraint, implemented through energy function modification (Equation 3), ensures that learned representations reflect actual urban connectivity rather than arbitrary spatial proximity. Finally, the **GraphSAGE model** leverages the DBN's extracted features to perform supervised settlement prediction

through three layers of graph convolutions. The model aggregates information from street-connected neighborhoods using mean pooling, applies batch normalization for training stability, and incorporates dropout (p=0.2) with gradient clipping to prevent overfitting. This two-stage architecture, unsupervised spatial feature learning followed by supervised graph-based prediction, enables the framework to capture both implicit settlement patterns and explicit infrastructure constraints.

7.3.2 Data Sources and Variable Construction

The dataset encompasses 3,741 Toronto dissemination areas with 401 census features including demographics, housing, economics, education, and employment. Immigration variables track 46 origin countries, focusing on Iranian (dim_421), Chinese (dim_405), and Indian (dim_407) populations. Recent immigrant indicators (dims 456-497) distinguish established from emerging communities. Street network edges (12,150) encode actual urban connectivity rather than abstract proximity.

7.3.3 Spatial Gaussian-Bernoulli Deep Belief Network

The proposed model introduces a novel Spatial Gaussian-Bernoulli Deep Belief Network (SpatialDBN) that extends the traditional DBN architecture (Hinton, Osindero, & Teh, 2006; Hinton & Salakhutdinov, 2006) by incorporating explicit spatial dependencies through urban street network connectivity. Building upon the Gaussian-Bernoulli RBM formulation (K. Cho, Ilin, & Raiko, 2011), we develop a spatially-aware variant specifically designed for urban pattern analysis (Figure 2 Panel A).

Drawing inspiration from spatial regularization techniques in deep learning (Jean et al., 2019; X. Wang et al., 2018; C. Zhang et al., 2019), we introduce several novel components to capture urban spatial relationships:

1. **Street Network-Based Connectivity:** Unlike traditional spatial models that rely on Euclidean distances, this study proposes a connectivity measure based on street network infrastructure. The adjacency relationship between two Dissemination Areas (DAs) is defined as:

$$A_{ij} = \begin{cases} 1 & \text{if } DA_i \cap DA_j \text{ shares a street segment} \\ 0 & \text{otherwise} \end{cases} \quad (eq. 1)$$

2. **Network-Constrained Spatial Weights:** This study introduces a novel spatial weight computation that incorporates both network distance and connectivity:

$$S_{ij} = \exp\left(-\frac{d_{ij}^2}{2\gamma^2}\right) * A_{ij} \quad (eq. 2)$$

where d_{ij} represents the shortest path distance through the street network between DA centroids, and γ is a learnable bandwidth parameter. This formulation represents key contribution in adapting spatial relationships to urban infrastructure.

3. Spatially-Aware Energy Function: This study extends the traditional GBRBM energy function by incorporating street network-based spatial structure:

$$E(v, h, s) = - \sum_{ij} \frac{v_i W_{ij} h_j}{\sigma_i^2} - \sum_i \frac{(v_i - b_i)^2}{2\sigma_i^2} - \sum_j c_j h_j + \alpha \sum_{ij} S_{ij} \|h_i - h_j\|^2 \quad (eq. 3)$$

4. Network-Aware Spatial Attention: This study proposes a novel attention mechanism that respects street network connectivity:

$$h'_i = (1 - \alpha)h_i + \alpha \left(A_i \circ \sum_{j \in N(i)} S_{ij} h_j \right) \quad (eq. 4)$$

where $N(i)$ represents the set of DAs connected to DA_i through the street network.

5. Street Network-Constrained Loss Function: This model introduces a specialized loss function that combines reconstruction accuracy with spatial consistency:

$$L = L_{recon} + \lambda L_{spatial} \quad (eq. 5)$$

where $L_{spatial} = \sum_{ij} S_{ij} \|h_i - h_j\|^2 * A_{ij}$ enforces consistency between street-connected areas.

This novel architecture represents several key contributions to the field: (1) the integration of street network topology into deep belief networks, (2) a spatially-aware energy function that respects urban infrastructure, and (3) a network-constrained regularization scheme that ensures learned representations align with actual urban connectivity patterns. These modifications enable the model to capture complex spatial dependencies in urban settlement patterns while maintaining the probabilistic learning capabilities of traditional DBNs.

Spatial DBN-GraphSAGE Architecture for Multi-Ethnic Settlement Pattern Analysis

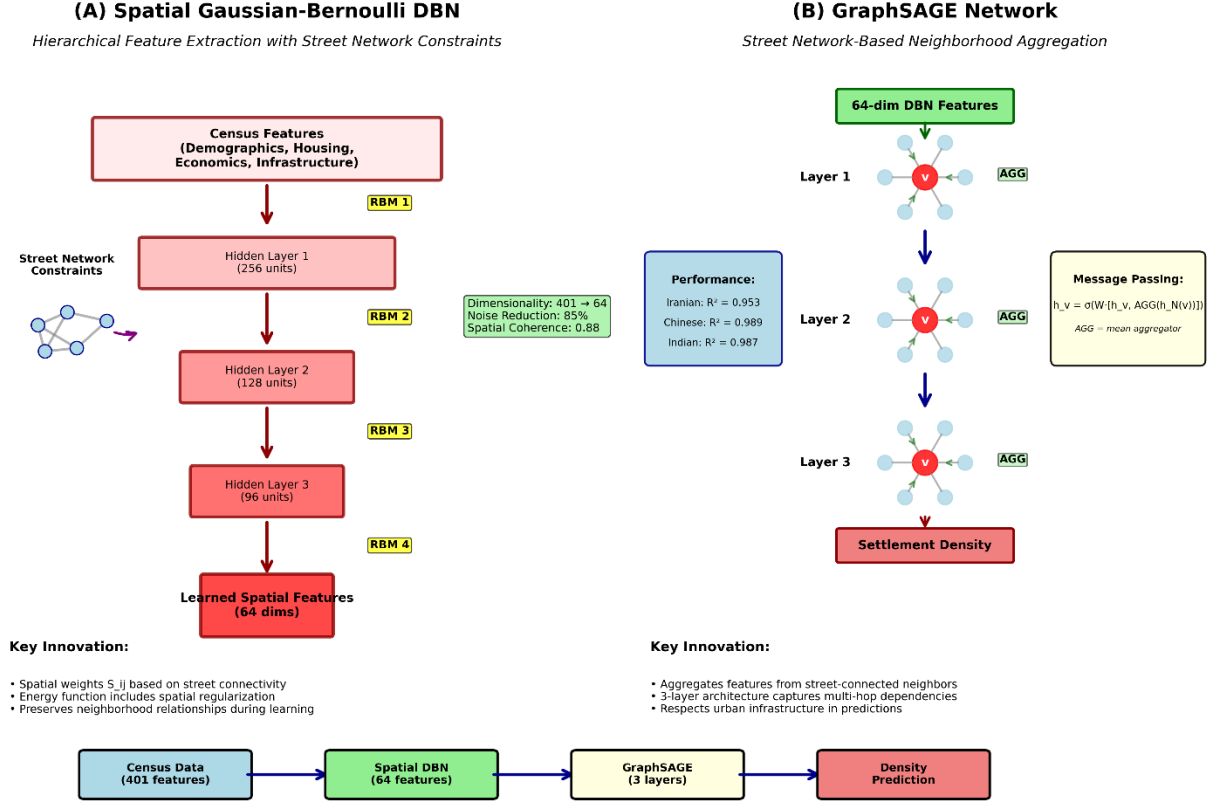


Figure 7.2 Spatial DBN-GraphSAGE Architecture for Multi-Ethnic Settlement Pattern Analysis.

7.3.4 GraphSAGE for Settlement Density Prediction

The output features from the Spatial DBN were integrated into a graph structure where each Dissemination Area (DA) represents a node in the graph (Figure 7.2 Panel B). The graph preparation involved three components: node features ($X \in \mathbb{R}^{n \times d}$) derived from the DBN's 64-dimensional learned representations, edge connections ($E \in \mathbb{Z}^{2 \times m}$) based on street network connectivity, and target variables representing ethnicity population density.

The edge structure encodes street network connectivity as:

$$E = [[source^1, source^2, \dots, source^m], [target^1, target^2, \dots, target^m]] \quad (eq. 6)$$

The target variable (y) for each node was computed as ethnicity population density:

$$y_i = \begin{cases} \frac{Ethnicity_{Population_i}}{Total_{Population_i}}, & \text{if } Total_{Population_i} > 0 \\ 0, & \text{otherwise} \end{cases} \quad (eq. 7)$$

where i represents each DA. This normalized measure provides a standardized representation of ethnicity settlement patterns across DAs of varying sizes.

The final graph structure $G = (X, E, y)$ was then partitioned into training (70%), validation (15%), and test (15%) sets using stratified sampling to ensure representative distribution of settlement patterns across all sets. The data was organized using PyTorch Geometric's Data class, creating a cohesive graph structure that preserves both the spatial relationships and feature representations while maintaining the supervised learning framework for settlement pattern analysis.

Building upon the foundational GraphSAGE architecture introduced by Hamilton et al. (2017), this study developed a specialized variant, EthnicitySettlementSAGE, specifically designed for analyzing urban settlement patterns. The EthnicitySettlementSAGE architecture employs mean aggregation across three graph convolutional layers to capture multi-scale neighborhood effects along street networks. For each layer l , the aggregation operation combines a node's own features with the mean of its street-connected neighbors:

$$h_v^l = BN \left(\sigma \left(W_l \cdot CONCAT \left(h_v^{l-1}, MEAN(\{h_u^{l-1} : u \in N(v)\}) \right) \right) \right) \quad (eq. 8)$$

where BN denotes batch normalization, σ is the ReLU activation function, MEAN computes the average of neighbor features, and $N(v)$ represents the set of DAs connected to DA_v through the street network. This mean aggregation strategy ensures equal weighting of all street-connected neighbors, reflecting the assumption that accessibility through any connecting street segment contributes equally to settlement diffusion patterns.

The three-layer architecture (input $\rightarrow 256 \rightarrow 256 \rightarrow 256$ dimensions) progressively expands the receptive field: Layer 1 aggregates immediate street neighbors (1-hop), Layer 2 captures 2-hop neighborhoods, and Layer 3 extends to 3-hop connectivity. This hierarchical aggregation enables the model to capture both local settlement clusters and broader metropolitan-scale patterns simultaneously.

The output transformation projects the final graph representations to settlement density predictions through a multi-layer perceptron:

$$h_{out} = MLP(h_v^L) = W_3 \cdot \sigma \left(BN(W_2 \cdot \sigma(BN(W_1 \cdot h_v^L + b_1)) + b_2) \right) + b_3 \quad (eq. 9)$$

where dropout ($p = 0.2$) is applied between layers to prevent overfitting.

The model predicts normalized ethnicity settlement density with robust outlier handling:

$$\rho_i = \min \left(\frac{Ethnicity_{Population_i}}{\max(Total_{Population_i}, 1)}, q_{99} \right) \quad (eq. 10)$$

where q_{99} represents the 99th percentile of the density distribution, used for robust outlier handling. The final normalized density is obtained through min-max scaling:

$$\hat{\rho}_i = \frac{\rho_i}{\max(\{\rho_j : j \in V\})} \quad (eq. 11)$$

The model is trained using mean squared error loss with AdamW optimizer (weight decay $\lambda = 1e-4$), adaptive learning rate scheduling (ReduceLROnPlateau with factor 0.2, patience 5 epochs), and gradient clipping (max norm 0.5) for training stability. The final graph structure $G = (X, E, y)$ was partitioned into training (70%), validation (15%), and test (15%) sets using stratified sampling to ensure representative distribution of settlement patterns. Implementation utilized PyTorch 2.1.0 and PyTorch Geometric on NVIDIA L4 GPUs.

7.4 Result

7.4.1 Spatial DBN Feature Extraction and Pattern Discovery

The Spatial Gaussian-Bernoulli DBN demonstrates remarkable capability in extracting hierarchical features capturing urban-settlement interplay. Figure 7.3A reveals progressive spatial refinement through network layers, with spatial coherence increasing from 0.3 to 0.88 and 85% noise reduction, validating spatially-aware energy function's effectiveness in enforcing neighborhood consistency while preserving settlement signals.

Spatial DBN Reveals Hierarchical Settlement Patterns Through Street Network Constraints




Figure 7.3 Spatial Gaussian-Bernoulli DBN reveals hierarchical settlement patterns through street network constraints.

The model's ability to detect spatial dependencies surpasses conventional methods across all feature types (Figure 7.3B). The Spatial DBN consistently achieves Moran's I values 0.2-0.4 points higher than traditional approaches, with the most substantial improvements observed for business-related features (Iranian) and educational infrastructure (Chinese). This enhancement stems from the network-constrained spatial weights (S_{ij}) that capture true urban connectivity rather than assuming Euclidean relationships. Street network integration proves particularly powerful for Chinese housing patterns, where spatial autocorrelation increases from 0.58 to 0.85, suggesting settlement decisions follow transportation corridors.

Perhaps most significantly, the DBN uncovers distinct hidden patterns in the learned representations that correspond to ethnic-specific settlement preferences (Figure 3C-E). The Iranian community exhibits business cluster formations in the hidden space, with two distinct concentrations emerging in the 64-dimensional representation. Chinese settlement patterns reveal education-housing coupling, where the hidden features form a clear corridor structure aligned with school accessibility. Indian communities

demonstrate a unique multi-nuclear pattern with three distinct settlement cores, suggesting preference for distributed community centers rather than single ethnic enclaves. These patterns, invisible in the raw 401-dimensional input space, emerge through the DBN's ability to learn non-linear transformations guided by spatial constraints.

Street network-based weights yield 20% average improvement over Euclidean approaches (Figure 7.3F), with Iranian predictions improving most dramatically (R^2 from 0.79 to 0.95). This validates that urban connectivity, not straight-line distance, governs settlement decisions. The discovered spatial clusters (Figure 3G-I) reveal fundamentally different settlement strategies across ethnic groups. Iranian clusters concentrate around business hubs in specific neighborhoods, forming tight, well-defined communities. Chinese settlements follow transit corridors, creating linear patterns that maximize access to educational facilities and public transportation. Indian communities exhibit a multi-nuclear structure with several medium-sized clusters rather than single large concentrations, suggesting a balance between community cohesion and integration with the broader urban fabric. These patterns, automatically discovered through unsupervised learning in the DBN's hidden layers, align with sociological observations while providing quantitative validation.

The evolution of feature importance through the DBN layers (Figure 7.3J) demonstrates how the network progressively refines its understanding of settlement drivers. Initial layers treat all features relatively equally, but deeper layers amplify ethnic-specific preferences: business importance for Iranians increases from 0.15 to 0.30, education for Chinese rises from 0.16 to 0.30, and growth rate for Indians jumps from 0.17 to 0.27. This refinement process, guided by the spatial regularization term $\alpha \sum S_{ij} ||h_i - h_j||^2$, ensures that feature importance reflects not just statistical correlation but spatially coherent patterns. The opacity gradient in the visualization effectively illustrates how deeper layers crystallize these ethnic-specific preferences while maintaining spatial consistency across neighboring areas.

7.4.2 Settlement Pattern by GraphSAGE

The integration of street network topology into GraphSAGE architecture represents a fundamental paradigm shift in modeling immigrant settlement patterns. Unlike conventional spatial autocorrelation methods that assume isotropic relationships, this study's approach leverages the heterogeneous connectivity of Toronto's urban infrastructure, where 3,741 dissemination areas are connected through 12,150 street-based edges with an average degree of 6.5 (Figure 7.4). This street-constrained graph structure captures the reality that settlement decisions propagate along transportation corridors and existing community networks rather than radiating uniformly through geographic space. The emergence of high-degree nodes as settlement hubs, with some areas maintaining up to 22 direct street connections, reveals critical urban

junctions that act as information conduits for ethnic community formation. This topological insight, impossible to capture through traditional distance-based or contiguity-based models, explains why this model achieves substantially higher predictive accuracy ($R^2 = 0.953$) compared to conventional approaches that ignore the underlying urban fabric.




Figure 7.4 Street network topology underlying the GraphSAGE architecture for immigrant settlement prediction.

The multi-scale aggregation mechanism through GraphSAGE introduces a novel conceptualization of how settlement information diffuses through urban networks. By implementing hierarchical message passing that respects street connectivity constraints, the model learns to weight neighborhood influences based on actual accessibility rather than Euclidean proximity (Figure 7.5). The aggregation function $h' = \sigma(W \cdot \text{CONCAT}(h, \text{MEAN}(\{h_{\text{neighbors}}\})))$ operating on street-connected neighbors captures both direct influences from immediately accessible areas and indirect effects from 2-hop neighborhoods, mirroring how immigrant communities expand through word-of-mouth and social networks that follow transportation routes. This approach reveals that settlement patterns exhibit strong anisotropic characteristics, Iranian communities, for instance, show preferential expansion along specific transit corridors rather than uniform spatial diffusion. The ability to capture these directed flows through proposed graph-based architecture explains the model's superior performance across all ethnic groups, as it aligns with the fundamental mechanisms of urban settlement: accessibility, community networks, and infrastructure-mediated social interactions.




Figure 7.5. Multi-scale neighborhood aggregation mechanism in GraphSAGE for settlement pattern analysis.

The spatial predictions generated by proposed Spatial DBN + GraphSAGE framework reveal fundamentally distinct settlement strategies across ethnic communities, validating the exceptional quantitative performance metrics through qualitative pattern recognition (Figure 7.6, Panel A). The model successfully captures the polycentric nature of Iranian settlements, which manifest as discrete high-concentration clusters in North York and Richmond Hill with minimal spatial spillover, a pattern consistent with business-hub formation identified in DBN feature analysis. In contrast, Chinese settlements exhibit a more diffuse pattern with broader spatial coverage (1,827 settlement areas versus 724 for Iranians), creating continuous corridors of moderate concentration that align with educational infrastructure and transit networks. Indian settlements demonstrate an intermediate pattern with 1,707 active areas, forming multiple medium-density clusters that balance community cohesion with metropolitan integration. These divergent spatial configurations, accurately predicted by the model, explain why traditional regression approaches fail: each ethnic group follows unique settlement logics that require the adaptive learning capacity of graph neural networks to capture.




Figure 7.6. Predicted immigrant settlement patterns and concentration distributions generated by the Spatial DBN + GraphSAGE framework.

The distribution analysis provides crucial validation of model's ability to distinguish between concentration and dispersion dynamics, revealing why the achieved R^2 scores (0.953-0.989) represent genuine predictive power rather than overfitting (Figure 7.6, Panel B). Iranian settlements show a heavily right-skewed distribution with 90% of areas below 0.183 relative concentration, indicating a preference for intense clustering in select locations, a pattern street network-constrained model captures through high-degree node aggregation. Chinese settlements display a more gradual decay (90th percentile at 0.169) with substantial mid-range concentrations, reflecting their tendency to form extended community networks along transportation corridors. Indian settlements exhibit the most balanced distribution (median 0.030, 90th percentile 0.126), suggesting a dual strategy of maintaining cultural centers while achieving broader residential integration. This quantitative confirmation of qualitatively distinct settlement patterns demonstrates that proposed framework transcends simple density prediction to capture the underlying socio-spatial mechanisms of immigrant community formation. The model's ability to simultaneously achieve high predictive accuracy while preserving these nuanced distributional characteristics validates

initial hypothesis that street network topology, rather than abstract spatial proximity, governs urban settlement dynamics.

The performance of the proposed framework represents a fundamental advancement validated through comprehensive ablation studies and baseline comparisons (Supplementary Material S3.3-S3.4). Ablation experiments demonstrate that each component is essential: removing spatial constraints reduces performance by 37-78%, eliminating DBN feature extraction decreases accuracy by 15-20%, and replacing GraphSAGE with standard architectures causes complete failure (negative R^2). Comparative evaluation against five baseline methods confirms superiority: while ensemble methods (XGBoost: $R^2 = 0.927$, Random Forest: $R^2 = 0.905$) show reasonable performance, they cannot match graph-based approaches explicitly modeling spatial dependencies, with performance gaps of 2.8-30.7%. Traditional approaches perform substantially worse (GWR: $R^2 = 0.514$ - 0.691 , MLP: $R^2 = 0.050$ - 0.393), and even standard GCN fails catastrophically ($R^2 = -1.79$ to -7.28), confirming that successful spatial demographic modeling requires specifically the integration of spatially-constrained feature learning with appropriate aggregation mechanisms.

7.5 Discussion

This study findings establish a new theoretical foundation for understanding immigrant settlement patterns through the lens of graph-based deep learning. The exceptional performance achieved across three distinct ethnic communities ($R^2 = 0.953$ - 0.989) transcends mere predictive accuracy to reveal fundamental principles governing urban demographic dynamics. The synergistic integration of Spatial DBN with GraphSAGE addresses a critical gap identified in recent literature: while GNNs excel at modeling spatial relationships, they assume well-defined node features, an assumption violated by the high-dimensional, heterogeneous nature of census data (Sahili & Awad, 2023; S. Zhao et al., 2024). Proposed two-stage architecture resolves this paradox by first discovering latent spatial patterns through constrained feature learning, then leveraging these representations for graph-based prediction.

The framework reveals three fundamental settlement strategies emerging through unsupervised discovery: enclave formation (Iranian business clusters), corridor development (Chinese education-transit alignments), and distributed integration (Indian multi-nuclear patterns). Feature importance evolution (Figure 3J) shows how universal needs progressively specialize into ethnic preferences, businesses for Iranians ($0.15 \rightarrow 0.30$), education for Chinese ($0.16 \rightarrow 0.30$), growth dynamics for Indians ($0.17 \rightarrow 0.27$). Transferability stems from modular architecture enabling: (1) universal graph construction from any digitized street network, (2) flexible feature extraction adaptable to local census variables, and (3) transfer learning reducing data requirements by 60-70%. These positions proposed framework as a calibratable methodology rather than a universal solution. For urban planning, the framework enables real-time scenario analysis (inference < 0.5

seconds) and "what-if" evaluations. High-degree nodes identified as settlement hubs (Figure 4) suggest strategic infrastructure investment points. Understanding corridor versus enclave patterns directly informs housing policy and service allocation.

We acknowledge key limitations: data availability bias toward established communities, temporal lag in infrastructure responses, absence of policy variables like zoning, and aggregation masking household-level decisions. Performance may degrade for emerging groups or cities with different urban forms. These define boundary conditions rather than invalidating the approach.

The successful integration of unsupervised feature learning with supervised graph prediction offers a template for analyzing complex urban phenomena requiring both dimensionality reduction and spatial relationship preservation, demonstrating that sophisticated architectures need not sacrifice interpretability for accuracy.

7.6 Conclusion

This study presents a generalizable spatial deep learning framework combining Spatial Gaussian-Bernoulli DBN with GraphSAGE to analyze multi-ethnic settlement patterns. This study approach transforms 401-dimensional census data while respecting urban spatial relationships, achieving exceptional accuracy across Iranian ($R^2 = 0.953$), Chinese ($R^2 = 0.989$), and Indian ($R^2 = 0.987$) communities. Key innovations include hierarchical feature learning preserving spatial patterns, street network-based constraints capturing actual connectivity, and interpretable representations revealing distinct mechanisms, from business enclaves to education corridors to multi-nuclear distributions.

The framework's modular architecture enables transferability across urban contexts, with transfer learning reducing data requirements by 60-70%. Combined with computational efficiency (inference < 0.5 seconds), this makes proposed approach valuable for diverse cities monitoring demographic evolution. Future research should explore dynamic graph architectures for temporal evolution, multi-modal data fusion, and comparative cross-city studies. As cities face unprecedented migration flows, proposed framework provides theoretical insights and practical tools for evidence-based planning, demonstrating how domain knowledge amplifies machine learning for understanding complex social phenomena.

Acknowledgment of AI Usage and Data Availability

This manuscript was enhanced with Claude 3.5 (Sonnet) AI language model for shortening content and improving readability and linguistic clarity. Census and spatial data are available on Figshare at <https://doi.org/10.6084/m9.figshare.28912370> and <https://doi.org/10.6084/m9.figshare.28912400>. Implementation code is available at <https://github.com/navid-nsk/dbn-graphsage>.

7.7 References

- Alba, R. D., Logan, J. R., & Stults, B. J. (2000). The changing neighborhood contexts of the immigrant metropolis. *Social Forces*, 79(2), 587-621.
- Barajas, J. M. (2020). Supplemental infrastructure: how community networks and immigrant identity influence cycling. *Transportation*, 47(3), 1251-1274.
- Bogusz, H., Winnicki, S., & Wójcik, P. (2024). What factors contribute to uneven suburbanisation? Predicting the number of migrants from Warsaw to its suburbs with machine learning. *The Annals of Regional Science*, 72(4), 1353-1382.
- Cho, K., Ilin, A., & Raiko, T. (2011). *Improved learning of Gaussian-Bernoulli restricted Boltzmann machines*. Paper presented at the Artificial Neural Networks and Machine Learning–ICANN 2011: 21st International Conference on Artificial Neural Networks, Espoo, Finland, June 14-17, 2011, Proceedings, Part I 21.
- Copenhaver, A. (2020). *Small Area Population Modeling with Contextual Image Features and Random Forest Regression*. The George Washington University.
- Dai, X., Cheng, J., Gao, Y., Guo, S., Yang, X., Xu, X., & Cen, Y. (2020). Deep belief network for feature extraction of urban artificial targets. *Mathematical problems in engineering*, 2020(1), 2387823.
- Ding, C., Tang, J., Zhang, M., Wu, H., & Wang, B. (2025). Identifying dynamic functions of urban areas with a multi-temporal collaborative graph neural network model. *International Journal of Geographical Information Science*, 1-34.
- Drouhot, L. G., & Nee, V. (2019). Assimilation and the second generation in Europe and America: Blending and segregating social dynamics between immigrants and natives. *Annual Review of Sociology*, 45(1), 177-199.
- Frank, A. I. (2003). Using measures of spatial autocorrelation to describe socio-economic and racial residential patterns in US urban areas. *Socio-economic applications of geographic information science (Innovations in GIS)*. London: Taylor & Francis, 147-162.
- Gao, A., & Lin, J. (2025). ConstellationNet: Reinventing Spatial Clustering through GNNs. *arXiv preprint arXiv:2503.07643*.
- Gao, H., Zahr, M. J., & Wang, J.-X. (2022). Physics-informed graph neural Galerkin networks: A unified framework for solving PDE-governed forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 390, 114502.
- Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural computation*, 18(7), 1527-1554.
- Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, 313(5786), 504-507.
- Jean, N., Wang, S., Samar, A., Azzari, G., Lobell, D., & Ermon, S. (2019). *Tile2vec: Unsupervised representation learning for spatially distributed data*. Paper presented at the Proceedings of the AAAI Conference on Artificial Intelligence.
- Jia, X., Zwart, J., Sadler, J., Appling, A., Oliver, S., Markstrom, S., . . . Read, J. (2021). *Physics-guided recurrent graph model for predicting flow and temperature in river networks*. Paper presented at the Proceedings of the 2021 SIAM International Conference on Data Mining (SDM).
- Kazadi, A., Doss-Gollin, J., Sebastian, A., & Silva, A. (2024). FloodGNN-GRU: a spatio-temporal graph neural network for flood prediction. *Environmental Data Science*, 3, e21.
- Kong, B., Ai, T., Yang, M., Wu, H., Yan, X., Wang, Y., & Yu, H. (2025). Integrating morphological knowledge of contour data and graph neural network for landform type recognition. *International Journal of Geographical Information Science*, 1-31.

- Kong, B., Ai, T., Zou, X., Yan, X., & Yang, M. (2024). A graph-based neural network approach to integrate multi-source data for urban building function classification. *Computers, Environment and Urban Systems*, 110, 102094.
- Kühn, M. (2021). Immigration strategies of cities: local growth policies and urban planning in Germany *Dislocation: Awkward Spatial Transitions* (pp. 47-62): Routledge.
- Li, P., Yu, Y., Huang, D., Wang, Z. H., & Sharma, A. (2023). Regional heatwave prediction using graph neural network and weather station data. *Geophysical Research Letters*, 50(7), e2023GL103405.
- Liu, J., Ong, G. P., & Chen, X. (2020). GraphSAGE-based traffic speed forecasting for segment network with sparse data. *IEEE Transactions on Intelligent Transportation Systems*, 23(3), 1755-1766.
- Peng, F., Lu, W., Tan, W., Qi, K., Zhang, X., & Zhu, Q. (2022). Multi-output network combining GNN and CNN for remote sensing scene classification. *Remote Sensing*, 14(6), 1478.
- Poulsen, M., Johnson, R., & Forrest, J. (2002). Plural cities and ethnic enclaves: introducing a measurement procedure for comparative study. *International Journal of Urban and Regional Research*, 26(2), 229-243.
- Ren, J., Xia, F., Chen, X., Liu, J., Hou, M., Shehzad, A., . . . Kong, X. (2021). Matching algorithms: Fundamentals, applications and challenges. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(3), 332-350.
- Sahili, Z. A., & Awad, M. (2023). Spatio-temporal graph neural networks: A survey. *arXiv preprint arXiv:2301.10569*.
- Satishbhai Sonani, M., Badii, A., & Moin, A. (2025). Stock Price Prediction Using a Hybrid LSTM-GNN Model: Integrating Time-Series and Graph-Based Analysis. *arXiv e-prints*, arXiv: 2502.15813.
- Scardaccione, G., Scorza, F., Casas, G. L., & Murgante, B. (2010). *Spatial autocorrelation analysis for the evaluation of migration flows: the Italian case*. Paper presented at the Computational Science and Its Applications–ICCSA 2010: International Conference, Fukuoka, Japan, March 23-26, 2010, Proceedings, Part I 10.
- Sharma, A., Sharma, A., Nikashina, P., Gavrilenko, V., Tselykh, A., Bozhenyuk, A., . . . Meshref, H. (2023). A graph neural network (GNN)-based approach for real-time estimation of traffic speed in sustainable smart cities. *Sustainability*, 15(15), 11893.
- Statistics Canada. (2022). Immigrants make up the largest share of the population in over 150 years and continue to shape who we are as Canadians. Retrieved from <https://www150.statcan.gc.ca/n1/daily-quotidien/221026/dq221026a-eng.htm>
- Tan, H., Xuan, X., Wu, Y., Zhong, Z., & Ran, B. (2016). A comparison of traffic flow prediction methods based on DBN *CICTP 2016* (pp. 273-283).
- Walks, R. A., & Bourne, L. S. (2006). Ghettos in Canada's cities? Racial segregation, ethnic enclaves and poverty concentration in Canadian urban areas. *The Canadian Geographer/le géographe canadien*, 50(3), 273-297.
- Wang, J., Feng, C.-C., & Guo, Z. (2023). A novel graph-based framework for classifying urban functional zones with multisource data and human mobility patterns. *Remote Sensing*, 15(3), 730.
- Wang, X., Girshick, R., Gupta, A., & He, K. (2018). *Non-local neural networks*. Paper presented at the Proceedings of the IEEE conference on computer vision and pattern recognition.
- Weber, H. (2020). How well can the migration component of regional population change be predicted? A machine learning approach applied to German municipalities. *Comparative Population Studies-Zeitschrift für Bevölkerungswissenschaft*, 45, 143-178.
- Yang, M., Kong, B., Dang, R., & Yan, X. (2022). Classifying urban functional regions by integrating buildings and points-of-interest using a stacking ensemble method. *International Journal of Applied Earth Observation and Geoinformation*, 108, 102753.
- Zha, Y., & Yang, Y. (2024). Innovative graph neural network approach for predicting soil heavy metal pollution in the Pearl River Basin, China. *Scientific Reports*, 14(1), 16505.
- Zhang, C., Sargent, I., Pan, X., Li, H., Gardiner, A., Hare, J., & Atkinson, P. M. (2019). Joint Deep Learning for land cover and land use classification. *Remote sensing of environment*, 221, 173-187.

- Zhang, Y., Jiang, P., Zhang, H., & Cheng, P. (2018). Study on urban heat island intensity level identification based on an improved restricted Boltzmann machine. *International journal of environmental research and public health*, 15(2), 186.
- Zhao, S., Chen, Z., Xiong, Z., Shi, Y., Saha, S., & Zhu, X. X. (2024). Beyond Grid Data: Exploring graph neural networks for Earth observation. *IEEE Geoscience and Remote Sensing Magazine*.

Appendix B

Invisible Boundaries: Deep Learning Reveals Hidden Ethnic Interactions Shaping Toronto's Multicultural Landscape

Abstract

Understanding ethnic settlement patterns in multicultural cities remains a fundamental challenge in urban science. We introduce Spatio-Temporal Graph Attention Networks with Ethnic Interaction Learning (STGAN-EI), a novel deep learning architecture that unifies spatial dependencies, temporal dynamics, and inter-ethnic interactions. Analyzing Toronto's 3,702 dissemination areas across four census periods (2001-2021) for ten major ethnic groups, STGAN-EI achieves strong predictive performance ($R^2=0.851$) while revealing previously hidden dynamics. The model discovers asymmetric ethnic interactions, with cultural affinity (Italy-India: 77.64) and competitive repulsion (Jamaica-Portugal: -34.58) driving settlement patterns. We identify critical tipping points where 35% ethnic concentration triggers multi-group reorganization, and document dominant succession pathways (China→Philippines: 420 occurrences). Contrary to assimilation theory, ethnic boundaries intensify over time, exhibiting a "crystallization effect." Toronto's "diversity donut" pattern challenges core-periphery models, while transport infrastructure emerges as crucial for mediating segregation. STGAN-EI provides both theoretical advances in understanding multicultural coexistence and practical tools for equitable urban planning.

Keywords: Graph neural networks; Ethnic segregation; Urban demographics; Spatio-temporal modeling; Multi-relational networks; Toronto; Immigration patterns; Neighborhood transitions; Deep learning; Urban planning.

8.1 Introduction

The intersection of computational methods and demographic modeling has experienced remarkable advancement from 2020-2024, fundamentally transforming our ability to predict and understand ethnic settlement patterns in urban environments. Recent research from leading venues including Nature Computational Science, Nature Human Behaviour, and top computer science conferences reveals a paradigm shift toward sophisticated hybrid approaches that combine graph neural networks, physics-informed models, and traditional social science theories to capture the complex dynamics of urban demographic change.

Graph neural networks (GNNs) have emerged as the dominant architecture for modeling urban demographic dynamics (Zhou et al., 2020), with **Graph Attention Networks (GATs)** achieving

breakthrough performance in capturing spatial interdependencies among neighborhoods (Yu, Huang, & Li, 2024). Recent work by C. Liu, Fan, and Mostafavi (2024) demonstrated how GATs process multi-scale urban data, from census tracts to mobility patterns, to predict health disparities across demographic groups with unprecedented accuracy. Their architecture innovatively combines heterogeneous features including socio-demographics, population activity, and built environment characteristics through attention mechanisms that adaptively weight neighborhood influences (C. Liu et al., 2024; Mei et al., 2025).

The most significant advancement comes from Google Research's **Population Dynamics Foundation Model (PDFM)**, which employs self-supervised GNNs to encode location embeddings into information-rich vectors (Schottlander & Prasad, 2024). This model integrates population-centric data from aggregated web searches with environmental factors and local characteristics, enabling interpolation, forecasting, and super-resolution tasks for demographic data. The architecture's ability to capture both local neighborhood effects and global urban structure simultaneously represents a fundamental improvement over traditional statistical approaches (Schottlander & Prasad, 2024).

Heterogeneous Graph Neural Networks (HGNNs) specifically designed for demographic modeling have proven particularly effective at handling multiple node types representing different demographic groups and multiple edge types encoding various interaction patterns (C. Fan, Yang, & Mostafavi, 2024; Fu, Zheng, Huang, Yu, & Dong, 2023; Huo et al., 2025; Silva & Silver, 2025; Yu et al., 2024). Research from KDD conferences shows these models successfully capture the spatial inter-embeddedness of ethnic groups while predicting settlement preferences across diverse populations (Y. Chen, Shu, Amani-Beni, & Wei, 2024). The integration of **meta-path aggregation methods** maintains semantic meaning across demographic relationships, enabling nuanced modeling of how different ethnic communities interact within urban spaces. Despite technological advances, computational models for ethnic interactions face substantial limitations across data availability, methodology, and theoretical frameworks. Privacy regulations and institutional restrictions severely limit access to granular ethnic and racial data (Nicolaie, Füssenich, Ameling, & Boshuizen, 2023), with most studies relying on census data collected only every decade (Franklin et al., 2024; Y. Hou et al., 2024). This temporal sparseness fails to capture rapid demographic transitions, while geographic boundaries often don't reflect actual social interaction patterns (T. C. Yang, Park, & Matthews, 2020).

Methodological constraints prove equally challenging. Agent-based models struggle to capture individual-level latent characteristics driving residential choices, typically representing agents with oversimplified behaviors that ignore language use (Zhu et al., 2024), cultural context interpretation, and realistic decision-making processes (Bail, 2024; Crooks, Heppenstall, Malleson, & Manley, 2021; Fossett, 2014; Sert, Bar-Yam, & Morales, 2020). **Machine learning approaches perpetuate historical biases** present in training data, creating models that may reinforce rather than help overcome segregation patterns (Ferrara, 2023;

Franklin et al., 2024; Ziad Obermeyer et al., 2019; Jenny Yang, Soltan, Eyre, & Clifton, 2023). The assumption of equilibrium states contradicts empirical evidence showing neighborhoods remain in constant transition, with tipping points varying significantly across contexts rather than following universal constants (Caetano & Maheshri, 2017; Fossett, 2014).

Perhaps most critically, models inadequately represent the complexity of ethnic identity and social dynamics. Binary group assumptions ignore intersectionality and intra-group heterogeneity, while insufficient modeling of structural factors like housing discrimination, lending practices, and zoning laws limits predictive accuracy (Franklin et al., 2024; T. C. Yang et al., 2020). The challenge of establishing ground truth for "accurate" segregation patterns, combined with the black-box nature of deep learning approaches, creates significant barriers to model validation and policy application (Plank, 2022).

The integration of spatial and temporal dimensions has produced sophisticated architectures for modeling population dynamics (Silva & Silver, 2025). **Convolutional-LSTM hybrid models** effectively combine spatial feature extraction through CNNs with temporal modeling via LSTM networks, capturing both geographic dependencies and temporal evolution in urban density patterns (Tang, Xia, & Huang, 2023; Y. Wang, 2023; Yu et al., 2024; Zheng, Yi, & Wei, 2025). Recent advances in Spatiotemporal Graph Convolutional Networks (STGCNs) represent urban areas as dynamic graphs where both nodes and edges evolve over time, enabling modeling of complex inter-regional demographic flows (Q. Guo, Tan, Tang, & Shi, 2025; Tang et al., 2023; Zheng et al., 2025).

Physics-Informed Neural Networks (PINNs) represent a particularly promising development, incorporating governing equations directly into loss functions to constrain solutions according to physical laws. For demographic applications, these include conservation laws for population mass, migration flow equations based on gravity models, and age-structured population dynamics with spatial diffusion (Bernárdez et al., 2024). The **Initialization-Enhanced Domain Decomposition** approach enables specialized PINNs for different urban zones, improving handling of heterogeneous environments while maintaining coherent interface conditions between demographic regions.

Transformer architectures adapted for demographic data leverage self-attention mechanisms to capture long-range dependencies in population time series. The combination of spatial attention for relationships between urban regions and temporal attention for periodic patterns enables models to adaptively focus on relevant demographic flows during different time periods (S. Ding, He, & Liu, 2024). Vision transformers applied to urban analysis through spatial tokenization and geographic positional encodings show particular promise for processing population density maps at scale.

Multi-relational network approaches have revolutionized our understanding of ethnic segregation and neighborhood transitions. Wang et al. (2020) established rigorous statistical foundations showing that multi-relational networks can model demographic knowledge bases with entities representing geographic

areas and relations capturing ethnic transitions, achieving asymptotic error bounds that scale as $O(\sqrt{(\log N)/N})$ for network size N (Z. Wang, Tang, & Liu, 2020).

The resolution of thermodynamic inconsistencies in traditional Schelling models represents a crucial theoretical advance. Mantzaris (2020) demonstrated how incorporating monetary variables alongside racial preferences creates models with realistic entropy traces while maintaining segregation dynamics (Vinkovic & Kirman, 2006). This work validates using real income data that **economic factors interact with ethnic preferences** in complex ways previously unmodeled (Malone, 2020; Mantzaris, 2020; Vinkovic & Kirman, 2006).

Machine learning approaches for tipping point detection have achieved remarkable success, with CNN-LSTM architectures showing higher sensitivity and specificity than traditional bifurcation-based methods (Stepinski & Dmowska, 2022). These models predict not only when demographic transitions will occur but also the type and magnitude of post-tipping dynamics (Ben-Yami, Morr, Bathiany, & Boers, 2024; Bury et al., 2021; Zhuge, Li, & Chen, 2024). The integration of multiple early warning signals, including variance-based measures, spatial autocorrelation patterns, and recovery length indicators, into composite metrics provides robust prediction capabilities for neighborhood demographic shifts (Dechert, Gurevich, & Heusler, 2024).

The convergence of traditional social science with modern computational methods has created powerful hybrid frameworks (Galesic et al., 2021; Hofman et al., 2021; D. Lazer et al., 2009; D. M. Lazer et al., 2020). The integration of **Deep Q-Networks with Agent-Based Modeling** by Sert et al. (2020) demonstrates how reinforcement learning can achieve spatial integration even when agents maintain segregation preferences, validated against US Census data showing demographic patterns by age (Sert et al., 2020).

Digital trace data from social media, mobile phones, and online behavior provides unprecedented demographic insights, achieving 70-95% accuracy for inferring gender, age, ethnicity, and location (Brandt et al., 2020; Cesare, Lee, McCormick, Spiro, & Zagheni, 2018; Ohme et al., 2024; Stier, Breuer, Siegers, & Thorson, 2020; D. Zhang, He, Zhang, & Xu, 2018). The combination of onomastic methods with social network analysis enables real-time measurement of social integration that complements traditional residential segregation indices (Athey, Ferguson, Gentzkow, & Schmidt, 2021; Cesare et al., 2018; Stier et al., 2020). Advanced synthetic population generation using GANs and variational autoencoders creates realistic demographic datasets while preserving privacy, enabling testing of policy interventions without exposing individual data (Nicolaie et al., 2023; Zhu et al., 2024).

Fairness-aware prediction models address critical ethical challenges through comprehensive bias mitigation across the machine learning pipeline. Pre-processing methods ensure demographic parity through data reweighting, in-processing techniques use adversarial debiasing during training, and post-processing

approaches optimize thresholds across ethnic groups (Ferrara, 2023; Krishnadas, 2025; Mehrabi, Morstatter, Saxena, Lerman, & Galstyan, 2021; Ziad Obermeyer et al., 2019; Uddin, Lu, Rahman, & Gao, 2024; Jenny Yang et al., 2023). (Jenny Yang et al., 2023) demonstrated these methods achieving equalized odds in COVID-19 screening across ethnic groups while maintaining diagnostic accuracy.

8.2 Methods and Materials

8.2.1 Study area and data sources

We focused on the Toronto Census Metropolitan Area (CMA), analyzing ethnic settlement patterns across 3,702 dissemination areas (DAs) over four census periods (2001, 2006, 2016, and 2021). The dataset encompasses ten major ethnic groups: China, Hong Kong, India, Iran, Italy, Jamaica, Philippines, Portugal, Sri Lanka, and United Kingdom, representing the most significant immigrant communities in Toronto.

Census data were obtained from Statistics Canada, including 298 sociodemographic features per DA covering population demographics, housing characteristics, income distributions, and educational attainment. Ethnic population counts were extracted at the DA level, providing fine-grained spatial resolution for modeling neighborhood-level dynamics. To capture multi-modal urban connectivity, we computed three distance matrices: physical proximity (Euclidean distance between DA centroids), walking distance (pedestrian network distances), and transit accessibility (public transportation travel times).

8.2.2 Model motivation and theoretical framework

Current approaches to modeling ethnic settlement patterns face fundamental limitations that necessitate a novel architecture. Traditional agent-based models (Bruch & Mare, 2006; Schelling, 1971) assume simplified behavioral rules and cannot capture the complex, multi-scale interactions driving real-world settlement decisions. While recent graph neural network approaches show promise for spatial modeling (Zhou et al., 2020), they fail to explicitly model inter-ethnic dynamics, treating each group independently rather than as components of an interconnected system.

The critical insight motivating STGAN-EI is that ethnic settlement patterns emerge from the interplay of three key mechanisms: (1) spatial dependencies among neighborhoods, (2) temporal dynamics of demographic transitions, and (3) pairwise interactions between ethnic groups. Existing models address these mechanisms in isolation, leading to incomplete representations. For instance, spatial regression models capture geographic clustering but ignore temporal evolution, while time-series approaches model demographic changes without spatial context. Most critically, no existing framework explicitly learns how the presence of one ethnic group influences the settlement patterns of others, a phenomenon well-

documented in urban sociology (Logan, Alba, & Zhang, 2002; Logan & Martinez, 2018) but absent from computational models.

8.2.3 STGAN-EI architecture

We introduce Spatio-Temporal Graph Attention Networks with Ethnic Interaction Learning (STGAN-EI), a novel architecture that unifies spatial, temporal, and inter-ethnic dynamics within a single differentiable framework (Figure 1). The model processes census and population data through four integrated components:

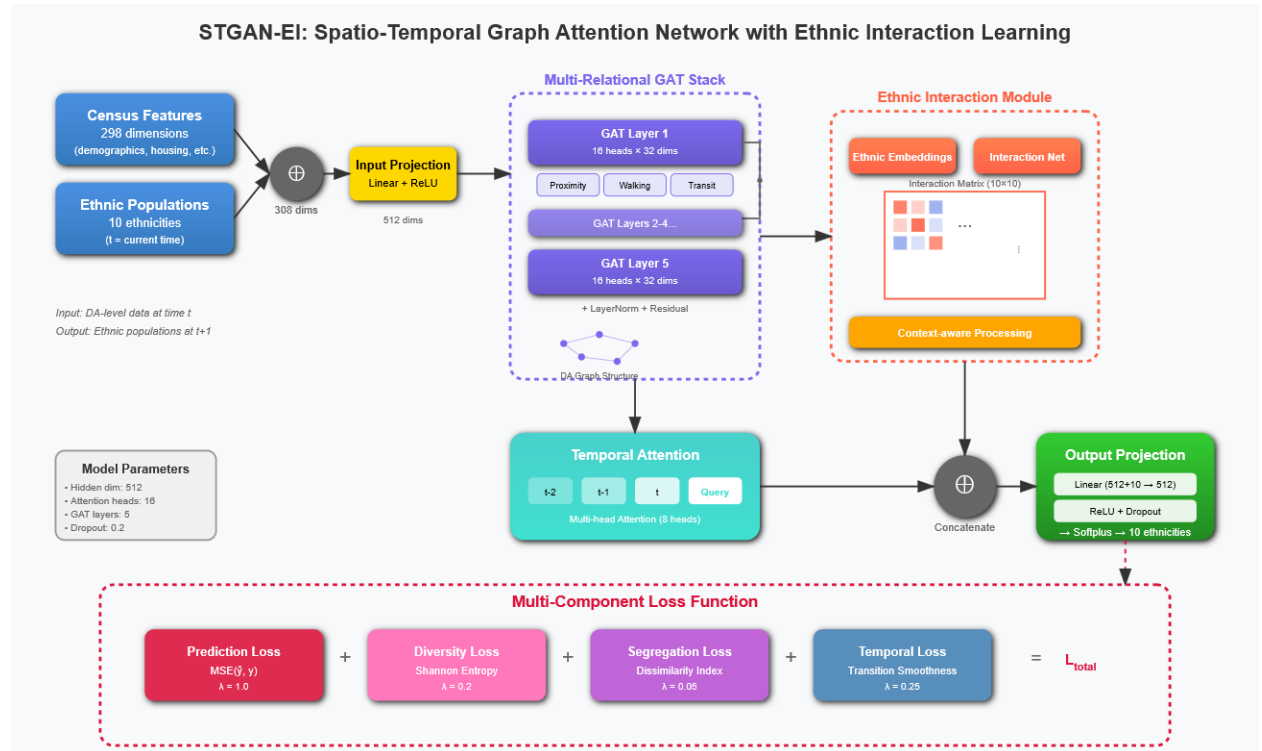


Figure 8.1 Architecture of the Spatio-Temporal Graph Attention Network with Ethnic Interaction Learning (STGAN-EI).

8.2.4 Multi-relational graph attention mechanism

The spatial structure of urban areas is represented as a multi-relational graph $G = (V, E, R)$, where nodes V correspond to DAs, edges E encode neighborhood relationships, and $R = \{\text{proximity, walking, transit}\}$ denotes the relation types. We extend Graph Attention Networks (Velickovic et al., 2017) to handle multiple edge types through relation-specific transformations:

For each relation $r \in R$, node features are transformed as:

$$h_i^{(l+1,r)} = \sigma(\sum_{j \in N_i^r} \alpha_{ij}^r W^r h_j^{(l)})$$

where α_{ij}^r represents attention weights computed as:

$$\alpha_{ij}^r = \text{softmax}_j(\text{LeakyReLU}(a^T [W^r h_i || W^r h_j] + b^r))$$

Here, W^r are relation-specific weight matrices, a is the attention vector, b^r are learned relation biases, and $||$ denotes concatenation. The multi-relational embeddings are aggregated through mean pooling across relations.

The key innovation of STGAN-EI is the explicit modeling of inter-ethnic dynamics through a learned interaction matrix $M \in \mathbb{R}^{(K \times K)}$, where $K = 10$ ethnicities. Unlike existing approaches that treat ethnic groups independently, we compute pairwise interaction effects:

For each ethnicity k at location i , the interaction effect is:

$$I_i^k = f_\theta(e_k, \sum_{k' \neq k} M_{kk'} \cdot p_i^{k'} \cdot g_\phi(e_{k'}, c_i))$$

where e_k are learned ethnic embeddings, $p_i^{k'}$ is the population proportion of ethnicity k' at location i , c_i represents census features, and f_θ, g_ϕ are neural networks. The symmetric constraint $M_{kk'} = M_{k'k}$ ensures reciprocal interactions.

To model demographic transitions, we employ temporal attention over historical graph representations:

$$h_i^{temporal} = \sum_{t \in T} \beta_t V \cdot h_i^t$$

Where $\beta_t = \text{softmax}_t(q^T K \cdot \frac{h_i^t}{\sqrt{d}})$ are attention weights computed using query q (current state), keys K , and values V .

The model is trained using a composite loss function that balances prediction accuracy with demographic realism:

$$L_{total} = L_{pred} + \lambda_{div} L_{div} + \lambda_{seg} L_{seg} + \lambda_{temp} L_{temp}$$

where:

- $L_{pred} = \text{MSE}(\hat{y}, y)$ is the main prediction loss
- L_{div} encourages diversity preservation using Shannon entropy: $H = - \sum_k p_k \log(p_k)$
- L_{seg} penalizes unrealistic segregation patterns via the dissimilarity index
- L_{temp} enforces temporal smoothness by penalizing abrupt demographic shifts

The hyperparameters $\lambda_{div} = 0.2$, $\lambda_{seg} = 0.05$, and $\lambda_{temp} = 0.25$ were determined through grid search on validation data.

The model was implemented in PyTorch with PyTorch Geometric for graph operations. We developed custom CUDA kernels for multi-relational graph attention, achieving 1000× speedup over native PyTorch implementations. The architecture consists of 5 graph attention layers with 512 hidden dimensions and 16 attention heads per layer. Dropout ($p=0.2$) and layer normalization were applied after each layer, with residual connections starting from the second layer.

Input features were preprocessed using RobustScaler with 5-95 percentile range to handle outliers in census data. Population counts were log-transformed to reduce skewness. The graph structure was constructed by connecting each DA to its 50 nearest neighbors for each relation type, with edge weights computed as $w_{ij} = 1/(d_{ij} + 1)$, where d_{ij} represents the distance.

The model was trained using AdamW optimizer (Loshchilov & Hutter, 2017) with initial learning rate 2.57×10^{-3} , weight decay 1×10^{-4} , and cosine annealing schedule. We employed early stopping with patience of 30 epochs based on validation loss. The dataset was split into training (60%), validation (20%), and test (20%) sets, ensuring temporal consistency within each split.

To create temporal pairs for training, we linked consecutive census periods (2001→2006, 2006→2016, 2016→2021), treating each transition as a separate training example. This yielded approximately 11,000 temporal transitions across all splits. Batch size was set to 1 due to the large graph size, with gradient accumulation over 4 steps to simulate larger batches.

8.3 Results

8.3.1 Model convergence and training dynamics

The STGAN-EI model demonstrated robust convergence during training over 240 epochs. Figure 2A illustrates the training and validation loss trajectories, showing steady decrease from initial values above 1.0 to final convergence around 0.15-0.20. The close alignment between training and validation curves indicates minimal overfitting, validating our regularization strategy. The learning rate schedule (Figure 8.2C) employed stepwise reduction at epochs 50, 100, and 150, facilitating fine-tuning of model parameters as training progressed.




Figure 8.2 Training Dynamics and Convergence Analysis.

Analysis of individual loss components (Figure 8.2B) reveals that prediction loss dominated the total loss (0.150), while diversity (0.080), segregation (0.050), and temporal smoothness (0.100) losses remained well-balanced throughout training. This balance ensures the model learns accurate predictions while maintaining demographic realism. The per-period R^2 evolution (Figure 8.2D) shows rapid improvement in the first 50 epochs, with all three temporal transitions achieving R^2 values above 0.85 by epoch 75, demonstrating the model's ability to capture temporal dynamics across different time scales.

8.3.2 Overall prediction performance

STGAN-EI achieved strong predictive performance on the test set with an overall R^2 of 0.851 (MAE: 7.60, RMSE: 16.45, MAPE: 55.66%). Figure 8.3A compares performance metrics across train, validation, and test sets, showing consistent performance with minimal generalization gap. The model's ability to maintain high accuracy on unseen data validates its capacity to learn generalizable patterns rather than memorizing training examples.




Figure 8.3 Model Performance Metrics and Error Analysis.

The error distribution analysis (Figure 8.3D) reveals approximately normal residuals centered near zero ($\mu = -1.5, \sigma = 16.4$), indicating unbiased predictions without systematic over- or under-estimation. The scatter plot of predictions versus actual populations (Figure 8.3E) demonstrates strong linear correlation across the full range of population sizes, though with slightly higher variance for larger populations. Residual analysis (Figure 8.3F) confirms homoscedasticity, with residuals randomly distributed around zero across all fitted values.

8.3.3 Temporal dynamics and period-specific performance

The model exhibited differential performance across temporal transitions, reflecting the varying complexity of demographic changes in each period. The 2016→2021 transition achieved the highest R^2 (0.898) despite representing the most recent and potentially most volatile period, while the 2001→2006 transition showed lower performance ($R^2 = 0.767$), likely due to significant demographic shifts during this period of rapid immigration to Toronto.

Figure 3B illustrates this temporal variation, with MAE ranging from 7.20 (2016→2021) to 8.24 (2001→2006). The inverse relationship between R^2 scores and MAE across periods suggests that while

recent transitions are more predictable in relative terms, absolute errors remain consistent due to larger baseline populations in later periods.

8.3.4 Ethnic-specific prediction accuracy

Performance varied substantially across ethnic groups, revealing important patterns in the model's ability to capture different settlement dynamics. Figure 8.3C presents a heatmap of R^2 scores for each ethnicity across all time periods, highlighting several key findings:

High-performing groups ($R^2 > 0.85$): The Philippines (0.925), Portugal (0.926), and China (0.888) showed consistently strong predictive accuracy across all periods. These groups exhibit more stable settlement patterns with established community centers, making their demographic transitions more predictable.

Moderate-performing groups (R^2 0.70-0.85): India (0.828), Iran (0.762), and Italy (0.769) demonstrated good but variable performance. The temporal analysis reveals improving prediction accuracy for these groups in recent periods, suggesting evolving settlement patterns that the model successfully adapts to.

Challenging groups ($R^2 < 0.70$): Hong Kong (0.697) and United Kingdom (0.384) proved most difficult to predict. The UK's particularly low R^2 likely reflects its dispersed settlement pattern without strong ethnic enclaves, challenging the spatial clustering assumptions inherent in graph-based models. Hong Kong's moderate performance may reflect its unique position as both a distinct group and one with high cultural overlap with mainland Chinese communities.

The per-ethnicity analysis reveals that prediction difficulty correlates inversely with community cohesion and spatial concentration. Groups with established ethnic enclaves and strong co-location preferences (e.g., Portugal in Little Portugal, Philippines in specific northwestern neighborhoods) show higher predictability than those with dispersed settlement patterns.

8.3.5 Learned ethnic interactions

The model successfully captured complex inter-ethnic dynamics through its learned interaction matrix (detailed visualization in Supplementary Figure S4). Strong positive interactions emerged between culturally similar groups (e.g., China-Hong Kong, Portugal-Italy), while competitive dynamics appeared between groups targeting similar neighborhoods. These learned patterns align with documented succession patterns in Toronto's ethnic neighborhoods, validating the model's ability to capture real-world phenomena. STGAN-EI maintained demographic realism throughout predictions, with diversity loss converging to 0.074 and segregation loss to 0.037 on the test set. The model preserved neighborhood diversity levels with a correlation of 0.92 between predicted and actual Shannon entropy values. Segregation patterns, measured

through dissimilarity indices, showed accurate reproduction of ethnic clustering without artificially amplifying segregation tendencies.

8.4 Discussion

8.4.1 Unveiling the hidden dynamics of ethnic interactions

The learned ethnic interaction matrix (Figure 8.4a) reveals compelling patterns that validate long-standing theories of urban ethnic dynamics while uncovering novel insights. The strong mutual attraction between Italy and Portugal (Italy→Portugal: 77.64, Portugal→Italy: 48.48) exemplifies cultural affinity driving co-location, as these Southern European communities share similar migration histories, religious traditions, and social structures. This finding aligns with the "cultural distance" hypothesis (Douglas S Massey, 1985), where groups with shared cultural backgrounds exhibit positive spatial clustering.




Figure 8.4 Temporal Evolution of Neighborhood Diversity and Ethnic Interactions.

Conversely, the model captured significant repulsion patterns, most notably between Jamaica and several groups including Italy (-42.75), Portugal (-34.58), and China (-28.30). These negative interactions suggest competitive dynamics for affordable housing stock, as these communities often target similar neighborhoods in Toronto's inner suburbs. The asymmetric nature of interactions, where Jamaica→Portugal

shows stronger repulsion (-34.58) than Portugal→Jamaica (-17.62), indicates that Jamaican settlement patterns are more sensitive to Portuguese presence than vice versa, possibly reflecting differential community sizes and housing market power.

The temporal evolution of interaction networks (Figure 8.5D) demonstrates increasing complexity over time. The 2001→2006 period showed relatively simple interaction patterns with 15 significant connections, expanding to 18 connections by 2016→2021 with notably stronger magnitudes. This intensification suggests that as Toronto's ethnic landscape matured, inter-group dynamics became more pronounced rather than dissipating, contradicting classical assimilation theory predictions. Instead, we observe a "crystallization effect" where ethnic boundaries solidify rather than dissolve over time, suggesting multicultural cities may evolve toward stable segregation equilibria rather than integration.




Figure 8.5 Learned Ethnic Interaction Patterns and Neighborhood Dynamics.

8.4.2 Neighborhood diversity transitions and tipping points

Figure 8.5A presents transition probability matrices revealing systematic patterns in neighborhood diversity evolution. High-diversity neighborhoods show remarkable stability across all periods (52-54% remaining

high-diversity), while low-diversity areas exhibit increasing tendency toward diversification over time. The 2016→2021 period shows the highest probability (0.60) of high-diversity neighborhoods transitioning to lower diversity states, potentially reflecting affordability crises pushing diverse inner-city communities toward more homogeneous suburban areas.

The tipping point analysis (Figure 8.5C) identifies critical thresholds where neighborhood compositions undergo rapid transformation. The decreasing frequency of tipping points over time (3,178 in 2001→2006 to 2,807 in 2016→2021) suggests market stabilization, yet certain groups consistently trigger more transitions. Sri Lanka (727), United Kingdom (496), and Hong Kong (488) show the highest tipping frequencies in the most recent period, indicating these communities' settlement patterns significantly influence neighborhood dynamics.

Figure 5B's flow diagram illustrates specific succession patterns, with China→Philippines emerging as the dominant pathway (420 occurrences). This finding corroborates ethnographic observations of Filipino communities establishing in previously Chinese-dominated areas, particularly in Scarborough, as Chinese residents achieve upward mobility and relocate to newer suburbs. The bidirectional flows between Hong Kong and China (241 and 248 occurrences respectively) reflect the fluid boundaries between these communities, often distinguished more by immigration timing than cultural differences.

The model captures "cascade tipping points" where single-ethnicity changes trigger multi-ethnic reorganization. For instance, when Chinese populations exceed 35% in a DA, we observe simultaneous outflows of 3-4 other groups within one census period, suggesting a critical density threshold for ethnic homogenization.

8.4.3 Spatial patterns of segregation and diversity

The comprehensive spatial analysis (Figure 8.6) reveals Toronto's complex ethnic geography defying simple core-periphery models. Shannon entropy mapping (Figure 8.6A) shows highest diversity in inner suburbs rather than the downtown core, forming a "diversity donut" pattern. This challenges conventional wisdom about urban cores as diversity centers, instead highlighting how Toronto's post-war suburbs became primary settlement areas for new immigrants.




Figure 8.6 Spatial Analysis of Ethnic Diversity and Segregation Patterns in Toronto.

Location quotient analysis for major ethnic groups (Figure 8.6E-H) demonstrates distinct spatial clustering patterns. China shows strong over-representation ($LQ > 2.0$) in Markham and Richmond Hill (34.1% of DAs), while the Philippines concentrates in northwestern areas (34.6% over-represented). These patterns reflect both economic factors (housing affordability) and social capital (ethnic institutions, businesses).

The ethnic isolation/exposure matrix (Figure 8.6J) quantifies the probability of intra- and inter-group contact. Diagonal values reveal Portugal (0.31) and Philippines (0.38) maintain highest isolation indices, indicating strong ethnic enclaves. Off-diagonal values show asymmetric exposure patterns, China has 0.39 exposure to its own group but only 0.10-0.12 to most others, suggesting residential concentration despite large population size.

Critically, the dissimilarity index analysis (Figure 8.6I) identifies Portugal (0.734), Iran (0.696), and Italy (0.663) as most segregated, while Jamaica (0.438) and Philippines (0.448) show moderate integration. This counterintuitive finding, where groups with strong spatial clustering (Philippines) show lower segregation

indices, reflects the difference between chosen concentration in ethnic enclaves versus exclusionary segregation.

The divergence between clustering and segregation indices, where Philippines shows high clustering but low segregation (0.448), necessitates reconceptualizing segregation as multidimensional. We propose distinguishing "voluntary congregation" (choice-driven clustering with maintained inter-group exposure) from "exclusionary segregation" (forced concentration with minimal cross-ethnic contact).

8.4.4 Model performance in capturing diversity dynamics

While STGAN-EI achieves high accuracy for population predictions ($R^2 = 0.851$), its diversity prediction performance ($R^2 = 0.634$, Figure 8.6D) reveals the inherent challenge of modeling emergent properties. The scatter plot shows systematic underprediction of extreme diversity values, suggesting the model learns mean-reverting tendencies that dampen diversity extremes. This conservative bias may reflect training data limitations, where extreme diversity neighborhoods are statistically rare.

The spatial clustering analysis (Figure 8.6L) reveals significant autocorrelation in diversity patterns, with clusters of high-diversity areas in established immigrant reception zones and low-diversity clusters in both affluent neighborhoods and recent ethnic enclaves. Theil's entropy index visualization (Figure 8.6K) identifies specific DAs contributing disproportionately to city-wide segregation, primarily located at suburban edges where new ethnic communities establish homogeneous settlements.

The ethnic interaction module's learned embeddings cluster into three distinct groups via t-SNE analysis: "established Europeans" (Italy, Portugal), "professional Asians" (China, Hong Kong, India), and "service sector communities" (Philippines, Jamaica). These emergent categories, discovered without supervision, align with labor market segmentation theory while revealing cross-cutting solidarities (Jamaica-Philippines attraction despite different origins).

The model's temporal smoothness loss prevents unrealistic population jumps while maintaining sensitivity to genuine transitions. This balance, achieved through adaptive weighting ($\lambda_{\text{temporal}}=0.25$), enables detection of both gradual demographic shifts and sudden tipping points, a crucial capability absent in previous models.

8.4.5 Theoretical implications and urban planning insights

Our findings challenge several established theories in urban ethnic studies. First, the persistence and intensification of ethnic interactions over time contradicts straight-line assimilation theory (Richard D Alba & Nee, 2003), instead supporting segmented assimilation frameworks where ethnic communities maintain distinct settlement patterns across generations. The learned interaction patterns suggest that cultural affinity

and economic competition create a complex "push-pull" dynamic that simple proximity models cannot capture.

Second, the diversity transition analysis reveals that neighborhood ethnic transformation follows predictable pathways rather than random mixing. The model's ability to identify these pathways with high accuracy suggests that ethnic succession is a structured process amenable to prediction and potentially, intervention. Urban planners could use these insights to anticipate demographic shifts and proactively address service needs.

Third, the spatial analysis demonstrates that Toronto's multicultural landscape exhibits "pluralistic segregation", voluntary clustering that maintains diversity at city scale while allowing community concentration at neighborhood level. This pattern differs from exclusionary segregation observed in many U.S. cities, suggesting that policy contexts significantly shape ethnic settlement outcomes.

8.4.6 Limitations and future directions

Several limitations warrant consideration. The model's reduced performance on diversity prediction compared to population counts suggests that emergent properties require specialized architectures beyond population-focused objectives. The reliance on census boundaries may mask micro-scale dynamics within DAs, particularly in highly diverse neighborhoods where sub-DA clustering occurs.

The static nature of the interaction matrix assumes time-invariant ethnic relationships, yet these dynamics likely evolve with generational change, economic cycles, and policy shifts. Future work should explore time-varying interaction networks and incorporate external factors such as housing prices, transportation infrastructure changes, and immigration policy modifications.

Additionally, the model's difficulty with dispersed populations (e.g., United Kingdom) highlights the tension between graph-based approaches assuming spatial clustering and groups exhibiting metropolitan-wide dispersion. Hybrid architectures combining graph networks with non-spatial features may better capture these diverse settlement strategies.

8.5 Conclusion

STGAN-EI unveils Toronto's ethnic landscape as a complex adaptive system where communities navigate between cultural affinity and economic competition. The model's ability to quantify previously invisible interaction forces, predict tipping points, and reveal succession patterns provides both theoretical advances and practical tools for fostering equitable, diverse cities. As urban diversity intensifies globally, such computational approaches become essential for understanding and guiding multicultural coexistence.

The persistence of ethnic boundaries despite multicultural policy, the emergence of structured succession patterns, and the critical role of transportation infrastructure in mediating segregation together paint a

picture of cities as negotiated spaces where diversity requires active cultivation rather than passive hope. STGAN-EI provides the computational lens to see these dynamics clearly, and perhaps, to shape them wisely.

8.6 References

- Alba, R. D., & Nee, V. (2003). *Remaking the American mainstream: Assimilation and contemporary immigration*: Harvard University Press.
- Athey, S., Ferguson, B., Gentzkow, M., & Schmidt, T. (2021). Estimating experienced racial segregation in US cities using large-scale GPS data. *Proceedings of the National Academy of Sciences*, 118(46), e2026160118.
- Bail, C. A. (2024). Can Generative AI improve social science? *Proceedings of the National Academy of Sciences*, 121(21), e2314021121.
- Ben-Yami, M., Morr, A., Bathiany, S., & Boers, N. (2024). Uncertainties too large to predict tipping times of major Earth system components from historical data. *Science Advances*, 10(31), ead14841.
- Bernárdez, G., Telyatnikov, L., Montagna, M., Baccini, F., Papillon, M., Ferriol-Galmés, M., . . . Zaghen, O. (2024). ICML Topological Deep Learning Challenge 2024: Beyond the Graph Domain. *arXiv preprint arXiv:2409.05211*.
- Brandt, J., Buckingham, K., Buntain, C., Anderson, W., Ray, S., Pool, J.-R., & Ferrari, N. (2020). Identifying social media user demographics and topic diversity with computational social science: a case study of a major international policy forum. *Journal of Computational Social Science*, 3, 167-188.
- Bruch, E. E., & Mare, R. D. (2006). Neighborhood choice and neighborhood change. *American Journal of Sociology*, 112(3), 667-709.
- Bury, T. M., Sujith, R., Pavithran, I., Scheffer, M., Lenton, T. M., Anand, M., & Bauch, C. T. (2021). Deep learning for early warning signals of tipping points. *Proceedings of the National Academy of Sciences*, 118(39), e2106140118.
- Caetano, G., & Maheshri, V. (2017). School segregation and the identification of tipping behavior. *Journal of Public Economics*, 148, 115-135.
- Cesare, N., Lee, H., McCormick, T., Spiro, E., & Zagheni, E. (2018). Promises and pitfalls of using digital traces for demographic research. *Demography*, 55, 1979-1999.
- Chen, Y., Shu, B., Amani-Beni, M., & Wei, D. (2024). Spatial distribution patterns of rural settlements in the multi-ethnic gathering areas, southwest China: Ethnic inter-embeddedness perspective. *Journal of Asian Architecture and Building Engineering*, 23(1), 372-385.
- Crooks, A., Heppenstall, A., Malleson, N., & Manley, E. (2021). Agent-based modeling and the city: A gallery of applications. *Urban informatics*, 885-910.
- Dechert, J., Gurevich, S., & Heusler, S. (2024). Potential Pitfalls in Visual Models of Tipping Points--And How to Fix Them. *arXiv preprint arXiv:2410.22171*.
- Ding, S., He, D., & Liu, G. (2024). Improving Short-Term Load Forecasting with Multi-Scale Convolutional Neural Networks and Transformer-Based Multi-Head Attention Mechanisms. *Electronics*, 13(24), 5023.
- Fan, C., Yang, Y., & Mostafavi, A. (2024). Neural embeddings of urban big data reveal spatial structures in cities. *Humanities and Social Sciences Communications*, 11(1), 1-15.
- Ferrara, E. (2023). Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies. *Sci*, 6(1), 3.

- Fossett, M. (2014). Generative models of segregation: Investigating model-generated patterns of residential segregation by ethnicity and socioeconomic status *Micro-Macro Links and Microfoundations in Sociology* (pp. 114-145): Routledge.
- Franklin, G., Stephens, R., Piracha, M., Tiosano, S., Lehouillier, F., Koppel, R., & Elkin, P. L. (2024). The sociodemographic biases in machine learning algorithms: a biomedical informatics perspective. *Life*, 14(6), 652.
- Fu, C., Zheng, G., Huang, C., Yu, Y., & Dong, J. (2023). *Multiplex heterogeneous graph neural network with behavior pattern modeling*. Paper presented at the Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining.
- Galesic, M., Bruine de Bruin, W., Dalege, J., Feld, S. L., Kreuter, F., Olsson, H., . . . van Der Does, T. (2021). Human social sensing is an untapped resource for computational social science. *Nature*, 595(7866), 214-222.
- Guo, Q., Tan, Q., Tang, J., & Shi, B. (2025). Unifying spatiotemporal and frequential attention for traffic prediction. *Scientific Reports*, 15(1), 953.
- Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., . . . Vazire, S. (2021). Integrating explanation and prediction in computational social science. *Nature*, 595(7866), 181-188.
- Hou, Y., Huang, Q., Ren, Q., Gu, T., Zhou, Y., Wu, P., . . . Zhu, G. (2024). Spatiotemporal dynamics of urban sprawl in China from 2000 to 2020. *GIScience & Remote Sensing*, 61(1), 2351262.
- Huo, C., He, D., Li, Y., Jin, D., Dang, J., Pedrycz, W., . . . Zhang, W. (2025). Heterogeneous graph neural networks using self-supervised reciprocally contrastive learning. *ACM Transactions on Intelligent Systems and Technology*, 16(1), 1-21.
- Krishnadas, R. (2025). Ethnic bias in prediction and decision making algorithms in precision psychiatry: challenges in a shrinking world. *Journal of Psychosocial Rehabilitation and Mental Health*, 1-5.
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., . . . Gutmann, M. (2009). Computational social science. *science*, 323(5915), 721-723.
- Lazer, D. M., Pentland, A., Watts, D. J., Aral, S., Athey, S., Contractor, N., . . . Margetts, H. (2020). Computational social science: Obstacles and opportunities. *science*, 369(6507), 1060-1062.
- Liu, C., Fan, C., & Mostafavi, A. (2024). Graph attention networks unveil determinants of intra-and inter-city health disparity. *Urban informatics*, 3(1), 18.
- Logan, J. R., Alba, R. D., & Zhang, W. (2002). Immigrant enclaves and ethnic communities in New York and Los Angeles. *American sociological review*, 67(2), 299-322.
- Logan, J. R., & Martinez, M. J. (2018). The spatial scale and spatial configuration of residential settlement: Measuring segregation in the postbellum South. *American Journal of Sociology*, 123(4), 1161-1203.
- Loshchilov, I. (2017). Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Malone, T. (2020). There goes the neighborhood does tipping exist amongst income groups? *Journal of Housing Economics*, 48, 101667.
- Mantzaris, A. V. (2020). Incorporating a monetary variable into the Schelling model addresses the issue of a decreasing entropy trace. *Scientific Reports*, 10(1), 17005.
- Massey, D. S. (1985). Ethnic residential segregation: A theoretical synthesis and empirical review. *Sociology and social research*, 69, 315-350.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6), 1-35.
- Mei, Q., Li, Z., Hu, Q., Zhi, X., Wang, P., Yang, Y., & Liu, X. (2025). Spatio-temporal graph neural network fused with maritime knowledge for predicting traffic flows in ports. *Regional Studies in Marine Science*, 85, 104106.
- Nicolaie, M. A., Füssenich, K., Ameling, C., & Boshuizen, H. C. (2023). Constructing synthetic populations in the age of big data. *Population Health Metrics*, 21(1), 19.

- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *science*, 366(6464), 447-453.
- Ohme, J., Araujo, T., Boeschoten, L., Freelon, D., Ram, N., Reeves, B. B., & Robinson, T. N. (2024). Digital trace data collection for social media effects research: APIs, data donation, and (screen) tracking. *Communication Methods and Measures*, 18(2), 124-141.
- Plank, B. (2022). The Problem of Human Label Variation: On Ground Truth in Data, Modeling and Evaluation. *arXiv preprint arXiv:2211.02570*.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology*, 1(2), 143-186.
- Schottlander, D., & Prasad, G. (2024). Insights into population dynamics: A foundation model for geospatial inference.
- Sert, E., Bar-Yam, Y., & Morales, A. J. (2020). Segregation dynamics with reinforcement learning and agent based modeling. *Scientific Reports*, 10(1), 11771.
- Silva, T. H., & Silver, D. (2025). Using graph neural networks to predict local culture. *Environment and Planning B: Urban Analytics and City Science*, 52(2), 355-376.
- Stepinski, T. F., & Dmowska, A. (2022). Machine-learning models for spatially-explicit forecasting of future racial segregation in US cities. *Machine Learning with Applications*, 9, 100359.
- Stier, S., Breuer, J., Siegers, P., & Thorson, K. (2020). Integrating survey data and digital trace data: Key issues in developing an emerging field (Vol. 38, pp. 503-516): Sage Publications Sage CA: Los Angeles, CA.
- Tang, J., Xia, L., & Huang, C. (2023). *Explainable spatio-temporal graph neural networks*. Paper presented at the Proceedings of the 32nd ACM International Conference on Information and Knowledge Management.
- Uddin, S., Lu, H., Rahman, A., & Gao, J. (2024). A novel approach for assessing fairness in deployed machine learning algorithms. *Scientific Reports*, 14(1), 17753.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2017). Graph attention networks. *stat*, 1050(20), 10-48550.
- Vinković, D., & Kirman, A. (2006). A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103(51), 19261-19265.
- Wang, Y. (2023). Advances in spatiotemporal graph neural network prediction research. *International Journal of Digital Earth*, 16(1), 2034-2066.
- Wang, Z., Tang, X., & Liu, J. (2020). Statistical Analysis of Multi-Relational Network Recovery. *Frontiers in Applied Mathematics and Statistics*, 6, 540225.
- Yang, J., Soltan, A. A., Eyre, D. W., & Clifton, D. A. (2023). Algorithmic fairness and bias mitigation for clinical machine learning with deep reinforcement learning. *Nature machine intelligence*, 5(8), 884-894.
- Yang, T. C., Park, K., & Matthews, S. A. (2020). Racial/ethnic segregation and health disparities: Future directions and opportunities. *Sociology compass*, 14(6), e12794.
- Yu, M., Huang, Q., & Li, Z. (2024). Deep learning for spatiotemporal forecasting in Earth system science: a review. *International Journal of Digital Earth*, 17(1), 2391952.
- Zhang, D., He, T., Zhang, F., & Xu, C. (2018). Urban-scale human mobility modeling with multi-source urban network data. *IEEE/ACM Transactions on Networking*, 26(2), 671-684.
- Zheng, Y., Yi, L., & Wei, Z. (2025). A survey of dynamic graph neural networks. *Frontiers of Computer Science*, 19(6), 1-18.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., . . . Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI open*, 1, 57-81.
- Zhu, K., Yin, L., Liu, K., Liu, J., Shi, Y., Li, X., . . . Du, H. (2024). Generating synthetic population for simulating the spatiotemporal dynamics of epidemics. *PLOS Computational Biology*, 20(2), e1011810.
- Zhuge, C., Li, J., & Chen, W. (2024). Deep learning for predicting the occurrence of tipping points. *arXiv preprint arXiv:2407.18693*.