

NLP4PBM: A Systematic Review on Process Extraction using Natural Language Processing with Rule-based, Machine and Deep Learning Methods

William Van Woensel^a[0000-0002-7049-8735] and Soroor Motie^a[0000-0002-3675-4774]

^aTelfer School of Management, University of Ottawa, 55 Laurier E., Ottawa, K1N 6N5, Ontario, Canada

ARTICLE HISTORY

Compiled September 10, 2024

ABSTRACT

This literature review studies the field of automated process extraction, i.e., transforming textual descriptions into structured processes using Natural Language Processing (NLP). We found that Machine Learning (ML) / Deep Learning (DL) methods are being increasingly used for the NLP component. In some cases, they were chosen for their suitability towards process extraction, and results show that they can outperform classic rule-based methods. We also found a paucity of gold-standard, scalable annotated datasets, which currently hinders objective evaluations as well as the training or fine-tuning of ML / DL methods. Finally, we discuss preliminary work on the application of LLMs for automated process extraction, as well as promising developments in this field.

KEYWORDS

Business Process Management; Process Extraction; Natural Language Processing; Machine Learning; Deep Learning; Language Models

1. Introduction

Business processes are structured as series of tasks that collectively accomplish a business objective [1]. As such, they can provide an objective framework for an organization's day-to-day operations [1]. Business Process Management (BPM) is devoted to the analysis, design, implementation and management of business processes [1]. BPM commonly relies on process models, i.e., structured representations of the processes in question [2]. Process models may represent control flow (e.g., sequence, concurrency) or decisional (e.g., criteria and their requirements) aspects of processes. Moreover, in the former category, process models can be imperative, i.e., exhaustively prescribing the occurrence and ordering of activities; or declarative, i.e., declaring high-level constraints on the occurrence and ordering of specific activities¹ [3].

Currently, most organizations rely solely on natural text to describe business processes and their requirements [4]. Vast amounts of textual data relevant to business processes is constantly being generated, such as reports, emails, and ad-hoc documentation. To implement BPM in such organizations, textual sources could be used as a

CONTACT William Van Woensel. Email: wavanwoen@uottawa.ca

¹E.g., a *precedence* constraint on (A, B) states that B only occurs if preceded by A.

source for structured process models. The extraction of process models from natural language text, albeit manually or automatically, is referred to as *process extraction* [2].

Nevertheless, according to a 2019 report from Deloitte [5], only a few organizations (18%) analyze unstructured data such as natural text to gain business insights, including as the processes described therein, since it is more challenging to interpret. Indeed, manual process extraction in particular is known to be time-consuming and error-prone [6; 7; 2]; Herbst reported that more than half of the time in BPM is often spent on manual process extraction [8]. Hence, automated support for process extraction, if done accurately and reliably, can improve the efficiency of process extraction. To that end, a number of studies have used Natural Language Processing (NLP) methods to implement process extraction: we refer to this as *automated process extraction*, or simply *process extraction* if the automated aspect is clear from the context. By providing structured process models, automated process extraction thus has the potential to improve the effectiveness of BPM in organizations currently relying on textual data. The availability of extracted process models, i.e., their management, storage and retrieval, will further have implications for information management within enterprises.

Traditionally, process extraction methods have involved rule-based and/or Machine Learning (ML)-based methods to implement two steps: (1) *NLP*, which categorises the fundamental building blocks of the text; such as labeling nouns, verbs, and adjectives, recognizing unique entities (e.g., people, events), and tagging their semantic concepts (e.g., actor, activity); and (2) *process generation*, which maps the NLP output to a process model that captures the control-flow or decisional elements from the text. With the emergence of Deep Learning (DL) in NLP, the last 5 years have seen a paradigm shift: authors are increasingly using DL models such as Transformers (e.g., BERT [9]) and Long Short-Term Memory (LSTM) [10] in process extraction. Even more recently, the advent of Large Language Models (LLM) has led to a surge of interest in LLM-powered NLP for a number of fields. We expect LLM to also impact the process extraction field, given their capabilities to generate code and models [11]; we have already seen preliminary work on this topic [12; 13; 14; 15]. We discuss this initial work, and the promise of recent innovations, in our Discussion section.

With these promising natural language applications on the horizon, we believe it is opportune to present a systematic review of process extraction from the “pre-LLM” era. For researchers who want to contribute to this field, we aim to acquaint them with state-of-the-art methods and pipelines, applied evaluation methods and datasets, and preliminary work on LLM in process extraction. For practitioners who want to apply process extraction, we present an overview of relevant methods, including publicly available NLP tools. We discuss both traditional approaches to process extraction, and emphasise recent work that relies on ML/DL. Hence, this paper will cover the following research questions:

- RQ1: Which NLP and process generation methods, including publicly available tools, have been studied and used for extracting process models?
- RQ2: To what extent have ML/DL methods been studied in process extraction?
- RQ3: Which evaluation methods and datasets have been utilised to evaluate process extraction?

The rest of the paper is structured as follows. Section 2 discusses other review papers on automated process extraction. Section 3 outlines our systematic review approach, including source databases, keywords, exclusion and inclusion criteria, and classification criteria. Sections 4-7 present our data synthesis results, including answers to the research questions. Section 8 summarises important observations from our study. Sec-

tion 9 discusses the limitations and threats to the validity of the study, and Section 10 concludes our paper.

2. Related Work

This section discusses other reviews on automated process extraction, and outlines the contributions of our systematic review to the literature. We differentiate between non-systematic and systematic literature reviews.

As a non-systematic review, Bellan et al. [2] conducted a qualitative analysis in 2020 on process extraction methods. The authors highlight multiple limitations when applied to *real-world* natural language text, and conclude that process extraction is a task that is far from solved. The same authors performed a non-systematic comparative analysis of the literature a year later [16], covering 13 papers published up to 2020, and categorised the used NLP techniques, supported process elements, and evaluations.

We identified systematic reviews on process extraction up to 2018. These reviews thus do not cover important DL advancements that took place after 2018, notably BERT (i.e., Transformers) [17]. The first relevant systematic review was conducted in 2016 by Riefer et al. [18], reviewing 5 papers in total. The authors analyzed 3 aspects, namely textual input, text analysis, and model generation. Bordignon et al. [7] performed a systematic review of works up to 2016 that apply NLP to the first 3 BPM lifecycle phases (i.e., identification, discovery and analysis), and categorised the NLP tools utilised. Maqbool et al. [6] covers process extraction up to 2018, focusing in particular on the extraction of Business Process Model and Notation (BPMN) [19] process models, categorising the used NLP tools and techniques, and BPMN modelling constructs and tools. Schüler et al. [20] performed a more general systematic review in 2022 (published in 2024) on the automated generation of process models from a range of input sources, including source code, business rules, event logs and natural text, among others. The authors only shortly discuss process extraction (2 paragraphs).

Compared to these studies, we present a systematic review of NLP-enabled process extraction literature up to 2023: we thus also include novel DL methods in our review. Further, we evaluate a more comprehensive set of review categories, including natural language analysis (text input, computational paradigm and tools), process generation (computational paradigm, intermediary representations, and target model), and evaluation (methods and dataset). Finally, our review pays special attention to works that rely on innovations in ML/DL-based NLP for process extraction.

3. Methods

We performed a systematic review of the literature based on the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines [21], which are widely recognised recommendations for conducting systematic reviews and meta-analyses. We searched 5 major scientific literature databases: Web of Science, Scopus, IEEE Xplore, ACM, and ScienceDirect, thus covering all disciplines relevant to our search. Our search queries focused on the intersection of two key concepts: NLP and BPM². Table 1 shows the synonyms and subfields (e.g., decision model extraction) considered for each concept and the corresponding search queries. We considered the title, abstract, and author’s keywords metadata fields in our search.

²Any work focusing on the use of ML/DL for these concepts would thus also be included.

Table 1.: Synonyms and search queries for NLP and BPM.

Concepts	Synonyms	Search Queries
Natural Language Processing	NLP	“NLP” OR
	Natural Language Processing	“Natural Language Processing” OR
	Text Mining	”Text mining” OR
	Text Analysis	“Text analysis” OR
	NL Processing	“NL processing” OR
	NLTK	“NLTK”
Business Process Management	Process Modelling	“Process Modelling” OR
	Business Process Management	“Business Process Management” OR
	BPM	“BPM” OR
	Business Process Model and Notation	“Business Process Model and Notation” OR
	BPMN	“BPMN” OR
	Process Extraction	“Process Extraction” OR
	Process Management	“Process Management” OR
	Business Process Elicitation	“Business Process Elicitation” OR
	Decision Model and Notation	“Decision Model and Notation” OR
	DMN	“DMN” OR
	Process Discovery	“Process Discovery” OR
	Decision Model Extraction	“Decision Model Extraction” OR
	Business decision management	“Business decision management” OR
	Decision Extraction	“Decision Extraction”

We removed duplicates from the search results and downloaded the remaining papers. The papers’ titles, abstracts and keywords were screened for eligibility based on the inclusion/exclusion criteria described in the subsections below. In case a paper was still considered inconclusive after screening, it was fully reviewed to determine its eligibility. We used the Covidence³ software to keep track of the review process.

3.1. Inclusion Criteria (IC)

Articles meeting all of the following criteria were included in our systematic review: **(IC.1)** full research articles published in a peer-reviewed conference or journal, where a full text was available, and written in English (there was no restriction on time period); **(IC.2)** primary research articles, i.e., presenting original research contributions⁴; **(IC.3)** articles specifically covering the use of NLP for process extraction from natural language text; **(IC.4)** articles that explicate a concrete method and describe experiments to empirically validate their method.

3.2. Exclusion Criteria (EC)

We excluded articles from our systematic review that did not comply with the IC as follows: **(EC.1)** articles that are not full research papers (such as posters, short papers, abstracts, reports, theses, or book chapters), not published in a conference

³<https://www.covidence.org/>

⁴In case of multiple papers related to the same contribution, we reviewed the most recent paper.

Table 2.: Themes and categories used for classification.

Theme	Category	Description	Possible Values
Natural Language Analysis	Text Input	Restrictions on textual data.	User/group stories, unrestricted or not mentioned.
	Computing paradigm	Computing paradigm for NLP.	Rule-based or ML/DL modules, end-to-end DL.
	NLP Tools	Off-the-shelf NLP tools used.	E.g., NLTK, spaCy, WordNet, CoreNLP.
Process Model Generation	Computing Paradigm	Computing paradigm for process generation from NLP output.	Knowledge-based, ML/DL.
	Intermediary represent.	Use of notation-agnostic, intermediary representations.	E.g., graph-based, dependency tuples.
	Target model	Type of process model and notation that was targeted.	Control flow (imperative / declarative) vs. decisional. E.g., BPMN, DCL, DMN.
Evaluation	Methods & Metrics	Methods and metrics used for evaluation.	Methods: component-based vs. holistic; systematic vs. expert-based. Metrics: e.g., P, R, F1.
	Evaluation Dataset	Dataset used for evaluation.	Type (real, synthetic), target model, size, domain, availability, link (if any).
Metadata	Year	Year of Publication.	From 2011 to 2023.

or journal (e.g., workshops, discussion forums, white papers, technical reports), not peer-reviewed (including preprints), where no full text was available, or that were not written in English; **(EC.2)** secondary research articles, i.e., using primary research to derive results (e.g., literature reviews, meta-analysis, comments); **(EC.3)** articles not specifically covering the use of NLP for process extraction. This includes studies on the use of NLP for process redesign, process matching, process prediction, or sentiment analysis, works that target individual labels instead of natural text, or studies on generating natural text from processes; **(EC.4)** articles that only superficially discuss an applied method, or do not describe experiments to validate their method.

3.3. Study Classification

We classified the remaining studies according to the orthogonal themes in Table 2. In line with our research questions, themes include NLP, process extraction, evaluation datasets and methods, and metadata to analyze publication trends. Based on the reviewed papers, we identified a set of categories per theme; in turn, these were used to classify the papers (i.e., qualitative inductive coding [22]).

4. Results

Figure 1 shows the results of the systematic review process using the PRISMA flowchart. A total number of 524 studies were identified by our database search, 405 of which were unique. Table 3 shows the breakdown of search results per database. Following title and abstract screening, 70 articles were retained for full-text screening.

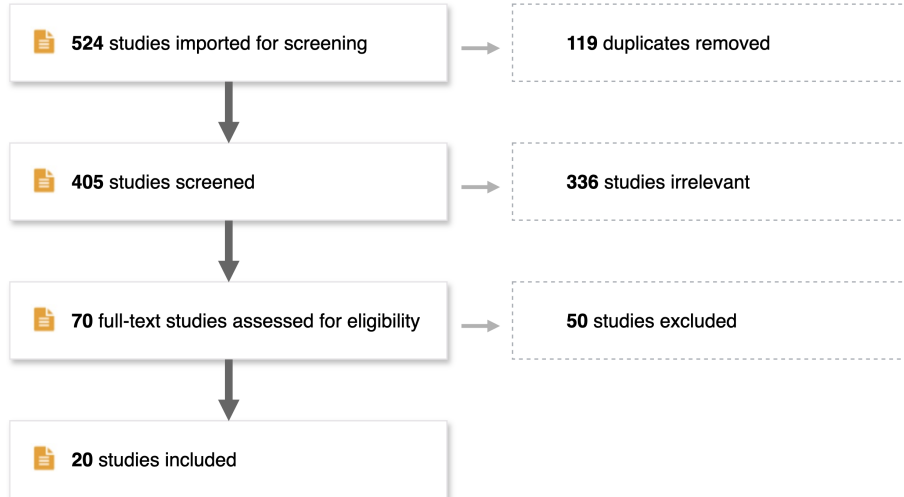


Figure 1.: PRISMA flow diagram with results for our systematic review.

In the end, 20 articles were found to fully meet our inclusion/exclusion criteria. Our search queries were executed in June 2023 and the found papers span a time period from 2011 to 2023.

Table 3.: Breakdown of search results per database.

Database	Web Address	Number of Results
Scopus	https://www.scopus.com/	303
Web of Science	https://www.webofknowledge.com/	90
IEEE Xplore	http://www.ieeexplore.ieee.org	33
ScienceDirect	https://www.sciencedirect.com	85
ACM Digital Library	http://dl.acm.org	13
Total with Duplicates		524
Number of Duplicates		119
Total without Duplicates		405

In the following sections, we answer our research questions RQ1-3 by discussing the reviewed papers in terms of the themes and categories from Table 2. Appendix A includes tables that categorize each paper along these themes in detail.

5. Natural Language Processing

This section discusses the NLP aspect of automated process generation, including restrictions on text input (Section 5.1), computational paradigms used to implement NLP (Section 5.2), and publicly available NLP tools that were used (Section 5.3).

5.1. Text Input

In agile software development, user stories are used as a natural language medium to describe software requirements [23]. Nasiri et al. [23] extracted Unified Modeling Language (UML) activity diagrams from user stories to capture software requirements

as process models. In an information systems context, Goncalves et al. [24] similarly focus on user stories for process extraction. User stories, gathered through group storytelling, are described as a natural way for people to outline their everyday activities, difficulties, and suggestions for solving problems. The authors note that writing stories in a free style poses obstacles for NLP, such as ambiguity and lack of clarity; hence, the authors restrict the textual input to a scenario structure.

Halioui et al. [25] focus on extracting bio-informatics workflows from scientific texts.

The remainder of the papers did not explicitly mention a restriction on text input. That said, approaches that target decisional models will clearly expect text input with decision-related information; idem for control flow models. Further, authors will evaluate their approach on datasets that, for instance, pertain to only one domain, or include sentences with minimal structural variability. Hence, a method’s (implicit) restrictions on text input may be reflected by the dataset used for evaluation. We discuss evaluation datasets in Section 7.1.

5.2. Computing Paradigm

Most reviewed papers involve a traditional “NLP pipeline”, which identifies and tags textual elements with increasing amounts of semantics—ranging from part-of-speech and grammar relations to high-level semantics concepts—until the output is sufficiently detailed to generate a process model. Although many customizations exist, a “traditional NLP pipeline” includes the tasks listed in Table 4.

Table 4.: Traditional NLP pipeline tasks.

NLP Tasks	Description
Tokenization	Splitting the raw text into chunks of words or sentences called <i>tokens</i> (e.g., based on whitespaces, periods, or regular expressions).
Stemming and Lemmatizing	Stemming reduces words to their <i>root form</i> by removing leading or trailing characters. Lemmatizing considers word meaning & context to more accurately find root forms.
Part-of-Speech (POS) Tagging	Assigning <i>grammatical categories</i> (e.g., nouns, verbs) to words.
Dependency Parsing	Analyzing grammatical relationships between words in a sentence to reveal <i>typed dependencies</i> (e.g., subject noun, verb object) [26].
Named Entity Recognition (NER)	Identifying <i>named entities</i> in text, such as names of persons, locations, and organizations. This term is often used interchangeably with Entity Extraction (EE) and Entity Recognition (ER).
Coreference Resolution	Linking pronouns (e.g., “they”) and noun phrases (e.g., “dog owner”) to the entities they refer to in the text (e.g., “Bob”). This term is often used interchangeably with Anaphora Resolution and Entity Resolution.
Semantic Analysis	Extracting and understanding the <i>meaning of words or phrases</i> within text (e.g., activities, actors). In this category, we include <u>Relation Extraction (RE)</u> , i.e., identifying semantic relations found in the text (e.g., sequence, parallel relations); <u>Word Sense Disambiguation (WSD)</u> , i.e., disambiguating word meanings based on context; and <u>word/sentence classification</u> , i.e., providing meaning to a word or sentence by assigning it a class.
Preprocessing	NLP pipelines that leverage ML and DL models often require pre-processing by generating word embeddings (e.g., using GloVe [27]) or generating feature-vectors (e.g., using bag-of-words or TF-IDF).

Table 5.: ML and DL models used for NLP tasks.

Task	ML / DL Models
Preprocessing	BERT (or possibly Word2Vec) for word embeddings [28]; Bag Of Words, TF-IDF, word embedding (GloVe) [9]; RNN and Bi-LSTM for sentence- & process-level encodings, ON-LSTM with process-level LM objective [10]
NER	Conditional Random Fields (CRF) [29]; Pre-Trained BERT [9; 30] and Bi-LSTM [9];
Coreference Resolution	Pre-trained DL model [29]; NeuralCoref4 [31; 32; 9]
<i>Semantic Analysis</i>	
RE	Catboost Gradient Boosting [29]; pre-trained BERT (binary, multi-class) [30]
WSD	PAUM (Perceptron Algorithm with Uneven Margins) [25]
Sentence/word classification	Pre-Trained BERT (DistilBERT), logistic regression, naïve bayes and support vector machines [9]; Bi-LSTM, CNN, MLP [28]; Pre-trained BERT [30]

Other works rely on “*non-traditional*” *DL pipelines*, as we discuss later.

To answer **RQ2**, we proceed by summarising the use of ML/DL models for NLP tasks from Table 4. Subsequently, we discuss performance comparisons between different computing paradigms; and describe non-traditional DL pipelines.

Figure 2 shows the evolution of computing paradigms used for NLP in reviewed papers from 2011-2023. In line with the advent of the DL BERT model in 2018, the first reviewed papers on DL-based process extraction appeared in 2020. In 2023, we find an even distribution across computing paradigms.

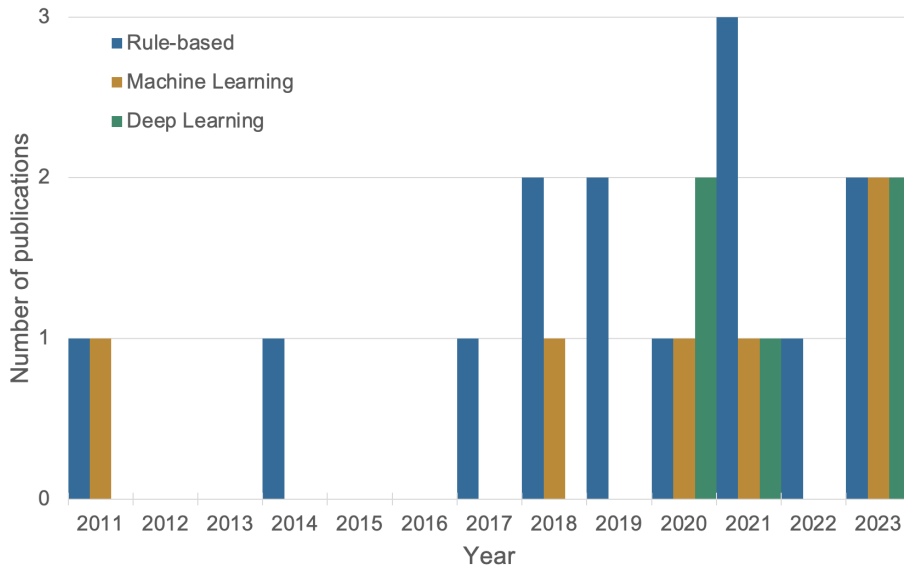


Figure 2.: Papers published over time with their computational paradigms for NLP⁵.

Table 5 summarises the ML/DL models used to implement NLP tasks.

⁵Note that papers covering multiple paradigms will be counted once for each paradigm.

Comparison between rule-, ML- and DL-based NLP methods.

Neuberger et al. [29] compared different NLP modules and pipelines:

- (1) An ML-based Relation Extraction (RE) module outperformed a rule-based RE module. The authors also noted that, while rule-based methods can be optimised for specific domains, they are hard to adapt to other applications.
- (2) Their NLP pipeline outperformed Jerex [33], an end-to-end DL method. The authors hypothesised that the used PET dataset [34] is not yet extensive enough for training state-of-the-art DL models.

Notably, their choice for Jerex was informed by the need for down-stream process generation: since Jerex’s is also able to identify the textual location of an entity, its surrounding text can be used to obtain richer BPMN activity labels.

Goossens et al. [9] make the following observations:

- (1) For sentence classification into dependency vs. logic sentences, a pre-trained BERT model (DistilBERT [35]) outperformed non-DL ML models (incl. logistic regression, naïve bayes and SVM).
- (2) For NER⁶, a pre-trained BERT model outperformed a trained-from-scratch Bi-LSTM-CRF model, although there were indications that the latter may work better on small datasets [36].

The authors also point out that the non-BERT models required extra pre-processing, respectively Bag of Words and TF-IDF, and word embeddings (GloVE [27]).

López et al. [30] compare ML-only, rule-only, and integrated ML + rule-based NER on roles, activities, and relations. The authors found the following:

- (1) Using only ML for NER yielded higher recall but lower precision;
- (2) Using only rules for NER yielded higher precision but lower recall;
- (3) Using ML for NER on roles and relations, and an integrated ML + rules for NER on activities, yielded the best F1 score.

Non-traditional, DL-based NLP pipelines.

Qian et al. [28] implemented the NLP aspect of process extraction using multi-grained text classification. In a course-grained phase, based on sentence encodings and features, sentences are classified (action vs. statement) and semantically tagged (e.g., sequential, concurrent). Next, based on the sentence-level encodings and features, together with word-level embeddings, a fine-grained phase semantically tags words and phrases with roles. The authors use a combination of Bi-LSTM, CNN, and multiple MLPs (Multi-Layer Perceptrons). Notably, the authors use LSTM as it was deemed suitable for the sequential nature of text. Han et al. [10] use RNN and Bi-LSTM to obtain sentence-level and process-level encodings; these are later passed to an Ordered Neurons LSTM (ON-LSTM) [37] for process model generation (Section 6.1).

In addition to a traditional NLP pipeline, Neuberger et al. [29] use an end-to-end DL model (Jerex [33]) to perform RE and NER.

Halioui et al. [25] combine ML with an ontology for semantic analysis. The authors map text elements to concepts from this ontology; precedence constraints from the ontology are then used to help reconstruct the process. The authors further consider the word’s context to perform WSD; i.e., POS tags and mapped concepts surrounding the word are used as context to train a ML (PAUM [38]) model. E.g., this allows mapping the term “MEGA” to different concepts based on surrounding context. No-

⁶In particular, extracting dependencies in terms of base, derived, and action tags.

Table 6.: Public NLP tools used for process extraction.

NLP Tools	Description	Papers
Ansj Tokeniser	Chinese word segmentation tool.	[39]
ELMo	A pre-trained model that generates embeddings for words based on their context across linguistic contexts.	[28]
FrameNet	A lexical database that annotates words in terms of their semantic frames and roles.	[40]
GloVe	An unsupervised learning algorithm for obtaining word vector representations based on word-word co-occurrence.	[28; 10]
MedPos Tagging	Assigns POS tags to words targeting <i>clinical text</i> .	[25]
NeuralCoref4	A spaCy extension that annotates and resolves coreference clusters using a neural network.	[31; 32; 9]
NLTK	Library with tools for symbolic and statistical NLP.	[24; 31]
OpenNLP Chunker	Identifies syntactic constituents (e.g., noun, verb phrases).	[25]
spaCy	Python NLP library.	[32; 31; 41; 42]
Stanford CoreNLP	Java suite of NLP tools.	[31; 23]
Stanford Lemmatiser	Converts words to their root form.	[43; 25]
Stanford Parser	Parses sentences to derive their grammatical structure.	[44; 40; 42]
Stanford POS Tagging	Assigns POS tags to words in an input text.	[32; 43; 45]
Stanford Tregex	Queries syntactic dependency trees based on patterns.	[46]
Stanza	Python NLP library developed by Stanford for multiple human languages, similar in functionality to CoreNLP.	[47]
Word2Vec	A technique that employs neural networks to learn word embeddings from a text corpus.	[28]
WordNet	A lexical database of English, grouping words into sets of synonyms and providing brief definitions.	[42; 40; 44; 43; 48; 45]

tably, PAUM was designed for imbalanced data and has successfully been applied to semantic annotation.

5.3. NLP Tools

Table 6 lists the NLP tools found in the reviewed papers. Many tools used in the reviewed papers have overlapping capabilities such as lemmatization, dependency parsing, POS tagging, and coreference resolution. Tools like Stanford POS Tagging, Stanford Parser, and spaCy can handle multiple tasks, from dependency tree generation to semantic role labeling. Specialised tools such as the NeuralCoref4 package or the Ansj Tokeniser target specific tasks, i.e., coreference resolution and word segmentation. Domain-specific tools such as MedPos focuses on POS tagging of clinical text.

6. Process Model Generation

This section discusses the generation of process models from NLP output: including the computational paradigm used (Section 6.1), intermediary representations, if any (Section 6.2), and targeted type of process model (Section 6.3).

6.1. Computing Paradigm

We observe that, despite the increasing popularity of ML/DL methods in NLP pipelines (Section 5), the majority of studies rely on knowledge-based methods for process generation—with the caveat that ML/DL enabled NLP pipelines may directly output a process model (see “Machine Learning / Deep Learning” below). We discuss these two approaches below.

Knowledge-based Methods.

Here, we consider the use of any knowledge-based artefact, such as concrete rules (e.g., Prolog), templates, or custom algorithms, to generate process elements based on NLP output. Below, we focus on the type of artefact and the NLP output involved.

- *Concrete Rules.* Nasiri et al. [23] provide a set of Prolog rules, which refer to Part-of-Speech (POS) tags and typed dependencies from NLP output, in order to generate activity diagrams. Halioui et al. [25] use JAPE (Java Annotation Patterns Engine) [49] pattern-based rules to, among other things, avoid incorrect entity recognitions. For instance, in case a noun, verb or adjective is in lowercase, a rule will ensure it is not matched to a software program⁷.
- *Custom algorithms.* Sonbol et al. [43] present a custom algorithm to generate an intermediary Text Graph (i.e., process model) based on semantic concepts, their tags and relations from the NLP step. Chen et al. [39] collect a “verb flow” from a given start word using a recursive algorithm, based on POS tags, grammar relations, and semantic tags (topics). E.g., for topic “House Cleaning”, this process would yield Clean→Dust Collection→Scrub→Wash→Wipe out. Van der Aa et al. [45] define an algorithm based on elements from NLP output that indicate modality (e.g., “must”), negation, connectors (e.g., “or”) and temporal relations (e.g., “before”); and consider verb-based (e.g., “create”) and noun-based (e.g., “creation”) activities. E.g., a sequential relation $rel(A,B)$, where activity A is mandatory but B is not, results in a *precedence* constraint. López et al. [30] apply an algorithm to integrate the results from RE and NER; when the binary RE classifier detects a sentence with a relation, the NER model extracts 2 events, with a relation as detected by the multi-class RE classifier.
- *Templates / patterns.* Etikala et al. [31] apply predefined patterns, such as “passive A<=B” (A follows from B) and relevant verbs (e.g., “determine”) to generate dependency tuples in the form of ⟨action, base, derived⟩ concepts (e.g., determine, height, Body Mass Index), based on typed dependencies and semantic concepts from the text. Later on, Etikala [32] additionally extracts decision logic by applying patterns with markers such as “if”, “then”, etc., which is ultimately converted into decision tables. Ferreira et al. [41] define high-level templates to identify BPMN elements (e.g., XOR gateway) based on POS tags and semantic tags (e.g., condition, conjunction). Quishpi et al. [46] use a query language (Tregex) to capture decision requirements (e.g., determine obesity from height)

⁷This could be considered part of the NLP pipeline; but, as they act on NLP output, we mention it here.

from a syntactic dependency tree with typed dependencies.

- *High-level Rules.* These include “if-then” rules narratively described in tables or text. (For instance, these could be implemented using concrete rules or a custom algorithm; see above.) Sholiq et al. [47] describe rules that, based on typed dependencies, extract “fact types” (FT); e.g., a binary FT would be “officer *accepts* document”. These fact types are mapped to BPMN elements. Goncalves et al. [24] describe rules to find activities, actors and other elements, based on POS tags and typed dependencies.
- *Custom.* Others narratively describe a series of steps in the paper text. Honkisz et al. [42] describe steps for identifying participants, subject-verb-object and BPMN gateway keywords, based on typed dependencies and semantic concepts from the text. Azevedo et al. [50] only shortly summarise their process generation approach; based on NLP output (e.g., POS tags), the authors generate process elements based on the sequential ordering of sentences. Friedrich et al. [44] describe an extensive process that includes extracting actors, actions and objects from text; merging actions, based on anaphora resolution; and generating process elements (e.g., exclusive, parallel) based on textual markers (e.g., “or”, “in parallel”). López et al. [48] suggest process elements from natural text for confirmation by the user, based on synonyms from Wordnet and a domain-specific R-A-R (role-activity-relation) dataset that is updated based on user interactions. So-called *highlights* keep the link between text and associated process elements.

Machine Learning / Deep Learning.

To the best of our knowledge, ML/DL is only used once to implement process generation in the reviewed papers [10]. Some authors cast process extraction as NLP problems (e.g., text classification, NER), which they then solve using ML/DL methods [28; 9]. In these cases, the NLP output directly yields the targeted output; however, extra work may still be needed to generate a complete process model from this output.

Qian et al. [28] cast process extraction as a multi-grained text classification problem; output includes semantic annotations of sentences and words/phrases. Goossens et al. [9] use the results from NER, i.e., a set of dependency tags (action, base, derived concepts) to generate dependency tuples; a tuple consists of dependency tags from 1 sentence. (In contrast, Etikala et al. [31] apply a separate knowledge-based step for extracting dependency tuples from NLP output.) Neuberger et al. [29] focus on RE and NER for the purpose of process extraction; the latter is considered out of scope.

In contrast to these works, Han et al. [10] pass the NLP output, i.e., sentence-level and process-level encodings, to a novel ML model called Ordered Neurons LSTM (ON-LSTM) [37], which separately implements process model generation. In a structure-retrieve step, the latent hierarchical process structure is inferred from the ON-LSTM model. A BPMN diagram can then be extracted using Jinja2⁸.

6.2. Intermediary Representation

Here, we include any method where a notation-agnostic, intermediary representation is separately generated. This representation can then be mapped to multiple different model notations. Note that we do not consider intermediate NLP output (e.g., dependency trees, semantic tags) here, nor internal model representations (e.g., latent representations in DL models).

⁸<https://jinja.palletsprojects.com/en/3.1.x/>

Table 7.: Target Process Models

Model Type	Notation	Papers
Control flow: Imperative	BPMN	[43; 47; 50; 41; 44; 40; 34; 10; 28; 42; 24]
	UML activity diagram [51]	[23]
	Custom / unspecified	[39; 29; 25]
Control flow: Declarative	Declare [3]	[45]
	DCR (Dynamic Condition Response) [52]	[48; 30]
Decisional	DMN [53]	[32; 9; 31; 46]

- *Graph-based Model.* Sonbol et al. [43] construct a Text Graph with sentences as vertices, which are connected based on their ordering in the text (i.e., 2 consecutive sentences as 2 connected vertices). They also construct a concept map that includes interrelated actors and resources. This concept map is leveraged to further connect sentence vertices, i.e., based on their shared concepts in the map. The Text Graph is converted into a BPMN model in an 8-step procedure. Friedrich et al. [44] generate a graph-based World Model that capture actors, resources, actions and flows as relations between them. Subsequently, a 9-step process generates a BPMN model from the World Model. We note that Azevedo et al. [50] mention an intermediary *Tree-based* Process Model, i.e., with sentences as tree nodes, but do not provide details.
- *Table.* Honkisz et al. [42] generate a table with columns including ordering, activity, condition and actor. The authors subsequently generate a BPMN diagram but do not detail the individual steps.
- *Dependency tuples.* Goossens et al. [9] and Etikala et al. [31; 32] generate intermediary dependency tuples that are ultimately converted into a Dependency Requirements Diagram (DRD).

6.3. Target Process Model

Regarding the type of process model that is targeted, we find a contrast between *control flow models* (BPMN, UML activity diagrams) and *decisional models* (Decision Model and Notation, DMN: decision requirements, tables). Some studies focus only on particular elements; e.g., Chen et al. [39] capture activity ordering (“verb-flow”).

In the control flow category, we further distinguish *imperative models*, i.e., which fully specify the activity occurrence and sequential ordering; and *declarative models*, i.e., which include high-level constraints on the occurrence and ordering of certain activities (e.g., precedence, response, succession [45]).

Table 7 shows a breakdown of the reviewed papers based on the type of target model and the particular notation targeted, if any. Overall, imperative control flow models are the most covered (15), followed by decisional models (4) and then declarative control flow models (3).

7. Process Extraction Evaluation

In this section, we review the datasets and methods commonly utilised to evaluate automated process extraction.

7.1. Evaluation Datasets

Table 8 shows the datasets used by the reviewed papers. Regarding *Type of Data*, *synthetic* means that the included sentences were manually written by the authors; *real* means that sentences were collected from real-world sources, such as academic papers, laws and regulations, interviews and Internet searches. We note that the *Dataset Size* column lists the level of detail provided by the authors; e.g., if the paper mentions X “example texts”, we repeat that in this column (as this does not necessarily constitute X separate cases). Below, we make several observations on these datasets; we also discuss multiple real datasets and how they were collected.

Main Observations on Evaluation Datasets.

- (1) There is a lack of consistently used evaluation datasets. Only the Friedrich [54] and PET [34] datasets are re-used, i.e., in 2 other studies (also, per dataset, only 1 of these 2 studies was not authored by its creators).
- (2) Many online datasets are no longer available. Some articles provide links to their evaluation dataset; however, Table 8 shows that 4/9 links no longer work⁹.
- (3) There is a paucity of gold-standard datasets. There is currently 1 reference dataset for control flow process extraction, i.e., the PET dataset (Section 7.1).
- (4) Most evaluation datasets are small in size. While a variety of datasets are being used, they are relatively small in size (i.e., tens or hundreds of process descriptions).

Real-world Datasets and Their Collection.

Friedrich et al. [54] introduced a dataset consisting of 47 text-model pairs, each pair including a textual description and a corresponding BPMN model created by a human modeler. It spans processes from academic sources, industry (e.g., BPM tool vendors), textbooks and the public sector. Hence, it can be considered a mix of synthetic (e.g., textbooks) and real (e.g., industry). Bellan et al. [34] introduced the PET dataset, which was constructed based on the above “Friedrich dataset”. It improves upon the granularity of the latter by annotating parts of the text with corresponding process elements. Also, since the textual dataset was annotated by 3 independent annotators, it can be considered more of a gold standard. The authors further provide baseline results for typical process extraction-related NLP tasks using the PET dataset, i.e., NER and RE using ML (CRF) and a custom rule-based approach.

Han et al. [10] collected 210 process descriptions from BluePrism RPA (Robotic Process Automation), and 50 process descriptions from the Internet such as whitepapers, user guides, and SAP support documents. Ferreira et al. [41] evaluated their system using 56 natural language text from the US immigrant visa process, Federal Network Agency of Germany, and others. Chen et al. [39] used a search engine to retrieve process descriptions on 5 topics (e.g., house cleaning).

To evaluate multi-grained classification for process extraction, Qian et al. [28] manually collected a dataset consisting of (a) Cooking Recipes (COR) with 200+ recipes from recipe.com; (b) Maintenance Manuals (MAM) with 160+ device maintenance descriptions from ifixit.com. Halioui et al. [25] conducted a search on the PubMed Central database for papers with “phylogenetic” and filtered on articles with suitable bio-informatics workflows (i.e., keywords “Data”, “InferenceProgram”).

Goncalves et al. [24] manually collected user stories (events, characters, and other elements) from people involved with course enrolment (26 story events) and managing

⁹We suggest readers to reach out to the paper authors to gain access to these datasets.

Table 8.: Evaluation datasets used in process extraction.

Type	Target Model	Dataset Size	Public	Domain	Notes & Papers
Synth./Real	Control (Imper.)	47 Processes, 432 Sentences	Public ¹	Academic, industry, textbook, public sector	‘Friedrich’ [54] [54; 44; 43]
Synth./Real	Control (Imper.)	45 Processes, 417 Sentences	Public ²	Cfr. Friedrich et al. [54]	‘PET’ [34] [34; 55; 29]
Real	Control (Imper.)	Ca. 260 Processes	Private	Business processes (RPA, BluePrism, SAP)	[10]
Synth.	Control (Imper.)	10 Cases	Public ³	Finance and HR	[47]
Real	Control (Imper.)	56 Texts	Private	US visa process, Federal Network Agency of Germany, others	[41]
Real	Control (Imper.)	-	Public ⁴	Internet searches; 5 topics (e.g., house cleaning)	[39]
Synth.	Control (Imper.)	-	Public ⁵	Computer repair	[42]
Real	Control (Imper.)	Sentences: 2636(COR), 2172(MAM)	Public ⁶	Cooking, maintenance	[28]
Real	Control (Imper.)	320 Texts	Public ⁷	Keywords: Phylogenetic, Data, InferenceProgram (PubMed PCM)	[25]
Synth.	Control (Imper.)	2 Cases (33, 17 Activities)	Public ⁸	Finance and Web	[23]
Real	Control (Imper.)	2 Cases (26, 15 Story Events)	Private	Course enrolment, manual process elicitation	[24]
Synth./Real	Control (Decl.)	103 Constraint Descriptions	Public ⁹	Industrial, academic (general, declarative)	[45]
Synth./Real	Control (Decl.)	107 Process Descriptions	Public ¹⁰	BPM academic initiative, students, social services	[30]
Synth.	Decis.	2 Cases	Public ¹¹	Finance and Health	[31]
Synth.	Decis.	20 Texts	Private	-	[32]
Real	Decis.	Sentences: 577 (train), 232 (test); Text: 6 (38-101 words)	Public ¹²	Sentences: papers, laws & regulations, Etikala et al. [31], others; Text: DM community, Internet	[9]
Synth.	Decis.	12 Texts	Public ¹³	Finance, healthcare	[46]

We could not find a description of the dataset in Azevedo et al. [50] and López et al. [48].

¹ <http://github.com/FabianFriedrich/Text2Process>

² <http://huggingface.co/datasets/patriziobellan/PET>

³ Listed in paper

⁴ <http://yun.baidu.com/share/link?shareid=3648609423&uk=1862989032> (not working at submission)

⁵ Listed in paper

⁶ <https://github.com/qianc62/MGTC>

⁷ <https://data.world/growler/datasets>

⁸ Listed in paper

⁹ <http://github.com/hanvandraa/declarextraction>

¹⁰ <http://github.com/RasmusIven/DCR-Textograph> (not working at submission)

¹¹ Listed in paper

¹² <http://github.com/Goossens496/Extracting-DMN-models-from-text> (not working at submission)

¹³ <http://github.com/ProjectTex2Dec/Text2Dec/tree/master/data/collected> (not working at submission)

process elicitation at a company (15 story events).

For declarative process extraction, Van der Aa [45] gathered constraint descriptions from industrial and academic sources, with both declarative and general (i.e., not necessarily declarative) processes. López et al. [30] gathered process descriptions from the BPM academic initiative, student interviews, and excerpts related to social services.

For decisional model extraction, Goossens et al. [9] collected a real-world dataset from the DM community, laws and regulations, Web searches, and their prior work [31]. Quishpi et al. [46] collected 12 text-model pairs (i.e., text paired with the corresponding DMN model created by a human) based on materials from other authors.

7.2. Evaluation Methods

We observed 2 orthogonal sets of evaluation methods for process extraction:

Component-based vs. Holistic Evaluation

Component-based evaluation evaluates the output of individual modules in the NLP and process model generation steps. This has the advantage of locating specific pain points, and avoids errors from propagating to subsequent modules [29]. It is typically combined with a holistic evaluation to evaluate overall pipeline performance.

Holistic evaluation evaluates the holistic extracted process model; e.g., by comparing its individual elements with those of a reference model (systematic), or evaluating their understandability and usefulness (expert-based). While this approach thus evaluates the overall pipeline performance, it may gloss over the contribution of (and issues with) individual modules.

Systematic vs. Expert-based Evaluation

Systematic comparison of generated and reference output. The generated output, albeit individual module output or the final process model, is systematically compared to a reference model. This can be an *element-by-element* comparison or based on an *aggregate* measure, such as processes' behavioral profiles [56]. Here, *graph-based* evaluations represent process models as graphs, and map the nodes between reference and extracted process graphs. A potential drawback is that the reference model may be only one of many valid process models for the same text [47].

Expert evaluation of generated output. Experts manually evaluate the generated output, providing their opinions on metrics such as understandability, usefulness, their agreement with the output, or the model's equivalency with the text. Here, the evaluation can again be *element-by-element* (i.e., evaluating individual elements), or focusing on properties of the output as a whole (*aggregate*). This avoids the aforementioned issue of multiple valid reference models, as it does not involve comparisons with a reference. However, it is also resource-intensive and relies on subjective judgments¹⁰.

Table 9 shows a breakdown of the reviewed papers along these categories. Below, we separately discuss systematic and expert-based evaluations from the papers.

Systematic Evaluations

The majority of evaluation methods are *holistic* and *systematic (element-by-element)*: individual elements from the extracted and reference process model are systematically compared. Clearly, the evaluated elements will depend on the target process model (control flow vs. decisional). Instead of *element-by-element*, Qian et al. [28] utilise

¹⁰Reference models (*systematic*) will typically be created by experts, and thus also include a degree of subjectivity. The differentiation between these types of evaluations is how the comparison is performed.

an *aggregate* measure: assessing the similarity of the extracted and reference process model based on their behavioral profiles [56].

Often in addition, several authors evaluate their approach in a *component-based* and *systematic (element-by-element)* way. Etikala [32] separately evaluates the decision logic and dependency tuple extractor components; individual output elements (e.g., dependency tuples) are compared to reference elements. Goossens et al. [9] evaluate their sentence classification and downstream decisional model extraction. To avoid errors early on in the pipeline from propagating, Neuberger et al. [29] evaluate modules' performance in isolation by also providing reference inputs. Qian et al. [28] evaluate the accuracy of the classifiers that are involved in NLP and process extraction.

In *graph-based* systematic evaluations, reference and extracted processes are represented as graphs. Then, based on a mapping of graph nodes between the reference and extracted process graph, authors calculate a similarity score. Friedrich et al. [44] and Sonbol et al. [43] rely on the Graph Edit Distance [57], which, based on the similarity of individual nodes, calculates the number of node and edge edits needed to get from one graph to another. Han et al. [10] similarly map identical nodes, in this case, based on their "step" descriptions and "parent" and "previous" nodes. A similarity score is calculated as the ratio of mapped nodes vs. the total number of model nodes.

Systematic evaluations typically evaluate correctness (false positives) and completeness (false negatives); hence, evaluation metrics tend to involve precision, recall, and F1-score. Sholiq et al. [47] calculate the Magnitude of Relative Error (MRE) as the ratio of differences between the extracted and reference model vs. the size of the reference model. Chen et al. [39], in line with their focus on activity orderings ("verb-flow"), calculate the Normalised Distance-based Performance Measure (NDPM) as a ratio based on incorrectly & correctly sorted and concurrent activities. Van der Aa [45] proposes a fine-grained "slot filling" approach, where the type of constraint and its arguments are "slots"; precision and recall are based on the number of matching slots between the extracted and reference constraints.

To systematically evaluate ML/DL approaches, the text dataset is typically split up into training and testing using different ratios (e.g., 2/3, 8/10) [9; 25], or an external (i.e., different from the training) dataset is used [30]. For nondeterministic ML/DL methods, average performances over a number of runs (typically 5 or 10) are presented [9]. Halioui et al. [25] perform a 10-fold cross validation to establish optimal parameters; Bellan et al. [55] and Neuberger et al. [29] perform 5-fold cross-validation and take the average performance.

Expert-based Evaluations

Sonbol et al. [43] asked experts to assess the understandability and usefulness of the extracted process model. Azevedo et al. [50] asked 48 participants to assess whether the knowledge represented by the extracted process model is equivalent to the textual description. Similarly, Ferreira et al. [41] asked experts whether they agreed with the extracted process models, using a 5-item Likert scale (from "Fully Agree" to "Disagree"). They proposed six process models to the experts; 2 modeled based on mapping rules, and 4 purposefully modeled incorrectly. All of the above represent *aggregate* evaluations, as they focus on the process model as a whole.

In *element-by-element* evaluations, authors ask experts their opinion on the individual elements. Goncalves et al. [24] evaluated the extracted processes using a human modeler who separated valid (useful) from noise (not useful) elements. In addition, Ferreira et al. [41] asked experts to manually identify process elements from text (3 sentences) and subsequently compare them with the extracted elements.

Table 9.: Evaluation methods used in process extraction.

Category 1	Category 2	Description	Papers
Component-based	Systematic	Components: Sentence classification	[9]
		Components: Entity, Relation Extraction	[55; 29]
		Components: Sentence classification, semantics recognition, semantic role labeling	[28]
Holistic	Systematic	Elements: decisional dependencies (dependency tuples and/or decision logic)	[9; 31; 32; 46]
		Elements: disambiguated words, extracted relations	[25]
		Elements: extracted entities, relations	[55; 29; 48]
		Elements: activity sorting, found activities	[39]
		Elements: UML activity/ BPMN / interm. repr. elements	[41; 23; 47; 42]
		Graph-based: Graph Edit Distance	[44; 43]
		Graph-based: Similarity score	[10]
		Aggregate measure: behavioral profiles [56]	[28]
		Elements: declarative constraints	[45; 30]
		Expert-based	Assess understandability and usefulness of extracted processes (aggregate)
	Assess knowledge equivalency between extracted, reference processes (aggregate)		[50]
	Gather opinions (fully-agree to disagree) on extracted processes (aggregate)		[41]
	Separate valid (useful) from noise (not useful) extracted process elements (element-by-element)		[24]
	Manually compare textual elements with extracted elements (element-by-element)		[41]

8. Discussion

We identified multiple challenges and trends related to automated process extraction. We further discuss preliminary work that has been done using LLMs in this field, as well as promising developments in LLM architectures.

8.1. Complexity of Natural Language

Clearly, the problem of complex natural language is not limited to process extraction, but rather any domain that relies on NLP. To curb this complexity, the text input can be restricted. We discuss works that rely on restricted text input in Section 5.1. For example, regarding user stories, Nasiri et al. [23] mention that a user story often uses the following format type: “As \langle Role \rangle , I want to \langle Action \rangle , so that \langle Benefit \rangle ”. To avoid ambiguity and lack of clarity of free text, Goncalves et al. [24] consider stories from input text that is structured as a set of scenarios. Some papers mention simplifying assumptions in the text: e.g., Text2Dec [31] assumes the description is sequential, contains no irrelevant information or redundancies, and contains only 1 main decision.

Quishpi et al. [46] make the point that language patterns expressing *decisions* is less variable than those for *control flow* descriptions. Hence, decisional models may be an easier target for process extraction. Van der Aa [45] pointed out additional challenges regarding *declarative* process extraction. For instance, subtle differences in text (e.g., modality; “can” instead of “must”) will lead to semantically different constraints (e.g., precedence vs. response); in imperative models, these would typically lead to the same sequential relation. Text may also indicate the negation of a constraint (e.g., “an invoice cannot be paid”), or logical relations between constraints (e.g., “can be approved or rejected”), which must similarly be reflected in declarative constraints.

8.2. Use of Deep Learning for NLP

We have observed a paradigm shift over the last 5 years that coincides with the emergence of DL in NLP; authors are increasingly using DL models such as Transformers (e.g., BERT [9]) and LSTM [10] in process extraction. Until now, 4 works have applied Large Language Models (LLM), as discussed in Section 8.4.

Quishpi et al. [46] note that classical, rule-based AI approaches can still have advantages over DL systems: (1) training DL systems from scratch requires huge amounts of annotated training data, and the cost of producing such datasets may be higher than the cost of encoding expert knowledge into rules; and (2) DL systems, in contrast to rule-based systems, operate as black boxes, and it is difficult to tailor their behaviour to improve output in case of wrong answers.

Indeed, regarding training data (1), Neuberger et al. [29] hypothesised that the reason the end-to-end DL model Jerex, which was trained from scratch, was outperformed, was that the training data (PET dataset) was not yet extensive enough. Regarding the tailoring of behavior (2), the same authors point out that rule-based systems are also hard to adapt to even minor changes in the data.

Goossens et al. [9] found that a *pre-trained* BERT model, fine-tuned on process extraction (i.e., not trained from scratch), outperforms ML methods and even a DL model trained from scratch. Hence, regarding training data (1), pre-trained models may obviate the need for large-scale training datasets. However, fine-tuning still requires training on a labeled dataset for adjusting parameters. Moreover, pre-trained

models may have difficulty with highly domain-specific terminology such as found in scientific texts [25].

Finally, authors have pointed out that certain DL models may be intrinsically more suited towards automated process extraction:

- LSTMs are adept at handling sequences, making them more suitable for modeling processes. Goossens et al. [9] evaluated a Bi-LSTM-CRF model, pointing out that their particular structure affords them better insights into sentences. Similarly, several authors leveraged Bi-LSTM to obtain sentence encodings [28; 10].
- Han et al. [10] leverage Ordered Neurons LSTM (ON-LSTM), pointing out that ordered neurons add a structure-oriented inductive bias that better supports hierarchical process models.
- Neuberger et al. [29] pointed out that Jerex [33], an end-to-end DL method, can identify the textual location of an entity; this allows obtaining its surrounding text for enriching BPMN activity labels.

8.3. *Lack of Consistently Used, Gold-standard, Accessible and Large-Scale Evaluation Datasets*

We previously pointed out a number of observations regarding evaluation datasets (Section 7.1, Main Observations). We discuss these further below:

(1) *There is a lack of consistently used evaluation datasets.* Almost every study reviewed in this paper utilised a different dataset in their evaluation. Only the Friedrich [54] and PET [34] datasets are re-used, but only in 2 other studies each time; also, per dataset, only 1 of these 2 studies was not authored by its creators. This greatly complicates an objective, third-party comparative evaluation. There are multiple potential reasons for this: the fact that many evaluation datasets can no longer be found online (issue (2)), or that there is lack of gold standard datasets (issue (3)).

(2) *Many online datasets are no longer available.* While multiple articles provide links to their evaluation dataset, 4/9 of these links no longer work¹¹, which prevents their re-use by third parties. We thus recommend making datasets available using immutable records, such as provided by Zenodo¹², as opposed to university resources or even code repositories such as github¹³. This could be made a requirement by journals and conferences; multiple journals already have a data availability requirement.

(3) *There is a paucity of gold-standard datasets.* There is currently 1 gold-standard dataset for control flow process extraction, namely the PET dataset (Section 7.1). However, this dataset has only been available since 2022; moreover, it seems geared towards *imperative* process models, as the manual annotations follow imperative model notations. For instance, annotated process elements include (AND/OR) gateways and associated conditions, branches, and flow (e.g., sequential) relations. It is thus an open question whether the PET dataset is amenable to declarative process extraction. Moreover, to the best of our knowledge, there is no gold standard dataset for extracting decisional models.

(4) *Most evaluation datasets are small in size, i.e., in the order of tens or hundreds of individual process descriptions.* There is a need for large-scale datasets, not just for

¹¹Of course, this is a known problem beyond the field of automated process extraction.

¹²<https://zenodo.org/>

¹³Unfortunately, the use of such repositories is not a solution, as 3/6 of those links no longer work either.

robust evaluations, but also for training/fine-tuning and testing DL models [29].

To resolve issues (3) and (4), Neuberger et al. [58] recently proposed the application of data augmentation techniques for automatically synthesizing natural text for process extraction. The authors found that simple data augmentation techniques already improved the accuracy of ML-based NER and RE models trained on the augmented data. Example techniques include the increasing of linguistic variability, introducing variations in span length, and changing the directionality of actor-performer-recipient relations. Nevertheless, data augmentation techniques still rely on an initial domain dataset to synthesize additional samples. To better evaluate manual data curation efforts, some conferences currently organize tracks that accept resources (datasets, benchmarks, software, ...) and their descriptions as publications. Finally, we note that LLMs (see below) may end up removing the need of large training/fine-tuning datasets, due to their in-context learning ability. Even then, however, there may be a need to compare their performance with prior ML/DL methods.

8.4. *Promise of Large Language Models*

Large Language Models (LLM) can perform a variety of natural language tasks without requiring a specialized training or fine-tuning step. Bellan et al. [12] note that LLMs “push pre-training to the extreme”: after being pre-trained on huge amounts of online textual data, they are able to carry out a variety of tasks by performing only small-scale “in-context learning”. The latter involves a **prompt** with task instructions, contextual knowledge, possibly accompanied by few-shot examples (i.e., input and corresponding output), and the natural language text. Hence, LLMs have a huge potential for automated process extraction, as scalable datasets for training and fine-tuning are currently lacking (Section 8.3).

At the same time, the fact that LLM re trained purely on Internet data may lead to inadvertently learning and reproducing biases therein [59]. In process extraction, for instance, this may lead to processes that perpetuate discrimination on gender, race, religion or age [60]. We note that this problem expands beyond LLM and apply to any pre-trained language models, such as BERT [61]; techniques have been studied to identify and mitigate such biases [60].

Below, we shortly summarize initial work on the use of LLM for process extraction (Section 8.4.1). Afterwards, we discuss promising developments such as Retrieval-Augmented Generation (RAG) and Multi-Agent Systems (Section 8.4.2).

8.4.1. *Current LLM-based Process Extraction*

At the time of writing (August 2024), to the best of our knowledge, 4 peer-reviewed papers [12; 13; 14; 15] present experimental results LLM-enabled process extraction. This relative paucity of work is not unexpected, considering their novelty; the “game-changing” ChatGPT was only released at the end of 2022. We discuss these works below along a salient set of dimensions.

Prompts. All reviewed approaches construct a prompt with the textual process description, task instructions, typically a few few-shot (positive and negative) examples, and context knowledge. Given the prompt, the LLM then generates a process model based on the natural text. Task instructions can range from relatively simple instructions [12; 13; 15] (e.g., list all activities in the text [12]) to detailed procedures (e.g., how to identify places and transitions for Petri nets [14]). Context knowledge

tends to depend on what output is targeted; since Bellan et al. [12] aim to extract participants, activities, and relations, they define these elements in the prompt, as well as the process domain. Grohs et al. [13] target both declarative and imperative control flow models; to cope with the LLM output limit, they define concise output formats that can later be mapped to e.g., Declare and BPMN, respectively. Forell et al [14] provide a comprehensive definition of a Petri net and an output format (JSON). Kourani et al. [15] target Partially Ordered Workflow Language (POWL) as output; they define POWL and its semantics, and a set of Python functions to generate the POWL (see below).

LLM Methodology. Most authors separate the task of generating a final output format. As mentioned, Grohs et al. [13] output a custom format that is later mapped to Declare or BPMN. Forell et al. [14] include a second phase, where the LLM is asked to output a concrete format for adaptability to different modeling tools. Kourani et al. [15] ask the LLM to output Python code, which is then executed for the safe generation of POWL models.

Bellan et al. [12] apply an incremental approach; an initial prompt asks the LLM to identify activities from the text; subsequent prompts include these activities and ask for participants in, and directly-follows relations between, the activities. Kourani et al. [15] have users engage with the LLM in an iterative error handling loop; refined LLM prompts detail the error and ask the LLM to fix it.

Suitability of LLM for process extraction. Bellan et al. [12] found that it is feasible to extract activities and their participants, but encountered challenges with directly-follows relations (avg. precision of 0.28). Grohs et al. [13] compare the performance of their approach with the pre-LLM work by Van der Aa et al. [45]; they found that GPT4 yields equal or higher F1 scores. Kourani et al. [15] found that, regarding choice of LLM, GPT-4 performed better than Gemini. While the evaluation showed the effectiveness of LLM, the authors also note that manual effort is needed during the error handling loop [15]. Forell et al. [14] performed a small scale evaluation that illustrated the ability of LLM to generate correct Petri nets; however, issues were reported with AND/OR split and joins. They also found that the non-deterministic nature of LLM caused inconsistent output across multiple runs, which can be an issue for robust process modeling.

A number of vision papers have illustrated the general promise of LLM in BPM [62; 63; 11]. For instance, Busch et al. [62] pointed out studies demonstrating that the performance of in-context learning can be comparable to, or even better than, fine-tuned models. The authors discuss a research agenda that includes the appropriate representation of processes in prompts; e.g., this is studied in detail by Berti et al. [64] for querying processes using natural language and LLM.

8.4.2. *Promising Developments in the Field of LLMs*

This section describes two recent innovations in the LLM field, namely Retrieval-Augmented Generation (RAG) and Multi-Agent systems (LLM-MA), and their potential to improve process extraction performance.

Retrieval-Augmented Generation (RAG)

In a RAG setup, given the LLM prompt, an extra step first queries an external source (typically, a vector database) for data relevant to the prompt. This data is then provided, together with the prompt, to the LLM [65]. By additionally providing the LLM with external, contextually relevant data, RAG aims to improve the accuracy and

up-to-dateness of LLM responses.

Relevant to process extraction, RAG has been studied in event extraction [66] and code generation from natural text [67; 68; 69]. Typically, RAG is applied to find external event/code samples with labels relevant to the given text; subsequently, these samples are provided as positive examples to the generative model.

Event extraction. Event extraction identifies sequences of events in text, including their “arguments” (related individuals, organizations, and locations) [66]. The authors note that the structure of event descriptions is often quite similar in terms of syntax and semantics. Hence, a RAG setup is proposed that retrieves relevant event samples, i.e., with labels similar to the given event description, and provide them to the generative model. Experimental results show that the proposed approach outperforms baseline approaches [66].

Code generation. Process extraction can be considered as a code generation problem, i.e., where “process code” is generated from input text¹⁴. Similar to event extraction, multiple works on code generation [67; 68; 69] use RAG to retrieve code samples with labels similar to the user prompt, and then provide them as positive examples. This follows human developers’ tendency to base themselves on code snippets that are relevant, but perhaps not identical, to their problem [68; 69]. In this vein, SkCoder [68] retrieves code snippets related to a textual description, and subsequently extracts a “sketch” (i.e., code skeleton) that retains only the code relevant to the given description. To cope with the maximum input length of LLM, AceCoder [69] applies a selector that filters out redundant code samples. Synchromesh [67] applies Target Similarity Tuning to select semantically relevant samples; Constrained Semantic Decoding is then used to generate only valid programs. Experiments show that the RAG setup can improve LLM accuracy and code validity [67], and can outperform baseline LLM and non-RAG approaches [68; 69].

We propose that similar setups can be used to improve process extraction. Such a setup could leverage a repository of process descriptions, i.e., structured processes labeled with a textual description, to provide relevant samples to the generative model. Moreover, an extra vector database can include word embeddings on organizational terminology; e.g., generated by a pre-trained BERT [17] model fine-tuned on such terminology for activities, roles and objects. Given an input text, an intermediate step can first extract individual terms and search the vector database for similar organizational terms; in the retrieval step, this terminology is then used to retrieve relevant samples from the process repository.

Multi-Agent Systems (LLM-MA)

In the code generation field, LLM-based Multi-Agent (LLM-MA) frameworks involve multiple LLM agents, each assigned a different role (e.g., analyst, coder, reviewer, tester). These agents then collaborate throughout the software generation lifecycle to generate the code output [70]. This setup draws inspiration from human software development, where complexity is managed by delegating tasks to team members, who then collaborate on producing the software. In the same vein, LLM-MA aim to improve the ability of LLM to manage the complexity of coding tasks.

Dong et al. [71] present an LLM-MA framework where LLM agents respectively operate as analyst, coder, and tester; agents are instructed on how to collaborate and interact to generate code. In a sequence of stages, agents pass on their outputs to other agents (e.g., analyst to coder; tester back to coder). Qian et al. [72] introduce ChatDev, an LLM-MA framework where agents assume e.g., programmer, reviewer, and tester

¹⁴E.g., types of process code includes JSON [14], custom formats [13], and Python [15], among others.

roles. The code generation process is broken into a sequence of phases and subtasks. To solve a subtask, in a multi-turn dialogue, an instructor agent (e.g., reviewer) instructs an assistant agent (e.g., programmer), who responds with a potential solution (e.g., fixing an endless loop). The solution from a subtask is passed on to the next subtask. A “de-hallucination” mechanism allows an assistant to seek clarification from the instructor before providing a solution. In both works, the authors found that their LLM-MA setup led to improvements in code generation performance.

We propose that a similar LLM-MA setup can be applied to process extraction. Such an LLM-MA would feature specialized agents, each dedicated to specific aspects of process extraction (e.g., identifying actors, activities, and objects, and relations between them). In a continuous feedback loop, a “critic” agent can iteratively evaluate their output, and, if needed, ask for refinements. An “integration” agent could then synthesize the outputs into a format such as e.g., BPMN, Petri nets, or Declare.

9. Threats to Validity and Limitations

We acknowledge the existence of the following limitations and threats to validity:

- *Search scope*: Our literature search’s effectiveness is invariably based on the chosen keywords (Table 1), which were selected to target both NLP methods and process extraction.
- *Non-peer-reviewed and non-English articles*: Our reliance on peer-reviewed sources written in English may exclude insights from grey literature, such as technical reports, theses or corporate white papers, and non-English sources. For instance, this may overlook advances made outside of academic circles.
- *Coding of papers*: both authors (WV, SM) read and coded (i.e., categorised) the papers. Any inconsistencies were resolved through internal discussions.
- *Rapidly evolving field*: the field of NLP is progressing rapidly, especially regarding LLM-related work. Recent developments may be reported in non-peer-reviewed channels (e.g., preprints from repositories like arXiv). Of course, publications after the preparation of this manuscript (August 2024) will also not be included.

10. Conclusion

We presented a systematic literature review, conducted using the PRISMA guidelines, on the field of NLP-enabled process extraction. We cover publications up to 2023 that include rule-based, ML and DL methods, and with targets including control flow (imperative and declarative) and decisional models. Review themes include natural language analysis, process model generation, and evaluation.

Below, we summarise our answers to our research questions:

RQ1: Which NLP and process generation methods, including publicly available tools, have been studied and used for extracting process models?

We presented a detailed landscape of process extraction methods, their NLP and process model generation components, and public tools (Sections 5- 7). As shown in Figure 2, historically (2011-2022), most approaches have used rule-based methods for NLP (14) compared to ML (6); DL methods started appearing after the advent of BERT in 2018 (5). At the time of submission, we found initial experimental results (4) on the use of LLM (Section 8.4.1). Finally, we point out an imbalance in the type of

target process model, with most works targeting imperative control flow models (15), followed at a distance by decisional (4) and declarative control flow models (3).

RQ2: To what extent have ML/DL methods been studied in process extraction?

Figure 2 shows that, at this point (2023), the interest in ML/DL is on par with rule-based methods. Table 5 summarises the ML/DL models used for NLP tasks. We further discussed experimental results that show performance benefits of ML/DL over rule-based methods for NLP (Section 5.2). We found that DL-based Transformers (BERT), which revolutionised NLP in 2018, were successfully applied to improve process extraction. We outlined “non-traditional” DL pipelines for NLP, and the DL models involved (Section 5.2). In our discussion, we summarize the observed use of DL for process extraction (Section 8.2); including the suitability of certain DL models for downstream process extraction (e.g., LSTM) highlighted by study authors.

RQ3: Which evaluation methods and datasets have been utilised to evaluate process extraction?

We provide a detailed listing of datasets used in the reviewed papers (Section 7.1). We hereby observed a barrier towards objective evaluation, namely the lack of accessible, scalable and gold-standard datasets for different types of process models. We further presented a breakdown of orthogonal sets of evaluation methods, and classified the reviewed papers accordingly (Section 7.2).

Finally, we discussed the more recent LLM revolution as an opportunity to further improve process extraction (Section 8.4). We shortly reviewed initial work here, which is promising but also points out current challenges; and discussed promising developments (RAG, LLM-MA) that may be amenable to process extraction. This systematic review of the “pre-LLM” era thus comes at an opportune moment: for practitioners who want to apply process extraction, we present a detailed overview of state-of-the-art methods that can be used, including public NLP tools. For researchers aiming to contribute to process extraction, we present the following:

- A review that describes a categorization of, and comparisons between, rule-based and ML/DL methods, non-traditional DL pipelines, and the use of LLM for process extraction. New studies can be inspired by this work, re-use their models, and/or compare their experimental results.
- A comprehensive analysis of evaluation methods and datasets for process extraction, which can be directly re-used in new research.

11. Acknowledgements

We would like to thank Daniel Amyot for a thorough proofreading of the paper, and Alireza Houshidari for providing insights on novel LLM setups.

12. Declaration of Interest

The authors report there are no competing interests to declare.

13. Funding

This work was supported by the Telfer School of Management, University of Ottawa, under a School of Management Research Grant (SMRG).

Appendix A. Detailed categorization of papers along classification themes.

Regarding the NLP theme, we note that the vast majority of papers did not mention a restriction on text input (Section 5.1). We refer to Table 4 for typical NLP pipeline tasks; Table 5 for the use of ML/DL for these NLP tasks (NLP Computing Paradigm theme); and Table 6 for public NLP tools utilized in the papers (NLP Tools theme).

Below, we categorize the papers along the remaining Metadata (Table A1), Process Model Generation (Table A2), and Evaluation (Table A3) themes.

Table A1.: Categorization of papers along Metadata themes.

Title	Authors	Year
A Machine Translation Like Approach to Generate Business Process Model from Textual Description	Sonbol et al. [43]	2023
Extracting decision model components from natural language text for automated business decision modelling	Etikala [32]	2021
Extracting Decision Model and Notation models from text using deep learning techniques	Goossens et al. [9]	2023
Text2Dec: Extracting Decision Dependencies from Natural Language Text for Automated DMN Decision Modelling	Etikala et al. [31]	2020
Generating BPMN diagram from textual requirements	Sholiq et al. [47]	2022
Automatic generation of business process models from user stories	Nasiri et al. [23]	2023
BPMN model and text instructions automatic synchronization	Azevedo et al. [50]	2018
A semi-automatic approach to identify business process elements in natural language texts	Ferreira et al. [41]	2017
Let me tell you a story - on how to build process models	Goncalves et al. [24]	2011
Process model generation from natural language text	Friedrich et al. [44]	2011
PEWP: Process extraction based on word position in documents	Chen et al. [39]	2014
An Approach for Process Model Extraction by Multi-grained Text Classification	Qian et al. [28]	2020
Extracting Decision Models from Textual Descriptions of Processes	Quishpi et al. [46]	2021
Bioinformatic Workflow Extraction from Scientific Texts based on Word Sense Disambiguation	Halioui et al. [25]	2018
A-BPS: automatic business process discovery service using ordered neurons LSTM	Han et al. [10]	2020
Beyond Rule-Based Named Entity Recognition and Relation Extraction for Process Model Generation from Natural Language Text	Neuberger et al. [29]	2023
A concept for generating business process models from natural language description	Honkisz et al. [42]	2018

Extracting Declarative Process Models from Natural Language	Van der Aa et al. [45]	2019
Declarative Process Discovery: Linking Process and Textual Views	López et al. [30]	2021
Assisted declarative process creation from natural language descriptions	López et al. [48]	2019

Table A2.: Categorization of papers along Process Model Generation themes.

Citation	Process Model Generation		
	<i>Computing paradigm*</i>	<i>Interm. repr.</i>	<i>Target model</i>
Sonbol et al. [43]	Knowledge-based - Custom algorithms	Graph-based	Control flow - Imperative (BPMN)
Etikala[32]	Knowledge-based - Templates / patterns	Dependency tuples	Decisional (DMN)
Goossens et al. [9]	ML / DL - (see Table 5) ¹⁵	Dependency tuples	Decisional (DMN)
Etikala et al. [31]	Knowledge-based - Templates / patterns	Dependency tuples	Decisional (DMN)
Sholiq et al. [47]	Knowledge-based - High-level rules	/	Control flow - Imperative (BPMN)
Nasiri et al. [23]	Knowledge-based - Concrete rules	/	Control flow - Imperative (UML)
Azevedo et al. [50]	Knowledge-based - Custom	Tree-based	Control flow - Imperative (BPMN)
Ferreira et al. [41]	Knowledge-based - Templates / patterns	/	Control flow - Imperative (BPMN)
Goncalves et al. [24]	Knowledge-based - High-level rules	/	Control flow - Imperative (BPMN)
Friedrich et al. [44]	Knowledge-based - Custom	Graph-based	Control flow - Imperative (BPMN)
Chen et al. [39]	Knowledge-based - Custom algorithms	/	Control flow - Imperative (Custom)
Qian et al. [28]	ML / DL - (see Table 5) ¹⁶	/	Control flow - Imperative (BPMN)
Quishpi et al. [46]	Knowledge-based - Templates / patterns	/	Decisional (DMN)
Halioui et al. [25]	Knowledge-based - Concrete rules	/	Control flow - Imperative (Custom)
Han et al. [10]	ML / DL - (see Table 5) ¹⁷	/	Control flow - Imperative (BPMN)
Neuberger et al. [29]	ML / DL - (see Table 5) ¹⁸	/	Control flow - Imperative (Unspec.)
Honkisz et al. [42]	Knowledge-based - Custom	Table	Control flow - Imperative (BPMN)
Van der Aa et al. [45]	Knowledge-based - Custom algorithms	/	Control flow - Declarative (Declare)
López et al. [30]	Knowledge-based - Custom algorithms	/	Control flow - Declarative (DCR)
López et al. [48]	Knowledge-based - Custom	/	Control flow - Declarative (DCR)

¹⁵Process extraction is cast as an NLP NER problem.¹⁶Process extraction is cast as an NLP text classification problem.¹⁷A process structure is extracted from the trained ON-LSTM model.¹⁸The work implements NER, RE, and coref. resol. for process extraction (not elaborated in paper).

Table A3.: Categorization of papers along Evaluation themes.

Citation	Evaluation*		
	Type**	Metrics†	Dataset‡
Sonbol et al. [43]	Holistic - Systematic (Graph-based); Holistic - Expert-based	GED	Friedrich [54]
Etikala [32]	Component-based - Systematic	P, R	N/A
Goossens et al. [9]	Component-based - Systematic; Holistic - Systematic	P, R, F1	N/A
Etikala et al. [31]	Holistic - Systematic	N/A	Listed in paper
Sholiq et al. [47]	Holistic - Systematic	MRE	Listed in paper
Nasiri et al. [23]	Holistic - Systematic	P, R, Acc	Listed in paper
Azevedo et al. [50]	Holistic - Expert-based	answers on 0-100 & Likert scale	N/A
Ferreira et al. [41]	Holistic - Systematic; Holistic - Expert-based	P, R, F1, Acc; Likert scale	N/A
Goncalves et al. [24]	Holistic - Expert-based	# invalid & total # elements extracted	N/A
Friedrich et al. [44]	Holistic - Systematic (Graph-based)	GED	Friedrich [54]
Chen et al. [39]	Holistic - Systematic	NDPM, R	N/A
Qian et al. [28]	Component-based - Systematic; Holistic - Systematic (Aggregate)	Acc	N/A
Quishpi et al. [46]	Holistic - Systematic	P, R, F1	N/A
Halioui et al. [25]	Holistic - Systematic	P, R, F1	https://data.world/growler/datasets
Han et al. [10]	Holistic - Systematic (Graph-based)	Graph similarity score	N/A
Neuberger et al. [29]	Component-based - Systematic; Holistic - Systematic	P, R, F1	http://huggingface.co/datasets/patriziobellan/PET
Honkisz et al. [42]	Holistic - Systematic	N/A	Listed in paper
Van der Aa et al. [45]	Holistic - Systematic	P, R, F1	https://github.com/hanvandraa/declareextraction
López et al. [30]	Holistic - Systematic	P, R, F1	N/A

López et al. [48]	Holistic - Systematic	N/A	N/A
----------------------	-----------------------	-----	-----

* P=Precision, R=Recall, Acc=Accuracy, F1=F1-score.

** ‘Systematic’ evaluations are ‘element-by-element’ unless otherwise specified.

† ‘N/A’ when evaluation results are narratively described.

‡ ‘N/A’ when the dataset is private *or* the link was not working at submission.

References

- [1] M. Weske, Business process management—concepts, languages, architectures, verlag, Berlin (2007).
- [2] P. Bellan, M. Dragoni, C. Ghidini, A qualitative analysis of the state of the art in process extraction from text., DP@ AI* IA (2020) 19–30.
- [3] W. M. P. van der Aalst, M. Pesic, H. Schonenberg, Declarative workflows: Balancing between flexibility and support, Computer Science - Research and Development 23 (2009) 99–113. doi:10.1007/s00450-009-0057-9.
- [4] H. Van der Aa, J. Carmona Vargas, H. Leopold, J. Mendling, L. Padró, Challenges and opportunities of applying natural language processing in business process management, in: COLING 2018: The 27th International Conference on Computational Linguistics: Proceedings of the Conference: August 20-26, 2018 Santa Fe, New Mexico, USA, Association for Computational Linguistics, 2018, pp. 2791–2801.
- [5] T. H. Davenport, J. Guszczca, T. Smith, B. Stiller, Insight-driven organization — deloitte insights (2019).
URL <https://www2.deloitte.com/us/en/insights/topics/analytics/insight-driven-organization.html>
- [6] B. Maqbool, F. Azam, M. W. Anwar, W. H. Butt, J. Zeb, I. Zafar, A. K. Nazir, Z. Umair, A comprehensive investigation of bpmn models generation from textual requirements—techniques, tools and trends, in: Information Science and Applications 2018: ICISA 2018, Springer, 2019, pp. 543–557.
- [7] A. C. d. A. Bordignon, L. H. Thom, T. S. Silva, V. S. Dani, M. Fantinato, R. C. B. Ferreira, Natural language processing in business process identification and modeling: a systematic literature review, Simpósio Brasileiro de Sistemas de Informação (SBSI) (2018).
- [8] J. Herbst, D. Karagiannis, An inductive approach to the acquisition and adaptation of workflow models, in: Proceedings of the IJCAI, 1999, pp. 52–57.
- [9] A. Goossens, J. De Smedt, J. Vanthienen, Extracting decision model and notation models from text using deep learning techniques, Expert Systems with Applications 211 (2023) 118667.
- [10] X. Han, L. Hu, L. Mei, Y. Dang, S. Agarwal, X. Zhou, P. Hu, A-bps: automatic business process discovery service using ordered neurons lstm, in: 2020 IEEE International Conference on Web Services (ICWS), IEEE, 2020, pp. 428–432.
- [11] T. Kampik, C. Warmuth, A. Rebmann, R. Agam, L. N. Egger, A. Gerber, J. Hoffart, J. Kolk, P. Herzig, G. Decker, et al., Large process models: Business process management in the age of generative ai, arXiv preprint arXiv:2309.00900 (2023).
- [12] P. Bellan, M. Dragoni, C. Ghidini, Extracting business process entities and relations from text using pre-trained language models and in-context learning, in: J. P. A. Almeida, D. Karastoyanova, G. Guizzardi, M. Montali, F. M. Maggi, C. M. Fonseca (Eds.), Enterprise Design, Operations, and Computing, Springer International Publishing, Cham, 2022, pp. 182–199.
- [13] M. Grohs, L. Abb, N. Elsayed, J.-R. Rehse, Large language models can accomplish business process management tasks, arXiv preprint arXiv:2307.09923 (2023).
- [14] M. Forell, S. Schüler, Modeling meets large language models, Modellierung 2024 Satellite

- Events (2024). doi:10.18420/modellierung2024-ws-003.
- [15] H. Kourani, A. Berti, D. Schuster, W. M. P. van der Aalst, Process modeling with large language models, in: H. van der Aa, D. Bork, R. Schmidt, A. Sturm (Eds.), *Enterprise, Business-Process and Information Systems Modeling*, Springer Nature Switzerland, Cham, 2024, pp. 229–244.
- [16] P. Bellan, M. Dragoni, C. Ghidini, Process extraction from text: state of the art and challenges for the future, arXiv preprint arXiv:2110.03754 (2021).
- [17] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, <http://arxiv.org/abs/1810.04805> (Oct. 2018). doi:10.48550/arxiv.1810.04805.
- [18] M. Riefer, S. F. Ternis, T. Thaler, Mining process models from natural language text: A state-of-the-art analysis, *Multikonferenz Wirtschaftsinformatik (MKWI-16)*, March (2016) 9–11.
- [19] Object Management Group (OMG), Business process model and notation (bpmn). URL <https://www.bpmn.org/>
- [20] S. Schüler, S. Alpers, State of the art: Automatic generation of business process models, in: J. De Weerd, L. Pufahl (Eds.), *Business Process Management Workshops*, Springer Nature Switzerland, Cham, 2024, pp. 161–173.
- [21] M. J. Page, J. E. McKenzie, P. M. Bossuyt, I. Boutron, T. C. Hoffmann, C. D. Mulrow, L. Shamseer, J. M. Tetzlaff, E. A. Akl, S. E. Brennan, R. Chou, J. Glanville, J. M. Grimshaw, A. Hróbjartsson, M. M. Lalu, T. Li, E. W. Loder, E. Mayo-Wilson, S. McDonald, L. A. McGuinness, L. A. Stewart, J. Thomas, A. C. Tricco, V. A. Welch, P. Whiting, D. Moher, The prisma 2020 statement: an updated guideline for reporting systematic reviews, *BMJ* 372 (2021). arXiv:<https://www.bmj.com/content/372/bmj.n71.full.pdf>, doi:10.1136/bmj.n71. URL <https://www.bmj.com/content/372/bmj.n71>
- [22] J. Fereday, E. Muir-Cochrane, Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development, *International Journal of Qualitative Methods* 5 (1) (2006) 80–92, publisher: SAGE PublicationsSage CA: Los Angeles, CA. doi:10.1177/160940690600500107. URL <http://journals.sagepub.com/doi/10.1177/160940690600500107>
- [23] S. Nasiri, A. Adadi, M. Lahmer, Automatic generation of business process models from user stories, *International Journal of Electrical and Computer Engineering* 13 (1) (2023) 809.
- [24] J. C. d. A. Goncalves, F. M. Santoro, F. A. Baião, Let me tell you a story-on how to build process models, *Journal of Universal Computer Science* 17 (2) (2011) 276–295.
- [25] A. Halioui, P. Valtchev, A. B. Diallo, Bioinformatic workflow extraction from scientific texts based on word sense disambiguation, *IEEE/ACM transactions on computational biology and bioinformatics* 15 (6) (2018) 1979–1990.
- [26] M.-C. De Marneffe, C. D. Manning, Stanford typed dependencies manual, Tech. rep., Technical report, Stanford University (2008).
- [27] J. Pennington, R. Socher, C. D. Manning, Glove: Global vectors for word representation, in: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [28] C. Qian, L. Wen, A. Kumar, L. Lin, L. Lin, Z. Zong, S. Li, J. Wang, An approach for process model extraction by multi-grained text classification, in: *Advanced Information Systems Engineering: 32nd International Conference, CAiSE 2020, Grenoble, France, June 8–12, 2020, Proceedings* 32, Springer, 2020, pp. 268–282.
- [29] J. Neuberger, L. Ackermann, S. Jablonski, Beyond rule-based named entity recognition and relation extraction for process model generation from natural language text, in: *Co-operative Information Systems*, Springer Nature Switzerland, Cham, 2023, pp. 179–197.
- [30] H. A. López, R. Strømsted, J.-M. Niyodusenga, M. Marquard, Declarative process discovery: Linking process and textual views, in: S. Nurcan, A. Korthaus (Eds.), *Intelligent Information Systems*, Springer International Publishing, Cham, 2021, pp. 109–117.

- [31] V. Etikala, Z. Van Veldhoven, J. Vanthienen, Text2dec: extracting decision dependencies from natural language text for automated dmn decision modelling, in: International Conference on Business Process Management, Springer, 2020, pp. 367–379.
- [32] V. Etikala, Extracting decision model components from natural language text for automated business decision modelling., in: RuleML+ RR (Supplement), 2021.
- [33] M. Eberts, A. Ulges, An end-to-end model for entity-level relation extraction using multi-instance learning, in: Conference of the European Chapter of the Association for Computational Linguistics, 2021.
URL <https://api.semanticscholar.org/CorpusID:231879840>
- [34] P. Bellan, C. Ghidini, M. Dragoni, S. P. Ponzetto, H. van der Aa, Process extraction from natural language text: the pet dataset and annotation guidelines, in: Proceedings of the Sixth Workshop on Natural Language for Artificial Intelligence (NL4AI 2022) co-located with 21th International Conference of the Italian Association for Artificial Intelligence (AI* IA 2022), Vol. 3287, CEUR-WS. org, 2022, pp. 177–191.
- [35] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter (2020). [arXiv:1910.01108](https://arxiv.org/abs/1910.01108).
- [36] A. Ezen-Can, A comparison of lstm and bert for small corpus (2020). [arXiv:2009.05451](https://arxiv.org/abs/2009.05451).
- [37] Y. Shen, S. Tan, A. Sordoni, A. Courville, Ordered neurons: Integrating tree structures into recurrent neural networks, [arXiv preprint arXiv:1810.09536](https://arxiv.org/abs/1810.09536) (2018).
- [38] Y. Li, K. Bontcheva, H. Cunningham, Using uneven margins svm and perceptron for information extraction, in: Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005), 2005, pp. 72–79.
- [39] Y. Chen, Z. Ding, H. Sun, Pwep: Process extraction based on word position in documents, in: Ninth International Conference on Digital Information Management (ICDIM 2014), IEEE, 2014, pp. 135–140.
- [40] K. Sintoris, K. Vergidis, Extracting business process models using natural language processing (nlp) techniques, in: 2017 IEEE 19th conference on business informatics (CBI), Vol. 1, IEEE, 2017, pp. 135–139.
- [41] R. C. B. Ferreira, L. H. Thom, M. Fantinato, A semi-automatic approach to identify business process elements in natural language texts, in: International Conference on Enterprise Information Systems, Vol. 2, SCITEPRESS, 2017, pp. 250–261.
- [42] K. Honkisz, K. Kluza, P. Wiśniewski, A concept for generating business process models from natural language description, in: Knowledge Science, Engineering and Management: 11th International Conference, KSEM 2018, Changchun, China, August 17–19, 2018, Proceedings, Part I 11, Springer, 2018, pp. 91–103.
- [43] R. Sonbol, G. Rebdawi, N. Ghneim, A machine translation like approach to generate business process model from textual description, *SN Computer Science* 4 (3) (2023) 291.
- [44] F. Friedrich, J. Mendling, F. Puhlmann, Process model generation from natural language text, in: Advanced Information Systems Engineering: 23rd International Conference, CAiSE 2011, London, UK, June 20–24, 2011. Proceedings 23, Springer, 2011, pp. 482–496.
- [45] H. van der Aa, C. Di Ciccio, H. Leopold, H. A. Reijers, Extracting declarative process models from natural language, in: P. Giorgini, B. Weber (Eds.), Advanced Information Systems Engineering, Springer International Publishing, Cham, 2019, pp. 365–382.
- [46] L. Quishpi, J. Carmona, L. Padró, Extracting decision models from textual descriptions of processes, in: International Conference on Business Process Management, Springer, 2021, pp. 85–102.
- [47] S. Sholih, R. Sarno, E. S. Astuti, Generating bpmn diagram from textual requirements, *Journal of King Saud University-Computer and Information Sciences* 34 (10) (2022) 10079–10093.
- [48] H. A. López, M. Marquard, L. Muttenthaler, R. I. Strømsted, Assisted declarative process creation from natural language descriptions, 2019 IEEE 23rd International Enterprise Distributed Object Computing Workshop (EDOCW) (2019) 96–99.
URL <https://api.semanticscholar.org/CorpusID:208207164>

- [49] H. Cunningham, V. Tablan, A. Roberts, K. Bontcheva, Getting more out of biomedical documents with gate’s full lifecycle open source text analytics, *PLoS computational biology* 9 (2) (2013) e1002854.
- [50] L. G. Azevedo, R. D. A. Rodrigues, K. Revoredo, Bpmn model and text instructions automatic synchronization., in: *ICEIS* (1), 2018, pp. 484–491.
- [51] Object Management Group (OMG), Unified modeling language (uml).
URL <https://www.omg.org/spec/UML/>
- [52] T. Hildebrandt, R. R. Mukkamala, Declarative event-based workflow as distributed dynamic condition response graphs, *PLACES* 69 (10 2011). doi:10.4204/EPTCS.69.5.
- [53] Object Management Group (OMG), Decision Model and Notation (DMN).
URL <https://www.omg.org/dmn/>
- [54] F. Friedrich, Automated generation of business process models from natural language input, M. Sc., School of Business and Economics. Humboldt-Universität zu Berli (2010).
- [55] P. Bellan, H. van der Aa, M. Dragoni, C. Ghidini, S. P. Ponzetto, Pet: An annotated dataset for process extraction from natural language text tasks, in: C. Cabanillas, N. F. Garmann-Johnsen, A. Koschmider (Eds.), *Business Process Management Workshops*, Springer International Publishing, Cham, 2023, pp. 315–321.
- [56] M. Weidlich, J. Mendling, M. Weske, Efficient consistency measurement based on behavioral profiles of process models, *IEEE Transactions on Software Engineering* 37 (3) (2011) 410–429. doi:10.1109/TSE.2010.96.
- [57] R. Dijkman, M. Dumas, B. van Dongen, R. Käärik, J. Mendling, Similarity of business process models: Metrics and evaluation, *Information Systems* 36 (2) (2011) 498–516, special Issue: Semantic Integration of Data, Multimedia, and Services. doi:<https://doi.org/10.1016/j.is.2010.09.006>.
URL <https://www.sciencedirect.com/science/article/pii/S0306437910001006>
- [58] J. Neuberger, L. Doll, B. Engelmann, L. Ackermann, S. Jablonski, Leveraging data augmentation for process information extraction, in: H. van der Aa, D. Bork, R. Schmidt, A. Sturm (Eds.), *Enterprise, Business-Process and Information Systems Modeling*, Springer Nature Switzerland, Cham, 2024, pp. 57–70.
- [59] A. G. Brij B. Gupta, V. Arya, Navigating the security landscape of large language models in enterprise information systems, *Enterprise Information Systems* 18 (4) (2024) 2310846. arXiv:<https://doi.org/10.1080/17517575.2024.2310846>, doi:10.1080/17517575.2024.2310846.
URL <https://doi.org/10.1080/17517575.2024.2310846>
- [60] I. Straw, C. Callison-Burch, Artificial intelligence in mental health and the biases of language based models, *PLOS ONE* 15 (12) (2020) 1–19. doi:10.1371/journal.pone.0240376.
URL <https://doi.org/10.1371/journal.pone.0240376>
- [61] R. Bhardwaj, N. Majumder, S. Poria, Investigating gender bias in bert, *Cognitive Computation* 13 (4) (2021) 1008–1018.
- [62] K. Busch, A. Rochlitzer, D. Sola, H. Leopold, Just tell me: Prompt engineering in business process management, in: *International Conference on Business Process Modeling, Development and Support*, Springer, 2023, pp. 3–11.
- [63] M. Vidgof, S. Bachhofner, J. Mendling, Large language models for business process management: Opportunities and challenges, arXiv preprint arXiv:2304.04309 (2023).
- [64] A. Berti, D. Schuster, W. M. van der Aalst, Abstractions, scenarios, and prompt definitions for process mining with llms: A case study, arXiv preprint arXiv:2307.02194 (2023).
- [65] P. Zhao, H. Zhang, Q. Yu, Z. Wang, Y. Geng, F. Fu, L. Yang, W. Zhang, J. Jiang, B. Cui, Retrieval-augmented generation for ai-generated content: A survey (2024). arXiv:2402.19473.
URL <https://arxiv.org/abs/2402.19473>
- [66] X. Du, H. Ji, Retrieval-augmented generative question answering for event argument extraction, in: Y. Goldberg, Z. Kozareva, Y. Zhang (Eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Association for

- Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 4649–4666. doi:10.18653/v1/2022.emnlp-main.307.
URL <https://aclanthology.org/2022.emnlp-main.307>
- [67] G. Poesia, O. Polozov, V. Le, A. Tiwari, G. Soares, C. Meek, S. Gulwani, Synchronesh: Reliable code generation from pre-trained language models (2022). arXiv:2201.11227.
URL <https://arxiv.org/abs/2201.11227>
- [68] J. Li, Y. Li, G. Li, Z. Jin, Y. Hao, X. Hu, Skcoder: A sketch-based approach for automatic code generation, in: Proceedings of the 45th International Conference on Software Engineering, ICSE '23, IEEE Press, 2023, p. 2124–2135. doi:10.1109/ICSE48619.2023.00179.
URL <https://doi.org/10.1109/ICSE48619.2023.00179>
- [69] J. Li, Y. Zhao, Y. Li, G. Li, Z. Jin, Acecoder: Utilizing existing code to enhance code generation (2023). arXiv:2303.17780.
URL <https://arxiv.org/abs/2303.17780>
- [70] T. Guo, X. Chen, Y. Wang, R. Chang, S. Pei, N. V. Chawla, O. Wiest, X. Zhang, Large Language Model based Multi-Agents: A Survey of Progress and Challenges, CoRR abs/2402.01680, arXiv: 2402.01680 (2024). doi:10.48550/ARXIV.2402.01680.
URL <https://doi.org/10.48550/arXiv.2402.01680>
- [71] Y. Dong, X. Jiang, Z. Jin, G. Li, Self-collaboration Code Generation via ChatGPT, ACM Trans. Softw. Eng. Methodol. Place: New York, NY, USA Publisher: Association for Computing Machinery (Jun. 2024). doi:10.1145/3672459.
URL <https://doi.org/10.1145/3672459>
- [72] C. Qian, W. Liu, H. Liu, N. Chen, Y. Dang, J. Li, C. Yang, W. Chen, Y. Su, X. Cong, J. Xu, D. Li, Z. Liu, M. Sun, ChatDev: Communicative Agents for Software Development, in: L.-W. Ku, A. Martins, V. Srikumar (Eds.), Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024, Association for Computational Linguistics, 2024, pp. 15174–15186.
URL <https://aclanthology.org/2024.acl-long.810>