

Neural computation through synaptic dynamics in serotonergic networks

Michael Benjamin Fernando Lynn

Thesis submitted to the University of Ottawa
in partial fulfillment of the requirements for the
Doctorate in Philosophy in Neuroscience

Department of Cellular & Molecular Medicine
Faculty of Medicine
University of Ottawa

TABLE OF CONTENTS

Abstract.....	iv
Acknowledgements.....	v
List of Figures.....	vii
List of Abbreviations.....	xi
Copyright authorizations.....	xiii
General introduction.....	1
1. Synaptic dynamics, composition of synaptic ensembles and neural computations	1
1.1 Plasticity at central synapses.....	2
1.2 Computational roles of short-term plasticity at single synapses.....	3
1.3 Short-term plasticity and neural computation in microcircuits.....	5
1.4 Synaptic ensembles and information storage at central synapses.....	7
1.5 The synaptic basis of computations at the population level.....	11
2. The dorsal raphe nucleus.....	13
2.1 Anatomical and cellular organization.....	14
2.2 Organization of afferent and efferent connections.....	17
2.3 Spiking properties and neural computations by serotonin neurons.....	22
2.4 Behavioral regulation by serotonin output.....	26
2.5 Serotonergic transmission in raphe and cortex.....	31
2.6 A dynamical systems view of the dorsal raphe nucleus.....	34
Preface to manuscripts.....	38
Manuscript I: A synthetic likelihood solution to the silent synapse estimation problem..	39
Manuscript II: Accurate silent synapse estimation from simulator- corrected	
electrophysiological data using the SilentMLE Python package.....	61
Manuscript III: A slow 5-HT1AR-mediated recurrent inhibitory network in raphe	
computes contextual value through synaptic facilitation.....	83
Manuscript IV: Multiplexing distinct inputs by regulation of spike synchrony in the dorsal	
raphe nucleus.....	130
Manuscript V: Acute stress shapes synaptic inhibition within an amygdala microcircuit.....	155
Manuscript VI: BK channels at dendritic spines: A mechanism for coupling morphology,	
plasticity and information storage?.....	163
General discussion.....	171
Summary of results.....	171
Characteristics of monoaminergic transmission.....	172
Neural coding strategies in the dorsal raphe nucleus.....	176
Silent synapses as a substrate for learning in cortical networks.....	186
Synthetic likelihood functions as a scalable approach for simulation-based	
inference.....	187
Concluding remarks.....	188
Bibliography.....	189
Appendices.....	211
Appendix A: Manuscript VII: Reconciling current theories of consciousness.....	212
Python packages	
Appendix B: Package I: SynapPy: Automated synaptic waveform analysis in Python.....	222
Appendix C: Package II: MouseBerry: Raspberry Pi-based interface for controlling	
behavioral experiments.....	230
Appendix D: Package III: SilentMLE Python package for silent synapse estimation...	240
Appendix E: Supplemental for Manuscript I.....	246

Appendix F: Supplemental for Manuscript III.....	260
Appendix G: Supplemental for Manuscript IV.....	273

Abstract

Synapses are a fundamental unit of computation in the brain. Far from being passive connections between spiking neurons, synapses display striking short-term dynamics, undergo long-term changes in strength, and sculpt network-level processes in a complex manner. These synaptic dynamics, both in time and across space, may be a fundamental determinant of population-level computations and behavioral output of the brain, yet their role in neuromodulatory circuits is relatively under-explored. *First*, I developed and validated a set of likelihood-based inference tools to quantify the dynamics of synaptic ensemble composition throughout development. *Second*, I examined network computations in the serotonergic dorsal raphe nucleus through a dynamical lens, exploring the role of short-term synaptic dynamics at sparse recurrent connections, and of distinct long-range synaptic inputs, in shaping the output of spiking populations. *1. Simulation-based inference of synaptic ensembles.* Functional features of synapses are typically inferred by sampling small ensembles of synapses, yet it is unclear if such subsamples exhibit biases. I developed a statistical framework to address this question, using it to demonstrate that common bulk electrical stimulation methods for characterizing the fraction of silent synapses exhibit high bias and variance, and using typical sample sizes, possess insufficient statistical power for accurate inference. I developed and validated a novel synthetic likelihood-based inference approach based on a simulator of the underlying experimental methodology. This new estimator, made available in an object-oriented Python toolbox, reduces bias and variance compared to previously reported methods, and provides a scalable method for examining synaptic dynamics throughout development. These tools were validated by targeted recording from hippocampal CA1 neurons in juvenile mice, where they reveal fundamental tradeoffs between release probability, number of synapses sampled, and statistical power. *2. Synaptic dynamics and population computations in the serotonin system.* This part is comprised of two manuscripts. First, in the dorsal raphe nucleus, I uncovered slow, inhibitory recurrent interactions between serotonin neurons that are generated by local serotonin release. These connections were probabilistic, displayed striking short-term facilitation, gated the spiking output of serotonin neurons, and could be activated by long-range excitatory input from lateral habenula, representing threat signals. Targeted physiology and modeling revealed that these recurrent short-term facilitation features generated paradoxical excitation-driven inhibition in response to high-frequency habenula input. These facilitation rules additionally supported winner-take-all dynamics at the population level, providing a contrastive operation between functionally distinct serotonergic ensembles. Behaviorally, activating long-range lateral habenula input to dorsal raphe nucleus generated a transient, frequency-dependent suppression of reward anticipation consistent with these recurrent dynamics, without modulating the underlying reward association itself. These dynamics, we suggest, support sharp behavioral state transitions in changing environments. In a second manuscript, I explored the multiplexing of distinct long-range inputs in serotonergic circuits through spike synchrony. I demonstrated that a population of serotonergic neurons receives input from both lateral habenula and prefrontal cortex. These inputs produced similar subthreshold events, but prefrontal cortex triggered spikes with much higher latencies, supporting a population synchrony code for input identity. These input-specific spike timing patterns could be read out by simple linear decoders with high accuracy, suggesting they could be demultiplexed by downstream circuits receiving sparse innervation by serotonergic axons. We uncovered a novel intracellular calcium conductance in serotonergic neurons that altered the spectral characteristics of membrane voltage in a manner sufficient to generate long-latency, power law-distributed spike times, suggesting a simple dynamical origin for the production of synchronous or asynchronous spiking. This work indicates that serotonergic circuits can multiplex distinct informational streams through population spike synchrony mechanisms. Together, these investigations reveal that the dynamics of short-term facilitation and synaptic ensemble composition can act as the fundamental substrate for flexible computation by spiking networks across the brain.

Acknowledgements

Graduate school has been one of the great formative periods in my life. It has been a period characterized by joyful curiosity, increasingly fierce intellectual independence, intense emotions, pleasantly surprising sources of camaraderie and companionship, and idle daydreaming and manic work in equal measure. There are so many people to thank along this journey.

I will start by thanking my supervisor, Dr. Jean-Claude Béique. Jean-Claude has allowed me to have a large degree of freedom while choosing the eclectic mixture of personally compelling neuroscience problems that form the investigations in this thesis. I have been given the intellectual breathing room required to develop tools and techniques in the lab, some obscure (simulation-based inference techniques applied to synaptic ensembles), some practical (the development of head-fixed behavioral setups that now form a major component of the lab's work). It is rare indeed to have a supervisor who trusts the vision and skills of their trainees to this level, and I am greatly indebted to Jean-Claude for fostering such a training environment. Most of all, Jean-Claude is a fundamentally decent person who cares about the happiness of the trainees in the lab, and I can express my strong appreciation for his mentorship and guidance.

I would also like to thank my thesis advisory committee, Dr. André Longtin, Dr. Diane Lagace, and Dr. Tuan Bui, for guiding me throughout my studies. A special thank you to close collaborator Dr. Richard Naud for insightful and exciting discussions, a deep knowledge of the literature, and practical guidance and mentorship. I would also like to express my deep gratitude to co-supervisor Dr. Leonard Maler, who made all of this possible by taking me in as an eager undergraduate student, conveying the excitement of cellular physiology and neural coding principles, and being a source of guidance, brilliant scientific ideas, and unwavering support at many key points throughout my career.

I would also like to thank my parents, Mary and Brian, for their endless support, deep love, boundless faith in my abilities, sharp humour, practical advice, and ability to always be

there for me when I need them. I love you both very much. It is no exaggeration to say that I would probably not be doing a PhD without your support and guidance. Emma, my sibling - it's been great to grow up in this crazy world with you, and I feel very lucky to have such a supportive and amazing sister. Finally, to my grandparents, Dr. Aggie Fernando and Dr. Herbert Fernando, thank you for showering me with microscopes, bugs and books as a child, and teaching me the joy of doing science.

To scientific friends that have become personal friends - Sébastien Maillé, Emerson Harkin, Candice Lee, Léa Caya-Bissonnette, Rishi Rajalingham, Diego Vargas, and many more. The journey wouldn't have been possible without your friendship, long lunches, supportive chats, shared knowledge, and collective joy and misery in this endeavour to understand the brain.

Finally, to Éloïse, the love of my love - thanks for being there and supporting me through all of this. We'll get through this together.

List of Figures

Manuscript I

Figure 1: The FRA technique returns high-variance estimates under conditions with no silent synapses present.

Figure 2: Experimentally bounded simulations reveal a fundamental bias inherent in bulk electrical stimulation.

Figure 3: Electrophysiological recordings at CA3-CA1 synapses are consistent with simulated FRA distributions.

Figure 4: Statistical properties of FRA-MLE estimator compared with other methods.

Manuscript II

Figure 1: Maximum-likelihood estimation on two datasets with synthetic data.

Manuscript III

Figure 1: Habenula input to raphe triggers direct excitation and heterosynaptic inhibition in 5-HT neurons.

Figure 2: 5-HT neurons in raphe are organized in a recurrent inhibitory network.

Figure 3: Serotonin release is probabilistic and strongly facilitating at inhibitory recurrent connections.

Figure 4: A quantitative model of short-term release dynamics captures patterned inputs.

Figure 5: Habenula input to raphe generates a nonmonotonic network response through recurrent short-term plasticity.

Figure 6: Winner-take-all dynamics generated by recurrent short-term plasticity.

Figure 7: Recurrent inhibition gates spiking activity of 5-HT neurons.

Figure 8: Habenula input to raphe disrupts the expression of learned reward associations in a value- and frequency-dependent manner.

Manuscript IV

Figure 1: Subthreshold and spike timing characteristics of long-range PFC and LHb input to 5-HT neurons.

Figure 2: Decoding long-range inputs by spike-timing and subthreshold features in 5-HT populations.

Figure 3: An inward calcium conductance in 5-HT neurons shapes EPSPs and contributes to membrane noise.

Figure 4: The statistics of membrane noise shapes stochastic spike timing in 5-HT neurons.

Manuscript V

Figure 1: Glucocorticoids interact with the endocannabinoid system following acute stress exposure to shift excitation/inhibition balance in the BLA.

Manuscript VI

Figure 1: A model of BK channel regulation of synaptic transmission and plasticity

Appendix E: Supplemental for Manuscript I

Supplemental Figure 1: The poor estimation characteristics of the FRA equation do not vary significantly over the parameter space of release probabilities and synapse numbers. Related to Figure 1.

Supplemental Figure 2: Simulations exploring low synaptic release probabilities. Related to Figure 2.

Supplemental Figure 3: Simulations exploring an FRA assumption violation with unequal release probabilities between silent and active synapses. Related to Figure 2.

Supplemental Figure 4: A maximum likelihood estimator for silent synapse fractions. Related to Figure 4.

Appendix F: Supplemental for Manuscript III

Supplemental Figure 1: Validation of SERT-CRE-tdTomato mouse line and LHb injection site.

Supplemental Figure 2: LHb sends axonal projections to the DRN.

Supplemental Figure 3: Voltage-clamp experiments provide evidence that the afterhyperpolarization of serotonin neurons is mediated by SK channels, not 5-HT_{1A}Rs.

Supplemental Figure 4: Fluorescent dextran injections in DRN reveal local 5-HT projections.

Supplemental Figure 5: Validation of local CHETA optogenetic strategy in DRN.

Supplemental Figure 6: Postsynaptic dynamics of recurrent 5-HT_{1A} connections indicate non-desensitizing receptors.

Supplemental Figure 7: Frequency dependence of 5-HT_{1A} connections at suprathreshold stimulation frequencies.

Supplemental Figure 8: Nonlinear transformation of habenulo-raphé afferents generated by recurrent plasticity rules, with tenfold decrease in recurrent inhibitory connection probability (sum 0.25nS conductance received by each serotonin neuron).

Supplemental Figure 9: Plasticity-dependent winner-take-all dynamics in raphe, with tenfold decrease in recurrent inhibitory connection probability (sum 0.25nS conductance received by each serotonin neuron).

Supplemental Figure 10: LHb-driven, 5HT_{1A}R mediated, recurrent feedforward modulation of 5-HT firing activity in the DRN.

Supplemental Figure 11: The habenulo-raphé pathway disrupts reward-conditioned responses in a value-dependent manner and does not affect unconditioned response.

Appendix G: Supplemental for Manuscript IV

Supplemental Figure 1: Dual-color opsin dissection of long-range inputs to 5-HT neurons.

Supplemental Figure 2: Further quantification of subthreshold and suprathreshold statistics of light-evoked events in 5-HT neurons.

Supplemental Figure 3: Low-threshold calcium channels modulate the kinetics of EPSPs in 5-HT neurons.

Supplemental Figure 4: Subthreshold noise properties are not affected by calcium channels in 5-HT neurons.

Supplemental Figure 5: Calcium channels affect spike timing in 5-HT neurons.

List of Abbreviations

2P	two photon
4-AP	4-aminopyridine
5-HT	serotonin
AHP	afterhyperpolarization
AMPA	α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid
APV	2-amino-5-phosphonopentanoic acid
BK	large conductance calcium-activated potassium channel
CA1, CA3	Cornu Ammonis areas 1 and 3 of the hippocampus
ChR2	channelrhodopsin-2
CNQX	6-cyano-7-nitroquinoxaline
DIC	differential interference contrast
DRN	dorsal raphe nucleus
EPSC	excitatory postsynaptic current
EPSP	excitatory postsynaptic potential
eYFP	enhanced yellow fluorescent protein
FRA	failure rate analysis
GABA	γ -amino butyric acid
GFP	green fluorescent protein
GIRK	G-protein coupled inwardly rectifying potassium channel
IPSC	inhibitory postsynaptic current
IPSP	inhibitory postsynaptic potential
LHb	lateral habenula
LTD	long-term depression
LTP	long-term potentiation
MLE	maximum likelihood estimation
NBQX	2,3-dioxo-6-nitro-1,2,3,4-tetrahydrobenzo[f]quinoxaline-7-sulfonamide
NMDA	N-methyl-D-aspartate

PFC	prefrontal cortex
PIC	picrotoxin
SERT	serotonin transporter
SK	small conductance calcium-activated potassium channel
STP	short-term plasticity
TTX	tetrodotoxin
VGLUT3	vesicular glutamate transporter 3

Copyright Authorizations

Manuscript I: This article was published in *Cell Reports* in 2020 under a Creative Commons license (CC BY-NC-ND 4.0): “The second option is the Creative Commons Attribution-Noncommercial-No Derivative Works License (CC BY NC ND), which allows users to copy and distribute the article, provided the work is attributed back to the original author and publisher. The article cannot be changed in any way or used commercially. This license gives the author the most control about how the work might be published.”

Manuscript II: This article was published in *STAR Methods* in 2020 under a Creative Commons license (CC BY-NC-ND 4.0): “The second option is the Creative Commons Attribution-Noncommercial-No Derivative Works License (CC BY NC ND), which allows users to copy and distribute the article, provided the work is attributed back to the original author and publisher. The article cannot be changed in any way or used commercially. This license gives the author the most control about how the work might be published.”

Manuscript III is under review and thus carries no copyright.

Manuscript IV is in preparation for submission and thus carries no copyright.

Manuscript V: This article was published in *J Neurosci* in 2017 under a Creative Commons license (CC-BY): “Authors grant JNeurosci a license to publish their work and copyright remains with the author. For articles published after 2014, the Society for Neuroscience (SfN) retains an exclusive license to publish the article for 6 months; after 6 months, the work becomes available to the public to copy, distribute, or display under the terms of the Creative Commons Attribution 4.0 International License (CC-BY). This license allows data and text

mining, use of figures in presentations, and posting the article online, provided that the original article is credited.”

Manuscript VI: This article was published in *J Physiol* in 2022 under an Exclusive License Agreement: “b. Re-use in other publications. The right to re-use the Final Published Version or parts thereof for any publication authored or edited by the Contributor (excluding journal articles) where such re-used material constitutes less than half of the total material in such publication. In such case, any modifications must be accurately noted.”

Manuscript VII: This article was published in *J Neurosci* in 2020 under a Creative Commons license (CC-BY): “Authors grant JNeurosci a license to publish their work and copyright remains with the author. For articles published after 2014, the Society for Neuroscience (SfN) retains an exclusive license to publish the article for 6 months; after 6 months, the work becomes available to the public to copy, distribute, or display under the terms of the Creative Commons Attribution 4.0 International License (CC-BY). This license allows data and text mining, use of figures in presentations, and posting the article online, provided that the original article is credited.”

Introduction

1. Synaptic dynamics, composition of synaptic ensembles and neural computations

Our brains consist of a complex, organized web of billions of neurons. These neurons act in a coordinated fashion to produce all of our thoughts, emotions, complex relationships, and ultimately the sophisticated behaviors that we engage in every day. What are the mechanisms and fundamental principles by which such a complex system generates the complex behaviors that define us as a human? One defining principle that has emerged over decades of work is the key contribution of an elementary component of neural networks: the chemical synapse. Synapses are small structural connections between neurons, acting as sites of specialized chemical transmission for signaling between neurons over timescales of milliseconds. There are trillions of synapses in the human brain, providing an effectively vast substrate for wiring and rewiring the brain for complex behavior over the lifetime of organisms. First described by Ramon y Cahal in the late 1800s, synapses consist of a presynaptic side that releases vesicle-packaged neurotransmitters such as glutamate or GABA upon sufficient depolarization, and a postsynaptic side in close apposition that expresses receptors that bind these neurotransmitters to elicit a change in membrane potential or intracellular processes in the receiver cell. These structures endow networks of neurons with considerable computational capabilities, by virtue of their spatial organization, their dynamics over short timescales, and long-term activity dependent changes that occur over long timescales. Below, I will discuss the synaptic basis of information transmission, how the dynamics of synapses over longer and shorter timescales can be harnessed by neural circuitry to sculpt activity patterns, and how ensembles of synapses with distinct features can provide a neural substrate for learning and memory.

1.1 Plasticity at central synapses

Far from being static, synapses can undergo dramatic structural and functional changes throughout development and learning, a general principle called synaptic plasticity. During learning, through a phenomenon termed *long-term plasticity*, long-lasting changes in the number, strength and/or functional makeup of these synapses can occur. Donald Hebb first proposed the central tenets of long-term (Hebbian) plasticity 1949, and the first experimentally reported evidence of long-term potentiation was in hippocampus (Bliss & Lomo, 1973). Long-term plasticity classically requires coincident activation of pre- and post-synaptic neurons to produce a long-lasting change in synaptic efficacy (for reviews, see Kandel, 2014; Magee & Grienberger, 2020) although the precise timing between pre- and post-synaptic activation can elicit either synaptic potentiation or depression (spike timing-dependent plasticity; Bi & Poo, 1998). Mechanistically, the classical form of long-term potentiation requires calcium entry at single spines through NMDA receptors during induction, and subsequent AMPA receptor insertion at the postsynaptic membrane to increase synaptic conductance. These long-term synaptic plasticity mechanisms are thought to reflect a fundamental substrate for learning and memory in the central nervous system.

Synapses can also exhibit dynamics over far shorter timescales. Indeed, the phenomenon of long-term plasticity coexists with forms of synaptic *short-term plasticity*, reflecting reversible, activity-dependent changes in the parameters of synaptic transmission that typically lasts seconds. Such short-timescale forms of plasticity were initially discovered during early experiments at the frog neuromuscular junction (Eccles et al., 1941; Feng et al., 1940). These dynamics can profoundly influence information transfer at single synapses. Some synapses undergo *short-term facilitation* when multiple presynaptic activation events occur in quick succession (tens of milliseconds), leading to a transient increase in the amplitude of postsynaptic responses. The most probable mechanism is presynaptic calcium accumulation that leads to an increased probability of vesicle release (Dobrunz & Stevens, 1997; Fisher et al., 1997, Fioravante & Regehr, 2011). Other synapses can undergo *short-term depression*, where

repeated presynaptic stimulation decreases the release probability of subsequent stimuli and the subsequent postsynaptic responses, likely due to a decrease in the size of the presynaptic readily releasable pool of neurotransmitter-filled vesicles (Dobrunz & Stevens, 1997; Fioravante & Regehr, 2011). There is substantial heterogeneity in the presence and strength of these two opposing processes, both between individual synapses and between neurons, and indeed they can interact in complex ways (Muller et al., 2010).

1.2 Computational roles of short-term plasticity at single synapses

Short-term plasticity can have substantial influences on information transfer and neural computations in the central nervous system. These naturally depend on the nature of firing statistics and on the form of short-term plasticity exhibited in the microcircuit in question. Cortical neurons can exhibit highly variable rates of discharge, from under 1Hz to over 200Hz; work over the past several decades has focused on how synaptic dynamics can modulate information transfer given these variable firing rates both within and between neurons (for reviews, see Abbott & Regehr, 2004; Zucker & Regehr, 2002; Fortune & Rose, 2001). Below I will discuss how facilitation and depression dynamics at individual synapses can modulate information transfer in complex ways.

A body of work has convincingly shown that one function of synaptic depression may be to expand the dynamic range of neural computations. Short-term depression at primary visual cortex synapses can provide an effective means of gain control for synaptic inputs with diverse discharge rates, such that neurons effectively respond to relative changes in rates of each input synapse rather than absolute changes (Abbott et al., 1997). Similar ideas about short-term depression expanding the dynamic range of information processing have been shown in cerebellum, although here depression at excitatory inputs effectively allows inhibition to control gain in a manner that does not depend on noise (Rothman et al., 2009). At thalamocortical synapses, strong synaptic depression regulates cortical responses to repeated sensory inputs to whiskers at ethologically relevant frequencies above 2Hz (Chung et al., 2002),

highlighting a role for depression in sensory adaptation mechanisms. Synaptic depression has also been shown to contribute to noise reduction for noisy input spiketrains from electrosensory afferents in electric fish (Khanbabaie et al., 2010). More generally, synaptic depression can act as a low-pass filter on information transfer at single synapses (Fortune & Rose, 2001), or a band-pass filter when the stochastic nature of neurotransmitter release is considered (Rosenbaum et al., 2012).

Synaptic facilitation dynamics, in contrast, can act as high-pass filter for information transfer at single synapses, with a variety of implications for neural computation (Jackman & Regehr, 2017; Rotman et al., 2011). This is especially true for synapses that have low initial release probabilities, as here facilitation can have strong consequences for the reliability and strength of transmission at central synapses. An excellent example of facilitation acting as a high-pass filter comes from work on cholinergic inhibitory synapses between medial olivocochlear neurons and cochlear outer hair cells, a feedback circuit that regulates the gain of auditory responses. Here, fast facilitation of these inhibitory synapses, combined with temporal summation, leads to a reduction in cochlear gain with firing intensity (Ballestro et al., 2011). Another example of an unreliable synapse being converted to strong and reliable synapse is at excitatory connections from medial superior colliculus to dorsal periaqueductal grey (Evans et al., 2018). Here, facilitation of these synapses provides a synaptic threshold mechanism for computing escape behavior to looming stimuli presented to mice, such that visual threat information processed by superior colliculus is nonlinearly converted into a rapid escape behavior.

An influential idea is that fast facilitation can endow postsynaptic neurons with burst detection capabilities by increasing the reliability of synaptic transmission or information transfer in these regimes (Izhikevich et al., 2003; Lisman, 1997; Rotman et al., 2011). This likely depends on a precise correspondence between presynaptic bursting statistics and the timescale of synaptic facilitation. Furthermore, network effects with short-term plasticity rules

can arguably provide a more comprehensive mechanism for burst detection (see Section 1.3 below for more information).

1.3 Short-term plasticity and neural computation in microcircuits

At the network level, distinct short-term plasticity rules exist at each possible pair of interacting cell types, leading to complex network interactions and functions. An outstanding question in cortex is whether the high-dimensional patterns of neural activity have structural features that support the representation of external stimuli as an internal state over long timescales. It has been suggested that short-term synaptic plasticity in networks of excitatory and inhibitory neurons could generate state-dependent computations of spatial and temporal patterns of inputs, or may participate in decoding of temporal patterns (Buonomano, 2000; Buonomano & Maass, 2009). In cerebellum, short-term synaptic dynamics displaying realistic heterogeneity were able to generate temporal patterns of firing in granular cells that supported the learning of sub-second representation of time (Barri et al., 2022). Thus, short-term synaptic dynamics can either generate or decode important temporal features of the activity of spiking networks over behavioral timescales.

Perhaps one of the best demonstrations of the potential computational power of short-term synaptic dynamics comes from work on short-term memory storage in cortex. Ever since the discovery of monkey prefrontal neurons exhibiting sustained delay-period spiking in a working memory task (Fuster & Alexander, 1971; Funahashi et al., 1989), work has focused on uncovering network or cellular mechanisms that can support this working memory trace. Influential theoretical work has posited that working memory in cortex could be generated by specific short-term synaptic plasticity rules. In this model, recurrent excitatory networks with both facilitation and depression components can encode a memory (an appropriately timed external signal) for approximately one second through presynaptic calcium accumulation at excitatory synapses, which acts as a ‘buffer’ by triggering short-term facilitation upon memory recall (Barak & Tsodyks, 2007, Mongillo et al., 2008; Barak & Tsodyks, 2014). As well as

enabling ‘activity-silent’ traces of memory, short-term facilitation can also explain irregular sustained firing observed in cortical neurons (Hansel & Mato, 2013) (but see Constantinidis et al., 2018; Masse et al., 2019). More generally, the powerful idea that latent synaptic traces can reflect past inputs to a microcircuit may be an important, generalizable principle for how groups of neurons communicate dynamically over longer timescales.

Single neurons may fire regular trains of action potentials defined by long inter-spike intervals, or bursts of spikes defined by very short inter-spike intervals. Whether or not these two spiking modes consist of distinct forms of information that are being multiplexed depends on whether downstream circuits can, in fact, decode them. Single neurons exhibiting synaptic facilitation can decode bursts in theory (Izhikevich et al., 2003; Lisman, 1997; Rotman et al., 2011). However, recent work has suggested that network motifs with biologically plausible short-term dynamics can efficiently and robustly parse out streams of bursts from events (Naud & Sprekeler, 2018). The motif for decoding burst probability consists of facilitating direct excitation combined with depressing disynaptic inhibition; to decode event rate, an excitatory depressing synapse is sufficient. These network motifs can act as the basis for a high-performing burst-dependent learning rule in hierarchical circuits (Payeur et al., 2021). These findings are immensely exciting as they indicate that multiplexing of distinct information sources is possible even within single neurons, with short-term plasticity as a crucial component of the readout mechanism.

Despite being extensively studied in cortex, cerebellum, and the sensory periphery, it is still relatively rare to see studies in neuromodulatory systems that rigorously quantify short-term synaptic dynamics and explore their role in network computations. This is surprising as neuromodulatory nuclei often send widespread efferent projections that target cortical and subcortical regions, positioning them as powerful regulators of brainwide activity. An important component of Manuscript III consists of an exploration of short-term synaptic dynamics at local serotonergic inhibitory connections between serotonin neurons. These connections are unreliable, exhibit successes and failures similar to classical glutamatergic or GABAergic

synapses, and inhibit their postsynaptic partners over timescales of hundreds of milliseconds. By providing a compact mathematical description of strong, fast facilitation and long-lasting depression that is scalable to simulations of thousands of neurons, I describe a suite of important network computations taking place through the interaction of fast, direct excitation with this novel slow lateral inhibition. A key finding is prominent winner-take-all (competitive) dynamics between connected ensembles of serotonin neurons that we show critically depends on short-term synaptic dynamics. Taken together with recent literature on the role of serotonin in behavioral tasks, we suggest that these dynamics support a contrastive value computation between submodules in raphe.

1.4 Synaptic ensembles and information storage at central synapses

A hallmark of mammalian intelligence is the capability to store sophisticated, overlapping memories of the world and recall them from incomplete inputs, a process that operates continuously and effectively for many decades in a seemingly effortless manner. A key component of our understanding of such learning and memory processes comes from work on underlying biological mechanisms of synaptic rewiring that occurs in mammalian brains. These experimental findings, I argue here, must be considered alongside rapid progress in deep learning architectures and learning rules in order to fundamentally advance our understanding of learning and memory.

Cortical neurons have tens of thousands of synaptic inputs, reflecting an effectively vast space for rewiring principles to operate over. The discovery that stimulation can produce long-term changes in synaptic efficacy lasting hours in hippocampus (Bliss & Lomo, 1973) provided a tantalizing putative link between synaptic changes and memories. Work over subsequent decades has focused on finding synaptic features that are sufficient to explain these long-term changes. A key finding was the presence of *silent synapses*, or individual synapses that contain NMDA receptors (NMDARs) yet lack AMPA receptors (AMPA receptors), in hippocampus (Isaac et al., 1995; Liao et al., 1995). These types of synapses, present largely during development of brain

circuitry (but see Vardalaki et al., 2022), are nonconducting (silent) at hyperpolarized membrane potentials, yet conduct current at depolarized potentials due to the opening of voltage-gated NMDARs. Thus, they do not typically contribute to depolarizing the postsynaptic neuron even when activated. However, activity-dependent plasticity can convert these so-called silent synapses into nonsilent synapses through AMPAR insertion into the postsynaptic membrane. This increases synaptic strength in a long-term manner, augmenting the effective number of synaptic inputs by ‘unsilencing’ a number of them at typical hyperpolarized membrane potentials (Isaac et al., 1995; Liao et al., 1995).

Silent synapses in hippocampus and visual cortex display strong developmental regulation, with the fraction of detectable silent synapses decreasing sharply with the age of animals (Soares et al., 2017; Lee et al., 2016; Durand et al., 1996; Rumpel et al., 1998). This presumably coincides with their pruning or conversion to nonsilent form during activity-dependent plasticity epochs that occur during learning. However, a body of recent work has argued that silent synapses may additionally be generated *de novo* during adulthood, providing a site for future plasticity of these neural networks. Addictive drugs such as morphine or cocaine can enhance silent synapse fraction in reward-related networks (Huang et al., 2009; Brown et al., 2011; Ma et al., 2014; Ma et al., 2016; Graziane et al., 2016). Furthermore, sensory deprivation, stress or even associative learning are additionally capable of inducing silent synapse formation in hippocampus or reward-related networks (Whitaker et al., 2016; Milshtein-Parush et al., 2017, Suvrathan et al., 2013). Together, these emerging results point to a revised model where silent synapses are formed and pruned during development, yet are capable of being generated again in adult circuits that may require their presence for learning-related plasticity or rewiring.

Theoretical work has shown that high fractions of silent synapses can, in principle, contribute to optimal information storage conditions (Brunel et al., 2004; Brunel, 2016; Clopath and Brunel, 2013). However, at present an open question is how these silent synapses

mechanistically contribute to the development of rich representations in biological networks throughout learning. How can we answer such a question with the tools available to us?

Machine learning advances have shown that deep neural networks can learn sophisticated representations of tasks, achieving near-human levels of performance. Furthermore, trained networks can have their weights frozen to examine task representations in hidden layers (ie neither input nor output), allowing a level of insight into fundamental mechanisms of network operation. An influential recent approach is to use deep learning principles to better understand how learning is coordinated in biological networks. This approach advocates focusing on the levels of abstraction appropriate for training deep neural networks: the architecture of the network, the learning rule and the objective function (Richards et al., 2019). Here, the fraction of silent synapses and their localization in networks would correspond to an architectural (connectivity) change to the network, and the contributions of these synapses to representations formed by neural ensembles could be quantified by training deep neural networks with such architectural changes and analyzing hidden representations.

Recent work has highlighted the difference between ‘rich’ and ‘lazy’ representations in trained deep neural networks, a feature that depends on the distribution of initialization weights in the network (Chizat et al., 2019; Flesch et al., 2022a; Flesch et al., 2022b). Lazy representations emerge when initialization weights are drawn from a broad distribution, and while such networks show fast, exponentially decaying error rates, hidden units do not learn structured representations of the task that are generalizable. In contrast, rich representations emerge from initialization weights drawn from a narrow distribution with a mean of 0. Networks with such an initial weight distribution learn more slowly, yet construct rich, structured representations of the world that allow for context-dependent decisions. It is possible that the presence of substantial fractions of silent synapses, which have effective weights of zero, could encourage rich representations of the external world to emerge in developing cortical networks throughout learning. The expression of such silent synapses, dynamically regulated

even in adulthood, could provide neural circuits with a powerful mechanism to optimize various tradeoffs in the learning process for a particular situation or task.

However, for more concrete links to be made between the machine learning literature and neuroscience, and thus to arrive at a deep and mechanistic understanding of how silent synapses contribute to network function, the field requires toolsets to accurately quantify the presence or absence of silent synapses in real networks. These toolsets are required both at the level of synaptic populations onto defined neural subtypes, and at the subcellular level of dendritic clustering of silent synapses. At present, the main approaches for quantifying silent synapses at scale include bulk electrical stimulation of synaptic inputs to a recorded neuron (Isaac et al., 1995; Liao et al., 1995), or, more recently, glutamate uncaging techniques (Béïque et al., 2006; Busetto et al., 2008; Lee et al., 2016; Soares et al., 2017) (see Introduction to Manuscript I for a more complete description of these approaches). It has been unclear whether these approaches are scalable, have appropriate statistical power, or exhibit biases - questions that are vital when quantifying differences in silent synapse fraction throughout development, or between cell types.

In Manuscript I, I address these questions directly with a combination of experiments, simulations and formal analysis. I show that common bulk electrical stimulation methods for quantifying silent synapse fraction are noticeably inaccurate, displaying poor bias and variance. Glutamate uncaging methods display slightly higher statistical power at low true silent synapse fractions (<10%) and much higher statistical power when performing constrained hypothesis testing. Furthermore, in an effort to improve the reliability of bulk electrical stimulation approaches, I develop a *synthetic likelihood*-based estimator, which uses constrained experimental simulators to infer silent synapse fraction with no bias and lower variance, leading to higher statistical power. This silent synapse estimator, and the underlying experimental simulator, is made freely available as a simple-to-use Python package which forms the basis of Manuscript II. We hope that these tools will allow the collection of comprehensive, accurate datasets of silent synapse fractions between cell types and throughout development in cortical

microcircuits. These datasets will be an essential component in bridging the gap between biological neuroscience and machine learning approaches, and thus in uncovering the mechanistic, functional roles for silent synapses in the context of learning within microcircuits.

1.5 The synaptic basis of computations at the population level

Many meaningful neural computations, in the behavioral or cognitive sense, are best thought of as emergent properties of large networks of neurons. Indeed, a trend over the last several decades of neuroscience research has been a realization that despite the diversity in single-cell responses in biological networks during complex tasks (e.g. Churchland & Shenoy, 2007), a clearer picture of the computation can emerge by simply observing the neural dynamics of a large population taking place in an appropriate low-dimensional subspace (Churchland et al., 2012; Mante et al., 2013; Vyas et al., 2020; Sylwestrak et al., 2022; Nair et al., 2023). These population trajectories, when plotted over time, can reveal the presence of dynamical elements such as line attractors, saddle points and oscillatory dynamics, which collectively can provide an explanation of features of the animal's behavioral response. For example, the choice behavior of monkeys on a visual task can be explained at the neural population level as a line attractor and a contextual selection vector (Mante et al., 2013). Dynamical processes bearing strong similarities to those seen in real biological networks can be generated by appropriately trained recurrent networks (Mante et al., 2013; Vyas et al., 2020). This indicates that simple structural features are indeed sufficient to generate complex dynamics given proper learning rules. Similar approaches have been used to convincingly identify the key dynamical 'building blocks' that are re-used across diverse tasks in a flexible manner (Driscoll et al., 2022). This emerging body of work suggests that the spiking output of large neural populations is a useful level of abstraction to understand the mechanisms behind neural computation.

In biological networks, how do such striking population-level dynamics emerge from the collection of relatively simple elements (neurons and their connections) through learning and

development? The specific network architecture and value of synaptic weights play a significant causal role generating precise trajectories of network dynamics - for instance, standardized training of recurrent neural networks relies on adjusting synaptic weights through the now-ubiquitous backpropagation of error algorithm. However, a standing question is the role that the dynamics of synapses play in generating these complex population-level spiking processes. For example, as discussed above, short-term plasticity and the precise makeup of ensembles of synapses can alter information flow in a manner that is qualitatively distinct from simple synaptic weight changes. Work has begun to explore how dynamic synapses can imbue artificial networks with fundamentally different capabilities than standard recurrent networks with trained, fixed synaptic weights (Aitken & Mihalas, 2023). Despite this, the connection between short-term synaptic properties and population dynamics in real biological networks is still relatively unexplored. In Manuscript III, I explore the short-term dynamics at serotonergic recurrent connections in dorsal raphe nucleus. These synaptic dynamics, when inserted into a biophysically grounded network model, are necessary to generate strong winner-take-all competition between submodules receiving distinct long-range input - a dynamical process that evolves at the population level through recurrent interaction between principal serotonin neurons. These results suggest that synaptic facilitation can act in a switch-like manner to toggle between distinct modes of network dynamics, given spatially segregated synaptic input streams.

Additionally, the majority of these investigations on computations through population dynamics have focused on the mean firing rates of individual neurons, which are binned over small timesteps to explore low-dimensional subspaces of activity. While a powerful and practical approach, this neglects fine-timescale correlations between the timing of action potentials emitted by neurons, which may be an important aspect of neural computation in some brain regions. Temporal coding provides an alternate framework with which to consider neural activity, by contending that the timing of action potentials conveys information that is decoded by downstream regions. Indeed, population spike synchrony (the synchronous

production of spikes across a neural population) has been shown to be a viable coding strategy in olfactory bulb, somatosensory cortex, and other regions (Friedrich et al., 2004; Ratté et al., 2013; Lankarany et al., 2019; Kamaledin et al., 2022). Crucially, spiking synchrony allows populations to multiplex distinct sources of information or environmental variables. Thus, synchrony across populations or sub-populations may represent an additional axis on which a dynamical process evolves to represent information. In Manuscript IV, I compare two distinct sources of long-range synaptic inputs to serotonin neurons, finding differences in spike jitter that support a spike synchrony code for input identity at a population level. These spike timing differences appear to arise from cell-autonomous noise processes generated by stochastic ionic conductances, yet generate a coherent neural code at the population level. This suggests that population synchrony can represent a viable coding mechanism in midbrain circuits, and indeed can arise from the interaction of specific synaptic inputs with independent stochastic noise processes in individual neurons.

Thus, a significant portion of this thesis (Manuscripts III and IV) is concerned with forging links between meaningful population dynamics, on one hand, and the rules of synaptic transmission, on the other. These studies add to a body of literature showing that features of synapses, including short-term synaptic plasticity and input-specific synaptic properties, can impact computations and neural coding mechanisms at the population level in profound ways. Such structural principles may represent an important manner by which biological networks carry out functions.

2 The dorsal raphe nucleus

The dorsal raphe nucleus (DRN) is a highly interconnected, heterogeneous midbrain structure that contains the largest population of serotonin neurons in the mammalian brain (~9000 in the mouse). Far from being a passive relay structure, decades of work have revealed DRN to have substantial local interactions through complex microcircuitry, a highly organized anatomical structure, and heterogeneous genetically defined cell types. In contrast to other

monoamine systems which possess well-established theories enjoying decades of experimental support (such as the dopamine reward prediction error hypothesis), the serotonin field largely lacks a clear unifying theory of the behavioral and computational role played by this important neurotransmitter. This is made more complicated by reports of a diversity of task features communicated by spiking serotonergic neurons over multiple timescales, and of heterogeneous behaviors that have been shown to be impacted by manipulations of the serotonergic system.

Recent work has attempted to explain this diversity through the lens of multiple independent submodules in the DRN, defined through connectivity, anatomy and genetic markers, that mediate distinct behaviors and show unique responses during complex tasks (Commons et al., 2016; Ren et al., 2018; Okaty et al., 2019). Here, while considering the framework of heterogeneous submodules within the DRN, I will summarize the current state of knowledge about the anatomical and cellular organization of the DRN, its inputs and outputs, and computations performed by assemblies of spiking neurons during naturalistic behaviors. Additionally, I will discuss the nature of serotonergic release and transmission in central networks. Finally, I describe how fundamental progress in the serotonergic field can be made with a dynamical systems view of the DRN, and discuss two manuscripts contained within this thesis within this framework. This introductory section provides the framework for understanding the results in Manuscripts III and IV, which focus on novel coding properties within DRN.

2.1 Anatomical and cellular organization

Collectively, the 5-HT-producing nuclei in the midbrain and brainstem constitute the major source of serotonin in the mammalian brain. First described over half a century ago (Dahlstroem & Fuxe, 1964), the serotonergic nuclei consist of nine distinct structures spanning the brainstem up to the midbrain at the rostral end, traditionally distinguished by their developmental origin (rhombomere B1-B9; Jacobs & Azmitia, 1992). Of these, the DRN is the

nucleus which contains the majority of serotonergic neurons, consisting of B6 and B7 in the classical nomenclature. The DRN is located just below the periaqueductal gray and cerebral aqueduct, and extends in the caudal and ventral direction from the aqueduct, displaying substantial heterogeneity along these axes. The DRN is also located rostrally and dorsally to the median raphe nucleus (MRN), another serotonergic nucleus which some have argued shares key connectivity characteristics with the DRN B6 caudal cell group (Commons et al., 2016). While it is the primary source of serotonin neurons in the brain, the DRN in fact consists of multiple intermingled genetically identified cell types, as well as sophisticated anatomical organization and local interactions between cell types. The DRN can be divided into the dorsomedial, ventromedial, and lateral wing segments, all of which extend in the rostrocaudal dimension. The composition of cell types within each of these segments differs slightly, and work has attempted to comprehensively compare the segments in terms of composition, morphology and anatomy.

The DRN consists of several anatomically organized cell types, including serotonin neurons, GABAergic (somatostatin) neurons, dopamine neurons, glutamate neurons, and several other smaller groups. While the cellular organization of the DRN has been investigated for many decades, recent work has focused on quantifying diversity within the serotonin system, using modern intersectional genetic and transcriptomic tools (Okaty et al., 2015; Huang et al., 2019; Okaty et al., 2020). These studies have confirmed the presence of a diversity of cell types arranged in an anatomically defined manner, including prominent co-expression of neurotransmitters such as glutamate and serotonin, and have revealed distinct serotonergic subtypes with characteristic anatomical localization and gene expression patterns.

Serotonin neurons are the primary output neurons of the DRN. In the rostral parts of the DRN, serotonin neurons are primarily located along the midline, while in caudal sections they are located along both the midline and lateral wings (Daszuta & Portalier, 1985; Fu et al., 2010). There are between five and 14 main subclasses of serotonin neurons that exhibit defined

anatomical organization within the DRN, overlap with defined projector subtypes, and express characteristic ion channels, transcription factors, kinases and other genetic markers (Okaty et al., 2015; Huang et al., 2019; Okaty et al., 2020). These subclasses likely reflect developmental lineage, and partially correspond to hindbrain rhombomeric segments (Okaty et al., 2015). Morphologically, serotonin neurons display substantial diversity, likely relating to their anatomical localization and other features (Calizo et al., 2011; Andrade & Haj-Dahmane, 2013; Okaty et al., 2019). Serotonin neurons are typically small, with round, ovoid or semi-pyramidal shapes. They exhibit several primary dendrites that branch sparsely and extend primarily along the dorsoventral axis for midline neurons, and with less stereotyped organization for neurons in the lateral wings (Calizo et al., 2011). Axonal projections from DRN innervate cortical and subcortical structures, with single axons collateralizing frequently and with the total axonal length from such single neurons on the order of centimeters (approximately 10 times longer than the span of the mouse brain along its longest dimension (Ren et al., 2019). The subtypes of serotonin neurons, their connectivity, and their functional properties are discussed further in Sections 2.2 and 2.3.

Dopaminergic neurons form the next largest genetically distinct cell class within DRN, and are found along the midline and wings (Huang et al., 2019; Fu et al., 2010). These DRN dopamine neurons show physiological characteristics closer to dopamine neurons in ventral tegmental area, and can be divided into vasoactive intestinal polypeptide (VIP)-expressing and non-VIP expressing subtypes (Dougalis et al., 2012). The activity of these DRN dopamine neurons, highly modulated by social contact, may be involved in mediating a loneliness-like affective state (Matthews et al., 2016). This population of dopamine neurons also responds strongly to salient cues regardless of their valence, and may use these functional properties to regulate wakefulness or arousal (Cho et al., 2017).

GABAergic neurons (predominantly somatostatin) are the third largest class. These GABAergic neurons form local disynaptic inhibitory microcircuits with serotonergic neurons (Geddes et al., 2016; Zhou et al., 2017; Hernández-Vázquez et al., 2019), and are mostly

localized to the lateral wings of the rostral DRN, and along the midline and wings as the DRN extends caudally (Huang et al., 2019; Fu et al., 2010). A subpopulation of GABAergic neurons send long-range projections to forebrain, although these are mostly located in and around the medial longitudinal fasciculus (Bang et al., 2012).

There is a subpopulation of glutamatergic neurons in DRN that are localized along the midline and are especially concentrated in the ventral caudal DRN (Okaty et al., 2020). A large fraction of these neurons coexpress both the serotonin transporter gene SERT and vesicular glutamate transporter VGLUT3, suggesting that they corelease both serotonin and glutamate neurotransmitters at target sites (Fu et al., 2010; Huang et al., 2019; Okaty et al., 2020). This transmitter corelease has been shown functionally in amygdala, a major output site of this DRN glutamatergic subpopulation (Sengupta et al., 2017), and in ventral tegmental area and nucleus accumbens, where corelease may play a role in signaling reward (Liu et al., 2014; Wang et al., 2019). Co-expression of glutamate and serotonin neuron markers varies by region, with the medial and caudal regions showing the strongest co-expression (approximately 80% of 5-HT neurons show co-expression), and the lateral wings showing sparse co-expression (Hioki et al., 2010; Calizo et al., 2011; Okaty et al., 2020).

Sparse populations of other neurons are present within DRN, including those that express cholecystokinin, substance P and other neuropeptides (Fu et al., 2010). These comprise the smallest class of neurons by number in the DRN (Huang et al., 2019).

2.2 Organization of afferent and efferent connections

Although it is anatomically small, the DRN is a hub-like network, both receiving afferents and sending efferents to a wide variety of cortical and subcortical structures. Since the mid-20th century, research using retrograde and anterograde tracer techniques has attempted to highlight major inputs and outputs of this region. These efforts accelerated over the past decade, where researchers have used modern intersectional genetic tools and rabies tracing techniques to assemble comprehensive, brainwide maps of monosynaptic inputs and outputs

to defined DRN populations. Simultaneously, these connectivity maps have been used in concert with the known anatomical and cellular heterogeneity of the DRN (discussed in Section 2.1) to define major subsystems within DRN, or groups of neurons that are distinct based on their connectivity, anatomy, activation in various behaviors, and other properties. An exciting idea emerging from this body of work is that the functional and behavioral heterogeneity of ensembles of serotonergic neurons within the DRN might be reconciled through a consideration of defined subsystems of serotonergic neurons that work together to perform related, but distinct, affective functions.

DRN serotonin neurons receive extensive innervation from a wide range of cortical and subcortical structures. Afferents to the DRN region have been characterized using classical retrograde tracer approaches for decades, highlighting strong innervation by subcortical structures, including from lateral habenula, as well as various cortical structures (e.g. Aghajanian & Wang, 1977; Peyron et al., 1998; Lee et al., 2003). These classical tracer injection techniques, although invaluable, cannot distinguish between inputs to serotonergic and nonserotonergic neurons, which are intermingled within DRN.

With the use of modern intersectional genetics and rabies tracing approaches, inputs to serotonergic neurons were comprehensively characterized in a trio of papers that were published almost simultaneously (Dorocic et al., 2014; Weissbourd et al., 2014; Ogawa et al., 2014). From this work, it has emerged that serotonin neurons receive slightly more inputs from subcortical structures than from cortical structures. From cortical structures, the largest fraction of inputs to serotonin neurons are from somatomotor and somatosensory areas, infralimbic and prelimbic cortices, and orbitofrontal cortex. Subcortical structures providing strong monosynaptic innervation to serotonergic neurons include amygdala, lateral habenula, lateral hypothalamus and striatum. In the midbrain, ventral tegmental area (VTA) and periaqueductal gray strongly innervate serotonin neurons. Hindbrain (medulla and pons) additionally provides moderate levels of innervation to serotonin neurons. Intriguingly, these inputs appear to target anatomically segregated regions of DRN, with further work showing

that rostral DRN is predominantly innervated with areas involved in motivated behavior and motor output, and caudal DRN is innervated by septal and hippocampal-associated areas (Commons et al., 2015; Commons et al., 2016). Together, these studies show that serotonergic neurons integrate inputs from a wide variety of sources signalling affective and cognitive variables, and outline the broad anatomical segregation of inputs to rostral and caudal areas.

Several lines of research have also attempted to comprehensively characterize monosynaptic inputs to GABAergic neurons in DRN. These GABAergic neurons appear to receive inputs from a broadly similar set of brain areas that innervate serotonin neurons, including cortical and subcortical areas. Cortical inputs make up a statistically smaller fraction of inputs to GABAergic neurons than to serotonergic neurons, especially for insular, motor, orbital and prelimbic areas (Weissbourd et al., 2014). In contrast, most subcortical areas appear to provide approximately equal numbers of inputs to serotonergic and GABAergic neurons in DRN, with the exception of central amygdala, which innervates DRN GABAergic neurons comparatively more than serotonergic neurons (Weissbourd et al., 2014).

Perhaps the strongest lines of evidence for functional inputs to serotonergic and GABAergic neurons come from work combining single-cell electrophysiology and long-range optogenetic activation techniques. Collectively, this work indicates that serotonergic neurons and GABAergic neurons receive strong monosynaptic excitatory (glutamatergic) input from many long-range structures, including from lateral habenula, prefrontal cortex and basal ganglia (Weissbourd et al., 2014; Dorocic et al., 2014; Geddes et al., 2016; Zhou et al., 2017). These excitatory inputs contact large numbers of serotonin neurons in many cases: prefrontal cortex monosynaptically excites over half of serotonergic neurons, and lateral habenula excites close to 75% of serotonergic neurons (Zhou et al., 2017). Pharmacological interventions reveal that these long-range synapses exhibit postsynaptic AMPA receptor and NMDA receptor expression. Additionally, some pathways, including prefrontal cortex and lateral habenula, may recruit feedforward (disynaptic) inhibition onto serotonergic neurons through coincident activation of GABAergic neurons (Geddes et al., 2016; Zhou et al., 2017). Finally, inputs from

hypothalamus may provide direct (monosynaptic) long-range inhibition onto serotonergic neurons, as well as monosynaptic excitation (Zhou et al., 2017). This work has also strongly suggested subregion-specific inputs, with ventromedial DRN receiving the majority of input from lateral habenula and prefrontal cortex, and the lateral wings receiving input from amygdala and preoptic area (Zhou et al., 2017).

Serotonergic neurons in DRN also display anatomically and genetically defined projection classes that together innervate a large number of cortical and subcortical structures. Investigations into the organization of these projectors have revealed strong anatomical organizing principles that complement the findings of anatomically defined inputs to serotonin neurons, together suggesting that at least two major subdomains exist within DRN that are defined by their connectivity, as well as being largely functionally distinct. One subdomain consists of mainly VGLUT3-negative serotonergic neurons in the dorsal region of DRN, and preferentially targets subcortical areas (including strong innervation of central amygdala); another subdomain consists of predominantly VGLUT3-positive serotonergic neurons in ventral DRN and targets cortical areas (primarily orbitofrontal cortex) (Ren et al., 2018). Monosynaptic inputs to each subdomain are strikingly different, with the cortical-projecting domain (ventral DRN) receiving inputs from lateral hypothalamus, lateral habenula and medulla, while the subcortical-projecting domain (dorsal DRN) receives strong innervation from central amygdala, its main output target, forming a feedback loop, as well as from ventral bed nucleus of stria terminalis (BNST) (Ren et al., 2018). Furthermore, neurons from these subdomains show distinct collateralization patterns. Both cortex-projecting and amygdala-projecting neurons are highly collateralized and target multiple regions within their subdomain in complex branching patterns (Ren et al., 2019). These results support and extend previous work on projection-defined subgroups of serotonin neurons (Prouty et al., 2017; Muzerelle et al., 2016), suggesting that serotonin neurons are organized in distinct subdomains based on anatomy, connectivity, genetic markers, and functional properties.

Much of this work outlining key organizational principles of serotonin neurons relies heavily on genetic mouse lines that allows researchers to target particular genetically defined cell types. However, it should be mentioned that recent work has shown substantial variability between genetic mouse lines that have been used to target similar serotonergic neuron populations. The SERT-Cre mouse line (targeting the promoter for the serotonin transporter gene) and the ePET-Cre mouse line (targeting the enhancer of the transcription factor *Pet1*) are both employed as a marker of serotonergic neurons, yet the SERT-Cre line is far more selective for serotonin neurons, and accordingly this line reveals slightly different projection patterns of labeled neurons (Pinto et al., 2019). Therefore, researchers must carefully consider genetic mouse lines in any evaluation of anatomical segregation patterns in dorsal raphe nucleus.

A critical set of questions emerge from this body of influential recent work on the anatomical and cellular organization of DRN microcircuits. First, it is unclear how the multiple convergent, mainly excitatory inputs to each subsystem are integrated by single neurons or microcircuits. Indeed, while work has characterized a variety of single inputs to serotonin neurons (e.g. Geddes et al., 2016; Zhou et al., 2017), it has not been clear whether or how these diverse input sources are integrated into a coherent neural code that shapes serotonin output in target regions. For example, lateral habenula and anterior cortex both provide strong monosynaptic input to the orbitofrontal-projecting ventral DRN subregion, yet appear to signal quite different variables. While lateral habenula neurons classically signal negative reward prediction errors through acute changes in firing rate (Matsumoto & Hikosaka, 2009), prefrontal cortex neurons show diverse responses that signal higher-order task variables through a slowly evolving rate code (e.g. Funahashi et al., 1989, Bari et al., 2019). Such qualitatively distinct signals, communicated through synaptic input, must be integrated by DRN neurons in a manner that generates coherent neural codes communicated to downstream areas. Second, the heterogeneity and strong anatomical organizing principles present for serotonin neurons and other intermingled subtypes suggests that local circuitry might perform distinct computations on long-range inputs through network mechanisms. These network-level

computations require a dynamical systems approach to unveil, as well as precise measurements of synaptic dynamics and input statistics (covered in Sections 1.3, 1.5). Before examining potential answers to these questions (which form the basis for Manuscript III and IV), I will first give an overview of the spiking properties, neural computations and behavioral output of DRN serotonin neurons.

2.3 Spiking properties and neural computations by serotonin neurons

Although the activity of serotonergic neurons is generally modulated by salient stimuli, reliably relating the activity patterns of DRN serotonergic neurons to a set of sensory, cognitive or motor variables experienced by animals has proved to be a tremendously challenging endeavor. Serotonin neurons emit action potentials at low rates at rest, and their firing rates are generally modulated by phasic, salient stimuli such as rewards, punishments, predictive cues and motor variables, as well as by more long-term variables such as recent reward history or the current phase in the sleep-wake cycle. Additionally, learning can change their response profiles. A general theme is the heterogeneity of responses between serotonergic neurons, which may arise from the highly structured anatomical and cellular organization of the DRN, or may reflect either differential long-range input or local processing by cellular ensembles (Sections 2.1-2.2).

Early studies investigated the basic firing properties of serotonergic neurons during *in vivo* and *in vitro* conditions, demonstrating that these neurons displayed regular pacemaker-like firing at rates of several hertz, and moreover that these reliable firing rates could be altered by application of drugs or stimulation of peripheral nerve stimulation (Aghajanian et al., 1968, Aghajanian et al., 1978; Vandermaelen & Aghajanian, 1983). This pacemaker-like activity appears to rely on a prominent low-threshold (T-type) calcium conductance present in these serotonergic neurons, which allows reliable rebound spiking after the strong post-spike afterhyperpolarization (Burlhis & Aghajanian, 1987), although high-threshold (L-type) calcium conductances may also play a role (Asaoka et al., 2017). The firing rates of individual serotonin

neurons does not appear to exhibit an anatomical gradient, and is not correlated with neuron size (Mlinar et al., 2016).

Serotonergic neurons additionally exhibit changes in tonic firing during the sleep-wake cycle. In a seminal electrophysiological study in cat DRN, neurons (putatively serotonergic) were found to be most active during wakefulness, less active during slow-wave sleep, and the least active during rapid eye movement sleep (McGinty et al., 1976). Further electrophysiology and fiber photometry recordings in rodents have supported the idea that serotonergic neurons show decreases in tonic activity during sleep (Urbain et al., 2006; Kato et al., 2022). This work has additionally highlighted the diversity of responses of individual neurons, a subpopulation of which show paradoxical increases in their firing rates during sleep, potentially reflecting non-serotonergic neurons within DRN (Urbain et al., 2006). Slice physiology work has demonstrated that the tonic firing rate of serotonergic neurons can be increased by noradrenergic input in a dose-dependent manner, implying that locus coeruleus input during wakefulness may play a similar role *in vivo and* providing a putative mechanism for sleep-wake cycle changes in serotonergic neuron activity (Mlinar et al., 2016).

Serotonin neurons signal information about the predicted value of cues, and about transient stimuli such as rewards and punishments, through changes in their firing rate. Much of the evidence from this comes from head-fixed behavioral tasks in rodents or monkeys combined with single-unit electrophysiological recordings (Nakamura et al., 2008; Ranade et al., 2009; Bromberg-Martin et al., 2010; Liu et al., 2014; Hayashi et al., 2015; Cohen et al., 2015; Li et al., 2016; Matias et al., 2017; Zhong et al., 2017; Grossman et al., 2022). In typical operant or classical conditioning paradigms where sensory cues are paired with delayed rewards or punishments, serotonin neurons generally show phasic responses to predictive cues, and to the rewards and punishments themselves (Nakamura et al., 2008; Bromberg-Martin et al., 2010; Liu et al., 2014; Cohen et al., 2015; Matias et al., 2017). Individual serotonin neurons show a diversity of responses, and can be either excited or inhibited by the cue or unconditioned stimulus (Nakamura et al., 2008; Cohen et al., 2015). However, responses to the

cue and unconditioned stimuli appear to be correlated at a single-neuron level (Bromberg-Martin et al., 2010), and individual neurons signal the graded value of unconditioned or conditioned stimuli faithfully, e.g. small and large rewards (Nakamura et al., 2008; Hayashi et al., 2015, Cohen et al., 2015) (but see Paquelet et al., 2022). This cue-related activity typically develops slowly throughout training, while the uncued (reward or punishment) response is stationary over training (Zhong et al., 2017). Stress can additionally provide a transient modulation of these serotonergic responses when delivered during task blocks (Zhong et al., 2017). Thus, the activity of serotonin neurons reflects the presence of environmental stimuli with positive or negative valences, as well as cues predicting these stimuli.

The net polarity of serotonergic population responses to rewards or punishments is presently unclear, possibly reflecting response heterogeneity within DRN subpopulations. Fiber photometry recordings in head-fixed mice report that overall serotonergic activity decreases following rewards and increases following punishments (Matias et al., 2017), largely in line with reported population statistics from single-unit recordings (Cohen et al., 2015). Other work reports that for freely moving animals, sucrose rewards or social rewards can produce strong excitation in serotonergic neurons while punishments have no effect (Li et al., 2016). In this last work, heterogeneity between serotonin neuron responses was also observed, perhaps indicating that anatomy and connectivity play a role (Li et al., 2016). More recent work, using miniaturized microscopy techniques, found that the majority of serotonin neurons are excited by emotionally salient stimuli with either positive or negative valence (eg rewards, conspecific interactions and footshocks), although a subpopulation is inhibited for each stimulus (Paquelet et al., 2022).

These phasic reward and punishment-related changes coexist with more complex correlates of serotonergic firing rate that reflect higher-order task variables. During dynamic foraging tasks with blocks exhibiting changing reward values, the tonic activity of serotonin neurons changes, signalling graded value (Cohen et al., 2015; Hayashi et al., 2015). Additionally, recent work has demonstrated that serotonergic neurons track multiple forms of

task uncertainty over slow and fast timescales, highlighting the complexity of the response profile of serotonin neurons (Grossman et al., 2022). Responses to emotionally salient stimuli such as punishments, rewards and social interactions appear to show moderate adaptation over periods of minutes in approximately half of serotonin neurons (Paquelet et al., 2022). Finally, while much work has focused on how single-neuron changes in firing rate encode task variables, there are also indications that computations occur at the population level, such as the presence of synchronized activity of ensembles during stimulus presentation (Paquelet et al., 2022).

This body of work, over the past decade and a half, highlights the challenge in ascribing a single computational role to serotonergic neurons: they show diverse polarities and strengths of responses, yet generally respond in a predictive manner to upcoming salient events as well as to the events themselves in complex behavioral tasks. These diverse responses can be compared with those of dopaminergic neurons, for which there is a well-supported theory of reward prediction error signaling within a reward learning framework (Montague et al., 1996; Schultz et al., 1997). At least two potential frameworks could help explain serotonin neuron responses. One idea is that serotonin neurons reflect an as-yet-undescribed latent internal or task variable (or possibly multiplex several latent variables). A second possibility is that the heterogeneity in serotonin neuron response actually reflects sampling from multiple submodules within DRN, defined by their connectivity and anatomical location, that each reflects consistent environmental or task variables. The first idea requires uncovering a set of hidden latent variables that serotonin neurons represent. While some theories have been proposed, such as the beneficialness hypothesis of serotonin neurons (Luo et al., 2016; Liu et al., 2020), the experimental evidence for such ideas are not presently clear. Some support for the second idea of response heterogeneity comes from recent work studying the response profile of defined projection classes of serotonergic neurons. One study, after showing that cortical-projecting and subcortical-projecting neurons formed distinct submodules, investigated the activity patterns of each submodule. They reported that the cortical-projecting

submodule contained neurons that were mostly excited by phasic reward delivery and inhibited by phasic punishments, while the subcortical-projecting submodule was excited by both rewards and punishments (Ren et al., 2018). Importantly, responses were more consistent when each submodule was considered separately than when all serotonin neurons were nonselectively imaged. Similarly, other work has shown that ventral tegmental area (VTA)-projecting serotonin neurons are more excited by sucrose consumption, while neurons that project to the bed nucleus of the stria terminalis (BNST) are preferentially excited by footshock and the first interaction with a conspecific (Paquelet et al., 2022). Additionally, distinct types of neuronal responses to the same stimuli (eg excited vs inhibited, adapting vs non-adapting) appear to reflect activity of spatially intermingled subpopulations of serotonin neurons (Paquelet et al., 2022), a confounding finding in light of other results appearing to indicate strong anatomical segregation of distinct subnetworks (e.g. Ren et al., 2018). Together, this indicates that heterogeneity in serotonin neuron responses may in fact reflect a set of cryptically organized subnetworks that each have more consistent preferred responses and communicate to a limited subset of downstream targets.

Slice physiology work characterizing single-cell properties of serotonergic neurons has highlighted the strong heterogeneity of this system, as well as several important cellular properties that support key coding principles. In one recent study (Harkin et al., 2023), by utilizing a simplified single-cell modeling framework, we demonstrated that genetically identified serotonergic neurons and especially somatostatin interneurons exhibit highly heterogeneous cellular properties. Additionally, strong adaptation mechanisms, feedforward inhibition and a prominent hyperpolarizing potassium current heavily limit the excitability of serotonergic neurons to external inputs. This heterogeneity, when combined with network mechanisms such as specific long-range connectivity patterns, may partially explain the heterogeneous responses of neurons to task variables. Likewise, the limited excitability of these serotonergic neurons could explain why tonic firing rate changes appear to be a prominent feature by which the network represents task information.

2.4 Behavioral regulation by serotonin output

Similar to the response heterogeneity of serotonergic neurons, manipulations of DRN circuits have shown a diversity of involvement in multiple distinct behaviors. Serotonin has been ascribed a role in the regulation of patience, reward persistence, cognitive flexibility, direct reinforcement of ongoing behavior, regulation of learning rate (meta-learning), behavioral inhibition, direct locomotor modulation, and other roles. It is likely that serotonin regulates multiple behaviors through widespread efferents to cortex and subcortex, and that defined output pathways from DRN perform discrete behavioral roles. Despite this diversity, a set of common themes have emerged from several decades of work. First, serotonin generally inhibits locomotor processes (although not in all contexts), and can promote a state of patience or waiting for future expected rewards. Second, defined projector classes of serotonin neurons, notably those that innervate the central amygdala, are highly anxiogenic. Third, while the majority of the literature indicates that serotonin itself is not directly reinforcing, it appears to regulate the rate at which animals learn from experiences. These themes, largely emerging from work on manipulation of the serotonin system, complement above findings of how the activity of serotonergic neurons represent behavioral and environmental variables over multiple timescales.

Early work detailed at least two key hypotheses for the behavioral role of serotonin. One influential hypothesis was that serotonin was involved in behavioral suppression - for example, tolerating a delay before acting to obtain a future reward (Soubrié, 1986). This work focuses on the idea that serotonin may inhibit or withhold behavioral responses when a reward is expected but cannot be immediately obtained. The fact that serotonin depletion or lesions can lead to increased startle responses has been seen as supporting this behavioral inhibition hypothesis (Davis & Sheard, 1974). Another key early hypothesis was that serotonin is possibly implicated in generating adaptive responses to aversive stimuli, consistent with anatomical inputs from areas processing such aversive stimuli such as lateral habenula (Deakin, 1983; Deakin & Graeff, 1991). These two ideas of behavioral suppression and generation of responses to aversive

stimuli can be seen as tightly linked, and are consistent with some clinical work and early manipulations of serotonin levels in rodents (reviewed in Cools et al., 2008). This set of ideas also supports a model where serotonin and dopamine can act in opponency to promote avoidance and approach, respectively. This opponency model, while proposed decades ago (Deakin & Graeff, 1991), has seen more sophisticated formulations and support in the form of strong theoretical links with reinforcement learning models (Daw et al., 2002; Boureau & Dayan, 2011; Wert-Carvajal et al., 2022). In particular, these models generally assume that serotonergic neurons respond to negative reward prediction errors, and directly gate aversive learning (Daw et al., 2002, Boureau & Dayan, 2011), yet there is conflicting experimental evidence for this (see Section 2.3; Cohen et al., 2015, Li et al., 2016).

Recent serotonin manipulation studies have attempted to determine whether the behavioral inhibition hypothesis of serotonin is experimentally supported. One interpretation of this hypothesis is that serotonin should promote locomotor inhibition. Indeed, serotonin has been shown to phasically modulate various locomotor or muscular outputs. Serotonin stimulation phasically increases pupil diameter in a manner that appears independent from locomotor effects (Cazettes et al., 2021). Additionally, activating serotonergic neurons can transiently and rapidly inhibit locomotion without generating place preference or aversion phenotypes (Correia et al., 2017). Other work has showed that this inhibitory locomotor effect can be observed for both orbitofrontal cortex-projecting and central amygdala-projecting serotonergic neurons, hinting that it may be a general principle of serotonergic function independent of output pathway (Ren et al., 2018). However, the simple hypothesis that serotonin directly inhibits locomotor behavior in all contexts has been challenged by other work. Activating inputs from prefrontal cortex to DRN promote behavioral activation (kicking) during a forced-swim test (Warden et al., 2012), and similarly, activating the subset of serotonin neurons that project to orbitofrontal cortex can also increase escape kicking during a forced-swim test (Ren et al., 2018). These studies provide evidence against the simple behavioral inhibition hypothesis and indicate that a higher-order account of behavioral goals is necessary

to explain the action of serotonin. Additionally, this body of work demonstrates that particular inputs or outputs of the serotonin system may perform qualitatively distinct behavioral roles - a finding that dovetails with work demonstrating that activity patterns of serotonin neurons are segregated by projection target (see Section 2.3).

Accumulating evidence indicates that activation of serotonergic neurons consistently produces waiting behavior (patience) during a variety of tasks (Miyazaki et al., 2011; Miyazaki et al., 2014; Fonseca et al., 2015; Xu et al., 2017; Miyazaki et al., 2020). This work, focusing on the role of serotonin in promoting patience for delayed rewards, could be considered a modern refinement of the behavioral inhibition hypothesis. However, rather than emphasizing motor suppression of an ongoing behavior, these studies argue that the effect of serotonin is primarily cognitive: it increases the subjective confidence that a reward will be delivered. Indeed, the effect of serotonin activation on waiting is contextual on reward probability and timing uncertainty, and can be coherently explained by a Bayesian model of decision-making where serotonin increases the prior probability of reward delivery (Miyazaki et al., 2018). Consistent with this idea, Lottem et al. (2018) found that serotonin activation can promote active persistence (continuous nose-poking) in a foraging task, inconsistent with a simple locomotor inhibition hypothesis, but which instead supports a cognitive effect of serotonin for increasing subjective reward confidence.

Recent work has also attempted to determine whether there is experimental support for the hypothesis that serotonin may generate adaptive responses to negative valence stimuli, including generating an anxiety-like state. Global activation of DRN serotonin neurons has been shown to have mixed effects on anxiety-like states. Activating DRN serotonergic neurons optogenetically does not appear to modulate anxiety in an elevated plus-maze assay, while activating median raphe serotonergic neurons increases anxiety (Ohmura et al., 2014). Similarly, in an open-field assay, optogenetic activation of DRN serotonin neurons does not modulate measures of anxiety (Correia et al., 2017). However, specific projection targets of serotonin neurons have been shown to be anxiogenic. DRN serotonin neurons that specifically project to

central amygdala have been shown to promote an anxiety-like state using both an open field test and elevated plus maze assay, while orbitofrontal cortex-projectors do not effect anxiety (Ren et al., 2018). These results indicate that globally, serotonin activation likely does not promote anxiety, but specific projection pathways, including to central amygdala, may be anxiogenic.

Most work demonstrates that stimulation of serotonergic neurons is not directly reinforcing in a variety of tasks (Miyazaki et al., 2014; Fonseca et al., 2015; Correia et al., 2017). However, some work has demonstrated that activity of serotonin neurons can be directly reinforcing in certain circumstances (Liu et al., 2014). This has been shown to reflect corelease of serotonin and glutamate, as the mouse line used displays enriched targeting of this glutamatergic population. Thus, it appears that serotonergic output to target areas transiently increases patience for delayed rewards, yet is not reinforcing in most scenarios outside of corelease with glutamate.

While serotonin does not appear to be directly reinforcing, it may instead modulate the learning rate of animals, reflecting a meta-learning function for this neuromodulator in addition to its effect on patience and locomotion. The role of serotonin in meta-learning appears to be complex and highly contextual. For mice performing a reversal learning task, inhibition of serotonin neurons leads to perseverative responding after a rule switch, consistent with a decrease in learning rate (Matias et al., 2017). Additional work has shown that rodents performing a decision-making task will typically follow a win-stay, lose-switch behavioural strategy following short intertrial intervals, yet after long intertrial intervals their choices reflect average outcome history. Choices during these long intertrial interval periods are modulated by optogenetic activation of serotonergic neurons, which can be explained by an increase in learning rate (Iigaya et al., 2018). Similarly, in rodents performing a dynamic foraging task, the activity of serotonin neurons reflects two forms of task uncertainty that alter learning rate in a meta-learning framework, and manipulations of serotonin are consistent with a causal role of this neurotransmitter in meta-learning (Grossman et al., 2022). In these studies, serotonin is not

directly reinforcing, yet can play a dynamic role in adjusting how animals learn from their environment.

2.5 Features of serotonergic transmission in raphe and cortex

Serotonin is released at widespread sites located throughout the brain. These include both local release events within the serotonergic nuclei, including DRN, and long-range release in cortical and subcortical networks. Serotonin can bind to at least 15 known types of serotonergic receptors which have diverse intracellular effects and show cell type-specific expression throughout the brain. Here I will briefly discuss several features of serotonergic release in DRN and downstream targets.

Serotonin receptors can be divided into 7 main groups (5-HT1-7) which are expressed both centrally and peripherally, and are distinguished from each other based on homology and their intracellular effects (for reviews, see Barnes & Sharpe, 1999; Hoyer & Martin, 2002). All serotonin receptors are G-protein coupled (metabotropic) except 5-HT3 receptors, which are ionotropic. The G-protein coupled subtypes generally span 7 transmembrane domains and either couple to phospholipase C (PLC) or adenylate cyclase (AC). 5-HT1 receptors, of particular importance in this thesis, couple negatively to AC via Gi proteins, resulting in a diverse set of intracellular cascades. In particular, the activation of 5-HT1A receptors leads to the opening of G-protein coupled inwardly rectifying potassium (GIRK) channels, resulting in neuronal hyperpolarization over long timescales. 5-HT2 receptors signal through Gq proteins, activating phospholipase C and increasing intracellular calcium, leading to neuronal depolarization. 5-HT3 receptors, unique among this receptor family, are ligand-gated ion channels that conduct cations, leading to a depolarizing neural response over shorter timescales than the other receptor types. 5-HT4-7 receptors couple to adenylate cyclase via the Gs (5-HT4, 6, 7 receptors) and Gi/o (5-HT5 receptor) proteins.

Serotonin is released throughout the brain, where it binds to receptors that orchestrate a diversity of downstream effects. In prefrontal cortex and cingulate cortex, pyramidal neurons

and interneurons express distinct subclasses of serotonin receptors, hinting at the complex nature of serotonergic modulation of cortical circuits. Pyramidal neurons express both 5-HT_{1A} (inhibitory) and 5-HT_{2A} (excitatory) receptors, with substantial fraction of neurons (~80%) expressing both receptors (reviewed in Puig & Gullledge, 2011 and Celada et al., 2013). 5-HT_{2A} receptors are primarily localized on the apical dendrite of cortical neurons (Jakab & Goldman-Rakic, 1998; Martin-Ruiz et al., 2001) indicating they may participate in dendritic filtering processes. Parvalbumin (fast-spiking) interneurons also express either 5-HT_{1A} or 5-HT_{2A} receptors, although these are segregated to two distinct populations of neurons (Puig et al., 2010). A subclass of GABAergic neurons located in layers 1-3 uniquely express 5-HT₃ receptors, which provide fast (ionotropic) excitation upon serotonin binding (Puig et al., 2004). While a comprehensive summary of the physiological effects of serotonin on cortical and subcortical circuitry is outside the scope of this thesis, it is clear that serotonin can have a multitude of complex effects on sensory responses (e.g. Azimi et al., 2020), plasticity (e.g. Cavaccini et al., 2018), oscillatory activity (e.g. Puig et al., 2010), and other aspects of circuit function (for review, see Celada et al., 2013).

Within the DRN, serotonergic axon terminals are densely packed, indicating that serotonin is released locally in DRN microcircuits (Descarries et al., 1982; Chazal & Ralston, 1987). Two key concepts, volume release and inhibitory serotonin autoreceptors, have played a dominant role in guiding the field's thinking about how serotonergic microcircuits function. Early ultrastructural studies began to demonstrate that within rat DRN, serotonergic axon terminals can rarely be visualized in close apposition to either labeled somas or dendrites, leading to the hypothesis of paracrine or '*volume*' release of serotonin (Descarries et al., 1982; Chazal & Ralston, 1987; Fuxe & Agnati, 1991). This hypothesis, which contrasts with the classic spatially constrained synaptic (or junctional) model of neurotransmission for both glutamate and GABA, obtained further support by later studies in rat spinal cord where similar nonjunctional ultrastructural features of serotonin axons were observed (Ridet et al., 1994). However, other work found strong evidence of mostly (>90%) junctional, non-volumetric

transmission from serotonin terminals in substantia nigra (Moukhles et al., 1997). Thus, while it is likely that serotonergic axons can exhibit either junctional or non-junctional connections in a target-specific manner, the dominant conception has been non-junctional, volumetric release of serotonin. Around the same time, work emerged showing that most, if not all, serotonin neurons within DRN express inhibitory 5-HT_{1A} receptors, leading to the description of these receptors as *autoreceptors* due to their localization on the same type of neuron that produces their endogenous ligand (Sprouse & Aghajanian, 1987; Soletto et al., 1990). This implied that serotonergic neurons can possibly influence their own activity through local serotonin release. Together, these two influential ideas of volume release and 5-HT_{1A} inhibitory autoreceptors has led to a dominant feedback inhibition model of DRN function, where serotonin is released diffusely in DRN by the bulk activity of serotonergic neurons, and plays an autocrine role in regulating the overall activity of serotonergic circuits by binding to inhibitory 5-HT_{1A} receptors (e.g. Jacobs & Azmitia, 1992; Pineyro & Blier, 1999; Dankoski & Wightman, 2013; Best et al., 2020). This model, essentially proposing that serotonin enacts a global gain modulation function within DRN, has remained unchallenged despite new results indicating strong heterogeneity and submodule-specific computations within DRN (described in sections 2.1-2.3). It is possible that serotonin may play a more spatially precise role in modulating neural activity within DRN, which would seem to be consistent with the strong anatomical and functional heterogeneity present in this region. This is especially timely given fresh work in the dopamine system, classically thought to be characterized by volume transmission as well, that indicates that spatiotemporally precise signaling occurs in many scenarios (Liu et al., 2021).

Serotonergic transmission has been shown to be highly complex within DRN, and is influenced by both release and reuptake dynamics. Bulk serotonin in the extracellular space within DRN, measured with fast-scan cyclic voltammetry, appears to be unaffected by blocking reuptake dynamics or the frequency at which serotonergic axons are stimulated (Bunin & Wightman, 1998). These results were interpreted as supporting the volumetric release hypothesis for serotonin within DRN circuits. However, when recording intracellularly from

serotonin neurons within DRN and using 5-HT_{1A} receptor-mediated inhibitory postsynaptic potentials as a reporter for serotonin binding, a different scenario emerges: the duration of inhibition is affected by both serotonin reuptake blockers and serotonin diffusion rates, and the amplitude of inhibition is highly frequency-dependent, showing strong synaptic depression over timescales of seconds (Courtney & Ford, 2016). This indicates that serotonergic transmission is a dynamic process that likely involves spatially localized release and binding, and is strongly history-dependent. However, the precise manner in which stimulus history and release dynamics work together to influence microcircuit excitability is largely unclear.

2.6 A dynamical systems view of the dorsal raphe nucleus

From the past several decades of work, a complex and often confusing portrait of the DRN serotonin system emerges. Anatomically, there are distinct gradients of inputs from a diverse array of areas, and two key output pathways to orbitofrontal cortex (and other cortical regions) as well as to central amygdala (and other subcortical regions). Serotonergic, GABAergic, dopaminergic and glutamatergic neurons are intermingled within DRN and show anatomical gradients of expression, forming local microcircuits that likely perform a set of computations on input spiketrains. Functionally, serotonin neurons generally respond to salient stimuli or cues that predict them, including a prominent reward cue response and punishment response, and display changes in tonic firing over long timescales that reflect mean reward rate and potentially reward uncertainty. However, single serotonin neurons largely fail to display the kinds of responses that can be interpreted in the framework of a quantitative computational theory, contrasting with the dopamine system and reward prediction error signals in a reinforcement learning framework. Furthermore, manipulations of the serotonin system have been mostly unsuccessful in leading to a unified interpretation of what role this important neuromodulatory system plays during naturalistic behavior. Serotonin appears to promote patience for delayed rewards, slow locomotion in some contexts, control learning rate while not being directly reinforcing, and potentially increase anxiety through defined output pathways. To

arrive at a principled understanding of the functions of the serotonin system, a new approach may be needed.

Here, I argue for a dynamical systems approach for understanding the serotonin system. Of course, it is necessary to continue to uncover new anatomical and physiological organizing principles within the DRN, as refinements of our understanding of the serotonin system in the past decade have largely been driven by such advances (e.g. Weissbourd et al., 2014; Dorocic et al., 2014; Ren et al, 2018; Harkin et al., 2023). However, it is also important to consider these new biological features in the context of DRN circuitry as a dynamical system. Here, the activity of individual neurons is not seen in isolation, but as the result of a complex system (for example, a microcircuit with particular synaptic dynamics, input filtering characteristics and spike generation processes) whose activity evolves throughout time based on external inputs and internally generated activity (see Section 1.5 for a brief summary of this framework in the context of synaptic dynamics). The general approach with such a framework is typically to consider the evolution of population activity over time in a reduced set of dimensions that capture most of the population variance. Such an approach, sometimes applied to areas that can be modeled as recurrent networks *in silico*, has yielded coherent explanations of contextual decision-making processes in prefrontal cortex (Mante et al., 2013), neural control of motor movements (Churchland et al., 2012; Shenoy et al., 2013), reward representations in habenula (Sylwestrak et al., 2022), and behavioral state in hypothalamus (Nair et al., 2023), to cover several prominent examples. In some cases, these population-level dynamics can explain the reported diversity of single-cell responses that otherwise are difficult to interpret.

This thesis takes the view that considering the DRN as a dynamical system evolving over time, with measured synaptic or cellular properties that imbue it with qualitatively unique emergent properties, is key to understanding its function. Thus far, much of the important work on the DRN on the measurement side has focused on how single serotonin neurons represent task variables through changes in mean firing rate, or on the manipulation side, how

synchronously driving the majority of DRN serotonergic neurons through phasic optogenetic stimulation (and largely independent of projection target or input statistics) can influence some measured aspect of behavior in an interpretable way. Such approaches are clearly invaluable and have led to substantial progress in our understanding of the functions and representational abilities of DRN serotonin neurons.

However, it is also important, I contend, to consider activity patterns that can emerge at a population level through dynamical interactions, for several reasons. First, given the modular architecture of DRN and complex anatomical subdivisions, dynamical interactions between these architectural elements may turn out to be an important consideration. These interactions could thus help explain elements of the heterogeneity of neural responses, or time-varying neural responses during a task that have network origins. Second, the focus on mean firing rate neglects temporal coding schemes - for instance, activity correlations between serotonin neurons, which could be a rich source by which DRN microcircuits represent internal or external information. In fact, recent cutting-edge work demonstrates that such activity correlations are a prominent feature of the response pattern of serotonin neurons (Paquelet et al., 2022). Dynamical analyses are uniquely suited to expose these temporal coding schemes in DRN. Third, taking a neuroethological perspective, sources of long-range synaptic input onto serotonin neurons each exhibit unique statistics which may influence how DRN microcircuits respond. The type of input received (eg synchronous or asynchronous, low or high frequency, etc.) may alter microcircuit dynamics in nonlinear ways given the anatomical complexity of the circuit. For example, some input regimes may trigger feedforward inhibition by GABAergic neurons onto serotonin neurons, while other input regimes may mostly result in direct excitation of serotonin neurons. By considering the DRN circuit as a complex system with intrinsic dynamics, we will be in a position to map defined input statistics to unique regimes of population activity.

Manuscripts III and IV utilize a dynamical perspective to infer key functions of DRN circuits. In particular, they first focus on quantifying key novel properties of the system (often

synaptic or intracellular), and next to use these as building blocks to a dynamical systems view of DRN. In Manuscript III, I uncover new aspects of the physiology of DRN (inhibitory recurrent connections between serotonin neurons, with strong short-term facilitation dynamics) and ask how these components change the response of DRN microcircuits to long-range excitatory inputs from lateral habenula through large-scale networks simulations. By focusing on population dynamics both within and between defined submodules, I demonstrate that the observed recurrent synaptic dynamics generate strong winner-take-all rules that depend on input statistics. These results make unique behavioral predictions for threat processing that I test behaviorally. In Manuscript IV, I ask how spike timing of serotonin neurons depends on long-range excitatory inputs from prefrontal cortex and lateral habenula. This investigation revealed significant spike timing differences between the two inputs that supports a synchrony code for input identity at the population level. By uncovering a novel intracellular calcium conductance and constructing a dynamical model of intracellular membrane voltage, this manuscript provides a mechanistic understanding of how synchrony can provide a viable coding mechanism in midbrain circuits.

Preface to manuscripts

Here, I present six manuscripts that were written during the course of my doctoral studies. The manuscripts were formatted based on the standards of the journals that they were submitted for, and therefore may exhibit small differences in structure.

Manuscripts I and II provide a validated solution for accurately estimating silent synapse fraction in central circuits through a novel synthetic likelihood method, which relies on an experimental simulator to correct estimates bias and variance. They were published in *Cell Reports* (2020) and *STAR Methods* (2020), respectively, and are presented here in their published form.

Manuscripts III and IV focus on a dynamical systems perspective of the dorsal raphe nucleus, with an emphasis on how synaptic dynamics generate population-level computations. Manuscript III was reviewed by *Nature Neuroscience* and we are planning to submit revisions in the upcoming three months. Manuscript IV is currently in preparation for submission.

Manuscripts V and VI are perspective pieces on two important new findings in synaptic physiology. They were published in *Journal of Neuroscience* (2017) and *Journal of Physiology* (2022) respectively, and are presented here in their published form.

Manuscript I**A Synthetic Likelihood Solution to the Silent Synapse Estimation Problem**

Michael B. Lynn, Kevin F.H. Lee, Cary Soares, Richard Naud*, Jean-Claude Béïque*

*: Corresponding authors.

This manuscript was published in *Cell Reports* in 2020:

Lynn MB, Lee KFH, Soares C, Naud R, Béïque JC. A Synthetic Likelihood Solution to the Silent Synapse Estimation Problem. *Cell Rep.* 2020 Jul 21;32(3):107916. doi: 10.1016/j.celrep.2020.107916. PMID: 32697998.

Statement of contribution

For this work, I performed all simulations and experiments, performed formal analysis, assembled all figures, wrote all manuscripts drafts, and edited all drafts prior to submission in conjunction with the other co-authors. Jean-Claude Béïque and Richard Naud guided the project, edited drafts, and provided essential high-level input for the paper. Richard Naud gave input and guided the formal analysis.

I am especially indebted to Kevin Lee and Cary Soares, who were senior graduate students in the lab at the time of writing this paper. They were instrumental in the early planning phase of this project and had already identified several key issues with failure rate-based quantification of silent synapses before this manuscript was conceived. Their guidance and knowledge were essential for the creation of this manuscript.

Summary

Functional features of synaptic populations are typically inferred from random electrophysiological sampling of small subsets of synapses. Are these samples unbiased? Here, we developed a biophysically constrained statistical framework for addressing this question and applied it to assess the performance of a widely used method based on a failure-rate analysis to quantify the occurrence of silent (AMPA-lacking) synapses. We simulated this method *in silico* and found that it is characterized by strong and systematic biases, poor reliability and weak statistical power. Key conclusions were validated by whole-cell recordings from hippocampal neurons. To address these shortcomings, we developed a simulator of the experimental protocol and used it to compute a synthetic likelihood. By maximizing the likelihood, we inferred silent synapse fraction with no bias, low variance and superior statistical power over alternatives. Together, this generalizable approach highlights how a simulator of experimental methodologies can substantially improve the estimation of physiological properties.

Keywords: silent synapses, plasticity, statistical inference, synthetic likelihood

Introduction

Activity-dependent synaptic plasticity has captured global research attention for well over four decades due to its hypothesized role in learning and memory (Bliss and Lomo, 1973). Silent synapses contain only NMDA receptors, not AMPA receptors, and represent preferential sites for receptor insertion during long-term potentiation (LTP; Liao et al., 1995, Isaac et al., 1995). The insertion of AMPARs is generally accompanied by an increase in spine size (structural plasticity; Soares et al., 2017, Lee et al., 2016, Matsuzaki et al., 2004, George et al., 2015, Harvey et al., 2008), linking structure and function (For mechanistic details, see Zito et al., 2009, Harvey and Svoboda, 2007). The relative proportion of silent synapses in neural circuits is believed to be a fundamental determinant of the plasticity potential of neural networks. Silent synapses are more prevalent in developing networks (Soares et al., 2017, Lee et al., 2016, Durand et al., 1996, Isaac et al., 1997) yet continue to play a role in adult brain circuits for instance in newborn neurons (Chancey et al., 2014), and mature synapses can at times undergo ‘silencing’ by removal of AMPARs (Abramsson et al., 2007). Furthermore, addictive drugs such as cocaine and morphine have been suggested to induce de novo formation of silent synapse in reward-related pathways of adult animals (Huang et al., 2009, Brown et al., 2011, Graziane et al., 2016, Ma et al., 2016). More generally, theoretical work has shown that high fractions of ‘potential’ synapses (e.g., silent) are associated with optimal information storage over a range of conditions (Brunel et al., 2004, Brunel, 2016, Clopath and Brunel, 2013). Together, whereas these observations point to the fundamental role of silent synapses in both developing and mature brain circuits, they also emphasize the corollary challenge of quantifying with precision their occurrence in synaptic populations with experimentally realistic throughput.

While sound methods exist to demonstrate the presence of individual silent synapses (Fig. 1Ai-ii) (Liao et al., 1995, Isaac et al., 1995), these binary single-synapse classification methods can only estimate the fraction of silent synapses in a population by pooling across repeated single-

synapse interrogations either using a minimum stimulation paradigm (Durand et al., 1996, Isaac et al., 1997), or glutamate uncaging techniques (Soares et al., 2017, Lee et al., 2016). In contrast, multi-synapse classification methods, such as the failure-rate analysis (FRA) method, have been developed to provide an estimate of the fraction of silent synapses in a larger synaptic population for every experimental attempt: here an experimenter records from a small, but unknown, number of synapses, and the FRA formalism estimates the proportion of these which are silent by means of a comparison of failure rates (Fig. 1Bi-ii) (methodology reviewed in Graziane and Dong, 2016). A number of studies have used this method to quantify changes in silent synapse fraction, notably in rodent models of drug addiction, but also during critical periods of plasticity in the visual cortex and hippocampus (Huang et al., 2009, Brown et al., 2011, Graziane et al., 2016, Ma et al., 2016, Neumann et al., 2016, Huang et al., 2015, Ma et al., 2014, Lee et al., 2013, Koya et al., 2012, Rumpel et al., 1998, Beroun et al., 2018, Wright et al., 2020). Yet, although FRA seeks to scale up the throughput of silent synapse estimation, its scalability hinges on its reliability and accuracy (especially compared with traditional binary classification methods), which have not been investigated with targeted experimental and statistical analyses.

Here, we first employ a combination of biophysically constrained numerical simulations and electrophysiological recordings to examine the performance of existing methods for silent synapse detection and quantification. Our analysis reveals substantial and general limits on the reliability and validity of the silent synapse fraction estimates obtained using the FRA formalism. To overcome these inherent limitations, we build on approaches developed in the dynamical systems field and employed our computational model to generate an approximate maximum-likelihood estimator. This approach provides dramatic improvements in estimator bias and variance. A power analysis reveals that the enhanced estimator can resolve fine-grained changes in silent fraction with experimentally feasible sample sizes. Our estimator does not require a novel set of observations and can thus be deployed on existing FRA

datasets. These tools and findings make large scale interrogation of changes in silent synapses possible and reliable.

Results

While minimum stimulation experiments have been critically important for demonstrating the existence of silent synapses, they do not formally allow estimation of the true silent synapse fraction (s_t) contained in a synaptic population. A simple experimental technique, the failure rate analysis (FRA), has been developed to provide a quantitative estimate (\hat{s}_{FRA}) of the true underlying silent fraction, s_t . In this method, the rate of synaptic failures at two holding potentials (-60 mV and +40 mV) are used as inputs to a mathematical equation (see STAR Methods for derivation), which returns an estimated fraction of silent synapses. Despite this formalism, it is unclear how biophysical variables such as stochastic neurotransmitter release and the inherent variability in the number of synapses stimulated between experiments shape the statistical power and discriminability of this method. Here, we use a set of increasingly constrained and experimentally realistic numerical simulations and statistical analysis to examine the performance and general usefulness of this estimator.

A. Nonsilent synaptic populations and FRA estimator performance

As a first step, we directly determined the extent to which the stochasticity of synaptic release affect the FRA estimator. We thus numerically simulated the simplest scenario where a single, non-silent synapse with probability of release $Pr = 0.5$ is electrically stimulated (Fig. 1Bii, inset). Henceforth we will refer to non-silent synapses as active synapses (i.e., synapses containing both AMPA and NMDARs). We simulated Bernoulli-distributed trials (failures and successes) in 50-sweep epochs. The FRA technique, applied to a single, active synapse (i.e., $s_t = 0$) should return $\hat{s}_{FRA} \approx 0$. With 50 sweeps, the distribution of \hat{s}_{FRA} across multiple experiments (Fig. 1C, blue) was strikingly broad (S.D. = 31.3%), non-Gaussian, had high

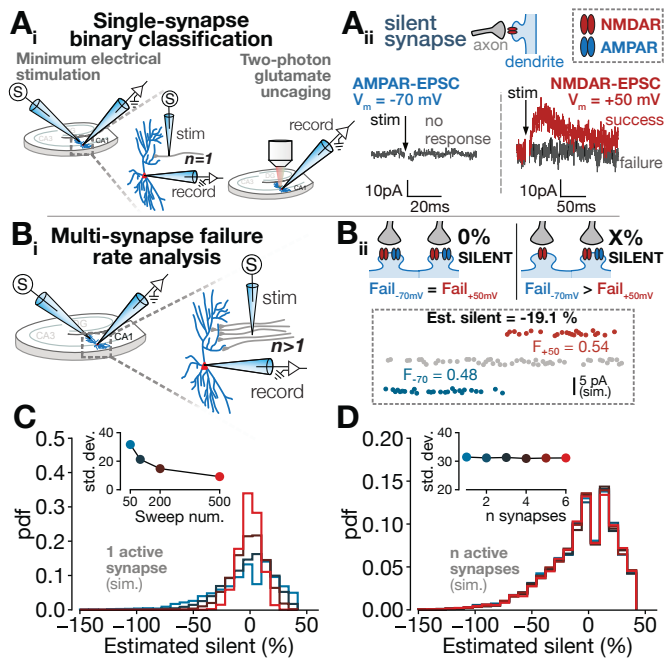


Figure 1: The FRA technique returns high-variance estimates under conditions with no silent synapses present. A_i: Single-synapse classification techniques rely on electrical (left) or optical (right) stimulation of a single synapse. A_{ii}: The responses (here averaged electrical traces), are classified as successes and failures at -70mV (left) and 50mV (right). B_i: In contrast, the multi-synapse failure rate analysis method relies on stimulating multiple synapses and (B_{ii}) comparing failure rates at two potentials to estimate silent fraction (inset shows example calculation). C: FRA estimate distributions for a single active neuron, as the number of sweeps per voltage condition is varied (inset, colors indicate sweep num.). D: FRA estimate distributions for varying numbers of active synapses (inset, colors indicate synapse num.), with mean failure rate held constant at 0.5.

kurtosis (kurtosis = 3.00) and a prominent left tail (skewness = -1.01). These attributes therefore make the FRA estimator unsuitable for the traditionally-used parametric tests that assume normality. Furthermore, 45.3% of the \hat{s}_{FRA} distribution falls below zero, which are biologically irrelevant estimates. Since negative silent synapse values are conspicuously absent from published datasets using FRA, it is possible that these negative values returned from the equation were arbitrarily set to 0% silent (i.e., systematically ‘zeroed’). This procedure would artificially inflate the mean silent synapse estimate.

We next asked whether the FRA estimator's accuracy could be improved by simply increasing the number of sweeps from the recommended value of 50 (Graziane and Dong, 2016) (Fig. 1C). Increasing the number of sweeps per epoch to 500 still produced a remarkably variable estimate (Fig. 1C inset; 500 sweeps: S.D. = 9.1%). Alternatively, recordings from multiple synapses may improve the reliability of \hat{s}_{FRA} , due to the intuitive averaging out of stochastic binary events. We thus numerically simulated FRA experiments on small numbers of synapses, first analytically calculating a fixed Pr for all synapses such that the mean failure rate was 50% (Fig. 1D). Surprisingly, \hat{s}_{FRA} variability was virtually identical over the entire range of active synapse numbers (Fig. 1D, inset). Therefore, in the simple case of non silent populations, the variability of the FRA estimates is not markedly improved when monitoring multiple synapses.

We considered the generalizability of our findings by asking how \hat{s}_{FRA} estimates varied over a wide Pr range and ensemble synapse number (Fig. S1). Our analysis revealed that the FRA estimator variance, surprisingly, is inversely proportional to synapse determinism (i.e., as $Pr \rightarrow 0^+$ and as $Pr \rightarrow 1^-$) and does not significantly decrease with any combination of Pr and synapse number. Taken together, these findings outline the general inability of the FRA method to return reliable (i.e., low variance) estimates of silent synapse fraction in simple conditions with no silent synapses present ($s_t = 0\%$).

B. Silent synapse-containing synaptic populations and FRA estimator performance

Under conditions where silent synapses are present ($s_t > 0\%$), the quantitative accuracy of the FRA method is naturally dependent on the particular arrangement of synapses sampled. It is not intuitively obvious whether traditional electrical stimulation selects an unbiased subsample of a larger synaptic population, either in terms of silent fraction or of individual synapse properties such as release probability. Therefore, before assessing the accuracy of the FRA equation, we first developed an obligatory intermediate method to simulate experimentally feasible ‘draws’ of synapse sets, with varying release probabilities, from larger mixed populations (Fig. 2A).

Briefly, we started with initial sets of 100 synapses with some true fraction silent (s_t). Release probabilities were uniformly distributed ($Pr \sim U(0,1)$). We next simulated as faithfully as possible a typical FRA electrophysiological experiment. Thus, the initial set of activated synapses was subjected to multiple rounds of stochastic synapse elimination (10% per round) corresponding to a decrease in electrical stimulation intensity that an experimenter typically performs during an FRA experiment. The rounds of elimination were ended when the failure rate of active synapses reached between 20% and 80%, as dictated by described experimental guidelines (Graziane and Dong, 2016). Synapse sets were discarded if they did not reach this criterion before the elimination of active synapses. The simulation thus returned stochastic, small subpopulations of silent and active synapses, each with an associated release probability, and with the failure rate range as the main constraint.

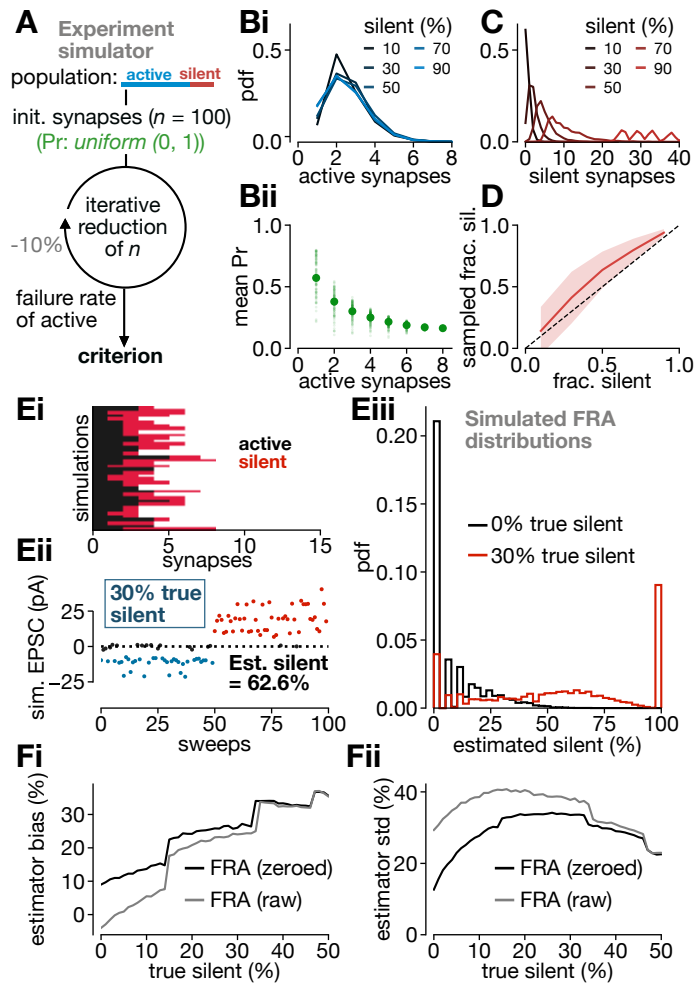


Figure 2: Experimentally bounded simulations reveal a fundamental bias inherent in bulk electrical stimulation. **A:** Schematic depicting experimental simulator approach. **Bi:** Probability distribution of number of active synapses obtained during each simulation. **Bii:** Release probability of active synapses. Small dots denote release probability for each simulation, while large dots denote the mean across simulations. **C:** Probability distribution of number of silent synapses obtained during each simulation. **D:** Sampled fraction silent is shown against the true fraction silent. **Ei:** Depiction of returned samples of silent/active synapses. **Eii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction. **Eiii:** Distribution of FRA estimates for two true silent values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii:** FRA estimator standard deviation for non-zeroed and zeroed distributions.

We found that the number of active synapses in each subpopulation reaching the selection criterion was very small - typically between 1 and 5 (Fig. 2Bi). As one may intuit, large numbers of synapses in a draw are rare because they require very low Pr for every synapse to achieve the desired failure rate (Fig. 2Bii). This experimentally-derived selection criterion therefore leads to a bias towards sampling synapses with low Pr . The number of silent synapses in the draws scaled with s_t , as expected (Fig. 2C). However, the sampled subsets were consistently enriched in silent synapses compared to the ground-truth population (Fig. 2D). This intriguing bias arises from the simple fact that the failure-rate criterion de facto solely depends on active synapse monitoring: therefore, the last elimination round ineluctably eliminates an active, rather than a silent, synapse. On average, the number of active synapses is therefore under-represented in the sample, which leads to a quasi-systematic overestimation of the silent fraction. Given that the small size of the sampled synaptic population (one to five synapses), this systematic bias is consequential. This appears to be a general feature of the FRA method when stimulation amplitude is adjusted to reach a particular failure rate.

We next considered the performance of the FRA estimator under this experimentally realistic scenario. Individual simulations returned highly variable synapse sets (Fig. 2Ei) and inaccurate estimates (\hat{s}_{FRA}) (Fig. 2Eii). Over multiple simulated experiments, the estimated silent synapse fraction showed distinctly non-normal distributions (Fig. 2Eiii). We next undertook a formal analysis of the FRA estimator's bias and variance. The bias of \hat{s}_{FRA} increased with the true silent fraction, s_t , and this occurred irrespective of whether negative estimates were erroneously set to zero ('zeroing') or not (Fig. 2Fi). The standard deviation of the estimate was high across all silent fractions (Fig. 2Fii), reaching a maximum of 40% with zeroing applied. We conclude that the FRA estimator displays a high positive bias and large variance, and thus is a flawed estimation technique.

While we drew the release probability of individual synapses from a uniform distribution following previous work assessing Pr from iGluSnFR-mediated single-synapse release events (Soares et al., 2017), we additionally asked whether different release probability distributions of a synaptic population would influence the FRA estimator's performance. First, we considered the gamma distribution (parameters: $\lambda = 5.8; n = 3$) reported from staining-destaining experiments (Fig. 1G in Branco et al., 2008). This distribution produced negligible changes in both the synapse selection bias, and in the bias and variance of the estimator (Fig. S2A-F). We additionally simulated the case with extremely small release probabilities ($\lambda = 5.8, n = 1$) (Fig. S2G-L). Although the mean number of active synapses rises from ~ 3 to ~ 10 (Fig. S2Hi), the estimator with zeroing applied still shows comparable bias (27% bias at 50% silent, compared to 35% bias at 50% silent for a uniform distribution) and variance (peak standard deviation of 30%) (Fig. S2Li-ii). We finally considered the hypothetical, but possible, scenario where release probabilities would be unequal between silent and active synapses (Fig. S3). These simulations show persistently high estimate variance (Fig. S3Fi; S3Mi), and an estimate bias whose magnitude depends on the release probability mismatch between active and silent synapses (Fig. S3Fii; S3Mii). However, it is unclear whether there is a strong mismatch between release properties at active and silent synapses, and furthermore, this scenario would violate a key assumption of the FRA formalism and thus invalidate its use. These results demonstrate the significant propensity of the FRA estimator towards large positive biases and large variances over a wide range of possible synaptic release probability distributions.

C. Electrophysiological verification of numerically simulated FRA estimate inaccuracies

We next sought to gain an experimental validation of the inherent propensity of the FRA approach to return variable estimates of silent synapse fraction. We thus performed voltage-clamp recordings from CA1 neurons from p14-15 mice while electrically stimulating Schaffer's collateral synapses (Fig. 3A; see STAR Methods). As per recommended practice for typical FRA

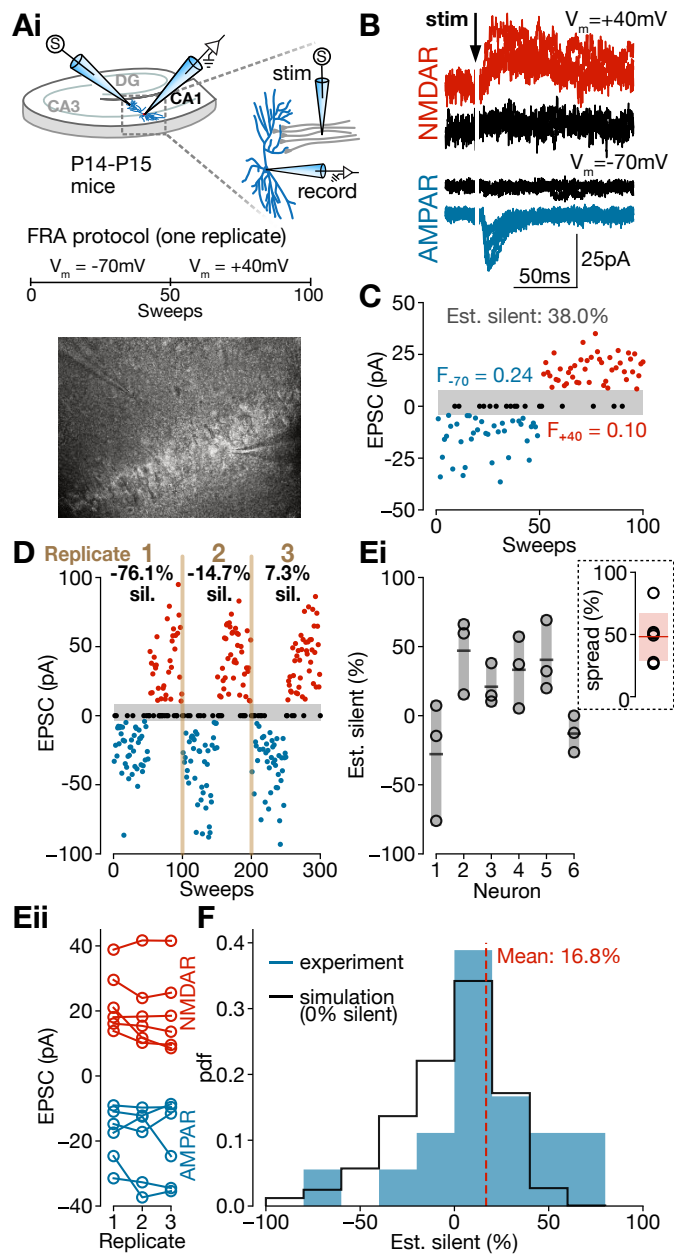


Figure 3: Electrophysiological recordings at CA3-CA1 synapses are consistent with simulated FRA distributions. **A:** Experimental schematic illustrating recording from CA1 neurons while stimulating CA3 afferents. **B:** Representative traces showing successes and failures at $V_m = +40\text{mV}$ (top) and $V_m = -70\text{mV}$ (bottom). **C:** Exemplar experiment showing failure rates and estimated silent fraction. **D:** Three replicates of an FRA experiment on a single neuron returns highly variable estimates consistent with biophysical limitations of the FRA method. **Ei:** Replicates for each neuron (circles) and the mean of these estimates (bar). Inset depicts estimate spread for each neuron. **Eii:** EPSC amplitude stability across replicates for each neuron. **F:** Histogram of all replicates (blue) overlaid with simulated histogram from null distribution (0% silent, black). Distributions are significantly different ($p = 0.049$, Kolmogorov-Smirnov two-sided test).

experiments, we tuned the stimulation intensity such that we obtained clear monosynaptic currents with ~50% failure rates (Fig. 3B,C). We performed FRA experiments in triplicate for each recording from individual CA1 neurons (Fig. 3D), and we verified that the evoked EPSC amplitudes at both -70mV and +40mV were stable between these successive iterations (Fig. 3Eii), confirming that we were stimulating the same subsets of synapses during each FRA iteration.

Since we performed the FRA experiments in triplicate for each recordings, we computed three independent \hat{s}_{FRA} estimates of the same synaptic population. As predicted from our simulations, these estimates were highly variable within the same neuron (Fig. 3Ei), and often returned biologically irrelevant negative values (Fig. 3F). Because we cannot directly measure s_i , we are unable to formally compare the bias and variance of the estimator with experimental ground truth, but these results nonetheless demonstrate the strong agreement between our computational model and experimental findings in documenting the fundamental limits of the FRA formalism.

D. An improved estimation technique based on synthetic likelihood functions

In principle, it should nonetheless be possible to extract reliable information about the true state of a synaptic population from unreliable and biased estimates returned by the FRA method. For instance, if we knew the distributions of failure rates at both hyperpolarized and depolarized potentials for each s_i , we could write expressions for likelihood functions and construct a maximum-likelihood estimator (MLE). The logic of this approach is to first determine the analytical function establishing the probability of an observation given some model parameters, and then to find the model parameter (here the silent synapse fraction) that maximizes this likelihood function. This type of estimator is of particular interest here since MLE minimizes the variance, at least asymptotically (Cramér, 1945).

In the present case, however, the likelihood functions are analytically intractable, as they result from a complex combination of the sampling protocol and the stochastic properties of synapses. Some simplifying assumptions could be made to reach a closed-form likelihood function, but this would require assuming that release probability distributions are independent of sampled synapse number, and that the sampled subsets themselves reflect unbiased samples from the population (both assumptions are shown by our model to be false). Alternatively, popular approaches for noisy dynamical systems sidestep analytical solutions and rather construct synthetic likelihood functions from Monte Carlo simulations (Wood, 2010, Sunnåker et al., 2013, Diggle and Gratton, 1984). Here, the dynamical system is iteratively simulated and summary statistics are extracted, allowing to likelihood-based estimation from a stochastic model. These likelihood function are approximations only in the sense that they are based on a finite number of stochastic simulations, but will converge to the true likelihood function. We use this formalism to construct an estimator, FRA-MLE, from synthetic likelihood functions (Fig. S4A-B) where the summary statistic is the biased FRA estimate itself (\hat{s}_{FRA}) (See STAR Methods for more details on why this summary statistic was chosen). Since our noisy sampling model already incorporates technical choices made by experimenters (e.g. the range of acceptable failure rates), this allows us the flexibility to construct estimators, $\hat{s}_{FRA-MLE}(\boldsymbol{\theta})$ which are a function of an experimental parameter vector $\boldsymbol{\theta}$ containing methodological variables and synaptic properties.

In contrast with the large bias found for the FRA estimator, we found that $\hat{s}_{FRA-MLE}$ is entirely unbiased (Fig. 4A, Fig. 2Fi). Also in contrast with the large variance found for the FRA estimator (Fig 2Fii), we found that although the variance of $\hat{s}_{FRA-MLE}$ increases with ground truth silent fraction, it remains lower than that of \hat{s}_{FRA} (Fig. 4B; Fig. S4C). Since the FRA-MLE method is based on a flexible experimental simulator, it can correct for assumption violations such as unequal release probabilities between silent and nonsilent synapses (Fig. S3). Applied to these

conditions, FRA-MLE still displays no measurable bias (Fig. S3Gi; S3Ni) and smaller variance (Fig. S3Gii; S3Nii) than the original FRA estimator. Therefore, the new FRA-MLE method circumvents the inherent sampling bias and high variance of the original \hat{s}_{FRA} .

As a final test, we applied FRA-MLE to the experimental values collected from hippocampus (Fig. S4D-F). We used FRA-MLE to arrive at a maximally likely estimate, $\hat{s}_{FRA-MLE}$, across all experimental observations. We found that the joint likelihood function was sharply tuned, peaking at a maximal likelihood estimate of 11.4% silent for our experiments (Fig. S4E). This was substantially lower than the estimate obtained by simply taking the mean of all experimental FRA values (16.8%), in agreement with our findings of a pervasive positive bias for FRA. We additionally verified our estimator and numerical simulation approach as follows: We employed our statistical sampling model (i.e., successive rounds of synapse elimination) to construct a hypothetical experimental distribution of FRA estimates for the maximally likely silent fraction $\hat{s}_{FRA-MLE} = 11.4\%$ obtained for our experiments. We found that this model-generated distribution was not significantly different from the experimental distribution (Fig. S4F; $p = 0.99$; KS-test, two-tailed). Together, these findings reiterate the general propensity of FRA to return positively biased estimates, and demonstrate the ability of FRA-MLE to extract not just locally optimal solutions, but solutions which are also in good quantitative correspondence with experimental distributions.

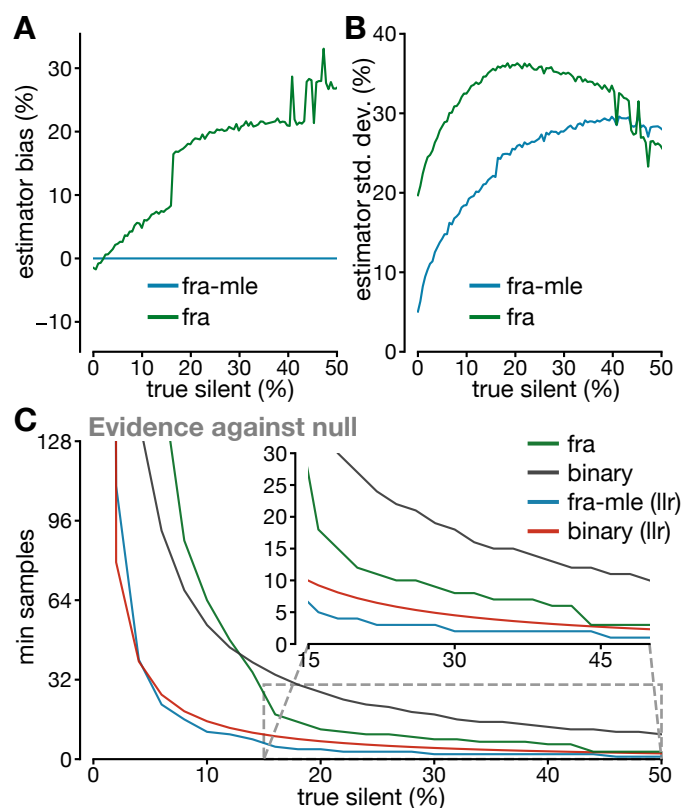


Figure 4: Statistical properties of FRA-MLE estimator compared with other methods. Shown are FRA -MLE (blue) and FRA (green) estimator bias (A) and standard deviation (B). C: Minimum samples required for each estimator to discriminate true silent fractions from a null population where no silent synapses are present. Inset depicts a zoomed view of 15-50% silent (for all cases).

E. A power analysis quantifies performance improvements of the FRA-MLE estimator

To achieve a high-throughput characterization, one needs to take two sets of samples from either of two separate conditions, and conclude with an experimentally realistic whether the two conditions do not have the same proportion of silent synapses (i.e. reject the null hypothesis). We used the framework of power analyses to ask whether our FRA-MLE estimator compares favorably with other methods, classic FRA as well as single-synapse binary classification methods, including minimum stimulation experiments as well as more modern 2-photon glutamate uncaging methods. We calculated the minimum sample size, n_{min} , required to discriminate silent-containing populations from an active-only population utilizing each technique. Depending on the estimator, calculating n_{min} is done by either analytically or

numerically calculating the minimum sample size to achieve a threshold statistical power (i.e. a threshold false negative rate; see STAR Methods). More specifically, since the FRA-MLE method calculates an approximate likelihood function explicitly, we could use the log-likelihood ratio (LLR) to test against our null hypothesis (Kass et al., 2014), an approach that is not possible with the FRA estimator. We can compare against the binary method (Fig. 1A) either using a LLR as done for the FRA-MLE estimator or with a paired sample test as done with the FRA method. Thus we compare four statistical protocols: FRA estimates using a paired test, FRA-MLE estimates using LLR, binary classification using a paired test and binary classification using LLR.

Figure 4C shows n_{min} against the true silent synapse fraction. As is expected, we found that more samples are required to discriminate smaller differences in silent synapse fractions from a fully active synapse population. Importantly, we found a dramatic improvement in statistical power afforded by the FRA-MLE technique (Fig. 4C) compared to FRA. While the FRA estimator required $n_{min} = 36$ to discriminate populations containing less than 15% silent, FRA-MLE required only $n_{min} = 8$. Averaged across all silent fractions, FRA-MLE achieved an identical statistical power to the FRA estimator with a 3.82-fold reduction in sample size. Binary classification required 1.88-fold more samples than FRA, although when LLR-based hypothesis testing was incorporated, the binary technique now required fewer samples than FRA such that its power almost matched that of the FRA-MLE method. Together, this power analysis shows that the FRA-MLE estimator renders more fine-grained silent fractions detectable with fewer samples than comparable methods.

Discussion

It has been long believed that the plasticity potential of a network is in part determined by the fraction of synapses which exist in a silent state, and the proportion of silent synapses is

regulated both throughout development (Durand et al., 1996, Isaac et al., 1997, Busetto et al., 2008, Ashby et al., 2011) and in response to diseases or drugs (Huang et al., 2009, Brown et al., 2011, Graziane et al., 2016, Ma et al., 2016, Neumann et al., 2016, Huang et al., 2015, Ma et al., 2014, Lee et al., 2013, Koya et al., 2012, Rumpel et al., 1998, Beroun et al., 2018, Wright et al., 2020). Despite a growing appreciation of the overall importance of silent synapses in mature neural networks, a systematic analysis of quantification methods for this parameter is, to our knowledge, lacking. Here, using statistical inference methods alongside electrophysiological recordings, we have demonstrated the significant bias and variance inherent to the popular failure-rate analysis method of silent synapse quantification. Through a set of numerically simulated likelihood functions relying on minimal assumptions, we have furthermore proposed and validated an alternate maximum-likelihood estimator which corrects for these flaws. This approach illustrates the advantages of a model-based statistical inference approach.

One assumption of our approach is that release probabilities across a synaptic population are drawn from a uniform distribution (Soares et al., 2017). It is possible that specific constellations of central synapses in the brain may deviate from the latter assumption. For instance, a population dominated by very low release probability synapses would inevitably result in the sampling of a greater number of synapses, which could in principle increase the statistical power of the method and potentially decrease bias. To address this possibility, we considered the release probabilities derived from staining-destaining experiments in hippocampus that did not follow a uniform distribution (Fig. S2; Branco et al., 2008) and found that it induced minimal changes in the mean estimates of sampled synapses. These considerations in effect highlight a strength of the MLE model-based approach in that experimenters can ensure that the statistical model is precisely calibrated to their experiment's methodology (e.g., here, to a specified release probability distribution).

We also note that any estimation of the failure rate itself is dependent on correctly binarizing a set of noisy analog traces into synaptic successes and failures. This is problematic given that accurate classification of small synaptic events relies on appropriate discrimination from noise levels that are inherently different between hyperpolarized and depolarized states. The generally unknown and variable presence of rectifying GluA2-lacking AMPA-Rs (that differentially contribute to the EPSC at depolarized and hyperpolarized states) may also contribute to this classification ambiguity. Our method works on the failure rates directly, and therefore it assumes that there are no systematic biases inherent in this signal classification. Examining and correcting such classification biases from the raw analog signals is beyond the scope of our MLE method, but future work is needed to address these concerns.

Our power analysis additionally provides key insights for more modern binary synapse classification tools, including 2-photon glutamate uncaging (Lee et al., 2016). Our results demonstrate that in these cases, a dramatic boost to statistical power (~3-fold) can be obtained through constrained hypothesis testing using analytical tools – for example, testing against the null hypothesis that there are no silent synapses by employing log-likelihood ratio tests. We suggest that such hypothesis-based testing constitutes a viable alternative to classical Chi-squared statistical tests when using modern binary classification techniques.

The FRA formalism relies on several rigid, core assumptions. First, FRA assumes that the release probability distributions for silent and nonsilent synapses are equivalent, since this is a core simplification needed to derive the FRA equation (see STAR Methods). However, it is entirely possible that silent synapses may have distinct release probabilities (Cabezas and Buno, 2006), which might bias the estimates. Second, FRA assumes that NMDA opening is the sole parameter that changes between recordings at hyperpolarized and depolarized potentials, and that other parameters such as synaptic release probability are constant at both holding potentials. However, several synapses have been shown to be regulated by depolarization-induced release of retrograde messengers, which can either increase (Carta et al., 2014) or decrease (Howlett et al., 2004, Chevaleyre et al., 2006) release probability. In these circuits, the

FRA approach may be fundamentally inadmissible, as the returned computation would reflect the effects of the retrograde messengers rather than the silent synapse fraction. Contrary to the FRA approach, synthetic likelihoods can incorporate the expected action of retrograde messengers. By comparing the likelihood with and without incorporating these changes in release probability, standard model validation methodologies (Kass et al., 2014) can assess whether the changes in failures rates truly reflect a change in silent synapse fraction. Altogether, the formalism developed herein provides not only an improved analytical tool to estimate silent synapse fraction from electrophysiological recordings, but also an independent means to assess assumption violations.

STAR METHODS

Please see Appendix E for methods.

Supplemental Information

Supplemental Information can be found online at <https://doi.org/10.1016/j.celrep.2020.107916>. This includes 4 supplemental figures.

Acknowledgements

We thank all members of the Beique and Naud Labs for helpful discussions. M.B.L. is thankful to have received graduate scholarships from the Government of Ontario. This work was supported by grants from the Canadian Institutes of Health Research, the Natural Sciences and Engineering Research Council of Canada, the Canada Foundation for Innovation, Brain Canada (Canadian Neurophonic Platform) and the Krembil Foundation.

Author Contributions

Conceptualization, M.B.L., K.F.H.L, C.S, R.N, J.-C.B; Methodology, M.B.L.; Software, M.B.L.; Formal Analysis, M.B.L.; Investigation, M.B.L.; Writing - Original Draft, M.B.L.; Writing - Review & Editing, M.B.L, R.N and J.-C.B.; Supervision, J.-C.B; Funding Acquisition, J.-C.B.

Declaration of Interests

The authors declare no competing interests.

Manuscript II**Accurate silent synapse estimation from simulator-corrected electrophysiological data using the SilentMLE Python Package**

Michael B. Lynn*, Richard Naud, Jean-Claude Béïque

*: Corresponding author.

This manuscript was published in *STAR Protocols* in 2020:

Lynn MB, Naud R, Béïque JC. Accurate silent synapse estimation from simulator-corrected electrophysiological data using the SilentMLE Python Package. *STAR Protocols*. 2020 Dec 18;1(3):100176. doi: 10.1016/j.xpro.2020.100176. PMID: 32697998.

Statement of contribution

For this work, I wrote all Python code, assembled all figures, wrote all manuscripts drafts, and edited all drafts prior to submission in conjunction with the other co-authors. Jean-Claude Béïque and Richard Naud guided the project, edited drafts, and provided essential high-level input for the paper.

Summary

The proportion of silent (AMPA-lacking) synapses are thought to be related to the plasticity potential of neural networks. We created a maximum-likelihood estimator of silent synapse fraction based on simulations of the underlying experimental methodology. We provide here a set of guidelines for running a Python package on compatible experimental synaptic data. Compared to traditional failure-rate approaches, this synthetic likelihood estimator improves validity and accuracy of the estimates of silent synapse fraction.

For complete details on the use and execution of this protocol (as well as a validation of the synthetic likelihood approach), please refer to Lynn et al. (2020).

Before you Begin

Synthetic likelihoods have been developed to improve the reliability and accuracy of estimated quantities in ecology (Woods et al. 2010), genetics (Beaumont et al., 2002) and more recently in neuroscience (Greenberg et al., 2019; Lueckmann et al., 2019, Goncalves et al., 2020). These approaches are computationally heavy, often requiring a simulator and extensive computer codes. Open-access packages for such simulation-based approaches makes these methods more accessible (Tejero-Cantero et al., 2020). Here we describe the use of a simulator of synaptic electrophysiology experiments along with a calculation of synthetic likelihoods used for the estimation of silent synapse fraction.

The SilentMLE package provides an estimate of silent synapse fraction from electrophysiological data collected using a well-known sampling paradigm known as the failure-rate protocol (see Graziane and Dong (2016) for more details). Stimulation at hyperpolarized potentials leads to synaptic responses which solely reflect the stochastic activation of active synapses (*i.e.*, both NMDAR- and AMPAR-containing), while stimulation at

depolarized potentials recruits both active synapses and silent synapses (*i.e.*, NMDAR-only containing synapses). The intuition behind this protocol is that by comparing the failure rates at depolarized and hyperpolarized holding potentials (which may recruit distinct but overlapping synapse populations), one should be able to quantitatively infer the makeup of the synapse population. Here, we describe best practices for performing these experiments and how to seamlessly and optimally use the SilentMLE synthetic likelihood estimation framework to quantify silent synapses.

One should also consider optical approaches to estimate silent synapse fraction (e.g., Lee et al., 2016; Soares et al., 2017), as these approaches are more direct, but require specialized equipment. A comparison between different approaches is presented in Lynn et al. (2020).

A. Data collection

SilentMLE operates on processed synaptic failure-rate data collected from electrophysiological recordings of neurons. In principle, compatible synaptic data can be acquired from any functional synaptic population probed by electrophysiological means, yet for a number of practical reasons this approach is typically restricted to either acute or organotypic brain slice recordings that largely maintain *in situ* synaptic organization. Detailed protocols have previously been described (Schwartzkroin, 1975; Schwartzkroin, 1981; Soares et al., 2014). Briefly, brain slices are prepared from an area of interest and neurons are patch clamped in the whole-cell configuration while a stimulating electrode is placed nearby in the slice. The raw electrophysiological data is typically obtained using specialized hardware (e.g. equipment produced by Molecular Devices, USA), digitized and stored.

The failure-rate protocol typically consists of 100 consecutive low frequency synaptic stimulation (ca., 0.1-0.06 Hz), where the effects of the first 50 electrical stimulations (referred to as ‘sweeps’) are recorded while voltage clamping at $V_m = -70$ mV while the subsequent 50

sweeps are recorded at $V_m = +40$ mV. See Graziane and Dong (2016) for more details on the failure-rate protocol. The sweep number should be chosen carefully. Increasing the number of sweeps will improve the accuracy of the estimates returned by SilentMLE (Lynn et al., 2020).

CRITICAL: Before data collection, it is important to decide on two parameters:

- An acceptable range of failure rates which all recordings should be uniformly subjected to (*i.e.*, the percentage of electrical stimuli which lead to no measurable post-synaptic currents). The failure rate range has an impact on the number of synapses recorded from, and on the subset bias (Lynn et al., 2020). Typically, a range of failure rates of 20% to 80% is sought. Before each recording, electrical stimulation intensity must be adjusted up or down such that the observed failure rates fall into the acceptable range at $V_m = -70$ mV. A key feature of SilentMLE is that the flexible experimental simulator can take into account different choices of failure rate ranges, but these choices must be made before data collection.
- A target sample size for sufficient statistical power, rather than relying on *post hoc* justification. A power analysis outlining the expected granularity of silent synapse estimates using this approach with different sample size is provided in the principal paper (Lynn et al., 2020; see Fig. 4).

B. Parsing of successes and failures

The analog electrophysiological data must be processed to parse individual events into successes and failures. This classification can be achieved using multiple methods, as long as there is a clear division in the sorted traces between successes and failures. For a principled approach, we recommend constructing a control distribution of mock event amplitudes from

the noise. Events can be classified as successes if their amplitude is greater than some threshold (e.g. 3 standard deviations) from the mean of the control distribution. While SilentMLE does not incorporate such an algorithm, it can be easily constructed in one's programming language of choice.

The control distribution should utilize the same amplitude detection protocol employed for detecting real events, and should apply to some equal span of time in the noise. This will produce a distribution of mock events from which the mean and standard deviation can be extracted.

CRITICAL: The traces should be visually inspected after any automated steps. In particular, the standard deviation threshold may need fine-tuning, and may differ between the hyperpolarized traces and the depolarized traces. The latter may exhibit more noise and thus event parsing should be carefully examined before moving onto the next step. Inaccurate parsing of successes and failures in this step can drastically reduce the accuracy of the SilentMLE method.

C. Calculation of failure rate

Estimated failure rates for the hyperpolarized condition (\hat{F}_h) and depolarized condition (\hat{F}_d) must be calculated for each neuron:

$$\hat{F}_h = \frac{n_{\text{failures}}}{n_{\text{failures}} + n_{\text{successes}}}$$

$$\hat{F}_d = \frac{n_{\text{failures}}}{n_{\text{failures}} + n_{\text{successes}}}$$

The SilentMLE package works on the calculated failure rates for each neuron. Alternately, for existing datasets where an estimated silent fraction is already calculated from the failure rates (see Graziane and Dong, 2016; for derivation and equation, see Lynn et al., 2020), SilentMLE can work on these processed estimates instead.

D. Python installation

It is necessary to ensure that you have a working, up-to-date installation of Python present on your computer, as well as an integrated development environment (IDE) or editor where code can be written and executed.

There are multiple ways of obtaining a working Python installation. For new users, we recommend an installation of Anaconda, as it provides a current installation of Python, a commonly used IDE (Spyder), as well as a full set of command-line tools.

Key Resources Table

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
Python	Python Software Foundation	Python ≥ 3.6
NumPy	NumPy	NumPy $\geq 1.18.1$
SciPy	SciPy	SciPy $\geq 1.4.1$
SilentMLE Python Package	This paper	https://github.com/miclynn/silentmle

Materials and Equipment

- Data, processed into failure-rates for each neuron (see Before You Begin)
- Python environment and required packages
 - Python >3.6
 - Numpy
 - Scipy
 - Matplotlib
 - h5py
 - Seaborn
- Hardware
 - No special requirements

Step-by-Step Method Details

Installing SilentMLE

Timing: 10 minutes

This step downloads the SilentMLE package from GitHub and installs it under your local Python installation. We provide two Variants for installation. Variant 1 is recommended for most users and Variant 2 provides guidelines for advanced users.

1. Variant 1: Simple installation of Python and SilentMLE using Anaconda

- a. Download and install the Anaconda python distribution for your operating system using the instructions at <https://docs.anaconda.com/anaconda/install/>
- b. Open an Anaconda terminal window using the instructions to “Open Anaconda Prompt” found by scrolling down on the page at <https://docs.anaconda.com/anaconda/user-guide/getting-started/>.
- c. In the terminal window, run the following commands, pressing Enter/Return after each line to execute:

```
conda install -y git
```

```
pip install git+https://github.com/micllynn/SilentMLE
python
```

- d. Verify the installation proceeded correctly
 - i. Run the following code from the running python instance:

```
import silentmle as sil
```

- ii. The import should run correctly. If not, see the CRITICAL section below.

2. Variant 2: Installation of SilentMLE for experienced users

- a. Experienced users are encouraged to initialize a separate virtual environment before installation:

```
conda create -n silentmle
conda activate silentmle
```

- b. Experienced users can install the package however best fits with their workflow. A setup.py file is provided for this purpose. The following installation command should be executed from the command line with the appropriate installation of Python:

```
python setup.py install
```

- c. Experienced users can also use their preferred text editor or IDE in place of Spyder (eg Emacs, vim, Atom, etc.), replacing the relevant sections below.

CRITICAL: It is important to install SilentMLE using the correct Python installation, or else the package will not be importable. See Troubleshooting if the installation verification does not proceed correctly.

Initializing an estimator from constrained experimental simulations

Timing: 5-20 minutes

In this step, we import the package and run the set of constrained experimental simulations which generate the mapping between the ground-truth fraction silent in some population, and

the biased measurement returned by electrical stimulation experiments. Two variants are provided: The first variant describes how to initialize a simple estimator when synaptic release probabilities are unknown (assumes a null case of a uniform distribution); the second variant describes how to initialize a more complex estimator with a known distribution of release probabilities.

3. *Variant 1: Initializing a simple estimator with no strong prior of synaptic release probability distribution.*

- a. To launch the Spyder IDE, open Anaconda Navigator, in the “Home” tab navigate to “Spyder”, and click Launch. Instructions for using Spyder can be found at <https://docs.spyder-ide.org/current/quickstart.html>.
- b. In Spyder, execute the following code to import the package.

```
import silentmle as sil
```

- c. To prepare the correct parameters, retrieve information from the slice physiology experiments on:
 - i. The number of sweeps of electrical stimulation each cell was subjected to in each voltage-clamp condition (hyperpolarized or depolarized potentials) (num_trials). For example, in a case of 50 depolarized sweeps and 50 hyperpolarized sweeps, num_trials=50.
 - ii. The lower and upper bounds of acceptable failure rate fractions, decided upon before running experiments (failrate_low and failrate_high). For example, a range of 20% failures to 80% failures would be represented as failrate_low=0.2, failrate_high=0.8.
 - iii. If using previously calculated silent fraction estimates, whether any negative silent fraction estimates were set to zero (zeroing=True) or not (zeroing=False).
- d. SilentMLE takes an object-oriented approach to the estimation process. An estimator object must be initialized first (which runs a set of simulations across the range of

possible silent fractions), and once initialized it can be used to perform maximum-likelihood estimation on data. To initialize the estimator, run the following code to initialize an estimator object, replacing the appropriate parameters with those specified in step 3c. (Note that `n_likelihood_points` refers to the number of silent fractions to simulate; lower values will result in more rapid estimator initialization at the expense of estimate precision.)

```
estimator = sil.Estimator(n_likelihood_points=250, num_trials=50,
                          failrate_low=0.2, failrate_high=0.8, zeroing=False)
```

4. *Variant 2: Initializing a more complex estimator with an experimentally supported synaptic release probability distribution.*

- a. To launch the Spyder IDE, open Anaconda Navigator, navigate to “Spyder”, and click Launch. Instructions for using Spyder can be found at <https://docs.spyder-ide.org/current/quickstart.html>.
- b. In Spyder, execute the following code to import the package.

```
import silentmle as sil
```

- c. To prepare the correct parameters, retrieve information from the slice physiology experiments on:
 - i. The number of sweeps of electrical stimulation each cell was subjected to in each voltage-clamp condition (hyperpolarized or depolarized potentials) (`num_trials`). For example, in a case of 50 depolarized sweeps and 50 hyperpolarized sweeps, `num_trials=50`.
 - ii. The lower and upper bounds of acceptable failure rate fractions, decided upon before running experiments (`failrate_low` and `failrate_high`). For example, a range of 20% failures to 80% failures would be represented as `failrate_low=0.2, failrate_high=0.8`.

- iii. If using previously calculated silent fraction estimates, whether any negative silent fraction estimates were set to zero (`zeroing=True`) or not (`zeroing=False`).
 - iv. A statistical distribution, with parameters, of synaptic release probabilities.
- d. Create a class instance representing the synapse release probability distribution retrieved in step 4civ,. Note that any distribution from `scipy.stats` is acceptable, and specified parameters for these distributions are required in the dictionary `args`.

```
import scipy.stats as sp_stats
```

```
pr_dist = sil.PrDist(dist=sp_stats.gamma,
                    args={'a':1, 'scale':0.3})
```

- e. `SilentMLE` takes an object-oriented approach to the estimation process. An estimator object must be initialized first (which runs a set of simulations across the range of possible silent fractions), and once initialized, the estimator object can then be used to perform maximum-likelihood estimation on data. To initialize the estimator, run the following code to initialize an estimator object, replacing the appropriate parameters with those specified in step 4c. (Note that `n_likelihood_points` refers to the number of silent fractions to simulate; lower values will result in more rapid estimator initialization at the expense of estimate precision.)

```
estimator = sil.Estimator(n_likelihood_points=250,
                          num_trials=50, failrate_low=0.2,
                          failrate_high=0.8, zeroing=False,
                          pr_dist_sil=pr_dist, pr_dist_nonsil=pr_dist)
```

CRITICAL: It is important to support the assertion of a particular synaptic release probability distribution with data and/or references from the literature. Otherwise, we recommend

assuming the null hypothesis of uniformly distributed synaptic release probabilities (Soares et al., 2017).

Note: `help(sil.Estimator)` provides comprehensive documentation of each input and its expected type for this class. Many parameters controlling the experimental simulations can be fine-tuned according to the information here.

Note: All specified simulation and experimental parameters are stored in the class instance as `estimator.params` for easy access. The observation space and hypothesis space employed are stored as `estimator.obs` and `estimator.hyp` respectively. The numerically simulated likelihood function is stored as `estimator.likelihood`.

Performing maximum-likelihood estimation on collected data

Timing: 5 minutes

In this step, we perform maximum-likelihood estimation using the initialized estimator. This can either be done using raw failure-rates for each neuron (Variant 1), or using existing datasets consisting of silent synapse estimates from the failure-rate analysis equation (Variant 2). See *Before You Begin, Section C* for more information on the two variants of processing data.

5. Variant 1: Data in raw failure-rate form

- a. Create a variable representing the data. Note that we create a list with each cell's failure rates expressed in the format $[\hat{F}_h, \hat{F}_d]$. Replace the values here with your own data.

Lists can be of arbitrary length.

```
data = [[0.2, 0.18], [0.4, 0.36], [0.38, 0.48], [0.58, 0.6], [0.46, 0.31], [0.7, 0.43],
        [0.58, 0.36]]
```

- b. Perform maximum-likelihood estimation using the initialized data variable and the existing estimator variable. Ensure that the correct datatype (dtype) is set.

```
likelihood = estimator.estimate(data, dtype='failrate')
```

- c. Maximum-likelihood estimation should be relatively rapid; wait for the estimation to complete.
- d. Examine the joint likelihood plot that was generated (see Figure 1 for examples). This depicts the likelihood of each hypothetical silent fraction being the correct value, given the data provided. One can conveniently read the maximally-likely estimate (MLE) for silent fraction indicated in red.
- e. The normalized joint likelihood function (shown on the y-axis of the returned plot) is stored in likelihood for further processing and analysis. The corresponding hypothesis space of silent synapse fractions (shown on the x-axis of the returned plot) is stored in estimator.hyp .

6. Variant 2: Data in processed silent synapse estimate form

- a. Create a variable representing the data (estimates of fraction silent synapses). Note that we create a list comprised of each cell's estimated silent synapse fraction. Replace the values here with your own data. Lists can be of arbitrary length.

```
data = [0.42, 0.32, 0.54, 0.43, 1.0, 1.0, -0.06, -0.18, 1, 0.42, 0.16]
```

- b. Perform maximum-likelihood estimation using the initialized data variable and the existing estimator variable. Ensure that the correct datatype (dtype) is set.

```
likelihood = estimator.estimate(data, dtype='est')
```

- c. Maximum-likelihood estimation should be relatively rapid; wait for the estimation to complete.
- d. Examine the joint likelihood plot that was generated (see Figure 1 for examples). This depicts the likelihood of each hypothetical silent fraction being the correct value, given

the data provided. One can conveniently read the maximally-likely estimate (MLE) for silent fraction indicated in red.

- e. The normalized joint likelihood function (shown on the y-axis of the returned plot) is stored in `likelihood` for further processing and analysis. The corresponding hypothesis space of silent synapse fractions (shown on the x-axis of the returned plot) is stored in `estimator.hyp` .

Note: `help(estimator.estimate)` provides comprehensive documentation of the expected type and format of the data, the parameters used for estimation, and the type and format of the returned likelihood vector.

CRITICAL: If previously obtained silent fraction estimates have been 'zeroed' (ie all negative values have been set to 0), the Estimator must have been initialized in steps 3-4 with the parameter `zeroing=True`.

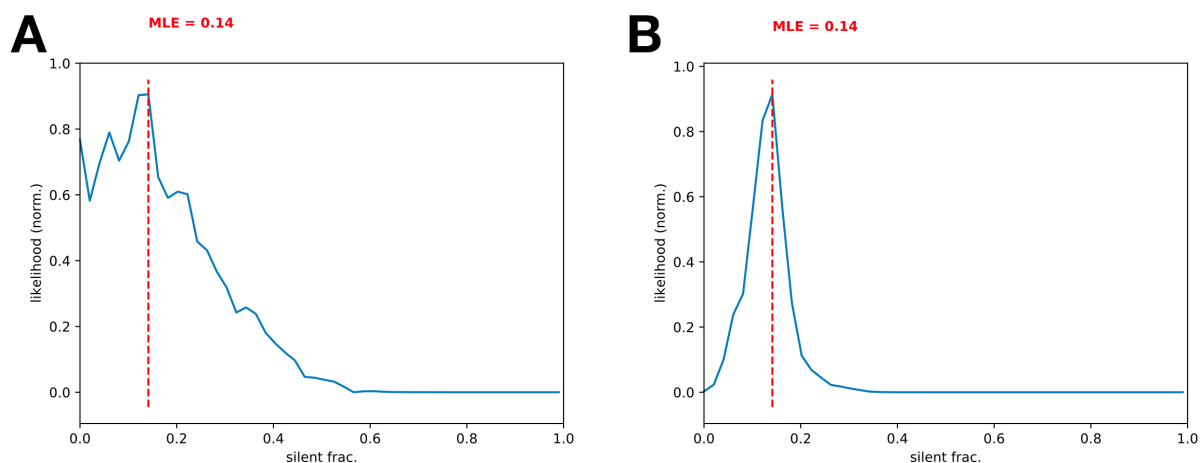


Figure 1: Maximum-likelihood estimation on two datasets with synthetic data.

Two example plots depicting the joint likelihood function returned by SilentMLE. **A:** An example likelihood plot for $n=2$ artificial data points. Note the broad likelihood distribution. **B:** An example likelihood plot for $n=20$ artificial data points. Note the narrower likelihood distribution despite the same maximum-likelihood estimate. The code to generate Figure 1B is available on the project's GitHub site at: <https://github.com/micillynn/SilentMLE#example-of-full-estimation-procedure-and-output>.

Expected Outcomes

Joint likelihood function

Figure 1 depicts the joint likelihood function associated with two simulated experimental datasets. In **Figure 1A**, a smaller dataset of $n=2$ is used as input to the maximum-likelihood estimator, generating a joint likelihood function with large variance and a maximally likely estimate of 0.14. In **Figure 1B**, a larger dataset of $n=20$ is used as input to the maximum-likelihood estimator. This generates an identical maximally likely estimate of 0.14, but the joint likelihood function has far smaller variance indicating a greater confidence in the estimate. This illustrates the importance of investigators deciding on a sufficient sample size for their experiments before using SilentMLE, as demonstrated in the power analysis in Lynn et al., 2020, to prevent estimates with low confidence. To aid users in verifying correct package

installation, and to give an example of a full step-by-step protocol, the full code to generate Figure 1B is available on the package's GitHub page (see figure legend).

Output of SilentMLE

Upon successful estimation, SilentMLE returns several important variables which can either be analyzed in the same session of Python, or exported for analysis in other software. `likelihood` represents the normalized joint likelihood function depicted in the y-axis of the plots in Figure 1. The corresponding hypothesis space of silent synapse fractions (shown on the x-axis of the returned plot) is stored in `estimator.hyp`. Together, these can be used to perform further analysis, including confidence intervals and log-likelihood ratio testing between datasets.

To export the variables in a neutral `.csv` format for use with other software, we recommend the Pandas package:

```
import pandas as pd
mle_output = pd.DataFrame({'likelihood': likelihood, 'silent': estimator.hyp})
mle_output.to_csv('mle_output.csv', index=False)
```

SilentMLE can additionally recreate complete figures from the original paper (Lynn et al., 2020). These take some time to run, as they recompute all relevant analysis steps and simulations. To generate the figures, the following functions can be run:

- `sil.figures.plot_fig1()`
- `sil.figures.plot_fig2()`
- `sil.figures.plot_fig4()`
- `sil.figures.plot_figS1()`
- `sil.figures.plot_figS2()`
- `sil.figures.plot_figS3()`
- `sil.figures.plot_figS4()`

Full documentation for all plotting functions can be found with `help(plotting-function)`. In addition to these functions, core functions are available for running the experimental simulator, performing the failure-rate analysis, and conducting power analyses. Full information on these functions can be found on the GitHub page (<https://github.com/miclynn/SilentMLE>), or by executing `help(sil.core)`.

Limitations

The accuracy of SilentMLE depends critically on several factors, including the quality of the electrophysiology data collected as well as its processing (summarized in the Before You Begin section), and the choice of parameters used to initialize the estimator. These assumptions and limitations include the following:

- Special attention should be paid to ensure that the parameters used as input to the computational model match those used during the experimental paradigm. All parameters, and the desired sample size, should be decided before experiments are started.
 - The failure-rate range provided to the model must match those in experiments.
 - The number of sweeps provided to the model must match the number in experiments.
 - Any specified synaptic release probability distribution must have strong support from either experiments or the literature. Otherwise, we recommend the null case of a uniform release probability distribution (Soares et al., 2017).
 - Any other parameters changed upon estimator initiation should be checked carefully and recorded, along with the maximum-likelihood estimate returned, for later reporting.
- A sufficient number of likelihood points are required when initializing the estimator (`n_likelihood_points`) to prevent stochasticity between estimation runs due to the Monte

Carlo nature of simulations of the experimental protocol. Ideally, this value should be set to >200.

- The underlying failure-rate analysis protocol (Graziane and Dong, 2016) relies on several key assumptions which, if violated, may limit the applicability of SilentMLE. These assumptions include:
 - Equivalent release probability distributions between silent and nonsilent synapses.
 - Random sampling of silent and nonsilent synapses through electrical stimulation. Our SilentMLE package does take into account the bias in sampling through gradual reduction in electrical stimulation intensity to reach the target failure rate, but it does rely on an initial, unbiased large sample from the population. Formally, any violation of this assumption (e.g. clustered silent or nonsilent synapses, causing large initial sampling biases before adjusting stimulation intensity) may invalidate the results of SilentMLE.
 - NMDA opening is the sole parameter which differs between hyperpolarized and depolarized voltage-clamp conditions. Synapses can be regulated by depolarization-induced release of retrograde messengers, which can either increase or decrease release probability (Carta et al., 2014; Howlett et al., 2004; Chevaleyre et al., 2006). In these circuits, SilentMLE may return unreliable results due to the effects of retrograde messengers.

Troubleshooting

Problem 1:

The SilentMLE package is not importable (steps 3b and 4b).

Potential Solution:

It is important to make sure that the SilentMLE package is installed under the correct installation of Python. For example, on macOS there is typically a built-in installation of Python which is separate from the Python installations installed with package managers such as conda. The latter are typically accessed for commonly used IDEs such as Spyder, and thus it is crucial to install the package under them, and not under the built-in Python installation.

To determine which Python installation is associated with a given binary (e.g. python3), the following steps may be necessary. To determine the location of the binary being referenced, one can use the terminal command “\$ which python-executable” on UNIX-based systems, or “where python-executable” on Windows systems. This should typically return the directory of a binary within the IDE or package manager folder (e.g. conda), and not within /usr/local/bin. If the wrong folder is associated with a Python binary you have used to install SilentMLE, you may need to reinstall the package using the correct binary. For more information on the correct binary, see your IDE or package manager of choice to determine which installation of Python is utilized.

Problem 2:

While initializing the estimator object (steps 3d and 4e), the execution hangs with the following output line : “Generating estimate distributions... 0.5%” (The precise number is not important, only that the simulation hangs at this point)

Potential Solution:

If the range of failure rates specified while initializing the estimator (steps 3d and 4e) is very narrow (e.g. 0.4 to 0.6), the experimental simulator may be unable to return the desired number of sets of synapses. This is because too many simulations will not reach a failure rate in the desired range, and thus will be discarded. By default, simulations are ‘oversampled’ by 8x from

0.9<silent_frac<1.0. To control the oversampling factor (default 8x) and threshold (default 0.9), the following kwargs can be passed to the `sil.estimate()` function: `sim_oversample_factor=8`, `sim_oversample_thresh=0.9`. These kwargs can be altered, but a higher oversampling factor or lower oversampling threshold will require more time to complete the simulations.

Problem 3:

While running the estimation on data, you receive the following message: “**ValueError:** operands could not be broadcast together...” (steps 5b and 6b).

Potential Solution:

Please ensure that the data initialized in the variable `data` is formatted correctly, and that it matches the `dtype` parameter of `estimator.estimate()`.

Resource Availability

Lead Contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Michael Benjamin Lynn.

Materials Availability

No Materials were generated in this protocol.

Data and Code Availability

All code associated with this Protocol is available at <https://github.com/micclynn/SilentMLE>

Acknowledgments

M.B.L. is thankful to have received graduate scholarships from the Government of Ontario. This work was supported by grants from the Canadian Institutes of Health Research, the Natural

Sciences and Engineering Research Council of Canada, the Canada Foundation for Innovation, Brain Canada (Canadian Neurophotonics Platform) and the Krembil Foundation.

We thank Emerson Harkin (University of Ottawa) and Dr. Philippe Vincent-Lamarre (University of Ottawa) for valuable help with package troubleshooting.

Author Contributions

Conceptualization, M.B.L., R.N, J.-C.B; Methodology, M.B.L.; Software, M.B.L.; Formal Analysis, M.B.L.; Investigation, M.B.L.; Writing - Original Draft, M.B.L.; Writing - Review & Editing, M.B.L, R.N and J.-C.B.; Supervision, J.-C.B; Funding Acquisition, J.-C.B.

Declaration of Interests

The authors declare that there are no competing interests.

Manuscript III**A slow 5-HT_{1A}R-mediated recurrent inhibitory network in raphe computes contextual value through synaptic facilitation**

Michael B. Lynn, Sean Geddes, Mohamad Chahrour, Sébastien Maillé, Emerson Harkin, Érik Harvey-Girard, Samir Haj-Dahmane, Richard Naud, Jean-Claude Béïque*

*: Corresponding author.

This manuscript is available on biorxiv:

Michael B. Lynn, Sean Geddes, Mohamad Chahrour, Sébastien Maillé, Emerson Harkin, Érik Harvey-Girard, Samir Haj-Dahmane, Richard Naud, Jean-Claude Béïque. A slow 5-HT_{1A}R-mediated recurrent inhibitory network in raphe computes contextual value through synaptic facilitation. Biorxiv, doi: <https://doi.org/10.1101/2022.08.31.506056>.

It was peer reviewed by *Nature Neuroscience* in November 2022 and we are planning to submit revisions in the next 3 months.

Statement of contribution

For this manuscript, I led project conceptualization and planning, and collected the majority of electrophysiological and anatomy data (Figures 1 and 2 involved data primarily collected by Sean Geddes, aided by Emerson Harkin and Sebastien Maillé and I). I performed all behavioral experiments, performed all simulations and analysis, assembled all figures, wrote all drafts of the manuscript, and participated in editing the drafts prior to submission.

I am greatly indebted to all co-authors for their contributions on this manuscript. Sean Geddes was instrumental, both for conceptualization and data collection, in early work characterizing recurrent inhibition and ruling out self-inhibition by serotonin neurons (Figures 1 and 2, as well as some data for Figure 7). Érik Harvey-Girard has helped with the anatomical characterization of recurrent connectivity with fluorescent dextran injections. Mohamad Chahrour participated in training mice for the behavioral task. Sébastien Maillé and Emerson Harkin helped with collecting electrophysiological data (Figures 1 and 2).

Abstract

Serotonin (5-HT) neurons in the dorsal raphe nucleus (DRN) receive a diverse constellation of long-range synaptic inputs, yet unifying principles of local circuitry and its dynamics are largely unknown - a crucial component of understanding how 5-HT output controls behavior. Here, we developed a formalism bridging optogenetic, electrophysiological, computational and behavioral strategies to reveal how the dynamics of local circuitry in DRN control the expression of reward associations. Using long-range input from lateral habenula (LHb) to interrogate functional DRN circuitry, we uncover 5-HT_{1A} receptor-mediated local recurrent connections between 5-HT neurons, refuting classical theories of autoinhibition by 5-HT_{1A} receptors. These inhibitory 5-HT connections were slow, stochastic, strongly facilitating, and gated spike output of 5-HT neurons. Targeted physiology and modeling approaches revealed that these functional connectivity features collectively support the emergence of a paradoxical excitation-driven inhibition in response to high frequency LHb activation, and of a winner-take-all computation over protracted timescales. In vivo, we found that optogenetic activation of LHb inputs to DRN transiently disrupted expression of a reward-conditioned response in an auditory conditioning task. In accordance with quantitative model predictions, this disruption occurred exclusively at the conjunction of high frequency LHb activation and high predicted reward value, and was not due to a modulation of the underlying reward association. Thus, we propose that recurrent dynamics in the DRN support a contextual value computation, where stable learned associations are integrated with dynamic environmental inputs to support sharp behavioral state transitions in changing environments.

Keywords: Serotonin; dorsal raphe nucleus; habenula; recurrent dynamics; short-term plasticity; winner-take-all; learned associations; value computation.

Introduction

Animals show a remarkable ability to adapt their behavioral state to environmental contexts. These adaptations can exhibit nonlinear relationships with environmental parameters that appear otherwise graded. For instance, various stress factors affect behavioral motivation positively at low levels but negatively at high levels (Yerkes-Dodson law; Yerkes and Dodson 1908; Broadhurst 1957 ; Salehi et al. 2010). While coordinated behavioral regulation has been associated with multiple neuromodulatory systems (Eriksson et al., 2017), it remains a challenge to identify how defined behavioral features emerge from specific network computations occurring in those neuromodulatory systems. Overall, neural computations implementing sharp behavioral state transitions from linear and graded neural representations of stimuli remain poorly characterized.

The serotonergic dorsal raphe nucleus (DRN) is a heterogeneous midbrain nucleus which regulates mood and behavior through its widespread efferent innervation (Dorocic et al., 2014; Weissbourd et al., 2014). The coding features of DRN 5-HT neurons still remain largely elusive, although recent studies have shown that they respond to environmental variables such as reward and threat (Cohen et al., 2015; Li et al., 2016; Zhong et al., 2017; Seo et al., 2019), making them suited to supervise behavioral policies based on internal representations of environmental contexts. Moreover, the DRN has been linked to a diverse range of mood-related states such as aggression (Takahashi et al., 2015) and patience (Miyazaki et al., 2014; Miyazaki et al., 2018; Miyazaki et al., 2020), as well as to the etiology in depressive and anxiety disorders (Graeff et al., 1996; Warden et al., 2012). DRN microcircuits are composed of multiple intermingled cell types, including dopaminergic, glutamatergic, serotonergic and GABAergic neurons, organized into functionally distinct modules (Ren et al., 2018). However, the precise rules by which these circuits process incoming information to ultimately provide instructive cues to regulate behavior are largely unknown.

One major source of input to the DRN is from the lateral habenula (LHb), which provides monosynaptic excitatory drive to 5-HT neurons (Zhou et al., 2017). LHb neurons fire in

response to punishments, reward omissions and threat by phasic, high-frequency discharges, providing a neural scaffold for valence-based behavioral modification and learning (Matsumoto & Hikosaka, 2009; Proulx et al., 2014; Bromberg-Martin & Hikosaka, 2011; Lecca et al., 2020; Amo et al., 2014; Yang et al., 2018). Recent work has suggested that both the raphe and habenula regulate behavioral policy over long timescales by modifying internal representations (Marques et al., 2020). How incoming synaptic inputs from LHb, carrying graded threat information over milliseconds, are processed by DRN microcircuits into an internal state representation that persists over longer timescales to influence behavioral policy is currently unknown.

Here, we developed a synthesis of optogenetic, electrophysiological, computational and behavioral strategies to elucidate key principles by which local DRN microcircuits enact computations. Using long-range inputs from LHb to interrogate DRN functional circuitry under ethologically meaningful drive, we uncovered an unexpected slow inhibitory conductance reflecting local recurrent connections between 5-HT neurons. These recurrent connections, mediated by inhibitory 5-HT_{1A} receptors (5-HT_{1A}Rs), were slow, highly facilitating, and modulated spiking output of the circuit. We provide a simple, unified model of the complex short-term dynamics of these recurrent serotonin connections, which, along with biophysically grounded modeling, predicts a sharp transition from net excitation to excitation-driven inhibition, as well as winner-take-all dynamics between long-range inputs to ensembles in DRN. These findings suggest operating principles for interactions between heterogeneous modules within DRN. The behavioral consequences of these computations were examined using a head-fixed auditory classical conditioning task, revealing that activating LHb inputs to DRN produced reward value-dependent effects on reward anticipation exclusively at high photostimulation frequencies quantitatively predicted by our network modeling. Our results provide a comprehensive description of a novel circuit motif in DRN, redefine the role of a key serotonin receptor subtype, and provide key algorithmic insights into how environmental

threats can be non-linearly processed and thresholded by a mood-related microcircuit to ultimately drive sharp transitions between competing behavioral policies.

Results

Excitatory monosynaptic LHb inputs to DRN trigger a frequency-dependent indirect inhibition

As a first step towards uncovering organizing principles of microcircuitry within DRN, we activated LHb, a strong long-range input to DRN, and probed neural responses. We utilized a SERT-CRE-tdTomato mouse line for whole-cell electrophysiological recordings from genetically identified 5-HT neurons, validating this approach with immunohistochemical staining (Fig. S1A-D; 98.4% overlap between TPH2 and tdTomato expression). Upon bilateral injection of AAV-GFP in LHb, we observed dense LHb axon density in DRN (Fig. S2A) consistent with recent work (Zhou et al., 2017). Next, we injected AAV-ChR2(H134R)-mCherry bilaterally into LHb of SERT-CRE-tdTomato mice (Fig. 1A and Fig. S1E, see Methods). Photostimulation induced strong and reliable action potentials in LHb neurons which were qualitatively similar to those following current step injections (Fig. S2B-C). Recording from identified DRN 5-HT neurons in the voltage-clamp configuration (Fig. 1C-D) showed that single brief (1-3 ms) pulses of blue light triggered synaptic release of glutamate that activated AMPARs (Fig. 1E; $V_m = -70$ mV) and NMDARs (Fig. 1F; $V_m = +40$ mV). LED-stimulated glutamate release was blocked by TTX and rescued by the K⁺ channel blocker 4-AP (Fig. 1G), indicating that LHb makes monosynaptic contacts with DRN 5-HT neurons. These excitatory inputs from LHb were capable of driving 5-HT neurons to spike (Fig. 1H). Thus, LHb afferents directly excite 5-HT neurons, consistent with previous reports (Zhou et al., 2017). While habenula-driven feedforward inhibition onto 5-HT neurons has been previously reported (Zhou et al., 2017), we observed no evidence of optically triggered IPSCs in our recording conditions (n=116 neurons). This stands in stark contrast to our previous finding of prominent feedforward inhibition onto 5-HT cells triggered by prefrontal cortex inputs to the DRN (Geddes et al., 2016).

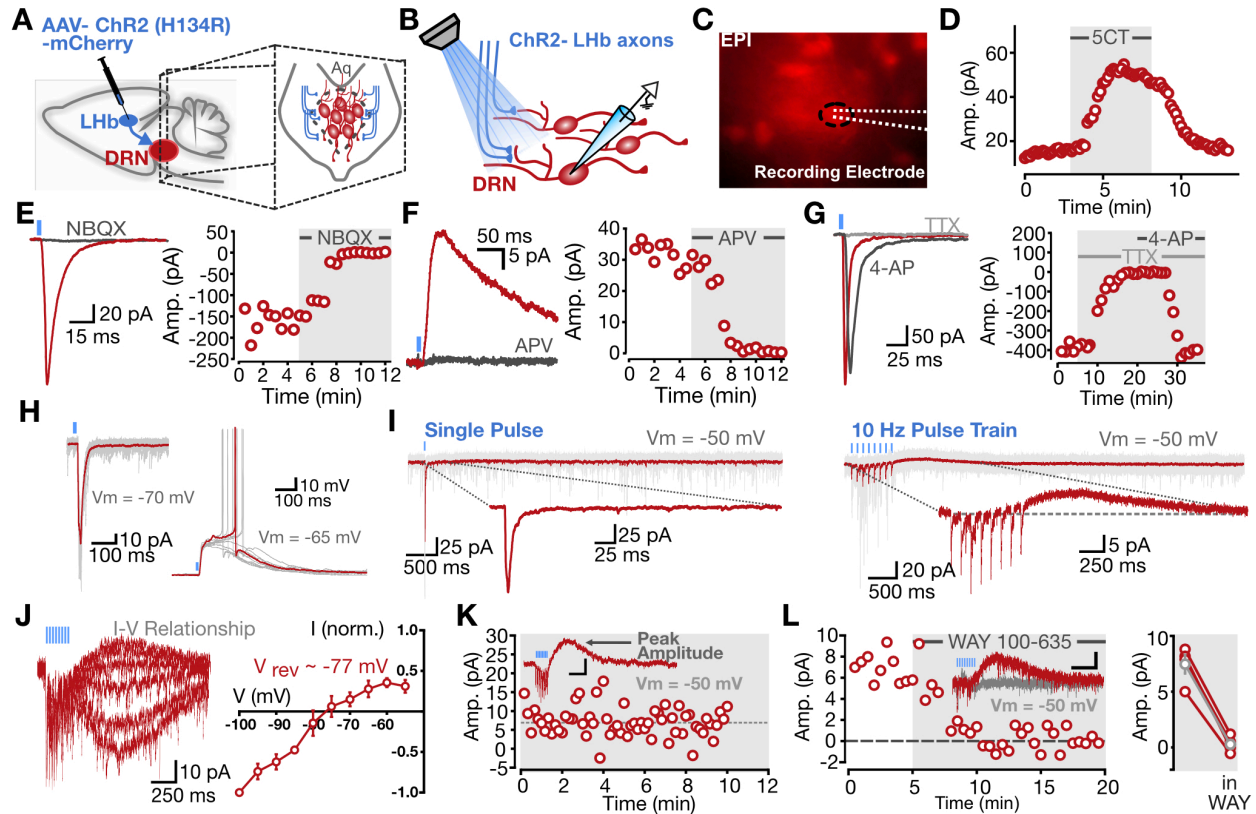


Figure 1: Habenula input to raphe triggers direct excitation and heterosynaptic inhibition in 5-HT neurons. **(A-B)** Schematic of experimental approach. **(C)** Epifluorescent image of tdTomato-positive 5-HT neurons. **(D)** Application of 5-CT (100nM) triggered a transient increase in holding potential. **(E)** Single pulses of light at $V_m = -70$ mV triggered EPSCs (left) which were abolished by application of NBQX (10uM, right). **(F)** At $V_m = +40$ mV, single pulses of light triggered EPSCs (left) which were abolished by application of APV (50uM, right). **(G)** Application of TTX (1uM) blocked EPSCs at $V_m = -70$ mV, while subsequent application of 4-AP (1mM) recovered the current. **(H)** Voltage-clamp recordings of ChR2-EPSCs ($V_m = -70$ mV) and spiking behavior in response to presynaptic activation. **(I)** Representative 5-HT neuron responses to single pulses of light (left) and pulse trains at 10Hz (right). **(J)** Representative traces of the current-voltage (I-V) relationship of the delayed outward current (left) and group data (right, $n=3$ neurons, mean \pm SE) **(K)** Representative timecourse of peak amplitude stability over time. **(L)** Representative traces and timecourse showing the blockade of the outward current by WAY 100-635 (400 nM; left). Population data presented in a scatter-line plot (right, Pre-WAY, 7.47 ± 0.8 ; Post-WAY, 0.28 ± 0.4 ; $n = 4$ neurons; $p = 0.001$; paired two-sided t-test).

LHb neurons can discharge either acutely at high rates following threatening stimuli, or at lower sustained levels (Matsumoto & Hikosaka, 2009; Yang et al., 2018, Dolzani et al., 2016, Amo et al., 2014). To explore how such variable input patterns, carrying graded threat information, are processed and potentially parsed by DRN microcircuits, we delivered trains of optical stimuli at a range of frequencies consistent with previously reported in vivo phasic habenular responses (Amo et al., 2014). Train stimulation at 10Hz or 20Hz, surprisingly, triggered the appearance of a delayed outward current lasting hundreds of milliseconds, following the initial monosynaptic EPSCs (Fig. 1I). The overall kinetics of this protracted current, as well as its underlying rectifying I-V curve displaying a reversal potential at $\sim -77\text{mV}$ (Fig. 1J), suggest the involvement of a G-protein inwardly rectifying potassium (GIRK) channel (Newberry & Nicoll, 1984). While several G-protein coupled receptor subtypes couple to GIRK channels, the inhibitory 5-HT receptor of the 5-HT_{1A} subtype are expressed at high level in DRN 5-HT neurons (Pazos & Palacios, 1985) and were thus reasonable candidates. Bath application of the 5-HT_{1A} antagonist WAY100-635 abolished the outward current triggered by burst activation of LHb axons in the DRN without influencing the light-induced EPSCs (Fig. 1L). Together, these results demonstrate that LHb input to DRN 5-HT neurons trigger a rapid monosynaptic glutamatergic-mediated excitation followed by a slower, 5-HT_{1A}-mediated inhibition.

DRN 5-HT neurons are organized in a recurrent inhibitory network

We next sought to identify the cellular and network underpinnings of this non-canonical mixed excitation/inhibition response, initially focusing on the protracted 5-HT_{1A}-mediated inhibition. A dominant and well-accepted model posits that 5-HT_{1A}Rs are autoreceptors. In this model, individual 5-HT neurons, upon firing, limit their own excitability through the dendritic release of 5-HT (Maxwell et al., 1983; Ridet et al., 1994; Bunin & Wightmann, 1998), which activates these so-called inhibitory somatodendritic 5-HT_{1A} autoreceptors in an autocrine fashion (Gozlan et al., 1983; Weissmann-Nanopoulos et al., 1985; Pineyro & Blier,

1999). This canonical autoreceptor model is nonetheless difficult to reconcile with our observations that train activation of LHb inputs triggered 5-HT_{1A}R-mediated outward currents in nonspiking, voltage-clamped 5-HT neurons (Fig. 1I,J), or in 5-HT neurons receiving no functional excitatory LHb inputs (Fig. S2D-E). We therefore next sought to experimentally disambiguate the autoreceptor hypothesis from the alternative hypothesis that DRN 5-HT neurons are organized in a recurrent inhibitory network architecture (Fig. 2A).

We first reasoned that if 5-HT_{1A}Rs were indeed autocrine inhibitory autoreceptors, evoked spikes from a single 5-HT neuron should elicit a 5-HT_{1A}R-mediated inhibition in the same isolated neuron, contributing to the post-spike afterhyperpolarization (AHP) commonly observed in 5-HT neurons (Vandermaelen & Aghajanian, 1983). Whole-cell recordings, however, revealed that the post-spiking AHP triggered by current injection was insensitive to application of the 5-HT_{1A}R antagonist WAY100-635 (Fig. 2B-D), yet fully abolished by application of the SK channel blocker apamin (Fig. 2E-H). The absence of autocrine effects on the AHP could be explained by the small number of action potentials emitted in current-clamp experiments. Voltage-clamp experiments showed that the robust outward current that followed long, constant, depolarizing voltage steps in 5-HT neurons was also sensitive to apamin, but not to WAY100-635 (Fig. S3). Together, these experiments failed to support the autocrine function for 5-HT_{1A}Rs on 5-HT neurons, instead implicating calcium-activated potassium channels of the SK subtype in generating the AHP in 5-HT neurons.

To begin exploring the alternative possibility of a recurrent network between 5-HT neurons in DRN, we injected a fluorescent dextran retrograde tracer into acute DRN slices to visualize possible local recurrent connectivity (Trinh et al., 2016; see Methods). These experiments revealed the presence of genetically-confirmed 5-HT neuron somas in DRN >150µm away from the purported site of injection (Fig S4). By suggesting the presence of local connectivity between DRN 5-HT neurons, these experiments prompted us to examine this possibility more rigorously. We thus designed a series of targeted optogenetic strategies based on viral transduction of 5-HT neurons with a Cre-dependent AAV ChETA-YFP in SERT-

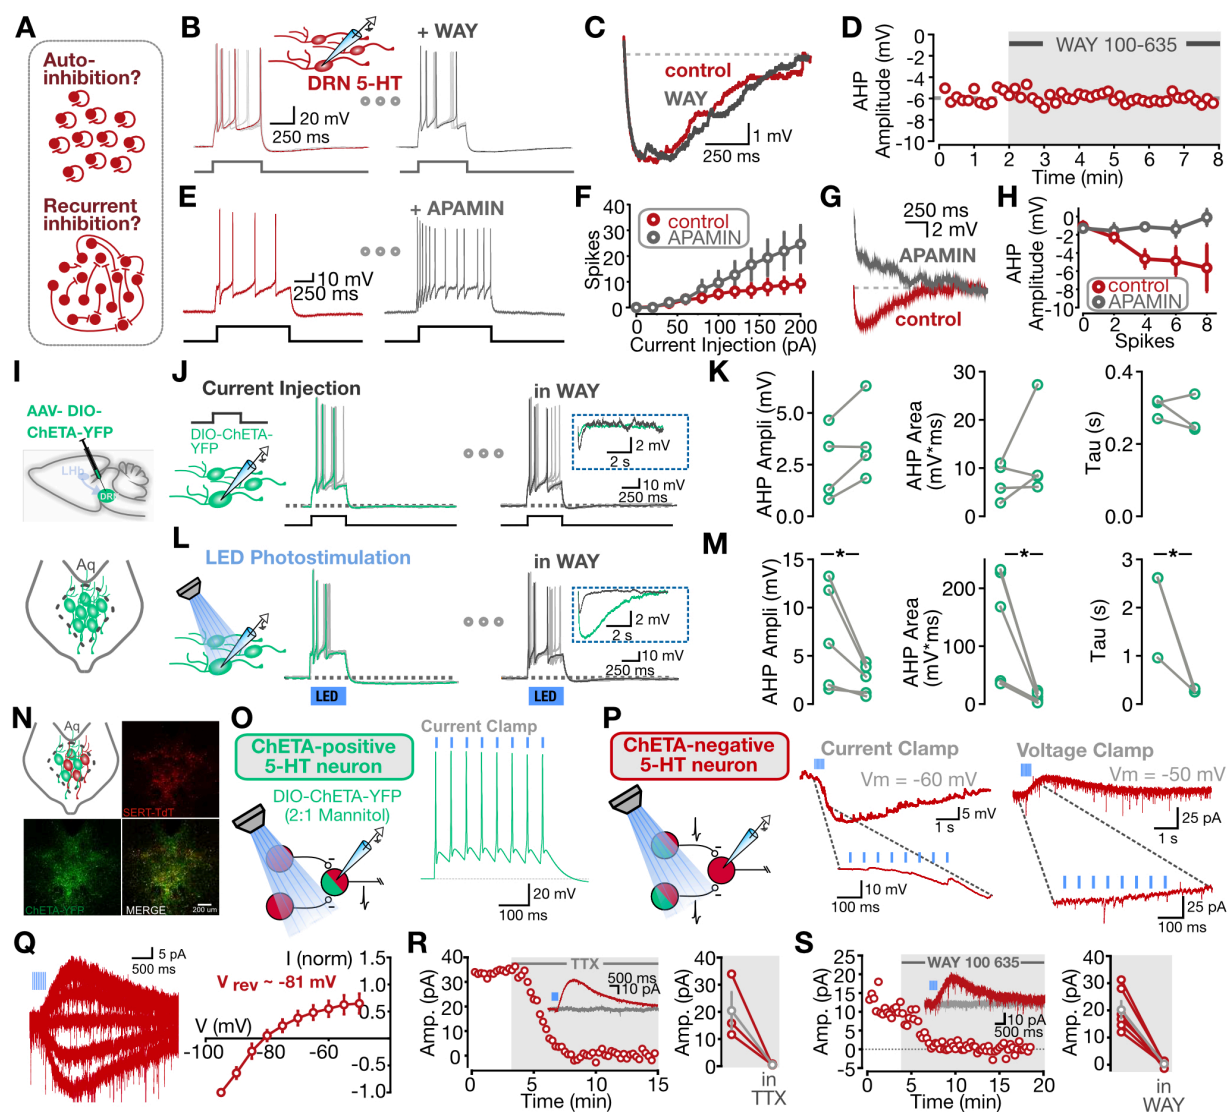


Figure 2: 5-HT neurons in raphe are organized in a recurrent inhibitory network. **(A)** Two hypothesized mechanisms of 5-HT_{1A}R-based signaling in raphe. **(B)** Representative traces of spiking in response to direct current injection before and after WAY 100-635 application. **(C)** Direct comparison of AHP before and after application of WAY 100-635. **(D)** Representative scatter plot showing no change in the AHP amplitude with bath application of WAY 100-635. **(E)** Representative traces of spiking in response to direct current injection before and after APAMIN application. **(F)** Relationship between current injection strength and evoked spikes, in the presence or absence of APAMIN ($n=3$ neurons, mean \pm SE) **(G)** Direct comparison of AHP before and after application of APAMIN. **(H)** Relationship between spike number and AHP amplitude, in the presence or absence of APAMIN ($n=3$ neurons, mean \pm SE). **(I)** Schematic of local optogenetic experimental approach. **(J)** Representative traces from CHETA-expressing 5-HT neurons showing current step-driven spiking before (green) and after (black) bath application of WAY100-635. **(K)** Quantification of AHP statistics for control and WAY100-635 conditions (AHP amplitude, $p=0.96$, $n=4$ neurons; AHP area, $p=0.85$, $n=4$ neurons; AHP tau,

p=0.09, n=3 neurons; paired two-sided t-tests). **(L)** Representative traces from CHETA-expressing 5-HT neurons showing photostimulation-driven spiking before (green) and after (black) bath application of WAY100-635. **(M)** Quantification of AHP statistics for control and WAY100-635 conditions (AHP amplitude, p=0.03, n=5 neurons; AHP area, p=0.03, n=5 neurons; AHP tau, p=0.02, n=3 neurons; paired two-sided t-tests). **(N)** Schematic and confocal images of sparse local optogenetic approach (AAV-DIO-CHETA-YFP diluted in mannitol 2:1). **(O)** While recording from an opsin-expressing neuron, brief optical stimulation could reliably elicit spiking behavior. **(P)** While recording from a non-opsin-expressing serotonergic neuron, photostimulation at 20Hz produced an inhibitory current; note the absence of direct photocurrents. **(Q)** Representative traces of the current-voltage (I-V) relationship of the inhibitory current (left) and group data (right, n=5 neurons, mean +/- SE) **(R)** Representative traces and timecourse of TTX (1uM) bath application (left) and group data (right; n=3 neurons; grey average trace: mean +/- SE). **(S)** Representative traces and timecourse of WAY100-635 (400nM) bath application (left) and group data (right; n=6 neurons; grey average trace: mean +/- SE).

Cre:tdTomato mice (Fig. S5A, see Methods). Ex vivo slice recordings from DRN 5-HT neurons co-expressing tdTomato and YFP-tagged ChETA opsin (Fig. S5B-C) showed that LED stimulation induced robust inward photocurrents/depolarizations that induced spiking reliably up to ~20 Hz (Fig. S5D-E). Train photostimulation of all transduced 5-HT neurons triggered a robust inhibitory 5-HT_{1A}R conductance (Fig. S5F-G). We then reasoned that neuronal spiking induced by photostimulating ChETA-expressing 5-HT neurons should reveal an additive 5-HT_{1A}R mediated component of the resultant AHP in the recorded neuron if DRN 5-HT neurons were indeed interconnected in a recurrent inhibitory network. To test this, we cross-calibrated single-cell direct current injections and LED intensity such that both methods independently triggered spiking of similar frequency (Fig. 2J and 2L). While the spiking triggered by direct current injections induced a strong and highly reliable AHP in these neurons (i.e., tdTomato and YFP-tagged ChETA-positive) that was not affected by the 5-HT_{1A}R antagonist WAY100-635 (Fig. 2K-L), the spiking triggered by LED photostimulation of the entire DRN network induced a substantially larger AHP, and the additional component was blocked by WAY100-635 (Fig. 2N). These results fail to support the autocrine hypothesis of 5-HT_{1A}R function, and rather provide evidence that 5-HT neurons are organized in a recurrent inhibitory network, communicating via 5-HT_{1A}R receptors.

In principle, however, it is possible that LED illumination depolarized poorly clamped distal 5-HT dendrites, which could release 5-HT to activate 5-HT_{1A}R in an autocrine fashion. To account for this possibility, we tested one final prediction of the recurrent network hypothesis; namely, that a 5-HT_{1A}R-dependent inhibitory conductance should be detected in non-photostimulated 5-HT neurons when other 5-HT neurons in the surrounding network are photostimulated. To directly test this, we implemented an AAV-ChETA-YFP transduction strategy (mannitol dilution, see Methods) in the DRN of SERT-CRE-tdTomato mice (Fig. 2O) to achieve sparse, mosaic-like expression of ChETA-YFP in 5-HT neurons. This allowed us to obtain visually targeted recordings either from ChETA-expressing 5-HT neurons (i.e., YFP and tdTomato expressing neurons; reliable optically driven spiking in Fig. 2P) or non ChETA-

expressing 5-HT neurons (i.e., YFP-negative, tdTomato-positive 5-HT neurons). Targeted recordings from non CHETA-expressing 5-HT neurons revealed that high-frequency LED photostimulation, which directly drives the surrounding network but not the recorded neuron, induced delayed strongly rectifying outward currents (Fig. 2Q) which were blocked by TTX and by WAY 100-635 (Fig. 2R-T). Taken together, these anatomical and optogenetic data strongly support the existence of an unsuspected recurrent inhibitory network in DRN, consisting of 5-HT neurons communicating via 5-HT1ARs over long timescales.

A quantitative description of release dynamics at 5-HT1AR-mediated recurrent connections

We next parametrized a number of key functional features of the 5-HT synapses between DRN 5-HT neurons using the 5-HT1AR – and GIRK-mediated hyperpolarization as a reporter of 5-HT release. Train optogenetic stimulation (20Hz, 8 pulses) of 5-HT neurons triggered a 5-HT1AR conductance that peaked approximately 400 ms following the end of the train (392.6 ± 106.9 ms) and that slowly decayed (2.1 ± 0.46 s). Varying the number of light pulses delivered at 20Hz revealed a relationship between pulse number and postsynaptic conductance that was almost linear, indicating that 5-HT1AR coupling to GIRK channels can follow multiple successive activation by endogenous 5-HT with limited desensitization (Fig. S6B-C). To infer presynaptic release dynamics, we next calculated the response variance as a function of pulse number, finding virtually no change in variance despite changes in amplitude (Fig. S6D). While such dynamics are consistent with deterministic neurotransmitter release (i.e., $P_{rel} \sim 1$; see Methods for full details), artefactual inflation of presynaptic release probability has previously been reported for optogenetic synaptic terminal activation (Zhang & Oertner, 2007). These results nonetheless outline the timecourse of a GIRK- and 5-HT1AR -mediated synapse, and indicate that postsynaptic receptor binding follows quasi-linearly presynaptic release.

To monitor 5-HT presynaptic release dynamics in the DRN circumventing the confound brought about by optogenetic stimulation, we next recorded DRN 5-HT neurons while electrically stimulating nearby serotonergic fibers with a bipolar concentric electrode (Fig. 3A).

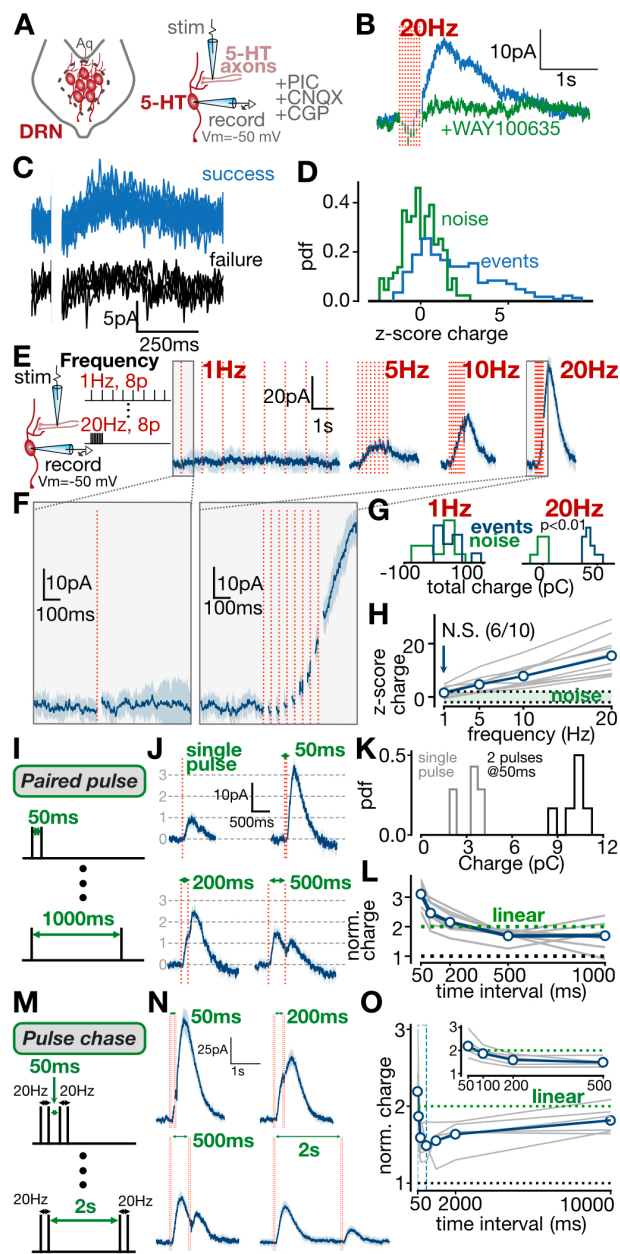


Figure 3: Serotonin release is probabilistic and strongly facilitating at inhibitory recurrent connections. **(A)** Experimental schematic. **(B)** Electrical stimulation at 20Hz evoked an outward current that is blocked by WAY100-635. **(C)** Representative traces depicting responses to low-amplitude electrical stimulation, showing sorted successes and failures from the same neuron. **(D)** Histograms of charge z-score, calculated for post-stimulus events (blue) and an equivalent period of non stimulus-driven noise from the same neuron (black) ($n=424$ events from 4 neurons). **(E)** Experimental setup for frequency-varying electrical stimulation experiments (left). Representative averaged traces from a single neuron depicting responses to electrical stimulation at various frequencies (right). **(F)** Expanded traces showing responses to 1Hz and 20Hz stimuli. **(G)** Charge histograms for stimulus-driven events and noise (one example neuron; 1Hz: $p>0.05$, 20Hz: $p<0.01$, two-sided t-tests). **(H)** Charge z-score as a function of frequency, with 95% noise C.I. shown in green ($n=10$ neurons, two-sided t-tests for each neuron as in G). **(I)** Paired-pulse electrical stimulation paradigm. **(J)** Representative averaged traces from a single neuron depicting responses to electrical stimulation at various paired-pulse intervals. **(K)** Charge histograms for single-pulse responses and paired-pulse responses at 50ms for one example neuron. **(L)** Post-stimulus charge as a function of paired-pulse time interval ($n=7$ neurons). Green line depicts linear summation. **(M)** Pulse-chase electrical stimulation paradigm. **(N)** Representative averaged traces from a single neuron depicting responses to pulse-chase electrical stimulation protocol at various temporal intervals. **(O)** Post-stimulus charge as a function of pulse-chase time interval ($n=6$ neurons). Green line depicts linear summation.

depicting responses to pulse-chase electrical stimulation protocol at various temporal intervals. **(O)** Post-stimulus charge as a function of pulse-chase time interval ($n=6$ neurons). Green line depicts linear summation.

A cocktail of synaptic blockers was bath-applied to isolate serotonergic transmission (Fig. 3A, see Methods). High-frequency electrical stimulation readily triggered a 5-HT_{1A}R-mediated outward current with a similar time course to that induced by LED stimulation (Fig. 3B). We next delivered isolated electrical pulses under a low-intensity minimal stimulation regime (i.e., activating small numbers of individual 5-HT synapses) while monitoring 5-HT_{1A}R-mediated currents. These recordings showed instances of clear successes (~5 pA) and failures (Fig. 3C) of 5-HT release, consistent with small numbers of 5-HT synapses each exhibiting a release probability < 1. Fig. 3D depicts the distribution of event charges plotted alongside neuron-specific control distributions (obtained from integrating an equivalent duration of noise; see Methods). The bimodal event distribution, exhibiting a left peak overlapping with the noise distribution, is consistent with probabilistic 5-HT transmission.

We next examined the short-term dynamics of these probabilistic 5-HT synapses by delivering electrical stimulation trains at higher (non-minimal) intensities with a fixed pulse number (8) and variable frequency (Fig. 3E). A synapse with no short-term dynamics should exhibit identical responses over a range of input frequencies if pulse number is fixed. However, while 1 Hz stimulation produced a notably weak response, higher frequencies of 10 Hz and 20 Hz elicited progressively larger inhibitory currents (Fig. 3E-F). To directly compare these protracted responses, which displayed rise-times and peak amplitudes that were highly dependent on stimulus parameters, we employed a z-score based metric of total charge. Briefly, we computed the total charge evoked by each stimulus train, and compared this with a surrogate noise distribution constructed by performing the same integration on an equivalent duration of baseline recording (pre-stimulation, computed within the same neuron, adjusted for each stimulation regime; Fig. 3G, See Methods). While low-frequency (1 Hz) stimuli did not produce responses significantly different from the surrogate noise distribution, high-frequency (20 Hz) stimuli produced progressively larger responses (Fig. 3G-H), consistent with short-term facilitation. To rule out subthreshold presynaptic summation as a possible confounding factor, these experiments were repeated with increased electrical stimulation strength, which returned

qualitatively similar short-term facilitation for these 5-HT connections (Fig. S7). We thus conclude that the protracted 5-HT_{1R}-dependent inhibition exhibits substantial short-term facilitation dynamics at frequencies above 5Hz.

To further quantify the timescale of these short-term synaptic dynamics, we delivered patterned electrical stimulation, starting with classic paired-pulse stimulation experiments. When paired-pulse stimuli were closely spaced (50-200 ms) we observed supralinear summation compared to isolated single pulses; in contrast, paired-pulse intervals of 500 ms and longer produced depressing responses (Fig. 3J-K). This short-term depression observed at timescales of seconds is consistent with previous reports on 5-HT_{1AR}-mediated inhibition (Courtney & Ford, 2016). By comparing the actual evoked charge with the predicted linear summation of isolated single-pulse responses, we were able to directly quantify the dynamics of these recurrent synapses across multiple timescales (Fig. 3L). Together, these results indicate that the short-term dynamics of these 5-HT synapses are dominated by strong short-term facilitation at short intervals, transitioning to mild depression at timescales > 500 ms.

To examine meta-dynamic features of these 5-HT synapses, we delivered a second set of patterned electrical stimuli consisting of a paired-pulse stimulation (20 Hz; ISI of 50 ms) separated from another paired-pulse stimulation by some varying time interval (pulse-chase; Fig. 3M). These experiments revealed, unexpectedly, that the paired-pulse facilitation at 20 Hz is highly dependent on stimulus history over a timescale of seconds. When paired-pulse sets were separated by short intervals, facilitation was observed, but when separated by intervals of >500 ms, the second paired-pulse stimulus elicited notably smaller responses, indicating release meta-dynamics, or a history-dependent change in facilitation (Fig. 3N). Charge quantification metrics, similar to above (see Methods), revealed that this dependence on stimulation history lasted up to 10 seconds (Fig. 3O).

A comprehensive picture of the network computations sustained by this recurrent inhibitory network requires, as an obligatory intermediate step, an accurate quantitative model of the synaptic dynamics described above. Thus, we next developed a mathematical model

capturing these ground-truth short-term dynamics of 5-HT synapses. We first fitted the experimentally-derived 5-HT_{1A}R-mediated inhibitory GIRK current with a simple summed exponential model (Fig. 4A). Next, to adequately capture the experimentally-derived short-term dynamics of 5-HT synapses, we fitted our experimental data using a previously described spike-response plasticity model (SRP; Rossbroich et al., 2021; see Methods). Briefly, a presynaptic spiketrain is convolved with a kernel indicating the sign and timescale of synaptic dynamics, passed through a sigmoidal readout, and convolved at each presynaptic spike-time with a postsynaptic event kernel (here, describing a GIRK-mediated current). We used a least-squares approach to fit this model to the experimental paired-pulse data, yielding an optimized short-term plasticity (STP) kernel that showed a short-timescale facilitation and long-timescale depression components (Fig. 4B). The resultant model effectively captured the dynamics of the paired-pulse experimental data across all timescales tested (Fig. 4C). Remarkably, the model was additionally able to describe the meta-facilitation dynamics observed in the pulse-chase experiments, despite not being exposed to this data during fitting – in effect acting as a test set for the model (Fig. 4D). To provide further intuition about this model, we finally conducted simple simulations where presynaptic trains were delivered to single neurons, and the model returned the kinetics and response magnitude of our earlier electrical stimulation experiments (Fig. 4E). Together, these results show that the complex dynamics of 5-HT release in the DRN can be faithfully captured by the SRP model.

Computations performed by a recurrent inhibitory network under excitatory drive

A natural question arising from the highly nonlinear short-term dynamics of 5-HT release observed above is how they contribute to the computations emerging from the DRN recurrent inhibitory network when receiving excitatory drive from long-range inputs. Our results so far show that the habenulo-raphé pathway can be characterized by a balance of direct excitation of 5-HT neurons on short timescales, and frequency-dependent recruitment of recurrent inhibition acting over protracted timescale. Since various connectivity motifs of

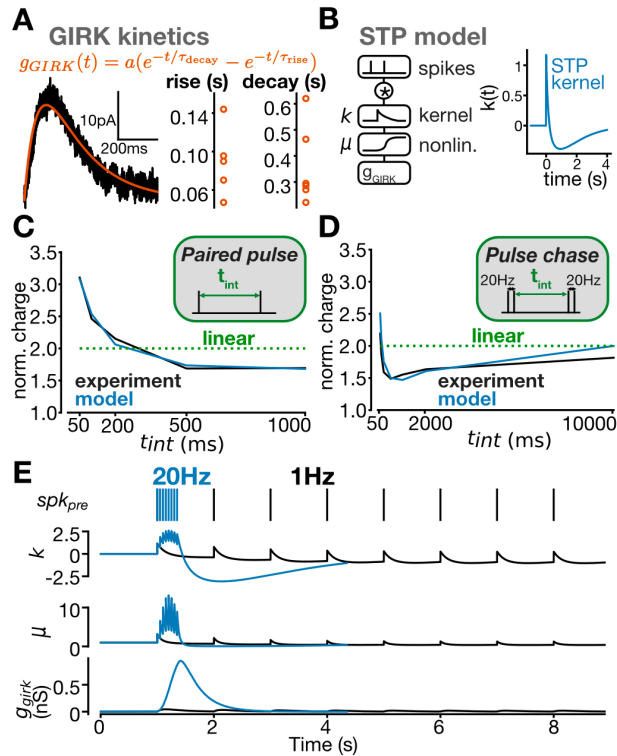


Figure 4: A quantitative model of short-term release dynamics captures patterned inputs. **(A)** Representative trace of kinetics fitted to an averaged electrically evoked GIRD current, and parameter distribution across the population ($n=6$ neurons). **(B)** Schematic depicting the linear-nonlinear short-term plasticity model and the plasticity kernel fitted using a least-squares approach to paired-pulse data. **(C)** Post-stimulus charge as a function of paired-pulse time interval, depicting experimental mean (black) and model predictions with fitted kernel and nonlinearity (blue). **(D)** Post-stimulus charge as a function of pulse-chase time interval, depicting experimental mean (black) and model predictions (blue). Note that model was fitted to paired-pulse, not pulse-chase dataset. **(E)** Simulations showing responses to 1Hz (black) and 20Hz (blue) presynaptic input, the model variables, and the resulting GIRD conductance.

excitation and inhibition have been ascribed a diverse set of network-level functions (Womelsdorf et al. 2014; Grienberger et al., 2017), we wondered about the potential computations supported by frequency-dependent feedback inhibition with fundamental timescale mismatch.

To investigate the computation carried out by DRN circuits under habenular drive (e.g., threat information), we leveraged our mathematical description of the short-term plasticity of 5-HT synapses in a network model (see Methods). Our model consisted of 5000 conductance-based leaky-integrate-and-fire 5-HT neurons, building on recent simplified single-cell modeling of 5-HT neurons (Harkin et al., 2021). 25% of 5-HT neurons received excitatory input from LHb ($\Sigma 7.5nS$), while all 5-HT neurons received background Poisson excitation as well as 5-HT_{1A}R-mediated inhibition from other 5-HT neurons ($\Sigma 2.5nS$) (Fig. 5A). The 5-HT synapse of recurrent inhibitory network were imparted with both the experimentally-fitted event kinetics and the short-term dynamics described above (i.e., from the SRP model).

To better understand how 5-HT neuron activity responded to LHb, we considered a range of LHb input frequencies and tracked 5-HT output over multiple timescales (Fig. 5). Responses to sustained 5 Hz or 20 Hz LHb activity suggested three periods (Fig. 5B-C): 1) an early phasic period (0-200 ms) followed by 2) a relaxation to a steady-state response (200-400 ms), and finally 3) a post-stimulus period. The early period resulted from 5-HT neurons responding to direct LHb excitation before 5-HT_{1A}R-mediated feedback inhibition could modulate the response. Consequently, the early period output rate was a rectified linear function of the input frequency (Fig. 5D). In contrast, the steady-state response was non-monotonic, displaying an inverted-U shape which peaked at input frequencies around 5 Hz and progressively decreased with higher LHb input frequencies (Fig. 5E). This atypical non-monotonic shape, essentially representing excitation-driven net inhibition, depended on the potent short-term facilitation at 5-HT synapses, as *in silico* removal of these short-term dynamics (while preserving the recurrent architecture) recovered a linear, monotonic, input-output function (Fig. 5E). In contrast, the early or post-stimulus period were not qualitatively

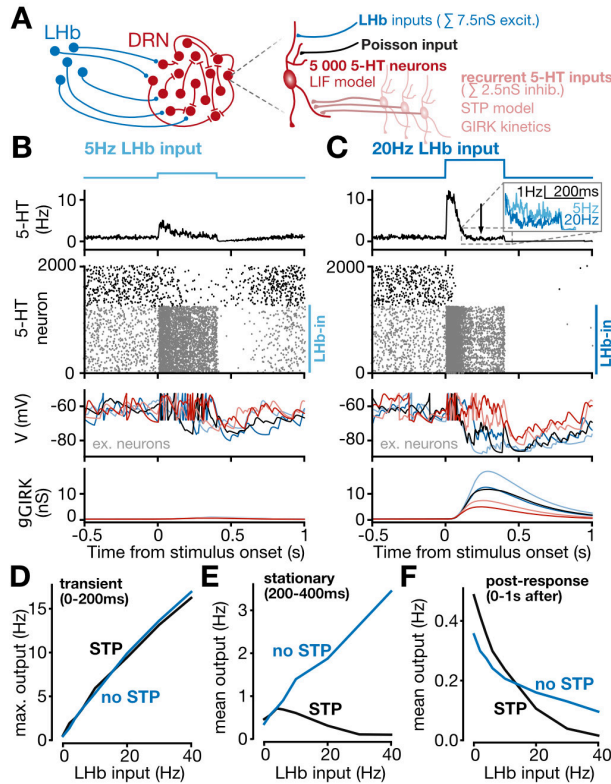


Figure 5: Habenula input to raphe generates a nonmonotonic network response through recurrent short-term plasticity. **(A)** Schematic of the LIF network model. **(B)** Results of a simulation where a 5Hz step input from LHB was delivered to a subset of 5-HT neurons. **(C)** Results of a simulation where a 20Hz step input from LHB was delivered to a subset of 5-HT neurons. **(D)** Input-output transformation for the transient (0-200ms) response phase, with (black) and without (blue) STP rules at recurrent serotonergic connections. **(E)** Input-output transformation for the stationary (200-400ms) response phase, with (black) and without (blue) STP rules at recurrent serotonergic connections. **(F)** Input-output transformation for the post-response (0-1s after stimulus) phase, with (black) and without (blue) STP rules at recurrent serotonergic connections.

affected by the removal of short-term recurrent dynamics (Fig. 5D-F). We repeated these simulations with connection probability decreased tenfold, which revealed qualitatively similar results, supporting the contention that these results may persist over a wide variety of possible network architectures (Fig S8). This indicated that the specific short-term dynamics, not the network architecture itself, was responsible for shaping excitation and inhibition. Together, these results suggests that synaptic dynamics imbue the serotonergic recurrent network with non-monotonic responses to long-range inputs, where response polarity is inverted at high input frequencies.

Recent work has indicated that there are multiple, largely non-overlapping submodules or cellular subgroups within DRN with distinct connectivity and functional rules (Ren et al., 2018; Huang et al., 2019; Paquelet et al., 2022), leading us to ask whether recurrent connections and their particular dynamics could gate interaction between functionally separate submodules or cellular assemblies. We modelled the canonical scenario of reward anticipation, where 5-HT neurons have been shown to display a ramping up of their activity from the time of cue to reward delivery (Zhong et al., 2017). We considered a simplified but realistic (based on Ren et al., 2018) model with two functionally distinct DRN submodules, one receiving step inputs from LHb, the other receiving ramping inputs from another generic area termed A (Fig. 6A). We first asked whether LHb inputs delivered to one submodule could modulate responses of the second, anatomically distinct submodule driven by area A. Ramping input from area A produced a corresponding activity ramp in A-recipient 5-HT neurons, qualitatively matching the 5-HT population responses observed in vivo during reward anticipation (Zhong et al., 2017). This population response was minimally influenced by low-frequency LHb inputs (Fig. 6B-C, left), yet high-intensity LHb drive (20 Hz) sharply disrupted A-driven 5-HT activity ramps (Fig. 6C-D, right) despite these LHb inputs not directly contacting A-recipient 5-HT neurons. To further characterize the interaction between 5-HT submodules, we computed the A-driven input-output transformation. The gain of this transformation was not altered by low-frequency LHb input (5 Hz), but was strongly depressed by high-frequency LHb input (Fig. 6D). Thus,

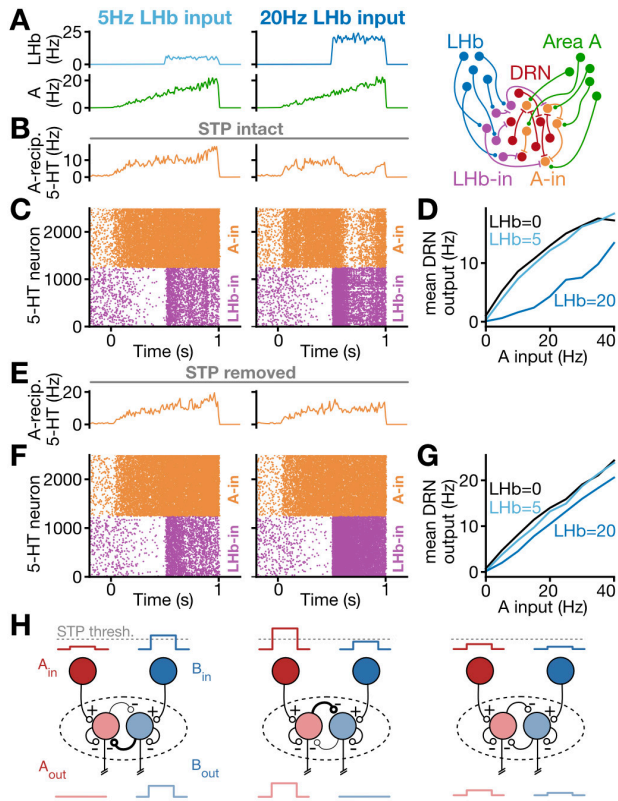


Figure 6: Winner-take-all dynamics generated by recurrent short-term plasticity. **(A)** Simulations with Lhb inputs of either 5Hz (left) and 20Hz (middle) were delivered to 5-HT neurons during ramp input from area A. Schematic of LIF network model with two inputs (right). **(B)** Mean responses of 5-HT neurons to 5Hz (left) and 20Hz (right) input with STP rules intact. **(C)** Raster plot showing 5-HT neuron responses to 5Hz (left) and 20Hz (right) input with STP rules intact. **(D)** Input-output function, for input A, with STP rules intact and coincident Lhb input at 0, 5 or 20Hz. **(E)** Mean responses of 5-HT neurons to 5Hz (left) and 20Hz (right) input with STP rules removed. **(F)** Raster plot showing 5-HT neuron responses to 5Hz (left) and 20Hz (right) input with STP rules removed. **(G)** Input-output function, for input A, with STP rules removed and coincident Lhb input at 0, 5 or 20Hz. **(H)** Schematic of winner-take-all computation with recurrently connected submodules. When one input crosses an STP threshold, output from that submodule dominates while other submodules are suppressed through recurrent inhibition (left and centre). When neither input crosses an STP threshold, outputs from both submodules persist (right).

strong LHb inputs to DRN have the ability to suppress anatomically separate submodules through recurrent inhibition, implementing a winner-take-all computation. Furthermore, additional simulations where we decreased the recurrent connection probability tenfold yielded qualitatively similar results (Fig. S9), indicating that these processing rules were not confined to narrow implementation of different network connectivity features.

To determine whether the short-term plasticity rules were necessary for the nonlinear computations, we performed additional simulations with *in silico* removal of the short-term dynamics (i.e., with linearly summing 5-HT synapses) while preserving recurrent architecture. These simulations revealed that without the STP dynamics, A-driven 5-HT ramps were largely unaffected by either weak or strong LHb input (Fig. 6E-F). Furthermore, without STP, the gain of the A-driven input-output transformation was not significantly modulated by LHb input strength (Fig. 6G). Together, these results indicate that the short-term dynamics of the recurrent 5-HT synapses are necessary to generate winner-take-all computations (Fig. 6H) where strong inputs (e.g. acute threat information from LHb) can inhibit weaker input-driven responses to dominate network output.

We next tested the biological plausibility of these winner-take-all rules through additional slice physiology experiments, where either habenula input or local recurrent connections were optogenetically stimulated to gate spiking output. First, a sparse ChETA expression paradigm was employed in DRN as described earlier (Fig. 2; Fig. 7A). In ChETA-negative 5-HT neurons, optical stimulation elicited outward 5-HT_{1A}R currents (Fig. 7B). To directly test whether recurrent connections could gate spiking output of (ChETA-negative) 5-HT neurons, we examined the ability of high-frequency (20 Hz) optical stimulation of DRN ChETA-expressing neurons to perturb the persistent spiking induced by direct current injection of 5-HT neurons (Fig. 7C-7D). Consistent with predictions from our computational model, photostimulation triggered a significant reduction in spiking (Fig. 7D-E) which was mediated by 5-HT_{1A}Rs (Fig. 7F-L). These data indicate that within DRN, recurrent inhibition can allow populations of 5-HT neurons to significantly shape spiking output of separate 5-HT

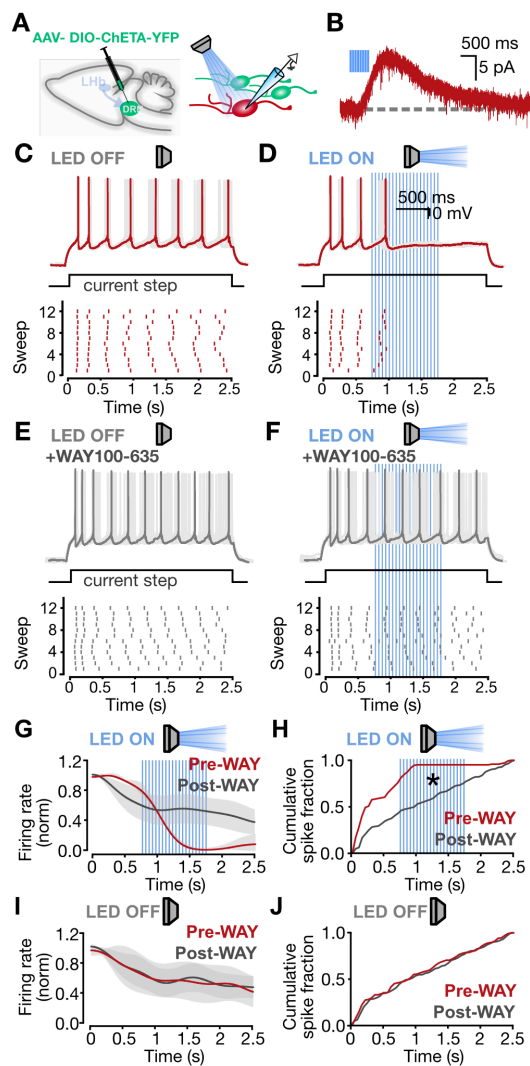


Figure 7: Recurrent inhibition gates spiking activity of 5-HT neurons. **(A)** Schematic of experimental design. Left, AAV-DIO-ChETA-YFP is injected into the DRN at a 2:1 ratio with Mannitol. Right, whole-cell recordings from ChETA-negative 5-HT neurons. **(B)** Recurrent 5HT1AR-mediated current in ChETA-negative 5-HT neuron. **(C)** Spiking of a 5-HT neuron in response to current injection. Below, raster plot showing the same neuron's activity across trials. **(D)** Spiking of a 5-HT neuron in response to current injection and 20Hz optical activation of CHETA-positive 5-HT neurons. Below, raster plot showing the same neuron's activity across trials. **(E)** Spiking of a 5-HT neuron in response to current injection in the presence of WAY100-635. Below, raster plot showing the same neuron's activity across trials. **(F)** Spiking of a 5-HT neuron in response to current injection and 20Hz optical activation of CHETA-positive 5-HT neurons, in the presence of WAY100-635. Below, raster plot showing the same neuron's activity across trials. **(G)** Smoothed firing rate (normalized) over time for pre- and post-WAY recordings with LED stimulation (mean \pm SE, $n=3$ neurons). **(H)** Cumulative spike fraction in pre- and post-WAY recordings with LED stimulation ($n=3$ neurons, 2-sample Kolmogorov-Smirnov test: $p = 7.4e-15$). **(I)** Smoothed firing rate (normalized) over time for pre- and post-WAY recordings without LED stimulation (mean \pm SE, $n=3$ neurons). **(J)** Cumulative spike fraction in pre- and post-WAY recordings with LED stimulation (3 neurons, 2-sample Kolmogorov-Smirnov test: $p = 0.85$).

populations. To ensure that the plausibility of these results extend to activation of long-range synaptic input, we found that high-frequency optical stimulation of LHb afferents in slice reliably inhibited spiking output of 5-HT neurons through 5-HT_{1A}Rs (Fig. S10).

Together, these findings indicate that the unique short-term dynamics present at DRN 5-HT synapses can facilitate key emergent network computations, including an input frequency-dependent polarity switch leading to network shutdown at high input rates, and winner-take-all dynamics where recurrent inhibition strongly suppresses submodules with weaker inputs.

Learned cue-reward associations are conditionally modulated by habenular drive to the DRN.

A synthesis of these local computations enacted by the DRN recurrent inhibitory motif makes circumscribed predictions about how the statistics of synaptic inputs to DRN should modulate 5-HT output and, by extension, DRN-dependent behaviors. Our ethological inspiration comes from recent work demonstrating that dynamics of DRN and LHb can represent behavioral state transitions (Marques et al., 2020; Amo et al., 2014; Lecca et al., 2020). Moreover, cue-reward learning has been shown to lead to a gradual increase in the population activity of 5-HT neurons during reward anticipation, peaking around reward delivery (Zhong et al., 2017). Because the internal state of reward anticipation can lead to measurable motor outputs such as post-cue anticipatory licking whose dynamics are modulated by 5-HT activity (Matias et al., 2017), we set out to develop a head-fixed classical conditioning task where we could optogenetically manipulate LHb inputs to DRN while monitoring licking behaviors as a metric of reward anticipation-linked internal state.

Our experimental and modeling work so far predicts that the statistics of incoming habenula input should determine whether DRN 5-HT neurons are excited or inhibited over protracted timescales: low-frequency LHb inputs to DRN should promote reward anticipation-linked activity ramps, while high-frequency LHb (e.g. severe threats), should suppress this activity by recruiting strong short-term facilitation of recurrent inhibition within DRN. Mice received a bilateral injection of AAV-ChR2-mCherry in LHb while an optic fiber was implanted in

DRN and a custom 3D-printed head bar affixed to the skull (Fig. 8A; see Methods). Mice underwent water deprivation before being trained on an auditory classical conditioning paradigm where one of three auditory cues (discriminable single frequency tones; de Hoz & Nelken, 2014) was randomly presented for each trial, followed by delayed delivery of either a small (5uL), a large reward (10uL), or no reward (Fig. 8B). Trained mice displayed robust cue-driven anticipatory licking, a conditioned response which indicates expression of a learned stimulus-reward association (Cohen et al., 2012). This conditioned licking response was strongest after large reward cues (Fig. 8C, left), and we also observed robust reward-driven licking (Fig. 8C, right). Animals were considered trained if they reached a criterion performance of anticipatory licking during large reward trials (see Methods).

To assess the behavioral impact of optogenetically activating the habenula input to the raphe, we devised a trial structure interleaving optogenetically stimulated and non-stimulated trials to permit within-animal behavioral comparisons for each discrete reward value (Fig. 8D, top). During stimulated trials, LHb afferents in the DRN were optogenetically activated by brief pulses of blue light during the period between cue and reward presentation (see Methods). During non-stimulated large reward trials, animals displayed robust conditioned anticipatory licking after the auditory cue (Fig. 8E, black). High frequency (20Hz) photostimulation of LHb afferents induced a striking and rapid decrease in anticipatory licking which recovered rapidly after stimulation ceased (Fig. 8E, blue). In contrast, in agreement with our quantitative model predictions, low-frequency 5 Hz optogenetic stimulation had no effect on conditioned anticipatory licking (Fig. 8F). Anticipatory behavior was quantified by calculating the difference in mean cue-driven anticipatory licking between non-stimulated and interleaved stimulated trials for each animal; this analysis revealed a significant group difference between 20 Hz and 5 Hz stimulation on acute reward anticipation for large-reward trials (Fig. 8G).

As current theories postulate that 5-HT neurons represent valence information during associative learning (Liu et al., 2017), we next focused on the effect of photostimulation during small reward and no reward trials. We hypothesized that since these trials had lower expected

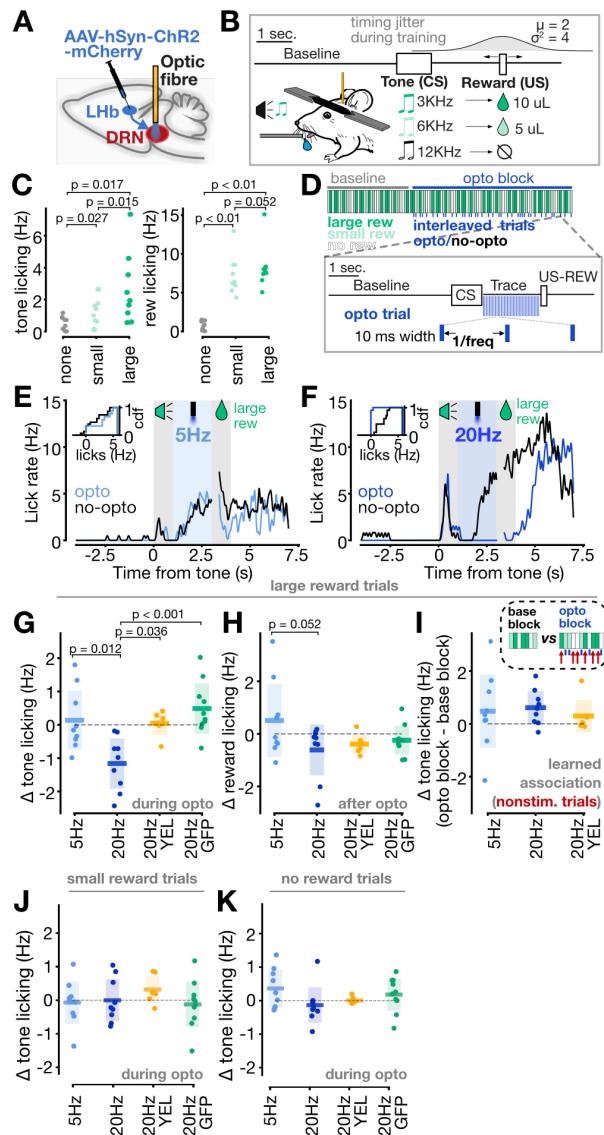


Figure 8: Habenula input to raphe disrupts the expression of learned reward associations in a value- and frequency-dependent manner. **(A)** Schematic illustrating experimental setup (left) and confocal image of Lhb injection site (right). **(B)** Schematic of head-fixed classical conditioning task. **(C)** Mice learned to lick in response to the small reward-predicting and large reward-predicting cues (left), and displayed elevated licking after reward delivery (right) ($n=9$ mice). **(D)** Schematic of trial structure for optogenetic manipulations. **(E)** Mean licking across all large reward trials for a single mouse depicting control trials (black) or 5Hz optogenetic stimulation trials (blue). Inset shows cumulative distribution of licking for all large-reward trials. **(F)** Mean licking across all large reward trials for a single mouse depicting control trials (black) or 20Hz optogenetic stimulation trials (blue). Inset shows cumulative distribution of licking for all large-reward trials. **(G-I)** Summary statistics for optogenetic manipulation during post-tone period of large-reward trials. **(G)** Change in conditioned licking response during tone presentation, between control and optogenetic stimulation trials during opto block (paired two-sided t-tests). **(H)** Mean reward licking change between control and optogenetic stimulation trials (paired two-sided t-tests). **(I)** Change in post-tone licking for nonstimulated trials between baseline block and optogenetic stimulation block. **(J)** Change in post-tone licking between control and optogenetic stimulation trials (small reward trials). **(K)** Change in post-tone licking between control and optogenetic stimulation trials (no reward trials).

reward values, activating LHb inputs to DRN would have a lesser impact on reward anticipation compared to large reward trial stimulation. During small reward trials, neither 20Hz or 5Hz photostimulation modulated anticipatory licking (Fig. 8J; Fig. S11B-C) despite the fact that mice showed anticipatory licking during interleaved non-photostimulated trials (Fig. S11B-C; Fig. 8C). For trials where no reward was delivered, photostimulating at 5Hz after the auditory tone triggered a slight, but at times clearly noticeable, increase in licking which was not seen during 20Hz photostimulation (Fig. 8K; Fig. S11D-E). These contextual effects on reward anticipation, whose sign and magnitude depend on both reward value and LHb input statistics, are broadly consistent with a model where recurrent inhibition leads to antagonistic competition between distinct valence representations spanning multiple timescales and submodules in DRN (Cohen et al., 2015; Ren et al., 2018).

To control for visual artifacts, we stimulated at 20 Hz with yellow light, which is beyond the excitation spectrum of ChR2 (Miyazaki et al., 2018; see Methods) and found no visual cue-induced effect on conditioned licking (Fig. 8G; yellow). As a separate control, in another cohort where AAV-GFP was injected bilaterally in LHb and an optic fiber implanted in DRN, 20 Hz stimulation with blue light did not elicit changes in conditioned licking (Fig. 8G, green). Finally, to rule out nonspecific stimulation-induced motor effects, we additionally photostimulated at high frequencies during the reward consumption phase and observed no significant modulation of licking in these conditions (Fig. S11F-G).

It is possible that the LHb inputs to the DRN may affect expression of conditioned responses or reward consumption over longer timescales, either within the same trial or across trials, possibly reflecting changes in cognitive flexibility or learning (Matias et al., 2017; Zhong et al., 2017). We assessed this in two ways. First, we examined licking behavior after acute stimulation had ended (i.e., during reward delivery) and compared stimulated and non-stimulated trials. 20Hz stimulation slightly but nonsignificantly reduced licking during reward consumption compared to 5Hz stimulation, likely reflecting a seconds-long persistent effect of stimulation consistent with the timescale of GIRK activation (Fig. 8H). Second, we considered

whether LHb activation updated the underlying cue-reward association within the session, which would manifest as changes in anticipation during interleaved, nonstimulated trials. We calculated the difference in conditioned licking response between the baseline block (first 40 trials) and nonstimulated trials during the opto block (last 80 trials), revealing no long-term changes in reward anticipation (Fig. 8I). From this, we conclude that at high input frequencies, the LHb input to the raphe acutely modulates behavior over timescales of hundreds of milliseconds, consistent with the timescales of GIRK inhibition uncovered in slice, yet does not modulate the conditioned association itself.

We additionally considered how behavioral variability between animals correlated with the effects of stimulation. We found that animals with more anticipatory licking experienced slightly greater suppression of licking upon high frequency stimulation (Fig. S11H). However, at low stimulation frequencies, there was a slight positive relationship between anticipation and stimulation effects (Fig. S11H). Finally, we observed no correlation between effects of high and low stimulation frequency across animals (Fig. S11I).

Together, these results indicate that LHb input to DRN suppresses reward anticipation in a frequency-dependent and temporally precise manner, exclusively during high reward value states. At high input frequencies, where strong STP-dependent recurrent inhibition is expected to dominate, anticipatory licking was suppressed; yet, at lower input frequencies where network excitation dominates, anticipatory licking was either not altered, or even enhanced slightly. During low reward value states, LHb input does not appear to substantively modulate reward anticipation. These behavioral results highlight how habenula inputs to raphe can alter behavioral state dynamics over protracted timescales in a manner contingent on input statistics and short-term recurrent dynamics.

Discussion

Here, using a combination of optogenetic, electrophysiological, computational and behavioral approaches, we provide evidence that DRN is organized in an unsuspected recurrent inhibitory

network, characterized by strong short-term release dynamics active over protracted timescales. We outline a core set of computations supported by this connectivity motif when activated by long-range excitatory LHB input, including a reversal of network response polarity at high input frequencies, and a winner-take-all computation where anatomically independent raphe subsystems can be antagonistically coupled via recurrent inhibition. Building on key predictions from this network modeling, we found behavioral manifestations of this non-linear network computation of the DRN in a frequency-dependent inhibition of reward anticipation triggered by LHB inputs to the DRN. These results add to the reported role of 5-HT in modulating cognitive processing, not direct appetitive or aversive proximal behaviors (Fonseca et al., 2015; Miyazaki et al., 2018) and links a local circuit's network architecture, and associated synaptic dynamics, to a core set of ethologically relevant computations.

We present several lines of experimental evidence supporting a connectivity model whereby 5-HT neurons in the DRN are organized in a local recurrent inhibitory network, including the observation that photostimulation of a subset of opsin-expressing 5-HT neurons induced robust 5-HT_{1A}R-mediated hyperpolarization/outward currents in non-opsin-expressing 5-HT neurons postsynaptic to the stimulated ensemble. This experimental architecture represents in effect an optogenetic equivalent of paired recordings from synaptically-connected neurons, which have provided decisive functional validation of the recurrent connectivity nature of several networks in the brain (Miles & Wong, 1986; Mason et al., 1991; Markram et al., 1997; Debanne et al., 2008). These experiments, unambiguously showing functional connections between 5-HT neurons, supports a model which is largely inconsistent with the widely and long-held idea that 5-HT neurons are regulated in an autocrine fashion by dendritically-released 5-HT acting on so-called 5-HT_{1A}R somatodendritic autoreceptors (Gozlan et al., 1983; Weissmann-Nanopoulos et al., 1985; Pineyro & Blier, 1999). This autoreceptor hypothesis makes several experimental predictions which we directly tested and failed to find support for. While our results do not formally allow to distinguish whether interconnected 5-HT neurons communicate through dendro-dendritic or canonical axo-

dendritic functional contacts, or whether release of 5-HT is spatially confined or consists of more spatially diffuse volume transmission (Maxwell et al., 1983; Ridet et al., 1994; Bunin & Wightmann, 1998), we nonetheless note that classic axonal release provides the most parsimonious and credible model accounting for our results. While a subset of 5-HT dendrites contains 5-HT containing vesicles (Descarries et al., 1982; Chazal and Ralston, 1987), 5-HT containing axon terminals are nonetheless widespread in the DRN and are structurally indistinguishable from those found elsewhere in the brain (Chazal and Ralston, 1987). Importantly, this view does not preclude the presence of NMDAR-dependent dendritic release of 5-HT that may operate in distinct physiological conditions (Colgan et al., 2012).

While there is compelling evidence that 5-HT is packaged in, and released from, vesicles (De-Miguel & Trueta, 1995), and mechanisms of 5-HT release in DRN have begun to be investigated (Courtney & Ford, 2016), an algorithmic description of its release dynamics over multiple timescales, and their underlying computations has been lacking. Through a series of minimal electrical stimulation experiments, we obtained strong evidence that 5-HT is released probabilistically, with complex facilitation and depression dynamics which could be precisely captured by a kernel-based synaptic release model. Previous work, relying on bulk monitoring of 5-HT release by fast-scan cyclic voltammetry, in contrast showed quasi-deterministic release of 5-HT and minimal short-term dynamics (Bunin & Wightman, 1998). Our approach rather relied on the non-desensitizing 5-HT_{1A}R-activated GIRK conductance as a reporter of 5-HT release (Fig. S6), and a statistical analysis of these events showed effective stimulation of a few 5-HT synapses (perhaps as little as one) was consistent with probabilistic release of 5-HT. It is likely that the distinct granularity of this approach compared to that afforded by voltammetry underlies some of these discrepancies. Irrespective of this point, it is still intriguing to emphasize that the nature of the short-term dynamics of 5-HT release stands in stark contrast to the robustly depressing dopamine release dynamics recently described in striatum (Liu et al., 2018, Liu et al., 2021). Especially considering that several features of neurotransmitter release are dictated by the postsynaptic target cell type (Geddes et al., 2016;

Koester & Johnston, 2005; Campagnola et al., 2022), future work will be required to determine whether the short-term synaptic dynamics reported here in the DRN are shared by other 5-HT synapses in the brain. These points of discussion nonetheless collectively point to our still rudimentary understanding of monoaminergic transmission dynamics and their underlying computations, especially when compared to those described for fast neurotransmitter release (Abbott & Regehr, 2004; Jackman & Regehr, 2017).

Through a series of constrained network models, we inferred several computations supported by the inhibitory recurrent network in DRN. Previous experimental and theoretical work has identified a number of possible computational functions of network recurrence, including winner-take-all competition (Coultrip et al., 1992), signal amplification and selectivity (Douglas et al., 1995) and probabilistic inference (Echeveste et al., 2020); however, the potential computations performed by recurrent inhibition over timescales of seconds were not immediately obvious. While canonical winner-take-all implementations require simple feedback inhibition from a shared interneuron (Coultrip et al., 1992), the implementation in DRN has several unique features. The extended timescale of the inhibitory GIRK conductance, lasting seconds, adds a history component to the winner-take-all computation. Additionally, the strong facilitation at recurrent inhibitory synapses amplifies inherent winner-take-all dynamics: past some facilitation threshold, as strongly activated groups of neurons inhibit less activated groups, they additionally receive less recurrent inhibition from these groups (Fig. 6H). Our model is experimentally grounded in several fundamental ways. First, individual 5-HT neurons were modeled with parameters extracted from an established integrate-and-fire framework (Harkin et al., 2021, Pozzorini et al., 2015). Second, the short-term dynamics at recurrent connections were modeled using a linear-nonlinear formalism (Rossbroich et al., 2021) which precisely captured responses to a test set of patterned input. Third, the core network computations persisted over a wide range of possible recurrent connectivity schemes, while *in silico* removal of the short-term recurrent dynamics themselves disrupted these computations. Finally, we tested key physiological predictions of a winner-take-all network, including that

serotonergic spiking can be suppressed by activation of recurrently connected neurons. By developing a formalism linking synaptic dynamics with emergent computations, we have provided a flexible architecture for understanding how local serotonergic circuitry is modulated by long-range inputs.

Our results show that the recurrent inhibitory dynamics lead to a non-monotonic input-output relationship in the DRN by supporting a strong excitation-driven inhibition, where larger excitatory inputs generate greater network suppression. A natural question is how this computation compares with gain control associated with classical feedforward inhibition motif. A range of distinct gain control regimes have been described with feedforward inhibition such as subtractive and divisive inhibition (Ayaz & Chance, 2009) and even non-monotonic inhibition (Mejias et al., 2014; Lewis et al., 2007). However, the recurrent inhibition motif described here is unique in that the phasic network response is dominated by direct excitation and linearly reports input intensity, while the delayed steady-state response exhibits non-monotonic behavior with increasing excitation. In principle, these temporally distinguishable epochs of 5-HT output dynamics could support multiplexed codes such as valence encoding over multiple timescales (Cohen et al., 2015). Our network model additionally uncovered an emergent winner-take-all-like computation in the DRN, where strong activation of connected modules of 5-HT neurons can inhibit activity of others. This builds on previous work describing anatomically segregated DRN subdomains (Ren et al., 2018), and implies that a winner-take-all computation governs behavioral state selection between submodules. While these simulations uncovered general algorithmic principles which were tested with a range of connectivity profiles in our network models, more precise quantitative information on the recurrent connectivity between submodules (Ren et al., 2018) or on the additional recruitment of feedforward GABAergic inhibition by some inputs (Zhou et al., 2017; Geddes et al., 2016) should bring about refinement of these general principles.

By combining patterned in vivo optogenetic stimulation of LHb inputs to the DRN with an auditory classical conditioning task, we were able to test key predictions from our network

modeling, and found that high-frequency, but not low-frequency, activation suppressed anticipatory licking before large rewards. While we optically stimulated LHb inputs in the DRN at 5Hz and 20Hz, these frequencies do not represent a one-to-one mapping onto habenular firing rates observed in vivo, since release probability is artificially inflated during optogenetic stimulation, and real population firing rates are likely highly heterogeneous and asynchronous over time. Thus, these frequencies are likely underestimates of the dynamic range of LHb firing rates within which the frequency-dependent inversion in the DRN occurs in situ. Nonetheless, a key finding of these experiments is that the behavioral effect of high-frequency activation was temporally precise, lasting several hundred milliseconds (in line with the timescale of a 5-HT_{1A}R-mediated GIRK activation) and did not extend beyond the running trial. Furthermore, the effect of optogenetic stimulation was specific to the anticipatory licking period, and did not reflect a non-specific, general inhibitory role on motor output since it was ineffective when administered during reward consumption. As a whole, these findings are consistent with previous work showing that optogenetically activating DRN promotes waiting, but is not rewarding and does not directly drive motor actions in and of itself (Fonseca et al., 2015; Miyazaki et al., 2018). More broadly, our behavioral results provide evidence that LHb inputs to the DRN rapidly and reversibly alter behavioral state, which is consistent with previous findings on the habenulo-raphe pathway (Amo et al., 2014; Marquez et al., 2020) and may represent an important role for this pathway in dynamic state selection.

Animals often exhibit sharp thresholds when transitioning between ethologically relevant behavioral states such as exploitation and exploration (Marques et al., 2020). These state transitions require brain circuits which continually evaluate graded representations of either internal or external cues, enacting thresholding operations on behavioral timescales. For example, threat may be best conjectured as being cautiously invigorating when small (go), but nefarious when strong (no-go). Building on recent theories (Liu et al., 2020; Luo et al., 2016; Cohen et al., 2015; Marques et al., 2020), the winner-take-all computation between ensembles of recurrently connected 5-HT neurons, operating over protracted timescales, satisfies many of

the conditions required for the largely exclusive, non-overlapping implementation of distinct naturalistic behaviors (Ren et al., 2018).

Supplemental Figures

Please see Appendix F for all supplemental information.

Acknowledgements

This work was supported by grants from the Canadian Institute for Health Research (JCB; RN), the uOttawa Brain and Mind Institute and the Canadian Heart and Stroke Foundation (JCB), Brain Canada (JCB), and the Canadian Foundation for Innovation (JCB). We thank members of the Beique, Naud and Maler lab for helpful discussions. We would also like to thank Drs Salah El Mestikawy and Sylvain Williams for providing us with the SERT-CRE mouse line.

Author contributions

Conceptualization, M.B.L., S.G., S.H-D., R.N., and J.-C.B.; Methodology, M.B.L., S.G. ; Software, M.B.L.; Formal Analysis, M.B.L.; Investigation, M.B.L., S.G., M.C., E.H., S.M., E.H.-G., ; Resources, M.B.L., E.H. and R.N. ; Writing – Original Draft, M.B.L., S.G.; Writing – Review & Editing, M.B.L., R.N., and J.-C.B.; Supervision, J.-C.B.; Funding Acquisition, J.-C.B.

Declaration of interests

The authors declare no competing interests.

Data Availability

The datasets generated during the current study are available from the corresponding author on reasonable request.

Code Availability

The code generated during the current study are available from the corresponding author on reasonable request.

Methods

Viral Constructs and Retrograde Tracer

Construct	Source	Titre (IU/ μ l)
AAV9.CAG.hChR2(H134R)-mCherry.WPRE.SV40	<i>(Penn Vector Core, AV-9-20938M)</i>	$\sim 10^{12}$
AAV9.Efla.DIO.ChETA(E123T/H134R)-eYFP.WPRE.hGH	<i>(Penn Vector Core, AV-9-26968P)</i>	$\sim 10^{12}$
rAAV2-retro-hSyn1-cre	<i>(Janelia Farm Research Campus)</i>	$\sim 10^{13}$
AAV2/9.emCBA.eGFP	<i>Canadian Neuorphotronics Platform, University of Laval</i>	$\sim 10^{13}$

Animal use

SERTcre: B6.Cg-Tg(Slc6a4- cre)ET33Gsat/Mmcd

R26R-TdT: Rosa26R-TdTomato fluorescent protein (TdT) reporter mice

All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services. For this study, SERTcre mice were crossed with the R26R line to generate a SERTcre-TdT mouse line. Mice were kept on a 12:12 hour light/dark cycle, with access to food and water ad libitum unless otherwise noted. A mix of male and female SERTcre-TdT mice (1-3 months for slice physiology experiments, 2-6 months for behavior experiments) were used.

For the behavior portion of the study, mice with a C57BL/6 background were utilized, either originating as mice with a wild-type genotype originating from the SERT-CRE-tdTomato line, or as mice purchased from Charles River Labs. Mice were kept on a 12:12 hour reverse light/dark cycle and were water restricted to 85-90% of their initial weight during behavioral testing.

Slice physiology

Stereotaxic Injections

All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services. For stereotaxic injections, mice were injected with 0.05 mg/kg buprenorphine and anesthetized by inhalation of isoflurane. Injections were performed using a 10 μ l Hamilton syringe with a 33 gauge needle and 0.4 - 0.5 μ l was injected per coordinate. The virus was injected at a rate of 0.1 - 0.2 μ l per minute. Stereotaxic co-ordinates are as follows: LHb (from Bregma skull landmark, in mm: AP -1.7, ML +/- 0.4, DV -3.0) and DRN (from Lambda skull landmark, in mm,

20 degree angle: AP +0.5, ML +1.0, DV -3.0). For retrograde tracing, 0.4 μ l of rAAV-retro-TdTomato was injected. For Chr2 (H134R) expression, 0.4 μ l of AAV9.CAG.hChr2(H134R)-mCherry.WPRE.SV40 (titre $\sim 10^{12}$ GC/mL) was injected bilaterally and animals were left to recover for 3 to 4 weeks prior to ex-vivo electrophysiological recordings. In anterograde tracing, AAV-GFP was injected in the Lhb and animals were perfused 2-3 weeks post-injections. For experiments requiring sparsely labeled cells in the DRN with ChETA-YFP, we further diluted virus in a 20 % mannitol/saline solution at a ratio of 2:1 (mannitol solution: virus solution). In these experiments, the virus solution was AAV9.Ef1a.DIO.ChETA(E123T/H134R)-eYFP.WPRE.hGH. The mixed mannitol and virus solution was then injected in the DRN at the coordinates above.

Slice preparation

Lhb and DRN containing brainstem slices were prepared from 8 to 10 week old mice (3 to 4 weeks post-stereotaxic injections). For electrical stimulation experiments, brainstem slices were instead prepared from P21-P30 mice. Acute slices were prepared similarly to previously described methods (Geddes et al., 2016). In brief, mice were anesthetized by isoflurane inhalation (Baxter Corporation, Canada) and sacrificed by decapitation. The brain was removed and brain slices were sectioned while immersed in ice-cold choline chloride-based cutting solution of the following composition (in mM): 119 choline-Cl, 2.5 KCl, 1 CaCl₂, 4.3 MgSO₄-7H₂O, 1 NaH₂PO₄, 1.3 sodium L-ascorbate, 26.2 NaHCO₃, and 11 glucose, and equilibrated with 95% O₂, 5% CO₂. Slices were recovered in a chamber containing standard Ringer's solution of the following composition (in mM): 119 NaCl, 2.5 CaCl₂, 1.3 MgSO₄-7H₂O, 1 NaH₂PO₄, 26.2 NaHCO₃, and 11 glucose, at a temperature of 37°C, continuously bubbled with a mixture of 95% O₂, 5% CO₂. Slices were left to recover for 1 hour in a recovery chamber and equilibrate to a temperature of approximately 21°C until recordings were performed.

Whole-Cell Electrophysiology

DRN neurons were visualized using an upright microscope: (1) Examiner D1 equipped with Dodt-contrast or differential-interference contrast (DIC) (Zeiss, Oberkochen, Germany; 40x/0.75NA objective) or, (2) BX51W1 equipped with DIC optics (Olympus, Center Valley, PA; 40x/0.80NA objective). 5-HT neurons were visually identified by TdTomato signal driven by the SERT promoter under EPI-fluorescence illumination (xCite Series 120pc). Whole-cell recordings were carried out using an Axon Multiclamp 700B amplifier, sampled at 10 kHz, digitized with an Axon Digidata 1440A (or 1550) digitizer and filtered at 2 kHz. Whole-cell recordings were performed using borosilicate glass patch electrodes (3-6 M Ω ; World Precision Instruments, Florida) pulled on a Narishige PC-10 pipette puller (Narishige, Japan). All experiments were performed at room temperature (except electrical stimulation experiments, conducted at a measured 30-32°C through the use of an in-line bath heater and stage heater.) Slices received constant bath application of Ringer containing (in mM): 119 NaCl, 2.5 KCl, 1.3 MgSO₄, 2.5 CaCl₂, 1.0 NaH₂PO₄, 11 glucose, 26.2 NaHCO₃, and 0.1 L-Tryptophan saturated with 95% O₂ and 5% CO₂ (pH = 7.3; 295-310 mOsm/L). Excitatory post-synaptic currents (EPSCs) were recorded using an intracellular solution of either of the following compositions: (1) 115 mM cesium methane-sulfonate, 5 mM tetraethylammonium-Cl, 10 mM sodium phosphocreatine, 20 mM HEPES, 2.8 mM NaCl, 5 mM QX-314, 0.4 mM EGTA, 3 mM ATP(Mg²⁺), and 0.5 mM GTP, pH 7.25 (adjusted with CsOH; osmolarity, 280–290 mOsmol/L). (2) For 5HT1AR current recordings, an internal solution of the following composition was used: 115 mM potassium gluconate, 20 mM KCl, 10 mM sodium phosphocreatine, 10 mM HEPES, 4 mM ATP(Mg²⁺), and 0.5 mM GTP, pH 7.25 (adjusted with KOH; osmolarity, 280–290 mOsmol/L). For all voltage clamp recordings, access resistance was continuously monitored using a 125 ms, 2 mV hyperpolarizing pulse, at least 245 ms prior to stimulation. Recordings were discarded when the access resistance changed by >30%. Liquid junction potential was not compensated for.

LED Photostimulation

LED photostimulation was delivered using a PlexBright LED module and controller (465 nm; Plexon, Texas). Light was delivered through a 200 μm patch fiber cable with a bare fiber tip with ~ 1 cm of glass exposed. For EPSC/EPSP stimulation, unless otherwise stated, optically evoked-PSCs were stimulated at a frequency of 0.1 or 0.067 Hz with light pulses of 1-3 ms in length and an irradiance of 792 mW/mm^2 . In driving 5-HT_{1A}-mediated currents, optical stimulation was accomplished using 10 ms pulses at 20 Hz for 400 ms at an irradiance of 792 mW/mm^2 .

Immunohistochemistry

Mice were perfused using a minimally invasive method developed by Kenneth D. Eichenbaum and colleague in 2005. Briefly, mice were put under general anaesthesia with 1.5% isoflurane. The animal was then placed on its back for injection (1 mL PFA4% in 0.1 M PBS). The injection was done with a 1 mL syringes equipped with a 26G needle over a 30 seconds course. The correct position of the needle was confirmed by the appearance of blood in the needle hub. Brains were then put in 10 mL of PFA4% overnight and then transferred in a 15 mL tube with 10 mL of PBS (0.1 M PBS, 0.1% Sodium azide). The fixed brains were then cut in 50 μm thick slices with a vibrotome (Vibrotome, 1500 sectioning system).

Free-floating sections were blocked in 1% BSA + 0.03% Triton X-100 (TX-100) for 45-60 min, and incubated in primary antibody in blocking solution at 4 °C. On the following morning, slices were washed three times, for 5 min each time, with 0.1% Tx-100 before incubation in the secondary antibody goat anti-rabbit-conjugated Alexa-647 (Life technologies, A21245) for 2 h at room temperature. Slices were washed with 0.1% Tx-100 three times, for 5 min each time, before mounting onto slides with Fluoromount-G (Southern Biothech, 0100-01). The primary antibody used for this paper were Tryptophan hydroxylase 2 (1:1,000; Millipore, ABN60), NeuN (1:1000; EnCor Biotechnology, RPCA-Fox3-AP) and Nissl (1:1000; NeuroTrace® 435/455 Blue Fluorescent Nissl Stain, Life technologies, N21479).

Fluorescent tracer injections in DRN

To examine the local connectivity and network structure in DRN, we performed visually guided injections of fluorescent tracer in DRN in acute slices (described fully in Trinh et al., 2015). Briefly, in these experiments, SERT-CRE-tdTomato animals (P30-P40) were sacrificed and acute slices prepared as described above (*slice preparation*). At 5x-10x magnification with a Zeiss upright microscope with an epifluorescent lamp, Alexa Fluor 488 dextran (suspended at 10 mg/mL in 0.1M phosphate-buffered saline (PBS) pH 7.2) was injected in DRN by iontophoresis (five to ten 400 ms pulses of ~ -100 V via a 7 mOhm electrode). After a fluorescent sphere of <50 μm diameter was observed, injection was stopped. Slices were incubated overnight in oxygenated ACSF to allow anterograde and retrograde transport, fixed in 4% paraformaldehyde and rendered transparent using the SeeDB procedure (see Trinh et al., 2015 for details). Slices were imaged with a Zeiss confocal microscope (LSM800 AxioObserver Z1 microscope; Zen software).

Electrical stimulation of 5-HT afferents

Electrical stimuli were delivered through a bipolar concentric stimulating electrode located >100 μm from the soma of the recorded neuron. Stimuli were 0.1 ms in duration, and consisted of frequencies provided in the text. For all electrical stimulation experiments, we bath-applied a cocktail of synaptic blockers to isolate the 5-HT_{1A}R response (10 μM NBQX, 100 μM picrotoxin, 200 nM CGP 55845). Additionally, for these experiments, a K+MeSO₄ internal solution with 10 μM BAPTA was utilized to prevent the after-hyperpolarization of action potentials from interfering with the recorded 5-HT_{1A}R responses. All 5-HT electrical stimulation experiments were performed at near-physiological temperature (30-32 °C).

Data Analysis

All electrophysiological recordings were analyzed offline using Clampfit 10.4 (Molecular Devices), Origin 8.5 software (OriginLab) and the Numpy and Scipy computational packages within Python. All data are presented as means \pm SEM. n refers to the number of cells recorded. Unless otherwise stated, statistical significances were derived using a standard Student's two-sided t-test ($p < 0.05$). P values of less than 0.05 were considered statistically significant and are indicated with an asterisk(*). Unless otherwise stated, all measurements were taken from distinct neurons.

Drugs and Chemicals

Picrotoxin, 2,3-Dioxo-6-nitro-1,2,3,4-tetrahydrobenzo[f]quinoxaline-7-sulfonamide di-sodium salt (NBQX), DL-2-Amino-5-phosphonopentanoic acid (D,L-APV), picrotoxin, CGP 55845 and WAY 100-635 were purchased from Abcam (Cambridge, MA). Tetrodotoxin (TTX) was purchased from Tocris Bioscience (Ellisville, MO). 4-Aminopuridine (4-AP) was purchased from Sigma Aldrich (St. Louis, MO).

Modeling of synaptic dynamics and network simulations

GIRK conductance

The GIRK conductance kinetics were modelled as a sum of exponentials:

$$g_{GIRK}(t | a, \tau_{decay}, \tau_{rise}) = a(e^{-t/\tau_{decay}} - e^{-t/\tau_{rise}}) \quad (1)$$

where $\tau_{decay} > \tau_{rise}$ and $t > 0$.

The parameters $a, \tau_{decay}, \tau_{rise}$ were fit to single-pulse electrical stimulation data ($n=6$ neurons) using a least-squares approach as follows: We first fit a monoexponential decay function $f(t | \tau_{decay}) = e^{-t/\tau_{decay}}$ to the normalized post-peak decay phase of $g_{GIRK}^{exp}(t)$. This yielded a constrained estimate of the GIRK decay, $\hat{\tau}_{decay}$. Next, fixing τ_{decay} , we found the values of $\hat{a}, \hat{\tau}_{rise}$ that minimized the mean-squared error over the whole observation period.

Each estimate was averaged across the population ($n=6$ neurons):

$$\begin{aligned} \hat{a} &= 0.725 \text{ nS} \\ \hat{\tau}_{decay} &= 308.04 \text{ ms} \\ \hat{\tau}_{rise} &= 85.97 \text{ ms} \end{aligned}$$

These spike triggered changes in conductances were simulated numerically using their equivalent in linear differential equations.

Quantifying 5-HT release dynamics through optogenetic stimulation

In some experiments, we optogenetically stimulated CHETA-expressing 5-HT neurons in slice, while recording 5-HT_{1A}R outward currents from identified 5-HT neurons. By photostimulating at 20Hz and varying the optical pulse number, we could infer basic properties of release dynamics at these serotonin-serotonin synapses.

We assume that the current generated by a single pulse is a random variable following a binomial distribution $B(n, p)$, which reflects the activity of n independent serotonin-serotonin synapses with release probability p . For a train of successive pulses, we assume p and n are fixed, and that each pulse is independent. Then $Q(n_{pulse} = a)$, the distribution of currents given by a pulses, is distributed as a sum of a binomial distributions:

$$Q(n_{pulse} = a) \sim B(a \cdot n, p)$$

Accordingly, the variance of Q as the pulse number a varies is as follows (Sigworth, 1980):

$$\text{var}[Q(n_{pulse} = a)] = (p(1 - p)n)a + \text{noise} \quad (2)$$

Since the variance of the current does not change as a function of pulse number (Fig. S6), we can conclude that the release probability is close to $p = 1$. This probably reflects the artefactual increase in release probability obtained when optogenetically activating synaptic terminals (Zhang & Oertner, 2007), necessitating other approaches to determine release probability under naturalistic conditions (see the section below).

Quantifying 5-HT release dynamics through electrical stimulation

We additionally performed experiments where 5-HT afferents were electrically stimulated in acute slice preparations while recording from identified 5-HT neurons. Electrical stimulation intensities were low (in the perithreshold response range), and stimulation frequency was 1Hz. To quantify successes and failures, we computed the total event charge triggered by each stimuli (by integrating current over time), which formed an *event* distribution containing both putative successes and failures. We constructed a surrogate *noise* distribution by performing an identical analysis on a nonstimulated portion of the same recording for each neuron (*ie* integrating equivalent durations of membrane noise where no stimuli were presented). Next, for each neuron we transformed the *event* and *noise* charge distributions into z-scored distributions (where the z-score was always computed from that neuron's *noise* distribution). Finally, we pooled z-scored distributions across neurons. This allows us to directly compare electrically evoked responses; for example, responses with a z-score > 3 likely represent successes, while responses with a z-score near 0 probably represent release failures. For electrical stimulation at higher frequencies (1-20Hz), we constructed surrogate noise distributions and event distributions in a manner identical to that described above.

For paired-pulse and pulse-chase stimuli, we quantified responses by first computing total charge elicited over the duration of each of the stimuli (integrating current over time). We next computed the total charge elicited by a single electrical stimulation delivered in isolation, and used this to calculate an expected value for purely linear summation of two pulses (*ie* by multiplying the single stimulation-induced charge by 2). Finally, for each neuron, we normalized the stimulus-evoked charges in the paired-pulse or pulse-chase condition by the expected value for linear summation of two individual pulses.

Short-term plasticity

Short-term plasticity dynamics were modeled using a linear-nonlinear approach called the spike-response plasticity (SRP; Rossbroich et al., 2021). It utilizes a possibly complex efficacy kernel to capture spike-triggered changes in short-term plasticity. Namely, each presynaptic spike is convolved with the following efficacy kernel:

$$\mathbf{k}_\mu(t) = a_s e^{-t/\tau_s} + a_l (e^{-t/\tau_l^+} - e^{-t/\tau_l^-}) \quad (3)$$

This kernel consists of multiple timescales of responses. The short component is controlled by τ_s and a_s , and dictates the rapid short-term facilitation dynamics. The long component is controlled by τ_l^+ and a_l , and dictates the more prolonged short-term depression dynamics. τ_l^- is a shaping term that reduces the effect of short-term depression in the early post-spike phases where short-term facilitation dominates. Since the kernel consists of a sum of exponential functions, it can be conveniently written as a dynamical system. Let k reflect the current efficacy state of the synapse ($k = \mathbf{k}_\mu * S$), which is transformed through a nonlinear operation into a spike efficacy, μ . In turn, μ dictates the magnitude of the postsynaptic response given a presynaptic spike, and is implemented as a nonlinear readout of k :

$$f(x) = \frac{1}{1 + e^{-x}} \quad (4)$$

$$\mu = \frac{1}{f(b)} f(k + b)$$

Here, $f(x)$ is the nonlinear function and b is a baseline term that controls whether the STP dynamics evolve sublinearly or supralinearly. In particular, b shifts the sigmoid $f(x)$ left or right, controlling the initial point to which synaptic input is added or subtracted, as well as providing a normalization of μ such that the first synaptic input leads to a unit change in μ . If b is negative (the sigmoid is right-shifted), the dynamics are supralinear, while if b is positive (the sigmoid is left-shifted), the dynamics are sublinear.

Fitting short-term plasticity dynamics to experimental data

We start with a set of paired-pulse electrical stimulation experiments where the time interval ranges from 50ms to 1sec. For our experimental dataset, we had computed, for each time interval t_i , a ratio of the total charge transfer for both pulses q_1 and q_2 to the charge transfer for a single pulse delivered separately:

$$R(t_i) = \frac{q_1 + q_2}{q_1} = \frac{\mu_1 + \mu_2(t_i)}{\mu_1} \quad (5)$$

Where μ_1 and μ_2 relative postsynaptic response amplitudes for the first and second pulse given by the SRP model. We wish to find optimal parameters for the efficacy kernel, $\mathbf{k}_\mu(t | a_s, a_l, \tau_s, \tau_l^+, \tau_l^-)$, such that we arrive at a local minimum in a cost function comparing the experimental and theoretical charges for various delays τ_i of the paired-pulse stimulation.

We start by deriving an analytical expression for (5) in terms of the underlying efficacy kernel. The first spike's efficacy is taken by evaluating the efficacy kernel at time 0:

$$\mu_1 = \frac{1}{f(b)} f(\mathbf{k}_\mu(t = 0) + b)$$

$$\mu_1 = \frac{1}{f(b)} f(a_s + b) \quad (6)$$

The second spike's efficacy is taken by evaluating the sum of an efficacy kernel at time t_i (efficacy from first spike at the time of spike two) and an efficacy kernel at time 0 (second spike's efficacy).

$$\mu_2 = \frac{1}{f(b)} f(\mathbf{k}_\mu(t = 0) + \mathbf{k}_\mu(t = t_i) + b)$$

$$\mu_2 = \frac{1}{f(b)} f(a_s(1 + e^{-t_i/\tau_s}) + a_l(e^{-t_i/\tau_l^+} - e^{-t_i/\tau_l^-}) + b) \quad (7)$$

Substituting (6) and (7) into (5):

$$R(t_i) = \frac{\frac{1}{f(b)} f(a_s + b) + \frac{1}{f(b)} f(a_s(1 + e^{-t_i/\tau_s}) + a_l(e^{-t_i/\tau_l^+} - e^{-t_i/\tau_l^-}) + b)}{\frac{1}{f(b)} f(a_s + b)}$$

$$R(t_i) = \frac{f(a_s + b) + f(a_s(1 + e^{-t_i/\tau_s}) + a_l(e^{-t_i/\tau_l^+} - e^{-t_i/\tau_l^-}) + b)}{f(a_s + b)} \quad (8)$$

Equation (8) and the corresponding experimental results (5) are used to perform least-squares optimization on the parameters $a_s, a_l, \tau_s, \tau_l^+, \tau_l^-, b$. This yields the following parameter estimates:

$$\begin{aligned} a_s &= 1.11 \\ a_l &= -1.80 \\ \tau_s &= 0.128 \text{ s} \\ \tau_l^+ &= 1.27 \text{ s} \\ \tau_l^- &= 0.698 \text{ s} \\ b &= -6.22 \end{aligned}$$

A general expression for spike efficacy given past spiking history

One can also write a solution to the generalized case of calculating the efficacy, μ_n , of a certain presynaptic spike n occurring at time t_n , given a past history of spikes with times $t_1, t_2, t_3, \dots, t_{n-1}$:

$$\mu_n = \frac{1}{f(b)} f\left(\sum_{s=1}^n \mathbf{k}_\mu(t = t_n - t_s)\right) \quad (9)$$

This equation is used to calculate the model's predictions for the matched pulse-chase experimental data.

Computational model of DRN network with recurrent GIRK-coupled 5-HT1AR connectivity, receiving LHb input.

To construct a simplified model of the DRN network, we took a modular approach where each individual component was fitted to experimental data where possible.

Individual 5-HT neurons were fit with a simplified leaky-integrate-and-fire model. The fitting was performed experimentally according to previously described methods (Pozzorini et al., 2015). Briefly, identified serotonin neurons from SERT-CRE-tdTomato mice were recorded in whole-cell current-clamp configuration, and a noisy input stimulus was presented through multiple test and training repetitions. These repetitions were used to train a generalized integrate-and-fire model to each individual serotonin neuron by maximizing the likelihood of the model (for full details, see Harkin et al., 2021).

In this work we focused on basic extracted parameters from this model: g_{leak} , E_{leak} , C_m , $t_{refractory}$, V_{thresh} . While Harkin et al., (2021) utilized a current-based generalized-integrate-and-fire model to represent serotonin neuron firing, here we consider a simplified, conductance-based, leaky-integrate-and-fire model of individual serotonin neurons:

$$\frac{dV}{dt} = \frac{-g_{leak}(V - E_{leak}) - g_{LHb}(V - E_{AMPA}) - g_{pois}(V - E_{AMPA}) - g_{GIRK}(V - E_K) - g_{AHP}^{tot}(V - E_K) - I_{const}}{C_m}$$

(10)

With the fitted parameters:

$$g_{leak} = 1.4 \text{ nS}$$

$$C_m = 0.0876 \text{ nF}$$

$$t_{refractory} = 6.5 \text{ ms}$$

$$V_{thresh} = -52.33 \text{ mV}$$

And the chosen ionic reversal potentials, which, for the purposes of this LIF model, are considered as fixed meta-parameters:

$$E_{leak} = -68.07 \text{ mV}$$

$$E_K = -90 \text{ mV}$$

$$E_{excit} = 0 \text{ mV}$$

Conductances for g_{LHb} , g_{pois} , g_{GIRK} , g_{AHP} were considered with the following dynamics, where s is a synapse state variable which instantaneously increments by 1 whenever a presynaptic neuron fires, and exponentially decays to 0 with a characteristic time-constant:

$$g_x = s_x \cdot g_x^{unitary}$$

$$\frac{ds_x}{dt} = \frac{-s_x}{\tau_x} \tag{11}$$

where x represents conductances of either *LHb*, *pois*, *GIRK* or *AHP*.

The AHP conductance is expressed as a sum of three unitary conductances with distinct timescales.

$$g_{AHP}^{tot} = s_{AHP}^{slow} g_{AHP}^{slow} + s_{AHP}^{med} g_{AHP}^{med} + s_{AHP}^{fast} g_{AHP}^{fast} \quad (12).$$

The following values were chosen for individual parameters of the model. Note that excitatory decays were chosen to reflect a mixture of both AMPA and NMDA currents. Note that these reflect unitary conductance values, and neurons receive multiple LHb (g_{LHb}) and recurrent GIRK (g_{GIRK}) inputs, as well as a constant background rate of Poisson inputs (g_{pois}).

$$\begin{aligned} g_{LHb} &\sim N(\mu = 1.5, \sigma = 0.2) \text{ nS} \\ g_{pois} &\sim N(\mu = 0.5, \sigma = 0.1) \text{ nS} \\ g_{GIRK} &= 0.01 \text{ nS} \\ I_{const} &\sim N(\mu = 10, \sigma = 5) \text{ pA} \\ \lambda_{pois} &= 25 \text{ Hz} \\ \tau_{A+N} &= 20 \text{ ms} \\ \tau_{GIRK} &= 50 \text{ ms} \end{aligned}$$

For the AHP dynamics, the following values were extracted from fits on identified 5-HT neurons using a previously described protocol (see Pozzorini et al., 2015; Harkin et al. 2021):

$$\begin{aligned} \tau_{AHP}^{fast} &= 3 \text{ ms} \\ \tau_{AHP}^{med} &= 30 \text{ ms} \\ \tau_{AHP}^{slow} &= 300 \text{ ms} \\ g_{AHP}^{fast} &= \frac{0.0281 \text{ nA}}{V_{thresh} - E_k} \\ g_{AHP}^{med} &= \frac{0.0333 \text{ nA}}{V_{thresh} - E_k} \\ g_{AHP}^{slow} &= \frac{0.0210 \text{ nA}}{V_{thresh} - E_k} \end{aligned}$$

Note that since the generalized integrate-and-fire model in Pozzorini et al. (2015) extracts a spike-triggered current for the AHP, we convert this to a conductance by dividing by the driving force at spike threshold.

The DRN network model consisted of 5000 5-HT neurons with the LIF dynamics described above. 5-HT neurons were recurrently connected in an unstructured manner with 0.5% random connection probability and 0.1 nS unitary GIRK synaptic conductance ($\sum 2.5 \text{ nS}$ net GIRK inhibition received by each neuron before facilitation/depression dynamics are considered). Short-term plasticity between 5-HT neuron connections was modelled using the linear-nonlinear description fit to data, described in the preceding section.

25% of DRN 5-HT neurons receive long-range input from a cluster of 500 Lhb neurons emitting Poisson spiketrains at rates specified in the text, where each Lhb neuron and 5-HT neuron is randomly connected with 1% probability (ie each 5-HT neuron receives an average of 5 Lhb synapses) with synapse strength specified above. This yields a mean Lhb conductance of $\sum 7.5$ nS for each 5-HT neuron.

Computational model of DRN network with recurrent GIRK-coupled 5-HT1AR connectivity, receiving Lhb and Area X input.

In some simulations (Fig. 6), we considered a DRN network where 25% of neurons received Lhb input, and 25% of neurons received input from a generic area termed Area X. In these simulations, the basic neural architecture and dynamics were identical to that described above, except that 25% of 5-HT neurons now received input from Area X (non-overlapping with the 25% of neurons already receiving Lhb input). Area X consisted of 500 neurons emitting Poisson spikes at rates specified in the text, where each X neuron and 5-HT neuron is randomly connected with 1% probability (ie each 5-HT neuron receives an average of 5 X synapses). This yields a mean X conductance of $\sum 7.5$ nS for each 5-HT neuron. Synapse strength is identical to that of Lhb-5-HT synapses, as are decay kinetics:

$$g_X \sim N(\mu = 1.5, \sigma = 0.2) \text{ nS}$$

$$\tau_{A+N} = 20 \text{ ms}$$

Auditory classical conditioning task

Surgery

All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services. For stereotaxic injections, mice were injected with 0.05 mg/kg buprenorphine 1 hour prior to surgery and anesthetized by inhalation of isoflurane (4% for induction and 1.5-2% for maintenance). The head was shaved, ointment was applied to each eye, and an incision was made above the skull. A dental drill was used to perform skull craniotomies at locations above the coordinates for Lhb and DRN described below.

For stereotaxic injections in Lhb, we utilized the virus AAV9.CAG.hChR2(H134R)-mCherry.WPRE.SV40 at a concentration of $\sim 10^{12}$ GC/mL. Injections were performed using a 10 μ l Hamilton Neuros syringe with point style 4 (12° beveled needle point). 0.4 - 0.5 μ l was injected per coordinate at a rate of 0.1 – 0.2 μ l per minute. Stereotaxic coordinates for injection are as follows for Lhb, from Bregma skull landmark, in mm: AP +1.7, ML +/- 0.4, DV -3.1 to -3.3. Following stereotaxic injections, we affixed a custom stainless steel or 3D-printed Acrylonitrile butadiene styrene (ABS) headbar to the anterior surface of the skull with C&B Metabond Quick Adhesive Cement (Parkell).

For optic fiber implantation in DRN, each animal received a fiber (Doric Mono Fibre-Optic Cannula, 4mm length, part number: MFC_100/125-0.37_4mm_ZF1.25_C45) which was implanted at the following coordinates, from Bregma skull landmark in mm: AP-4.6, ML 0, DV -2.9. In some animals, we used an angled injection to avoid the superior sagittal sinus (20° angle, from Bregma skull landmark in mm: AP-4.6, ML +1, DV -3.25). Optic fibers were slowly lowered to the desired dorsoventral location, and affixed with C&B Metabond Quick Adhesive Cement (Parkell) followed by Jet Denture repair cement (Lang) to provide structural stability to the implant.

After surgery, animals received topical bupivacaine (2%) and intraperitoneal (IP) carprofen (20 mg/kg). Post-operative care consisted of a further bupivacaine treatment at 4-6h post-surgery, and IP carprofen injection 4-6h post-surgery and for the next two days. Animals were allowed to recover from surgery for one week before water restriction was started. All animals were housed separately.

Behavior and water restriction

Animals were placed on water restriction for 7 days prior to the start of the task. During this time, they were habituated to the experimenters and the apparatus through handling and brief periods of head fixation while daily water (0.8 - 1 mL) was provided. After habituation, for 2-3 days each mouse was head-fixed for short bouts and 5-10 μ L apple juice was delivered to a lickport every 1 to 8 s without preceding auditory tones. Following this, mice underwent 1-2 weeks of training on the primary auditory classical conditioning task, described below.

The auditory classical conditioning task consisted of 120 randomly interleaved trials presented to each mouse per day, giving a total reward value of \sim 0.6 - 0.8 mL per mouse per day. Inter-trial intervals were drawn from an exponential distribution with $\lambda = 5$ s, where 2 s $<$ ITI $<$ 14 s to prevent excessively long or short ITIs. Exponentially distributed ITIs give the convenient property that trials follow a Poisson distribution in time, and therefore are memoryless (the mice cannot predict trial onset times). To prevent the mice from licking continuously during the baseline periods, new trial starts were contingent on baseline mouse licking of $<$ 10% (fractional time). Three trial types were randomly interleaved: a *large-reward trial*, consisting of a 3 kHz tone and a 10 μ L apple juice reward, a *small-reward trial*, consisting of a 6 kHz tone and 5 μ L apple juice reward, and a *no-reward trial*, consisting of a 12 kHz tone and 0 μ L apple juice reward. For each trial, the tone lasted 1 s, the trace period between tone and reward delivery was drawn from a Gaussian distribution with $\mu = 2$ s, $\sigma = 4$ s (values of $t_{trace} < 0.5$ s were prohibited), and the apple juice reward was concomitantly delivered with a 500 ms white noise tone to aid learning. Following reward delivery, licking was monitored for 4 s before the trial was terminated. This auditory classical conditioning task was presented daily to mice for 1-2 weeks until criterion performance had been achieved for $>$ 2 consecutive days. Criterion performance was defined as a mean lickrate of $>$ 0.5 Hz during the trace period of large-reward trials. While many mice displayed mean lickrates which were significantly higher than this (Fig. 8C), this was a convenient lower threshold for criterion performance.

After criterion performance was reached, mice experienced the auditory classical conditioning paradigm while we optogenetically activated the habenulo-raphé pathway in some trials. The trial structure consisted of 120 trials of interleaved large-reward, small-reward and no-reward trials as before. Each trial was composed of the same tone and reward associations and durations as described above, except the period between tone and reward was fixed at 2 s. For the first 40 trials, no optogenetic stimulation was delivered (*baseline block*). For the final 80 trials, optogenetic stimulation was delivered on 40 randomly chosen trials (*optogenetic stimulation block*). Optogenetic stimulation was composed of pulses of 10 ms width, at frequencies noted in the text (5Hz or 20Hz), delivered during the period between tone and reward. Optogenetic stimulation was delivered intracranially via an optic fibre implanted in DRN, to which a patch cable was connected before the session.

The behavioral apparatus consisted of a sound-insulated box containing a Raspberry Pi 3B, custom steel head fixation apparatus and equipment for delivering stimuli and acquiring signals. Lick signals were acquired via a custom 3D-printed combined infrared lickometer and stainless steel lickport. Lickport placement was adjusted daily for each mouse to ensure optimal performance and signal acquisition. A solenoid valve (NResearch 2-way NC isolation valve) gated apple juice delivery, and flow rate was calibrated weekly. Speakers connected to

the Raspberry Pi (Logitech Z150) delivered auditory stimuli. We visualized animal behavior with an infrared video camera connected to the Raspberry Pi (Raspberry Pi Camera Module 2 NoIR). In vivo optogenetic stimulation was delivered using a PlexBright LED module and controller (either blue light, 465; or yellow light, 590 nm; as noted in the manuscript; Plexon, Texas) and a PlexBright optical patch cable. A custom script (Python 3.8) on the Raspberry Pi registered lickometer data at 200 Hz while controlling trial structure, auditory stimulus delivery, reward delivery and optogenetic stimulation. Data was stored in hdf5 format.

After behavioral testing was completed for individual mice, they were perfused transcardially with 4% paraformaldehyde, and brains were cut in 50-100 μ m sections. Immunohistochemistry was performed as described above, and the injection sites and optic fiber tracts were verified.

Data analysis

Data was analyzed with custom Python 3.8 scripts, employing the numpy (van der Walt et al., 2011), scipy (Jones et al., 2001), h5py and matplotlib technical computing libraries. Briefly, raw licking data was binarized and converted into lick onset times. For visualizations (Fig. 8E-F), licking signals were calculated by convolving lick onset times with a gaussian kernel ($\sigma = 50$ ms) and averaging across all trials of that type. Optogenetic analyses are described below.

Optogenetically evoked changes in anticipatory (trace) licking were calculated as follows: we considered trials of a given type during the optogenetic stimulation block (last 80 trials), and calculated the mean licking frequency in the period between the tone and reward (anticipatory licking) for each trial. To calculate an optogenetically evoked change in anticipatory licking, the mean anticipatory licking frequency for optogenetically stimulated trials was subtracted from the mean anticipatory licking frequency for non-stimulated interleaved trials during the optogenetic stimulation block. Optogenetically evoked changes in reward licking were calculated in a similar way, but the period considered was the post-reward epoch.

We additionally considered whether optogenetic stimulation modulated behavior beyond the stimulated period (Fig. 8I). For these analyses, we compared nonstimulated trials during the baseline block with nonstimulated interleaved trials during the optogenetic stimulation block. For each trial of a given type, we first computed mean anticipatory licking frequency as described above. Next, to quantify long-term optogenetically evoked changes in anticipatory licking, we subtracted mean anticipatory licking frequency for nonstimulated trials during the optogenetic stimulation block from mean anticipatory licking frequency for trials during the baseline block.

Manuscript IV**Input multiplexing through cell-intrinsic regulation of spike synchrony in the dorsal raphe nucleus**

Michael B. Lynn, Leonard Maler, Jean-Claude Béïque

This manuscript is currently in preparation for submission.

Statement of contribution

For this manuscript, I led project conceptualization and planning, collected all electrophysiology data, performed all simulations and analysis, assembled all figures, and wrote the draft of the manuscript.

Abstract:

The serotonergic dorsal raphe nucleus (DRN) receives diverse long-range synaptic inputs, yet the relative contribution of each input to DRN output spiking patterns is unknown. Here, we use electrophysiological, optogenetic and computational tools to compare functional features of excitatory inputs from lateral habenula (LHb) and prefrontal cortex (PFC) onto DRN 5-HT neurons. Dual-color opsin strategies revealed that single 5-HT neurons receive functionally matched input from both PFC and LHb. Subthreshold features of excitatory post-synaptic potentials, including amplitude and decay, were largely overlapping. However, PFC inputs triggered spikes that displayed significantly higher latency and greater jitter than those from LHb. A support vector machine classifier demonstrated that input identity can be accurately decoded from the spike timing of under ten 5-HT neurons, suggesting that these timing differences can be robustly parsed by downstream circuits. By examining the intrinsic cellular mechanisms in 5-HT neurons underlying these transformations of EPSPs to spikes, we uncovered a prominent T-type calcium conductance which selectively boosts slow events, as well as subthreshold, voltage-dependent membrane noise which controls spike jitter and latency. Together, these results outline a mechanism by which intrinsic properties of 5-HT neurons functionally segregate LHb and PFC inputs into distinct spiking patterns which could, we hypothesize, be decoded by downstream areas innervated by 5-HT axons - a multiplexing operation.

Introduction

Neurons in the central nervous system transform noisy barrages of synaptic inputs into a spiketrain that conveys useful information to downstream circuits. While the nature of the information content of a spike train is unclear, a central question is whether populations of neurons can simultaneously represent independent features of their inputs (termed ‘multiplexing’), and how individual neurons might perform this computation. While time-division multiplexing is common in artificial neural networks (Rumelhart et al., 1986), it is unclear whether this coding strategy could be implemented in complex cortical networks. Several biologically plausible multiplexing schemes may exist in the brain, including burst multiplexing (Naud & Sprekeler, 2018; Payeur et al., 2021), synchrony-division multiplexing (Lankarany et al., 2019; Ratte et al., 2013; Friedrich et al., 2004), and rate coding in high-dimensional neural activity space using largely orthogonal subspaces (Stringer et al., 2019; Nogueira et al., 2023). An open question is how populations of neurons can generate multiplexed signals from their constituent synaptic inputs, and indeed whether multiplexing mechanisms can be generated from cell-intrinsic properties of individual neurons, or alternatively whether large-scale network interactions are required to generate multiplexed neural codes.

Theoretical work has been instrumental in linking intracellular noise processes to the statistics of spike generation. Membrane noise can dramatically enhance the reliability of spiketrains, although this depends on the statistics of the noise (Mainen & Sejnowski, 1995; Ermentraut et al., 2008). The activation of ionic receptors can also interact with membrane noise and synaptic input to shape spiking output or efficiently encode external signals (Schreiber et al., 2004; Marcoux et al., 2016). Intrinsic noise can additionally play a role in enhancing multiplexing of information at the population level, especially for temporal coding schemes (Kamaleddin et al., 2022). Thus, the generation of neural codes at the population level can be influenced by intrinsic, stochastic noise properties.

Here, we explore the mechanistic basis of multiplexing operations in the dorsal raphe nucleus (DRN), a midbrain hub that plays an essential role in controlling serotonin (5-HT) output

to cortical and subcortical structures. The anatomical and functional organization of DRN is heterogeneous and complex, with multiple submodules (Ren et al., 2018), convergent synaptic inputs and defined efferent targets (Weissbourd et al., 2014; Dorocic et al., 2014). DRN 5-HT neurons receive input from key structures representing cognitive and valence information, including strong inputs from medial prefrontal cortex (mPFC) and lateral habenula (LHb) (Geddes et al., 2016; Zhou et al., 2017). These two inputs target partially overlapping anatomical divisions (Zhou et al., 2017) including the ventral section of DRN, yet carry distinct information, with mPFC rate coding persistent cognitive states (Bari et al., 2019) and LHb representing negative reward prediction errors with high temporal precision spiking responses (Bromberg-Martin & Hikosaka, 2011). We undertook a quantitative comparison of these two inputs to 5-HT neurons to explore how such distinct information sources might be unified or multiplexed in the population neural code of DRN cells

Results

Single 5-HT neurons integrate multiple long-range inputs

We first sought to determine whether individual DRN 5-HT neurons receive input from both LHb and PFC, or alternatively whether these inputs contact distinct, non-overlapping populations of 5-HT DRN neurons. To probe the integration of PFC and LHb inputs at single 5-HT neurons, we employed a dual-color opsin approach (Klapoetke et al., 2014) wherein we injected adeno-associated viruses (AAVs) expressing Chronos (excited by blue light) in PFC and Chrimson (excited by red light) in LHb (Fig. 1A). Using an in vitro slice physiology preparation, we recorded in whole-cell configuration from genetically identified 5-HT neurons (SERT-CRE-tdTomato mouse line; see Methods) primarily in the ventromedial and dorsomedial regions of the DRN. In voltage clamp mode ($V_m = -70\text{mV}$), brief (1ms) pulses of blue (473nm) or red (550nm) light were delivered to the slice through an optic fiber. In agreement with previous work (Geddes et al., 2016; Zhou et al., 2017), we observed reliable inward photocurrents (excitatory post-synaptic potentials, EPSCs) which were time-locked to optical stimulation (Fig.

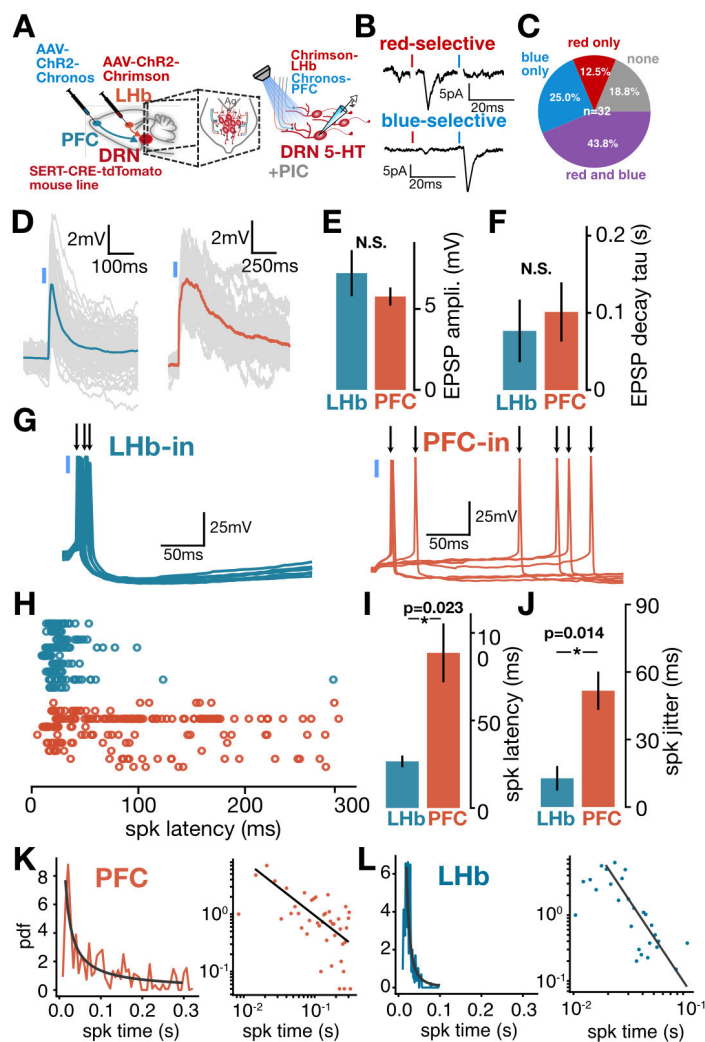


Figure 1: Subthreshold and spike timing characteristics of long-range PFC and LHb input to 5-HT neurons.

A: Experimental setup showing dual-color opsin strategy. **B:** Example voltage-clamp traces of individual 5-HT neurons responding only to red (top) and blue (bottom) photostimulation. **C:** Quantification of fraction of neurons responding to each light type (n=32). **D:** Example current-clamp traces of EPSPs of individual 5-HT neurons in response to photostimulation of PFC (red) or LHb (blue) inputs. **E-F:** Amplitude and decay (mean \pm SE) of EPSPs in individual neurons in response to PFC (red, n=9 neurons) or LHb (blue, n=9 neurons) input (paired T-test). **G:** Example current-clamp traces of spiking responses of individual 5-HT neurons in response to LHb (blue) or PFC (red) photostimulation. **H:** Spike timing of individual neurons (rows) in response to LHb (blue, n=9 neurons) or PFC (red, n=9 neurons) photostimulation. **I-J:** Spike latency and jitter (mean \pm SE) for individual neurons in response to LHb (blue, n=9 neurons) or PFC (red, n=9 neurons) photostimulation (paired T-test). **K:** Probability distribution of spike times

in response to PFC photostimulation across all neurons (left) and plotted as a log-log relationship with line of best fit for power-law scaling (right). **L:** Probability distribution of spike times in response to LHb photostimulation across all neurons (left) and plotted as a log-log relationship with line of best fit for power-law scaling (right).

1B), indicating the presence of long-range glutamatergic synaptic inputs from LHb or PFC. Single 5-HT neurons showed a diversity of responses, responding to neither, one, or both wavelengths of light. Overall, 43.8% of 5-HT neurons responded to both wavelengths of light, 18.8% of neurons responded to neither, and the remainder responded to only one wavelength (Fig. 1C).

To rule out overlap between the excitation spectra of these two opsins, we delivered a set of patterned dual-color stimuli taking advantage of opsin desensitization, similar to Christoffel et al (2021) (see Methods). This approach takes advantage of the fact that paired-pulse stimulation can lead to measurable short-term depression if the same opsin is activated twice (Fig. S1A, top), yet should not affect the other opsin if the excitation spectra are spectrally separated. Delivery of the same paired-pulse stimulation, but with sequential presentation of red and blue stimuli, revealed no short-term depression dynamics when compared with red-red paired-pulse stimuli (Fig. S1A, bottom; Fig. S1B). This indicates an effective separation of the excitation spectra of both opsins in this subset of dual-responder neurons. Single-opsin controls (Chrimson in LHb alone; Fig. S1C-D) revealed that while blue light could occasionally activate the red opsin, our patterned dual-color stimulation accurately discriminated this condition, with similar depression dynamics observed in red-red paired-pulse and red-blue paired-pulse stimulation as expected. These results indicate that a substantial fraction of single 5-HT neurons respond to inputs from both LHb and PFC.

A spike-timing code discriminates input identity from PFC and LHb in single 5-HT neurons

In principle, it is possible that 5-HT neurons combine long-range informational streams from PFC and LHb (and perhaps other areas) into a single low-dimensional rate code. However, populations of neurons can also represent multiple features of the environment simultaneously, whether through time-division multiplexing (Caruso et al., 2018), burst multiplexing (Naud et al., 2018), synchrony-division multiplexing (Lankarany et al., 2019) or other mechanisms. We therefore sought to quantitatively compare the subthreshold and suprathreshold features of 5-HT neuron responses to LHb and PFC inputs to gain insight into potential coding mechanisms. We injected an AAV expressing channelrhodopsin (ChR2) bilaterally into either LHb or PFC in two groups of mice, and conducted whole-cell recordings from identified 5-HT neurons. In voltage clamp ($V_m = -70\text{mV}$), brief pulses of blue light triggered an inward current, indicating monosynaptic release of glutamate from presynaptic terminals similar to the dual-opsin

experiments. In current clamp, photostimulation triggered time-locked excitatory post-synaptic potentials (EPSPs) that were qualitatively similar between LHb and PFC inputs. The mean amplitude and decay tau of EPSPs across cells were not statistically different for each input (Fig. 1D-F). However, we observed slight differences in the cumulative distribution across trials of amplitudes and decay taus; the amplitude cumulative distribution was right-shifted (higher) for EPSPs triggered by LHb inputs, while the decay tau cumulative distribution was right-shifted for EPSPs triggered by PFC inputs (Fig. S2C-F). Together, these results indicate that 5-HT neurons receive long-range inputs from LHb and PFC that are approximately matched in strength and amplitude at subthreshold membrane potentials.

How might these two long-range inputs be integrated into a coherent neural code? To answer this, we recorded in whole-cell configuration from 5-HT neurons in current clamp while optogenetically stimulating long-range inputs from either PFC or LHb to induce a sparse spiking response. We aimed to compare spiking statistics elicited by each input separately to infer how they might differentially contribute to the overall coding properties of 5-HT populations. Crucially, through direct current injection, we controlled resting membrane potential (Fig. S2E) and the fraction of trials eliciting a spike (Fig. S2F) such that there was no significant difference between LHb and PFC groups, and additionally calibrated current injection to ensure that spontaneous spikes (i.e. occurring before the optical stimulation) were not observed. LHb stimulation triggered low-latency spiking with very little variability (jitter) between trials, largely consistent with theories of spike times encoding the derivative of inputs - in this case, EPSPs with sharp rise-times (Mainen & Sejnowski, 1995). PFC stimulation, in contrast, produced much longer-latency spikes with high jitter (Fig. 1G-I). These PFC-induced high spike latencies, often lasting hundreds of milliseconds, were qualitatively similar across all 5-HT neurons measured (Fig. 1H), and suggested that the depolarization triggered by PFC input may drive cells towards a fluctuation-driven regime where action potentials are triggered by membrane potential fluctuations due to random synaptic input and/or intracellular conductances (Kuhn & Rotter, 2004).

Next, we assessed statistical properties of the spike-time distribution triggered by each input. PFC inputs yielded spike-times that were distributed with a heavy right tail that could be well fit with a power-law decay (Fig. 1K). LHb inputs triggered spike-times that could also be fit with a power-law decay with a larger negative exponent reflecting steeper decay (Fig. 1L). However, since spike-times spanned less than one order of magnitude we could not formally distinguish this from an exponential decay (Fig. S2G-J). This is consistent with theories of fluctuation-driven activity regimes, where spike timing depends on membrane potential and noise statistics (Gerstein & Mandelbrot, 1964; van Rossum et al., 2003).

Could differential stimulus-driven spiking for each input reflect an effective segregation of each informational stream within populations of 5-HT neurons? One compatible theory is that of synchrony-division multiplexing (Lankarany et al., 2019; Ratté et al., 2013; Friedrich et al., 2004) suggesting that synchronous firing of neurons in a population can carry information distinct from mean firing rate. Our data suggests that LHb inputs to DRN would trigger synchronous spiking in the 5-HT population, while PFC input would generate non-synchronous and potentially stochastic spiking of 5-HT neurons. These two firing patterns could, in principle, result in differential responses in downstream areas. To explore this possibility, we first applied linear classifiers to our spike timing data from each input, with the aim of determining whether the identity of the input (LHb or PFC) could be theoretically distinguished based on differential spike timing. We generated synthetic training and test datasets (5-fold cross-validation, see Methods) based on an unbiased sampling of real spike-times observed in single trials from each input (Fig. 2A). By training a support vector machine classifier and assessing performance on held-out data, we were able to quantify how spike timing-based decoding performance varied as a function of the number of spikes. Surprisingly, we found that spike timing information provided by as few as 6 neurons allowed good (>90%) decoding of the identity of the synaptic input (PFC or LHb) (Fig. 2B-C), indicating that synchrony/asynchrony-based decoding could be biologically plausible in downstream areas.

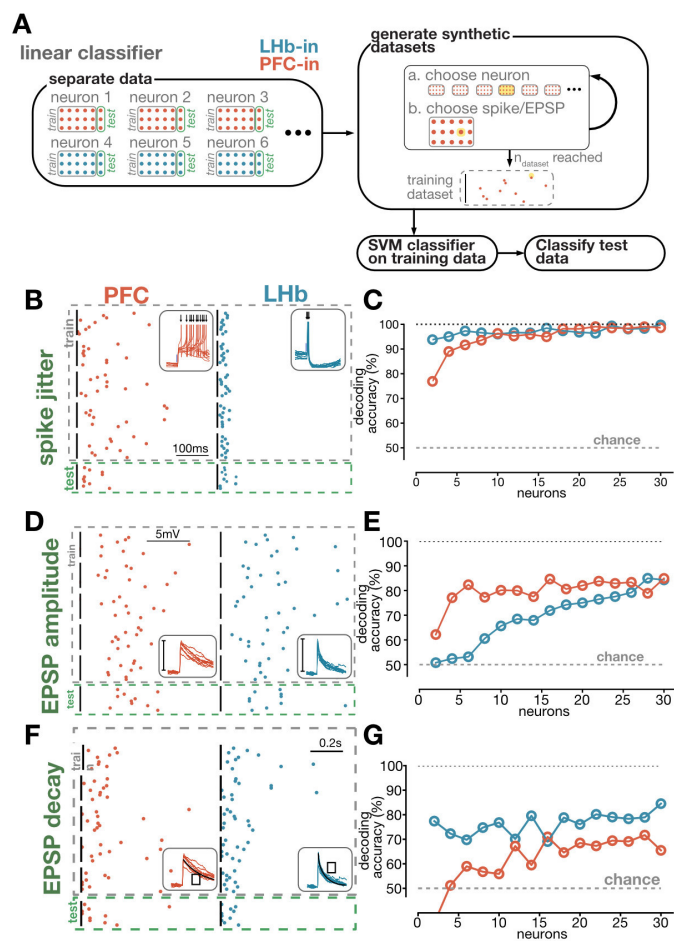


Figure 2: Decoding long-range inputs by spike-timing and subthreshold features in 5-HT populations. **A:** Schematic showing approach for generating synthetic datasets and applying a support vector machine linear classifier to data (5-fold cross-validation). **B:** Example synthetic datasets of spike times from PFC (red) or LHb (blue). Each vertical line represents one example dataset of $n=10$ spike times. **C:** Decoding accuracy of PFC (red) and LHb (blue) synthetic datasets from linear classifier as a function of dataset size (neurons). Dotted line indicates chance decoding performance (50%). **D:** Example synthetic datasets of EPSP amplitude from PFC (red) or LHb (blue). Each vertical line represents one synthetic dataset of $n=10$ EPSP amplitudes. **E:** Decoding accuracy of PFC (red) and LHb (blue) synthetic datasets from linear classifier as a function of dataset size (neurons). Dotted line indicates chance decoding performance (50%). **F:** Example synthetic datasets of EPSP decay from PFC (red) or LHb (blue). Each vertical line represents one synthetic dataset of $n=10$ EPSP

amplitudes. **G:** Decoding accuracy of PFC (red) and LHb (blue) synthetic datasets from linear classifier as a function of dataset size (neurons). Dotted line indicates chance decoding performance (50%).

Since LHb and PFC inputs could generate highly discriminable spike timing patterns, presumably by harnessing distinct ionic or synaptic mechanisms, we next asked whether inputs differed at the subthreshold level, providing a putative substrate for 5-HT neurons to generate these divergent spike timing properties. To determine this, we asked whether a linear classifier could correctly classify inputs based on subthreshold voltage statistics. We trained a support vector machine classifier on unbiased samples of subthreshold voltage amplitudes or decays, extracted from individual trials triggered by each input (Fig. 2D-G). When assessing performance using a held-out test set of observations, surprisingly, we found poor decoding for

both inputs, not exceeding 90% correct classification with >30 observations. This poor subthreshold decoding performance stands in stark contrast to the excellent spike timing discriminability, and is consistent with unobserved cellular features as the cause of differences in spike timing from each input.

Calcium channels expressed in 5-HT neurons modulate noise around spike threshold

We next sought to determine ionic and cellular mechanisms within 5-HT neurons which would be capable of generating reliable differences in spike timing. A potential ionic candidate is low-threshold (T-type) calcium channels, which can trigger rebound spiking in a subset of 5-HT neurons (Burlhis & Aghajanian, 1987) and are highly expressed in several key subgroups of DRN 5-HT neurons (Templin et al., 2012; Okati et al., 2020). We first confirmed their expression in genetically identified 5-HT neurons by recording in voltage-clamp configuration and delivering depolarizing step pulses. Voltage steps to $V_m = -40\text{mV}$ triggered delayed inward currents lasting several hundred milliseconds, which were blocked by application of NiCl, a low-voltage calcium channel blocker (Fig. 3A; Lee et al., 1999, see Methods). The inward current was voltage-dependent, with a protracted rise-time and decay consistent with reported values for a T-type calcium channel (Fig. 3A-B; Perez-Reyes, 2003). Together, these data indicate that T-type calcium channels are expressed in 5-HT neurons and are activated at depolarized voltages.

Next, we sought to determine how these T-type calcium channels modified incoming synaptic inputs with distinct kinetics potentially matching those of the PFC and LHb inputs. Recording from genetically identified 5-HT neurons in current-clamp configuration, we injected current waveforms which resembled incoming EPSCs with varying decay kinetics, which we refer to as synthetic EPSCs (Fig. 3C, left). Injecting synthetic EPSCs with slow decay kinetics led to voltage responses in 5-HT neurons which were occasionally prolonged and higher amplitude, which we hypothesized may be due to transient calcium channel activation (Fig. 3C, middle). Consistent with this hypothesis, bath application of low concentrations of NiCl

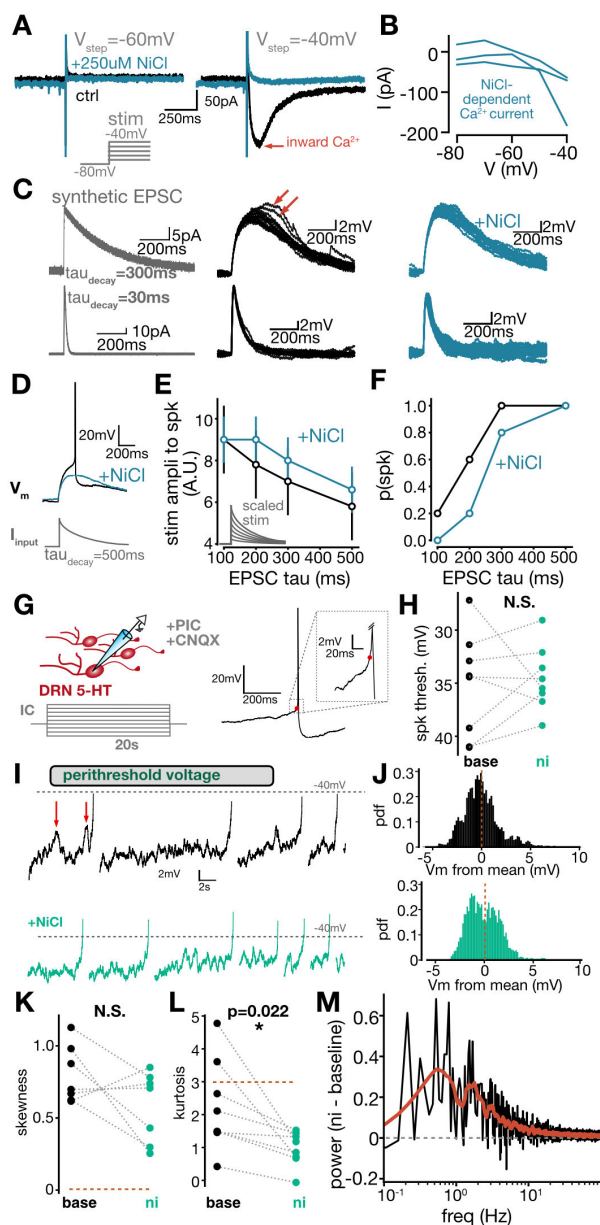


Figure 3: An inward calcium conductance in 5-HT neurons shapes EPSPs and contributes to membrane noise. **A:** Example traces in voltage-clamp showing an inward conductance at -40mV (right) but not -60mV (left) that is blocked by bath application of $NiCl_2$ (blue trace). **B:** Voltage-dependence of the putative inward calcium conductance ($n=3$ neurons). **C:** Synthetic EPSCs injected in current-clamp into 5-HT neurons (left) produce transient voltage excursions (middle) blocked by bath application of $NiCl_2$ (right). **D:** Example trace showing modulation of synthetic EPSC-triggered spiking by bath application of $NiCl_2$. **E:** Relationship between synthetic EPSC amplitude required to elicit a spike, and synthetic EPSC tau, in baseline conditions (black) and after bath application of $NiCl_2$ (blue) ($n=5$ neurons). **F:** Relationship between probability of obtaining an action potential across neurons, and synthetic EPSC tau, in baseline conditions (black) and after bath application of $NiCl_2$ (blue) ($n=5$ neurons).

G: Experimental setup for membrane noise experiments presented in **H-M** (left), and example of spike threshold estimation using fractional first-derivative method (right). **H:** Spike threshold is not modulated by bath application of $NiCl_2$ ($n=8$ neurons, paired T-test). **I:** Example traces of a single neuron recorded in current-clamp during injection of constant perithreshold depolarizing current, before (top) and after (bottom) application of $NiCl_2$. Red arrows depict large voltage excursions ('blips'). **J:** Probability distribution of membrane potentials for an example neuron before (top) and after (bottom) application of $NiCl_2$. **K:** Membrane potential distribution skewness is not significantly

affected by bath application of $NiCl_2$ ($n=7$, paired T-test). **L:** Membrane potential distribution kurtosis is significantly reduced by bath application of $NiCl_2$ ($n=8$, paired T-test). **M:** Relative spectral power of calcium conductance, calculated by subtracting baseline membrane potential spectral power from post- $NiCl_2$ washout membrane potential spectral power ($n=8$ neurons, mean shown.)

blocked these prolonged voltage responses to slow synthetic EPSCs (Fig. 3C, right; also see Supp. Fig. S3). In contrast, when charge-normalized synthetic EPSCs with faster decays were injected (see Methods), they produced voltage responses that lacked prolonged, high-amplitude events and were not measurably affected by application of NiCl (Fig. 3C; also see Supp. Fig. S3).

Since the activation of calcium channels and other ionic conductances can affect the production and timing of spikes (Cain & Snutch, 2013; van Rossum et al., 2003), we hypothesized that these channels might play a similar role in regulating spike timing in 5-HT neurons. Blocking ionic conductances should, in principle, increase the neuron's membrane resistance and make it easier for a fixed magnitude of current input to depolarize the neuron sufficiently to trigger spikes. However, we found, surprisingly, that bath application of NiCl could reduce or even eliminate spiking response for a fixed perithreshold current input (Fig. 3D). This may indicate complex interactions between voltage events and activation/inactivation kinetics of this calcium conductance, as well as spatial profiles of channel expression throughout the dendrite and soma. We used synthetic EPSCs to explore how the production of spikes is tuned by calcium channels across the parameter space of event kinetics. The decay kinetics and amplitudes of the injected synthetic EPSCs were systematically varied while we recorded the voltage response of 5-HT neurons. We found, consistent with these earlier findings (Fig. 3D), that NiCl application increased the amplitude of the synthetic EPSC required to elicit spiking behavior, especially at long decay tau values (Fig. 3E) and reduced spiking probability across all neurons at the same fixed EPSC amplitude (Fig. 3F). These results highlight an important role of 5-HT neuron calcium channels as regulators of spiking output.

We next aimed to explore how calcium channels might modulate spike timing evoked by PFC versus LHb inputs. In particular, this could provide a putative mechanism for controlling network synchrony through dendritic expression differences. We adopted a dynamical approach and focused on how voltage noise statistics are modulated by these channels. We conducted experiments where finely spaced, long (20s) current steps were injected into

identified 5-HT neurons before and after application of the calcium channel blocker NiCl (Fig. 3G). Throughout these experiments, a cocktail of synaptic blockers was applied, and the acute slices were kept at physiological temperature (30-32°C). Spikes were removed from the recordings using a fractional first-derivative method (Fig. 3G; see Methods; Trinh et al., 2019; Azouz and Gray, 2000). To determine how the dynamics of perithreshold voltage noise were modulated by T-type calcium channels, we first focused on the first ascending current step in each neuron that generated any action potentials. At these perithreshold potentials, we observed depolarizing non-spike events lasting ~1s, as well as prominent voltage noise (Fig. 3I, top). Furthermore, the membrane voltage distribution at these potentials was broad and long-tailed (Fig. 3J, top). We suspected that these slow depolarizing events, as well as fast voltage noise, were partially generated by stochastic activation and inactivation of these T-type calcium channels. Consistent with this, application of NiCl modulated the amplitude of slow events and fast voltage noise, and decreased the tail size of the membrane voltage distribution (Fig. 3I-J, bottom) without altering spike threshold (Fig. 3H).

To quantify these observations across the population of recorded 5-HT neurons, we measured the skewness (third central moment, indicating asymmetry) and kurtosis (fourth central moment, indicating tailedness) of the membrane potential distribution, before and after application of NiCl. Before NiCl application, the membrane voltage distribution had skewness values above 0, indicating a right tail, and NiCl application did not significantly affect this (Fig. 3K). The distribution's kurtosis was significantly reduced by NiCl application, consistent with calcium channels causally implicated in the generation of depolarized voltage excursions (Fig. 3L). To determine whether these effects were localized to perithreshold voltage regime, we considered the voltage dynamics during the first subthreshold current step (ie without spikes) and found no effect of NiCl bath application on kurtosis or skewness (Supp. Fig. S4). This indicates that the effects of T-type calcium activation on membrane dynamics is localized to perithreshold membrane potentials. Finally, we performed a fast Fourier transform of the time-varying perithreshold membrane potential (with spikes removed) to determine how calcium

channel blockade affected the frequency content of the noise. This analysis revealed a NiCl-dependent boost in frequencies between 0.1-10Hz (Fig. 3M). Collectively, these results indicate that T-type calcium channels expressed in 5-HT neurons increase membrane noise around spike threshold.

Given the above results, we hypothesized that this perithreshold calcium conductance was causally involved in generating the delayed spikes observed during photostimulation of long-range PFC input to 5-HT neurons. To test this hypothesis, we took two approaches. First, recording in current clamp from identified 5-HT neurons, we injected synthetic EPSPs with slow decay taus matching those of PFC-evoked EPSPs and monitored the timing of evoked spikes, before and after application of NiCl. We found that application of NiCl reduced the latency and jitter of evoked spikes (Fig. S5A-C). Second, we photostimulated long-range PFC inputs to 5-HT neurons before and after application of NiCl, finding a similar reduction in the latency and jitter of evoked spikes after NiCl application (Fig. S5D-E). Taken together, these results indicate that calcium channels expressed in 5-HT neurons play a critical role in controlling spike timing through a membrane noise mechanism.

Dynamical regimes for synchronous and asynchronous spike generation

We next sought a mechanistic understanding of how calcium channels could regulate the initiation time of action potentials and thus modulate spike synchrony. We conducted simulations to ascertain how the statistics of membrane potential fluctuations could control the timing of spike initiation. Specifically, we first considered a simplified scenario where the membrane voltage was initialized at a hyperpolarized value, and random noise was injected (modeled as Ornstein-Uhlenbeck noise to simulate barrages of synaptic input and/or intrinsic noise, see Methods; Destexhe et al., 2001). This produced a voltage signal that evolved over time in a pseudorandom fashion (Fig. 4A). The time at which the voltage signal first crossed a depolarized spike threshold (called the first-passage crossing time; yellow dot in Fig. 4A) was recorded over 20,000 stochastic, independent simulations, generating a probability distribution

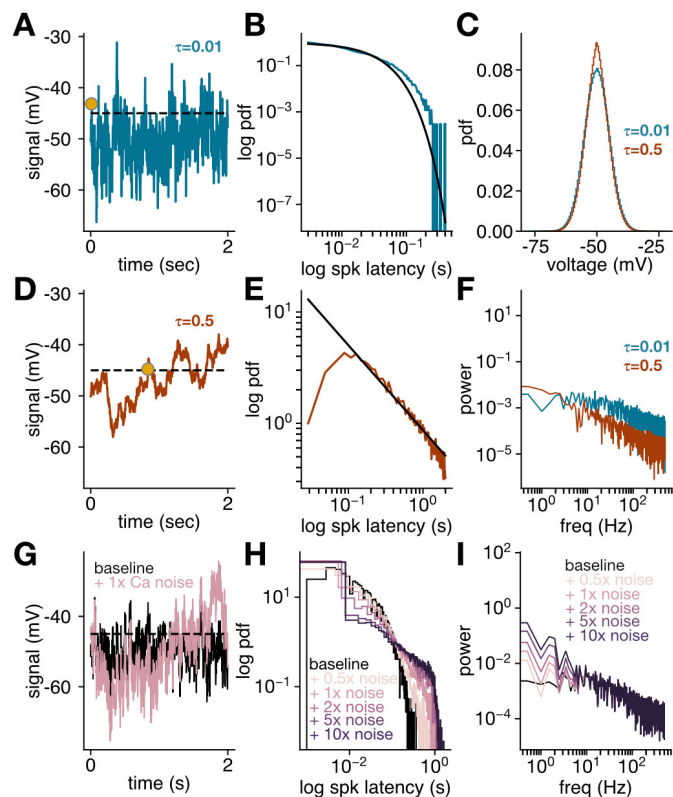


Figure 4: The statistics of membrane noise shapes stochastic spike timing in 5-HT neurons. A: Example simulation showing fast membrane potential noise represented by a fast ($\tau=0.01$ s) Ornstein-Uhlenbeck process. The dotted line represents spike threshold and the yellow dot represents the time of first-passage-crossing. B: Distribution of first-spike latencies across 20,000 independent initializations of OU noise. Black line represents an exponential fit. C: Distribution of simulated membrane voltages for fast and slow Ornstein-Uhlenbeck noise shown in A and D. D: Example simulation showing slow membrane potential noise represented by a slow ($\tau=0.5$ s) Ornstein-Uhlenbeck process. The dotted line represents spike threshold and the yellow dot represents the time of first-passage-crossing. E: Distribution of first-spike latencies across 20,000 independent initializations of OU noise. Black line represents a power-law fit. F: Spectral power of simulated membrane potential for fast and slow

noise processes. G: Example simulations showing fast membrane potential noise ($\tau=0.01$ s Ornstein-Uhlenbeck process, black) and fast membrane potential noise with calcium noise spectral power added in the frequency domain (pink). H: Distribution of first-spike latencies across 20,000 independent initializations of OU noise (black), and with increasing calcium noise spectral power added (pink). I: Spectral power of simulated membrane potential for fast noise process alone and with increasing calcium noise spectral power added.

of spike times (Fig. 4B). These simulations effectively correspond to a population of 5-HT neurons receiving a fixed synaptic input, and responding stochastically due to cell-intrinsic mechanisms. This approach allowed us to ask how the spike timing distributions we observed could be related to simple interactions between a stochastic noise process with defined characteristics, and a fixed spike threshold.

Guided by our empirical findings showing that calcium channels altered the temporal structure of membrane noise, we first used our *in silico* simulations to ask how the speed of noise fluctuations affected spike timing distributions. Ornstein-Uhlenbeck noise with a fast time

constant ($\tau=0.01$) generated rapidly fluctuating voltage signals that tended to reach spike initiation very quickly, with the spike times exponentially distributed and typically in the range of tens of milliseconds (Fig. 4A-B). Ornstein-Uhlenbeck noise with a slow time constant ($\tau=0.5$), and with a fixed membrane potential distribution compared to fast noise (Fig. 4C), generated time-varying signals that evolved more slowly and tended to show longer delays before reaching spike threshold. At the population level, these spike times were in the range of hundreds of milliseconds to seconds, and appeared to be power-law distributed (Fig. 4E). The spectral characteristics of these signals, revealed through a fast Fourier transform (Fig. 4F), showed that simulations with a slower noise time constant were specifically enriched in very low frequencies (0.5-2Hz). Together, these results provide a framework to understand how the distribution of spike initiation times depends on intracellular noise features. In particular, for a fixed distribution of membrane potentials over time, noise with fast fluctuations tends to produce short spike latencies which are exponentially distributed over time, while noise with slower fluctuations produces more protracted spike latencies which are power-law distributed.

To provide a more direct comparison to our T-type calcium noise, we asked whether the previously measured spectral components associated with this calcium noise (Fig. 3M) could alter spike latency distributions in a systematic manner. Starting with a fast-fluctuating Ornstein-Uhlenbeck noise process identical to that presented in Fig. 4A, we added the spectral components of the measured calcium-dependent noise in the frequency domain. Consistent with the calcium noise displaying a prominent low-frequency component, addition of spectral components of this noise generated slow fluctuations in the simulated signal (Fig. 4G, I). We added scaled components of the noise ranging from 0.5x to 10x, corresponding to weak or strong activation of calcium channels, and asked how this affected spike initiation times. Adding stronger versions of the calcium noise dramatically increased the average spike latency time, and prolonged the straight-line component of the spike latency distribution (in a log-log plot) to span two orders of magnitude, consistent with a power-law distribution of latencies (Fig. 4H). These results demonstrate that distinct spike latency distributions at the population

level can be readily explained by stochastic noise processes occurring in single neurons; additionally, this work demonstrates, somewhat surprisingly, that these individual neurons can toggle between distinct latency 'modes' through activation or inactivation of ionic conductances. Ultimately, this suggests a mechanism by which populations of 5-HT neurons receiving overlapping synaptic inputs can effectively segregate these inputs through dynamic generation of synchronous (low-latency) or asynchronous (high-latency) spiking, utilizing the simple building blocks of stochastic noise processes and ionic conductances with defined spectral characteristics.

Discussion

In this work, we demonstrate that a subpopulation of individual 5-HT neurons in DRN receives strong excitatory drive from both mPFC and LHb, and information from these two inputs can be multiplexed at the population level by cell-intrinsic regulation of spike synchrony. While synaptic inputs from LHb produce time-locked, low-latency spikes leading to population synchrony, mPFC inputs produce highly jittered, asynchronous spikes. We show that these two informational sources can be theoretically decoded by downstream circuits receiving sparse 5-HT afferents. Finally, we demonstrate that the delayed spiking pattern from mPFC inputs arise from a set of simple intracellular mechanism: perithreshold calcium channels increase membrane noise and trigger stochastic, delayed spike threshold crossings. Together, these findings indicate that long-range inputs from mPFC and LHb can be multiplexed in local serotonergic circuits by cell-intrinsic regulation of spike synchrony.

Our results showing multiplexed information in 5-HT populations are broadly consistent with recent reported in vivo recordings during complex behaviors. Paquelet et al. (2022) undertook miniscope calcium imaging of populations of 5-HT neurons during emotionally salient behaviors, finding highly correlated (synchronous) ensembles of neurons responding to sucrose or footshocks alongside non-selective ensembles. Li et al. (2016) demonstrated that 5-HT neurons encode reward signals, including a subpopulation that exhibits increased firing rate

with millisecond precision. Cohen et al. (2015) reported that 5-HT neurons exhibited tonic firing rates that tracked block reward size, alongside phasic responses to direct rewards and punishments. These findings, taken alongside the results reported here, are consistent with a model where information about ongoing environmental stimuli, such as rewards or punishments, is encoded as synchronous population activity in DRN microcircuits, while higher-order task features or environmental statistics may be represented as background, uncorrelated firing of 5-HT neurons. We expect that future work will help address the question of how synchronous and asynchronous activity evolves in larger populations of 5-HT neurons, and whether effective parsing of spiking in this manner provides meaningful information about task variables or internal variables.

In this work, we used optogenetic tools to drive defined long-range synaptic inputs to 5-HT neurons. This technique has several limitations which should be noted. First, optogenetics leads to an artefactual increase in release probability (Zhang & Oertner, 2007) as well as a synchronous release of neurotransmitter. It is possible that PFC or LHb synapses onto 5-HT neurons exhibit lower release probability, potentially increasing variability of the amplitude of postsynaptic potentials. Additionally, presynaptic terminals from PFC or LHb onto 5-HT neurons may exhibit release patterns that are more asynchronous during naturalistic behaviors, potentially altering spike timing. Finally, while we used dual-color optogenetics to uncover the presence of a subpopulation of 5-HT neurons that receive strong input from mPFC and LHb, further work will be needed to elucidate the precise location and size of this subpopulation, as well as to situate these results in the broader context of anatomical submodules within DRN (Zhou et al., 2017; Ren et al., 2018).

Our results demonstrate that 5-HT neurons exhibit calcium conductances that are activated in the perithreshold regime and add a broadband low-frequency component to ongoing membrane noise. How do synaptic inputs from mPFC preferentially activate these calcium conductances, while LHb inputs do not? Here we discuss two possibilities. First, a dendritic gradient of calcium channel expression, combined with localized targeting of calcium

channel-rich dendritic segments by mPFC inputs, might produce mPFC-specific activation of calcium conductances. Indeed, differential targeting of dendritic segments by long-range inputs can profoundly shape spike output and integration in other brain areas (MacAskill et al., 2012), indicating they are a plausible substrate for the effects reported here. A second possibility is that mPFC synaptic inputs to 5-HT dendrites could simply be more spatially clustered than LHB synaptic inputs, allowing mPFC inputs to exclusively trigger dendritic calcium events when activated synchronously. Work has shown that dendritic calcium amplification can play a functional role in coincidence detection (Lee et al., 2016; Higley & Sabatini, 2008). Future studies will be needed to disambiguate these and other possibilities.

These findings add to an emerging body of literature documenting how ionic conductances in 5-HT neurons can shape neural coding properties. Our work shows that calcium channels, in particular, control spike timing in a manner that may support synchrony division multiplexing from two distinct synaptic inputs. This could provide a putative mechanism for controlling network synchrony through an interaction between synaptic or dendritic features and these calcium conductances. Two striking features of our investigations were that calcium conductances in 5-HT neurons appeared to be preferentially activated by slow-decaying EPSPs, and that these conductances additionally increased voltage noise at low frequencies. These findings suggest a simple, dynamical model for the timing of spontaneous or evoked spikes in 5-HT neurons: a random walk of membrane voltage towards spike threshold, modulated by depolarizing synaptic inputs as well as ionic noise processes with certain statistics. This dynamical model may help explain features of random, uncorrelated spiking or synchrony of 5-HT neurons (Cohen et al., 2015; Paquelet et al., 2022).

These findings also add to recent work demonstrating that A-type potassium channels (I_A) are expressed in 5-HT neurons and play a role in regulating millisecond-level precision of spike timing after large membrane depolarizations (Harkin et al., 2023). It is likely that depending on the statistics of synaptic input, this A-type potassium conductance and the

calcium conductance described in the present work may be activated in a combinatorial manner to shape spike timing.

We used linear decoders to demonstrate that the observed spike timing patterns can theoretically be decoded by downstream circuits innervated by 5-HT axons. The biological implementation of this decoding could, we hypothesize, involve differential receptor subtype affinity in cortical microcircuits. In cortex, excitatory 5-HT_{2A} receptors are strongly expressed in apical dendrites of the majority of pyramidal neurons (Beique et al., 2007; Jakab et al., 1999). However, fast-spiking and somatostatin interneurons can either express excitatory 5-HT_{2A} or inhibitory 5-HT_{1A} receptors, but typically not both (for review, see Puig et al., 2011). 5-HT_{1A} receptors have a higher affinity for 5-HT than 5-HT_{2A} receptors (Peroutka et al., 1979; Nichols & Nichols, 2008), suggesting a decoding scheme where 5-HT_{1A}-expressing interneurons may be inhibited by both synchronous and asynchronous 5-HT axon release events, while 5-HT_{2A}-expressing interneurons and pyramidal neurons expressing 5-HT_{2A}Rs may be depolarized by only highly synchronous 5-HT release events due to lower receptor affinity. Such receptor binding rules in cortical microcircuitry could facilitate demultiplexing of information carried by 5-HT axon synchrony, although further work would be needed to establish the precise mechanisms by which such a decoding scheme would function.

Taken together, our results demonstrate that populations of 5-HT neurons receiving multiple sources of long-range synaptic input can effectively multiplex information, through predominantly cell-intrinsic mechanisms that regulate spike synchrony at the population level. These cell-intrinsic mechanisms involve stochastic membrane noise processes and differential calcium channel activation. More broadly, these results suggest that a complex population coding scheme such as synchrony-division multiplexing can plausibly arise from the autonomous, stochastic noise processes occurring in individual neurons, even in the absence of strong network-level interactions that are conventionally thought to shape information representations in cortex.

Methods

Animal use

SERTcre: B6.Cg-Tg(Slc6a4- cre)ET33Gsat/Mmcd

R26R-TdT: Rosa26R-TdTomato fluorescent protein (TdT) reporter mice

All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services. For this study, SERTcre mice were crossed with the R26R line to generate a SERTcre-TdT mouse line. Mice were kept on a 12:12 hour light/dark cycle, with access to food and water ad libitum unless otherwise noted. A mix of male and female SERTcre-TdT mice were used.

Slice physiology

Stereotaxic Injections

All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services. For stereotaxic injections, mice were injected with 0.05 mg/kg buprenorphine and anesthetized by inhalation of isoflurane. Injections were performed using either a 10 μ l Hamilton syringe with a 33 gauge needle, or a Nanoject machine with a glass micropipette. For all injections, 0.2 - 0.3 μ l of virus was injected per coordinate, at a rate of 0.1 – 0.2 μ l per minute. Stereotaxic coordinates are as follows: LHb (from Bregma skull landmark, in mm: AP -1.7, ML +/- 0.4, DV -2.7 to -3.0) and mPFC (from Bregma skull landmark, in mm: AP +2.5, ML +0.4, DV -1.5 to -1.8 from the surface of the brain).

For ChR2 (H134R) expression, 0.2-0.3 μ l of AAV9.hSyn.hChR2(H134R)- mCherry.WPRE.SV40 (titre ~10¹² GC/mL) was injected bilaterally into either LHb or mPFC and animals were left to recover for 3 to 4 weeks prior to ex-vivo electrophysiological recordings.

For the dual opsin experiments (Chrimson/Chronos), 0.2-0.3 μ l of AAV-Syn-ChrimsonR-TdTomato (Serotype:2/9) was injected bilaterally into LHb, and 0.2-0.3 μ l of AAV-Syn-Chronos-GFP (Serotype:1) was injected bilaterally into mPFC. In some control experiments (marked in text and figures), only AAV-Syn-ChrimsonR-TdTomato (Serotype:2/9) was injected bilaterally into LHb to rule out spectral overlap between opsins. Animals were left to recover for 3 to 4 weeks prior to ex-vivo electrophysiological recordings.

Slice preparation

LED Photostimulation experiments (Fig. 1). DRN containing brainstem slices were prepared from mice aged P40-P90. Acute slices were prepared similarly to previously described methods (Geddes et al., 2016). In brief, mice were anesthetized by isoflurane inhalation (Baxter Corporation, Canada) and sacrificed by decapitation. The brain was removed and brain slices were sectioned while immersed in ice-cold choline chloride-based cutting solution of the following composition (in mM): 119 choline-Cl, 2.5 KCl, 1 CaCl₂, 4.3 MgSO₄-7H₂O, 1 NaH₂PO₄, 1.3 sodium L-ascorbate, 26.2 NaHCO₃, and 11 glucose, and equilibrated with 95% O₂, 5% CO₂. Slices were recovered in a chamber containing standard Ringer's solution of the following composition (in mM): 119 NaCl, 2.5 CaCl₂, 1.3 MgSO₄-7H₂O, 1 NaH₂PO₄, 26.2 NaHCO₃, and 11 glucose, at a temperature of 37°C, continuously bubbled with a mixture of 95% O₂, 5% CO₂. Slices were left to recover for 1 hour in a recovery chamber and equilibrate to a temperature of approximately 21°C until recordings were performed.

Calcium conductance and membrane noise experiments (Fig. 3): DRN containing brainstem slices were prepared from mice aged P21-P40. To optimize the health of serotonin neurons for these sensitive experiments, acute slices were prepared using a modified NMDG two-bath cutting procedure, described in Ting et al. (2018). Briefly, mice were anesthetized by isoflurane inhalation (Baxter Corporation, Canada) and were transcidentally perfused with ice-cold NMDG-

based cutting solution with the following composition (in mM): 92 NMDG, 2.5 KCl, 1.2 NaH₂PO₄H, 30 NaHCO₄, 20 HEPES, 10 MgSO₄, 25 Glucose, 0.5 CaCl₂, 5 sodium ascorbate, 2 thiourea, 3 sodium pyruvate, with pH adjusted to 7.3–7.4 with HCl and osmolality adjusted to 300–310 mOsm/kg with sucrose, and equilibrated with 95% O₂, 5% CO₂. Following perfusion, the brain was removed and 300µm brain slices were sectioned in the same NMDG cutting solution. Brain slices were incubated for 0.5h in NMDG cutting solution at 37°C, with NaCl concentration gradually increasing to 92 mM over the incubation time (details in Ting et al., 2018), and then transferred to a HEPES holding ACSF bath, at room temperature, with the following composition (in mM): 92 NaCl, 2.5 KCl, 1.2 NaH₂PO₄H, 30 NaHCO₃, 2 MgSO₄·7H₂O, 2 CaCl₂·2H₂O, 20 HEPES, 25 Glucose, 5 Sodium ascorbate, 2 Thiourea, 3 Sodium pyruvate, equilibrated with 95% O₂, 5% CO₂. Slices were left to recover for 1 hour in a recovery chamber and equilibrate to a temperature of approximately 21°C until recordings were performed.

Whole-Cell Electrophysiology

DRN neurons were visualized using an upright microscope: (1) Examiner D1 equipped with Dodt-contrast or differential-interference contrast (DIC) (Zeiss, Oberkochen, Germany; 40x/0.75NA objective) or, (2) BX51W1 equipped with DIC optics (Olympus, Center Valley, PA; 40x/0.80NA objective). 5-HT neurons were visually identified by TdTomato signal driven by the SERT promoter under EPI-fluorescence illumination (xCite Series 120pc). Whole-cell recordings were carried out using an Axon Multiclamp 700B amplifier, sampled at 10 kHz, digitized with an Axon Digidata 1440A (or 1550) digitizer and filtered at 2 kHz. Whole-cell recordings were performed using borosilicate glass patch electrodes (3–6 MΩ; World Precision Instruments, Florida) pulled on a Narishige PC-10 pipette puller (Narishige, Japan).

All LED photostimulation experiments (Fig. 1) were performed at room temperature (20°C).

All calcium conductance and membrane dynamics experiments (Fig. 3) were conducted at a near-physiological temperature (measured 30–32°C) through the use of an in-line bath heater and stage heater. The purpose of these alterations were to replicate, as closely as possible, the membrane dynamics and channel kinetics that would be observed *in vivo*.

Slices received constant bath application of Ringer containing (in mM): 119 NaCl, 2.5 KCl, 1.3 MgSO₄, 2.5 CaCl₂, 1.0 NaH₂PO₄, 11 glucose, 26.2 NaHCO₃, and 0.1 L-Tryptophan saturated with 95% O₂ and 5% CO₂ (pH = 7.3; 295–310 mOsm/L). An intracellular solution of the following composition was utilized: 115 mM potassium gluconate, 20 mM KCl, 10 mM sodium phosphocreatine, 10 mM HEPES, 4 mM ATP(Mg²⁺), and 0.5 mM GTP, pH 7.25 (adjusted with KOH; osmolarity, 280–290 mOsmol/L). For all voltage clamp recordings, access resistance was continuously monitored using a 125 ms, 2 mV hyperpolarizing pulse, at least 245 ms prior to stimulation. Recordings were discarded when the access resistance changed by >30%. Liquid junction potential was not compensated for. For current clamp experiments, during LED photostimulation, direct current injection was used to maintain membrane voltage at approximately ~-50mV to allow a direct comparison between neurons in each stimulation condition (described in Fig. S2).

LED Photostimulation

LED photostimulation was delivered using a PlexBright LED module and controller (465 nm; Plexon, Texas). Light was delivered through a 200 µm patch fiber cable with a bare fiber tip with ~ 1 cm of glass exposed. For EPSC/EPSP stimulation, unless otherwise stated, optically evoked-PSCs were stimulated with single light pulses of 1–5 ms in length (individual trials 12–15 seconds apart) and an irradiance of 792 mW/mm².

Drugs and Chemicals

Picrotoxin, CNQX and nickel (ii) chloride were purchased from Abcam (Cambridge, MA), Tocris Bioscience (Bristol, UK) and Sigma-Aldrich (St. Louis, MO), respectively. The concentration of drugs employed was as follows: Picrotoxin 100uM, CNQX 10uM, NiCl 250uM.

Data Analysis

All electrophysiological recordings were analyzed offline in the Python computing environment using the Numpy and Scipy computational packages. All data are presented as means +/- SEM. n refers to the number of cells recorded. Unless otherwise stated, statistical significances were derived using a standard Student's t-test ($p < 0.05$). P values of less than 0.05 were considered statistically significant and are indicated with an asterisk(*).

Applying linear classifiers to sub-threshold and spiking statistics

The approach to apply linear classifiers to the LHB and PFC event data was as follows. We first describe methods to generate unbiased synthetic datasets of varying sample sizes from our physiological data, and second describe application of a linear classifier to this data.

Synthetic datasets. For a given number of samples and an event statistic (e.g. spike latency or EPSP amplitude), we generated 1000 synthetic datasets for both LHB and PFC inputs. These synthetic datasets were generated by iteratively performing the following: 1) randomly sampling a neuron, then 2) randomly sampling a trial from that neuron and storing the event statistic from that trial. This was iteratively performed until the desired number of samples was reached. This sampling procedure ensured that each neuron was equally represented in the dataset. Next, the 1000 synthetic datasets were partitioned into either *training* or *test* datasets, with 5-fold cross-validation.

Linear classifier. We first performed standardization of the data using a min-max scaler, and then trained a support vector machine linear classifier on the *training* dataset using stochastic gradient descent learning (employing the implementation in scikit-learn). The classifier was then used to predict labels (LHB or PFC) in the test dataset, with classification accuracy registered as the fraction of correct classifications across the entire test dataset. This linear classifier training and evaluation was then repeated for a total of 5 times (5-fold cross-validation) while rotating the held-out data in the test set. The classification accuracy was reported as the average accuracy across all 5 rotations of test and training data.

Generation and injection of synthetic EPSCs (Fig. 3)

Synthetic EPSCs were generated as a waveform with instantaneous rise-time, and a monoexponential decay with τ_{decay} that varied from 30ms to 500ms. These synthetic EPSCs were injected into the neuron in current clamp to elicit an evoked EPSP response. In Fig. 3C and Fig. S3, as various values of τ_{decay} were presented to the same neuron, the total charge transfer to the neuron was held constant by scaling the amplitude of each EPSC such that the integral of current over time was calculated to be identical.

Membrane noise analysis and spike threshold calculation (Fig. 3)

For the analysis of membrane potential dynamics, spikes were first identified using a fractional first-derivative method (Trinh et al., 2019; Azouz and Gray, 2000). Briefly, the maximum derivative of voltage over time, $\max(\frac{dV}{dt})$, was calculated for each spike, and the spike threshold was identified as the first point in the depolarization where the voltage derivative reached $0.01 \cdot \max(\frac{dV}{dt})$. A spike comprised all points from spike threshold to 0.4s after the

threshold; all such points were masked and not included in analysis of membrane voltage statistics or spectral characteristics. To calculate the spectral characteristics of membrane voltage, a 1-dimensional discrete fast Fourier transform of the voltage time series was computed, and the resulting spectrum was smoothed using a Savitzky-Golay filter with polynomial order 3.

Stochastic membrane noise simulations (Fig. 4)

To investigate how distinct spike timing patterns could arise from stochastic fluctuations in membrane voltage, we simulated membrane voltage as a dynamical process and solved the level-crossing problem to infer how spike timing distributions depended on membrane dynamics. We considered a simplified scenario where membrane voltage ($v_0 = -50mV$) was initialized close to spike threshold ($v_{thresh} = -45mV$). Membrane voltage evolved stochastically (described below) until spike threshold was reached. By repeating this over 20,000 simulations, we generated a distribution of level-crossing times (representing spike timing distributions), allowing us to assess how membrane potential dynamics were related to distinct regimes of spike timing.

We simulated the dynamics of membrane voltage as an Ornstein-Uhlenbeck noise with the following equation:

$$\frac{dv(t)}{dt} = -\frac{1}{\tau}(v(t) - v_0) + \sqrt{D} \cdot \eta(t) \quad (1)$$

where τ is a constant denoting the correlation timescale of the noise (indicated in text, ranging from 0.01 to 0.5), $v_0 = -50mV$ represents the resting membrane potential, and $\eta(t)$ is a white noise term. \sqrt{D} is a noise diffusion coefficient (scaling factor) that was varied across simulations with τ such that the quantity $\sigma^2 = D\tau/2$ was constant ($\sigma^2 = 100$); this produced fixed membrane potential distributions that were independent of τ as it was varied (Fig. 4B). Each individual simulation proceeded until v reached the simulated spike threshold $v_{thresh} = -45mV$, at which point the time from simulation onset to spike threshold was recorded as a simulated spike latency.

Numerically, the Ornstein-Uhlenbeck noise process was computed using the following constant definitions and update rule, following the well-known algorithm derived by Gillespie (1996) for exact numerical simulation of the Ornstein-Uhlenbeck Process:

$$A = \sqrt{\frac{D\tau}{2}(1 - e^{-\frac{2\Delta t}{\tau}})} \quad (2)$$

$$v(t + \Delta t) = v_0 + (v(t) - v_0) \cdot e^{-\Delta t/\tau} + A \cdot \mathcal{N}(0,1) \quad (3)$$

With a timestep of $\Delta t = 0.001s$.

For simulations in which the spectral features of the noise were altered, we conducted simulations as described above, then performed a fast Fourier transform of the simulated voltage, altered the signal in the frequency domain in the manner described in the text, then transformed back into the time domain.

Manuscript V**Acute stress shapes synaptic inhibition within an amygdala microcircuit**

Michael Lynn* and Sébastien Maillé*

Contributed equally to this work. Co-corresponding authors.*Review of Di et al. (2016)****This manuscript was published in the *Journal of Neuroscience* in 2017:**Maillé S & Lynn MB. Acute stress shapes synaptic inhibition within an amygdala microcircuit. *J Neurosci*. 2017 January 18;37(3):474-476. doi: 10.1523/JNEUROSCI.3145-16.2017.

Statement of contribution

This was a collaborative effort between myself and Sébastien Maillé, another graduate student in the Béïque Laboratory. We jointly led a review of the literature, guided the project, and wrote and edited all drafts.

We are grateful to Dr. Jean-Claude Béïque for providing an excellent training environment that allowed us the time and support to write such a review.

The amygdala is a limbic structure which regulates emotion-driven behaviors. Acute stress can modulate several amygdala-dependent functions, including fear conditioning and the expression of anxiety-like behavior. During acute stress, the hypothalamic–pituitary–adrenal (HPA) axis produces glucocorticoids, which interact with the noradrenergic system in the basolateral amygdala (BLA) to enhance fear memory storage (Roosendaal et al., 2006). This glucocorticoid stress response requires endocannabinoid signaling within BLA (Campolongo et al., 2009), but several questions regarding this signaling have remained unresolved. First, what molecular pathways link glucocorticoid binding to endocannabinoid release? Second, how does this stress response alter information processing within amygdala circuits and modulate functions such as anxiety-like behaviors or fear memory formation and expression?

In a recent issue of *The Journal of Neuroscience*, Di et al. (2016) identify a crucial piece of the molecular puzzle and provide convincing evidence showing these molecular pathways are sufficient to produce an anxiety-like phenotype. The authors recorded from BLA principal cells (PCs) in slices and showed that bath-application of glucocorticoids decreased the frequency of miniature inhibitory postsynaptic currents (mIPSCs). In addition, the authors electrically stimulated the external capsule to evoke inhibitory postsynaptic currents (eIPSCs), and showed that these responses did not decrease in amplitude upon glucocorticoid application. This surprising result argues against the straightforward interpretation that glucocorticoids decrease release probability at inhibitory synapses, since this would predict a reduction in both mIPSC frequency and eIPSC amplitude upon glucocorticoid application. Instead, the lack of correlated changes was interpreted as evidence that glucocorticoids may differentially modulate the two vesicle pools intended for spontaneous and evoked release.

Further experiments demonstrated that glucocorticoids suppress spontaneous inhibitory transmission by activating a G-protein-coupled, membrane-bound glucocorticoid receptor. This glucocorticoid receptor is located postsynaptically and triggers retrograde release of endocannabinoids (CB) upon activation, suppressing presynaptic GABA release. This glucocorticoid-endocannabinoid signaling pathway bears striking similarity to one already

uncovered in the hypothalamus (Di et al., 2003), indicating that this may represent a core molecular pathway through which glucocorticoids can tune synaptic inhibition.

The authors demonstrated the behavioral relevance of this molecular pathway in two ways. First, exposure to acute stress caused the emergence of an anxiety-like phenotype, measured as a decrease in time spent at the center of an open field. This phenotype was blocked by intra-BLA infusion of either a CB1 receptor antagonist or a 2-arachidonoylglycerol (2-AG) synthesis inhibitor, indicating that acute stress acts through an endocannabinoid signaling pathway to produce this anxiety-like phenotype. Second, the glucocorticoid-mediated suppression of inhibition observed in slices was occluded by prior exposure to behavioral stress, suggesting that CB1 receptors had been saturated by endogenous cannabinoids during stress. Interestingly, CB1R blockade in unstressed animals produces anxiety-like behavior, implying that tonic CB1R activation under basal conditions has an anxiolytic effect. Together, these data indicate that the glucocorticoid signaling pathway elucidated here may be sufficient to modify certain amygdala-dependent behaviors.

How does this glucocorticoid-induced suppression of inhibition fit within a broader picture of BLA circuit function? A highly simplified BLA circuit is outlined below. BLA PCs receive excitatory input from cortical and thalamic afferents passing through the amygdalar external capsule. Inhibitory drive onto these BLA PCs comes from two major sources: local inhibitory interneurons provide feedforward inhibition, which is enhanced by noradrenergic input (Kaneko et al., 2008); and paracapsular intercalated cells in the external and internal capsules provide inhibitory input under dopaminergic modulation (Marowski et al., 2005). With respect to the circuit's output, a population of BLA principal cells project to the central amygdala (CeA) to decrease anxiety-related behaviors when activated (Tye et al., 2011).

Di et al. (2016) report that acute stress decreased the frequency of spontaneous GABA release, which might shape circuit activity. Spontaneous inhibition has been shown to suppress action potential generation and reduce efficacy of coincident excitatory input (Carter et al., 2002). A decrease in spontaneous inhibition may therefore result in either an increase in basal

firing rate or more reliable responses to excitatory cortical input. If glucocorticoids increase activity selectively in CeA-projecting PCs, we would expect an anxiolytic effect (Tye et al., 2011) which is the opposite of results reported by Di et al (2016). It is therefore possible that PCs displaying glucocorticoid-induced activity changes do not directly project to CeA, and instead cause anxiety-like phenotypes through other pathways. This highlights the fact that BLA influences behavior in a highly projection-specific manner, and more work is required to link molecular modes of action with specific populations of BLA projection neurons.

Could glucocorticoids modulate evoked inhibitory transmission, and not simply spontaneous transmission? Di et al. (2016) found no effect of glucocorticoids on evoked transmission, but they evoked IPSCs by electrically stimulating within the external capsule, which likely produces strong activation of inhibitory intercalated cells. In support of this, recent work in which cortical afferents to amygdala were optogenetically stimulated in the absence or presence of the external capsule showed that the external capsule provides a prominent inhibitory input onto PCs likely arising from intercalated cells (Morozov et al., 2011). Furthermore, stimulation of cortical afferents in the external capsule would not produce excitatory synaptic currents in BLA interneurons due to the presence of glutamate antagonists DNQX and AP5, and therefore such stimulation would not produce feedforward inhibitory currents in BLA PCs. It is therefore possible that glucocorticoids reduce release probability from BLA interneurons onto BLA PCs, a mechanism which would decrease both spontaneous release frequency (detected by Di et al., 2016) and evoked event amplitude (undetected due to stimulating electrode placement) onto PCs.

Cholecystokinin (CCK) interneurons in BLA are a plausible target for presynaptic release probability modulations, because they express CB1 receptors and provide a distinct perisomatic inhibitory input onto PCs (Vereczki et al., 2016). During an acute stress response, the glucocorticoid cascade described by Di et al (2016) could act on these CCK interneurons, increasing BLA PC activity by relieving cortically evoked feedforward inhibition, and enhancing PC drive onto downstream targets to cause anxiogenic effects. However, evidence indicates

that a prominent class of PC projector neurons (CeA-projecting PC cells) in fact has an anxiolytic effect (Tye et al., 2011), implying again that greater circuit elucidation is required to fully link molecular pathways with behavior. Similar to the circuit hypothesis discussed here, recent work has shown that cortical input to dorsal raphe nucleus can be modulated in a switch-like manner due to endocannabinoid suppression of feedforward inhibition (Geddes et al., 2016).

Vesicle pools for spontaneous and evoked synaptic transmission are thought to represent distinct entities (Fredj et al., 2009), but it is unclear whether release probability can be modulated independently for each pool. Di et al (2016) reported that spontaneous, but not evoked transmission was glucocorticoid-sensitive, which was interpreted as a vesicle pool-specific modulation of release probability. This is an exciting and fundamentally novel concept, and has not been observed for kainate, another neuromodulator: Matthew et al. (2008) used a stryryl dye staining/destaining approach to show that kainate concomitantly enhances release probability of both spontaneous and evoked vesicle pools in GABA terminals in layers II/III of cortex. Further work is needed to clarify whether spontaneous and evoked vesicle pools have related release probabilities or whether they can be modulated independently.

Taken together, these observations add to an emerging model of how acute stress can recruit endocannabinoid machinery to dynamically alter the excitation/inhibition balance within BLA (Hill et al., 2010; Figure 1). Basal conditions are characterized by GABAergic inhibition of PCs, since glutamate release onto PCs is suppressed by tonic AEA release. Exposure to an acute stressor shifts the circuit's balance towards excitation. First, the suppression of excitatory drive onto PCs is relieved: corticotropin-releasing hormone (CRH) from the paraventricular nucleus of the hypothalamus (PVN) binds to PCs to increase FAAH activity and decrease tonic AEA release (Gray et al., 2015). Second, glucocorticoids synthesized by the HPA axis enhance glutamate release onto PCs by binding to presynaptic membrane-bound mineralocorticoid receptors (mbMRs; Karst et al., 2010). Third, these glucocorticoids also suppress inhibitory input to PCs by triggering the endocannabinoid signaling cascade

uncovered by Di et al (2016). Glucocorticoids may additionally disinhibit norepinephrine release through endocannabinoid actions (Campolongo et al., 2009), although more work is required to elucidate the behavioral and circuit-level implications of this disinhibition. Therefore, stress could lead to a constellation of synaptic alterations within BLA microcircuitry, shifting the excitation/inhibition balance of the network in a switch-like manner to influence behavior (Figure 1).

Di et al (2016) have elucidated a core molecular mechanism underlying the effects of the stress response on emotionally salient behavior. Clearly, a complex circuit in BLA shapes behavior, and stress-induced glucocorticoid release alters circuit function at several levels through mechanisms that are still not fully understood. Further work is now required to link expected activity changes in behaviorally defined pathways (such as anxiolytic CeA-projecting BLA PCs) with emerging molecular and synaptic details about circuit function. A deeper understanding of how networks in the brain are shaped by stress will help provide new insights into conditions such as anxiety and post-traumatic stress disorder.

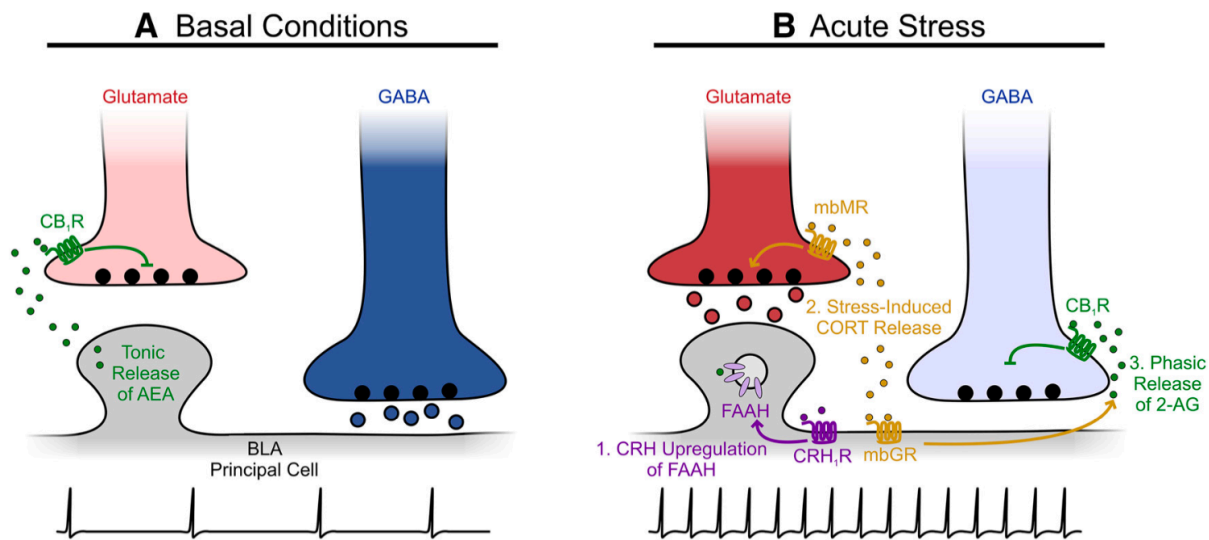


Figure 1. Glucocorticoids interact with the endocannabinoid system following acute stress exposure to shift excitation/inhibition balance in the BLA. A, Under basal conditions, tonic release of anandamide (AEA) from BLA principal neurons dampens glutamate release by binding to presynaptic endocannabinoid receptors (CB₁R) while GABA release is unrestrained. Thus, inhibition dominates the network and presumably suppresses principal neuron firing. B, Following exposure to acute stress, (1) Corticotropin releasing hormone (CRH) released from cells in PVN activates its receptor (CRH₁R) on BLA principal neurons, thus upregulating fatty acid amide hydrolase (FAAH) activity and halting the tonic release of AEA. (2) Glucocorticoids (CORT) produced by the HPA axis bind presynaptic mineralocorticoid receptors (mbMR) to enhance glutamate release, and simultaneously, CORT binds glucocorticoid receptors (mbGR) on principal cells to (3) trigger the retrograde release of 2-arachidonoylglycerol (2-AG), which activates presynaptic CB₁R to selectively suppress GABA release. Under these conditions, excitation dominates the network and presumably increases principal neuron firing rate.

Manuscript VI**BK channels at dendritic spines: A mechanism for coupling morphology, plasticity and information storage?**

Michael Lynn* and Sébastien Maillé*

***Contributed equally to this work. Co-corresponding authors.**

Review of Tazerart et al. (2022)

This manuscript was published in the *Journal of Physiology* in 2022:

Maillé S & Lynn MB. BK channels at dendritic spines: A mechanism for coupling morphology, plasticity and information storage? *J Physiol*. 2022 August;600(15):3399–3401. doi: 10.1113/JP283232.

Statement of contribution

This was a collaborative effort between myself and Sébastien Maillé, another graduate student in the Béïque Laboratory. We jointly led a review of the literature, guided the project, and wrote and edited all drafts.

We are grateful to Dr. Jean-Claude Béïque for providing an excellent training environment that allowed us the time and support to write such a review.

Cortical pyramidal neurons receive a constant barrage of temporally and spatially precise information in the form of local and long-range synaptic inputs. These inputs contact microscopic (<1 fL) protrusions called spines, connected to the neuron's extensive dendritic arborization by a slender neck. These spines express glutamate receptors (AMPA and NMDA) on their surface and act as electrical and biochemical compartments upon their excitation by presynaptic inputs. Work over recent decades has shown that the morphological and functional heterogeneity of spines enables them to perform sophisticated, combinatorial computations on incoming synaptic information. Spine morphology can influence synaptic plasticity, with spine neck shape and head volume controlling diffusion of important plasticity-related molecules such as calcium between the spine and the parent dendrite. Concurrently, plasticity can influence spine morphology, with learning-related changes over timescales of minutes to hours; long-term synaptic depression (LTD) can reduce spine head volume and lead to elimination, whereas long-term potentiation (LTP) can both increase spine head volume and shorten neck length (an in-depth review of these points is provided by Araya, 2014). An emerging body of work seeks to map out these complex relationships between morphology, plasticity and signal filtering, and connect them to core computations carried out by neural networks.

In a recent study reported in *The Journal of Physiology*, Tazerart, Blanchard et al. (2022) explored the expression and functional features of large conductance calcium-activated potassium channels (BK channels) within spines from basal dendrites of layer five pyramidal neurons in the mouse visual cortex. Previous work in hippocampal CA1 pyramidal neurons found that a related conductance generated by small conductance calcium-activated potassium channels (SK channels) expressed at spine heads modifies the waveform of synaptic potentials through a Ca^{2+} -mediated negative feedback loop (Ngo-Anh et al., 2005). To explore the subcellular expression of BK channels, Tazerart, Blanchard et al. (2022) combined western blotting on synaptoneurosomal fractions, immunofluorescence imaging and immunoelectron microscopy and found that BK channels are enriched at the postsynaptic

density for the majority (~50–90%) of these basal dendrite spines, across diverse spine morphologies and head sizes.

To probe the functional role of these spine-targeted BK channels, Tazerart, Blanchard et al. (2022) considered the effect of BK channels in shaping somatic spiking showing that pharmacological blockade of BK channels increased action potential width and decay time without impacting cell excitability. They next performed sophisticated glutamate uncaging experiments at single spines, combined with pharmacological manipulations, to show that BK channels suppress synaptic transmission exclusively at small-headed spines (despite functionally identical channel expression across spine sizes).

Focusing on these small-headed spines, Tazerart, Blanchard et al. (2022) explored mechanisms linking synaptic input with BK channel activation. The opening probability of BK channels depends on both large voltage depolarizations and high local calcium concentrations, which coincide within spine head compartments. By employing calcium imaging and pharmacological approaches, Tazerart, Blanchard et al. (2022) demonstrate that calcium influx through NMDARs is necessary for BK activation, and that these synaptically triggered spine calcium signals are inversely correlated with spine head size. This provides a putative explanation for the apparent spine volume-dependent effect of BK channels: smaller spines have higher calcium concentrations generated by synaptic inputs. These intuitions were confirmed with biophysical modelling, which demonstrated that, at small spines, NMDARs and voltage-gated calcium channels (VGCCs) can generate large voltage deflections and increased calcium concentrations, which are sufficient to activate BK channels.

These observations raise several questions about the mechanisms coupling spine calcium and ionic conductances. The initial increase of spine calcium required for BK channel activation could be provided by several different synaptic mechanisms, including direct influx through NMDARs or VGCCs, or mGluR/IP3R-dependent release from intracellular stores. In their study, Tazerart, Blanchard et al. (2022) failed to observe BK channel-dependent suppression of EPSPs after blocking NMDARs, strongly suggesting that, similar to SK channels

(Ngo-Anh et al., 2005), BK channels are primarily activated by Ca^{2+} provided by NMDARs. However, this does not rule out the possibility that secondary sources may further boost Ca^{2+} after initial NMDAR activation. Previous studies have shown that calcium-induced calcium release from intracellular stores can be triggered by NMDAR-dependent Ca^{2+} influx, thus significantly boosting Ca^{2+} signals in spines and dendrites; this facilitates co-operative plasticity between coactive synapses in developing networks (Lee et al., 2016). Therefore, it remains unclear whether supplementary calcium sources contribute in part to the observed spine calcium responses and to the observed selective activation of BK channels in spines, and whether these mechanisms might also operate in a spine size-dependent manner (for example, endoplasmic reticulum invasion in spines).

The work by Tazerart, Blanchard et al. (2022) significantly contributes to our understanding of the interactions between synaptic plasticity and spine morphology. The involvement of spine calcium in plasticity processes has received extensive experimental support; one influential hypothesis states that high concentrations of spine calcium trigger LTP, whereas low concentrations trigger LTD. Models of spine plasticity predict that the calcium concentration is undercompensated during spine growth: because spine volume grows as a cubic power of its radius, whereas spine surface area and associated expression of calcium channels only increase as a square power, then, as a spine undergoes LTP-dependent growth, synaptically evoked calcium concentrations will drop until a steady-state between LTP and LTD is reached (O'Donnell et al., 2011). Tazerart, Blanchard et al. (2022) confirmed that calcium influx is inversely proportional to spine head size, a key component of this undercompensation model. Next, modelling work predicted that, in small spines, BK channels down-regulate calcium concentrations, similar to the effect of SK channel activation (Ngo-Anh et al., 2005). Integrating this result with past work (O'Donnell et al., 2011) leads to the prediction that this BK-dependent calcium reduction in small spines shifts them further towards a compensating calcium regime, in which calcium concentrations become less dependent on spine head size (Fig. 1). The BK channels present at small-headed spines may therefore imbue them with a




Figure 1. A model of BK channel regulation of synaptic transmission and plasticity

Synaptic glutamate release activates AMPARs (blue/diagonal stripes), and NMDARs (green/vertical stripes), which depolarize the membrane and increase intracellular Ca²⁺. Large voltage deflections activate voltage-gated calcium channels (VGCCs; magenta/horizontal stripes) and high intracellular Ca²⁺ activates BK channels (yellow). In large spines (A), EPSPs are large but Ca²⁺ levels are not sufficiently elevated to activate BK channels. In small BK-lacking spines (B), EPSPs are smaller but spine Ca²⁺ levels are very elevated; in small BK-expressing spines (C), elevated Ca²⁺ levels activate BK channels, which lead to smaller, narrower EPSPs and lower Ca²⁺ concentrations. In this model, BK channel expression could facilitate pruning of some small spines (D).

greater capacity to undergo low calcium-dependent LTD and elimination (Fig. 1D), although we note that these spines may undergo a diversity of LTP or LTD responses depending on stimulation parameters, morphological factors and ionic channel expression.

If spine BK channels mediate this shift from undercompensating to compensating Ca^{2+} regimes, this could have several implications for network plasticity. First, this may explain the mechanism of small spine pruning during learning (Oh et al., 2013), which is not explained by simple undercompensating models of calcium. Indeed, these models predict that, because spine calcium increases exponentially as spine head size decreases, synaptic strength should show a unimodal, strictly positive distribution with no spine elimination (O'Donnell et al., 2011). A compensating calcium regime at small spines would shift this synaptic strength distribution to the left by including spine elimination dynamics. Second, theoretical work has shown that stimulus intensity can be stored as steady-state spine head volume (O'Donnell et al., 2011); calcium compensation may increase the gain of this stimulus–volume relationship at low stimulation frequencies, such that changes in these weak stimuli evoke proportionally larger changes in spine volume (O'Donnell et al., 2011).

Although some neural coding strategies rely on population firing rates to encode stimulus features (rate code), others depend on the precise timing of synaptic inputs or emitted spikes to represent information (temporal code). The exciting results reported by Tazerart, Blanchard et al. (2022) suggest that spine morphology and BK channel expression could be tuned to optimize synaptic input for either of these coding schemes. Classic work has demonstrated the computational power of timing-based (temporal) neural codes: for example, the brainstem nucleus laminaris of the barn owl encodes the azimuthal position of sounds through precise coincidence detection of synaptic inputs. In these temporal coding regimes, spines with small head volumes, exhibiting BK channel-mediated sharpening of EPSPs (Fig. 1C), may improve the precision of temporal discrimination. By contrast, association circuits often utilize a rate-based code, requiring neurons to perform spatiotemporal integration on synaptic inputs which can lack precise timing. In these coding regimes, large spines

(Fig. 1A) and those lacking BK channels (Fig. 1B) may optimize for spatiotemporal pooling of diverse synaptic inputs through protracted deflections in spine voltage and calcium, which could be boosted by active mechanisms including calcium-induced calcium release (Lee et al., 2016). Further work is needed to determine whether temporal and rate coding regimes are associated with distinct spine morphologies.

The first formal descriptions of dendritic spines were made in 1888 by Ramón y Cajal, who proposed that the thorny protrusions he saw on dendrites of bird cerebellar neurons could serve as contact sites between neurons. This assertion proved controversial not only because it challenged the dominant ‘reticular theory’ of neural circuits, but also because others had dismissed these protrusions as precipitated silver artefacts of the Golgi staining procedure. It was not clear why synaptic contacts, if they existed, should target these spines when they could directly contact the dendritic shaft, although Cajal proposed several innovative theories. Subsequent to these initial discoveries, technological advances have confirmed the importance of spine morphology for plasticity, information processing and storage. By identifying BK channels as a key component linking spine morphology and function, Tazerart, Blanchard et al. (2022) have suggested an exciting set of research avenues on plasticity, neural coding, and their subcellular implementations at spines of single neurons.

General discussion

Summary of results

Broadly, this thesis has focused on the role of synaptic dynamics and synaptic ensemble composition in generating population-level computations. In Manuscripts III and IV, I present work focusing on the synaptic basis of computations in DRN. Manuscript III demonstrates that serotonin neurons form a local recurrent inhibitory network within DRN, exhibiting striking synaptic facilitation rules that endow this inhibitory network with the ability to ‘ignite’ in response to high-frequency excitation to serotonin neurons. These synaptic dynamics sculpt nonlinear activity patterns within subnetworks of DRN, as well as appearing to underlie strong winner-take-all dynamics between DRN subnetworks. By optogenetically activating long-range excitatory input from lateral habenula to DRN, we demonstrate that these microcircuit dynamics are consistent with a contextual, phasic inhibition of a reward-conditioned response in a classical conditioning task. Manuscript IV focuses on comparing the spiking properties generated by long-range input from lateral habenula and prefrontal cortex to DRN serotonin neurons. This work demonstrates that information can be multiplexed in DRN circuits through a spike synchrony mechanism, extending this important temporal coding concept to midbrain circuits and providing a novel dynamical mechanism for regulating spike timing (here, T-type calcium conductances that generate delayed spikes). Manuscripts I and II focus on quantifying the makeup of synaptic ensembles in central networks. These two manuscripts demonstrate that while bulk electrical stimulation is a scalable method for quantifying the makeup of synaptic ensembles (specifically, the fraction of silent synapses in a microcircuit), a simulation-corrected estimator is necessary for high statistical power and low bias. These manuscripts provide tools to quantify the composition of synaptic ensembles at central synapses, which are necessary to accurately track how such compositions change throughout development and learning, as well as to forge links between machine learning and synaptic makeup in biological networks. Manuscripts V and VI provide commentaries on two exciting new papers that discuss features of central synapses. These include the role of

potassium channels in small spines for shaping postsynaptic potentials and contributing to temporal coding properties (Manuscript VI), and the ability of stress to shift the balance between excitation and inhibition within amygdala (Manuscript V). Together, the results presented in this thesis add to the body of knowledge about population-level neural computations that emerge from basic synaptic properties and ensemble makeup, and additionally suggest new perspectives on how midbrain circuits perform local computations on long-range input through recurrent dynamics and spike timing mechanisms.

In this General Discussion, I will discuss several specific findings in greater depth than what was possible in the discussion section of each manuscript. I will also explore links between the manuscripts, as well as proposing experimental and theoretical approaches to extend these findings.

Characteristics of monoaminergic transmission

Together, the findings in Manuscripts III and IV raise questions about the overall nature and dynamics of monoaminergic transmission, and contribute to emerging work indicating that monoaminergic signalling can be more spatiotemporally precise and complex than previously thought.

A key finding in Manuscript III is that DRN serotonin neurons form local recurrent connections, releasing serotonin in a manner that is probabilistic, highly facilitating, and history-dependent. This supports previous findings in the serotonin system indicating history-dependent release dynamics, including long-timescale depression (Courtney & Ford, 2016). By fitting these dynamics with a linear-nonlinear model over multiple timescales (Rossbroich et al., 2021), we are able to build on this in several ways. First, such quantitative models allow performance validation through separate experimental training and test sets, allowing us to obtain strong evidence that these models faithfully capture important aspects of the dynamics. Second, these linear models are easily expressed as differential equations, allowing them to scale to large population simulations. This allowed us to investigate the interplay between

these short-term plasticity dynamics and long-range inputs that may be arriving at distinct frequencies, in a manner that would be difficult without a dynamical model. Third, by comparing the response reliability and scaling with pulse number between electrical stimulation of serotonin axons and optogenetic stimulation of serotonin neurons themselves, our results strongly suggest that the locus of these short-term plasticity dynamics is presynaptic instead of postsynaptic. These results implicate a potential fast short-term facilitation mechanism consisting of presynaptic calcium accumulation and increased probability of vesicle release, rather than, for example, serotonin accumulation in the extracellular space which differentially acts on the binding curve of postsynaptic serotonin receptors. Our work additionally indicated short-term depression of serotonergic transmission over timescales of seconds, which may be due to a depletion of the presynaptic readily releasable pool of serotonin. Intriguingly, striatal dopamine shows distinct short-term dynamics from those reported here for serotonin. Dopamine release in striatum is characterized by high initial release probability and very strong short-term depression, supporting temporally precise coding by acting as a low-pass filter (Liu et al., 2018). This contrasts with our findings indicating that probabilistic transmission (*i.e.* lower presynaptic release probability), fast facilitation and slow depression characterize serotonergic transmission within DRN.

It is possible that phasic spatially confined serotonergic transmission coexists with more spatially diffuse serotonergic ‘tone’. While our findings do not, strictly speaking, support or disprove the volume transmission hypothesis (Descarries et al., 1982; Chazal & Ralston, 1987; Fuxe & Agnati, 1991), they nonetheless indicate that it is tractable to describe monoaminergic transmission with a similar mathematical formalism used to describe classical glutamatergic and GABAergic release (*i.e.* Npq, or number of axons, release probability, and quantal size). Several avenues of evidence support the idea of spatially constrained serotonin release. In Manuscript III, we find that direct current injection into individual serotonin neurons to elicit spiking does *not* trigger a 5-HT_{1A} receptor current in those same neurons. This indicates that whether serotonin is released dendritically or axonally from serotonin neurons,

the effective diffusion range must necessarily be less than the span of one serotonin neuron (soma and dendrites), or else we would have reliably witnessed a 5-HT_{1A} receptor-dependent postsynaptic current in activated serotonin neurons. Additionally, our work and that of Courtney & Ford (2016), using 5-HT_{1A} receptors as a postsynaptic reporter, finds significant history-dependent dynamics present in serotonergic signaling; this contrasts with reports using fast-scan cyclic voltammetry to measure bulk extracellular concentration (e.g. Bunin & Wightman, 1998) that report serotonin concentrations are not dependent on history or stimulation parameters. One manner of reconciling these reports is with a model where serotonin undergoes both phasic signaling with release zones in close apposition to receptors, and 'tonic' signaling through spatially diffuse serotonin binding (although the spatial scale must necessarily be limited). In fact, recent work in the dopamine system indicates that this neuromodulator, once thought to undergo predominantly spatially diffuse volume transmission, may exhibit mechanistically distinct spatiotemporal scales of signaling. This can be captured by a domain-overlap model where dopamine release and receptor binding are localized at spatially precise structures relative to one another (reviewed in Liu et al., 2021). Dopamine release in striatum shows both spatiotemporally precise dynamics at the axon level, and slower timescales of signalling in the extracellular space; these allow dopamine to encode aspects of both reward and locomotion (Howe et al., 2013; Howe et al., 2016). Thus, such varying scales of transmission may be a general feature shared between neuromodulatory systems.

To further investigate the mechanisms of serotonergic release at single sites, one could apply recent cutting-edge imaging techniques to the analysis of serotonergic circuits. For example, genetically encoded serotonin sensors (GRAB_{5-HT}) have recently been developed that permit spatiotemporally precise imaging of serotonin release by responding in a concentration-dependent manner to serotonin binding (Wan et al., 2021). By expressing this GRAB_{5-HT} sensor on the postsynaptic surface of genetically identified DRN serotonin neurons and electrically stimulating serotonin axons during two-photon imaging of the GRAB_{5-HT} sensor, one could investigate several aspects of serotonergic release. First, the subcellular location of activated

GRAB_{5-HT} sensor could provide evidence for volume release of serotonin (diffuse activation across the cell), or alternately for a more constrained somatic or dendritic binding of serotonin. Second, similar to techniques used to estimate the dynamics of glutamate release at single synapses (Soares et al., 2017), the response of small regions of activated sensor could be monitored in response to repetitive electrical stimulation of serotonin axons at different frequencies. In this way, probabilistic serotonin release could be visualized, and the dependence of release probability on stimulation frequency could be validated. In particular, using recently developed expectation-maximization techniques, quantal parameters of transmission could be estimated, including release probability at single serotonin release or binding sites (Soares et al., 2019). In the case that it is difficult to visualize the location of serotonin released from electrically stimulated axons, one could combine two-photon uncaging of ‘caged’ serotonin molecules (Cabrera et al., 2017) with electrophysiological recordings from identified serotonin neurons. This experiment would allow one to effectively ‘fix’ the presynaptic release probability at 1, revealing stimulus-dependent postsynaptic dynamics that may depend on receptor occupancy or binding. Our results in Manuscript III make the prediction of minimal postsynaptic dynamics, with the majority of short-term plasticity arising from presynaptic mechanisms, but further work such as the experiments described here will be necessary to determine these dynamics in further detail.

At the population level, Manuscript IV demonstrates that the firing patterns of serotonin neurons can support synchronous or asynchronous release at downstream targets, depending on input identity. Similar population-level synchrony patterns have been observed in DRN serotonin neuron during a diverse range of emotionally salient tasks (Paquet et al., 2022), and synchronous release from individual dopamine axons in striatum has also been reported (Howe et al., 2016), indicating that this may be a general feature of neuromodulatory systems. It remains to be seen how downstream areas respond to synchronous or asynchronous serotonin release, and indeed whether such release patterns can signal qualitatively distinct variables.

While the investigations in this thesis are confined to DRN serotonergic circuits, further work will be required to determine whether similar dynamics of serotonergic signaling are present in other serotonin circuits (for example, the median raphe nucleus) or indeed for serotonin signaling in cortex, where probabilistic release and presynaptic facilitation of serotonin axons could play important computational roles for shaping neuromodulatory tone in these areas. For example, such facilitation mechanisms could act as a high-pass filter to allow cortical circuitry to predominantly respond to high-frequency serotonin neuron activity. Since serotonergic axons in cortex can corelease glutamate (Liu et al., 2014), it is possible that release of these two transmitters might obey distinct dynamics of release that are dependent on stimulus statistics. Due to the diverse serotonin receptor subtypes expressed in a cell type-specific manner (see Introduction, Section 2.5), the overall effect on cortical circuitry would likely require population-level dynamical models to fully characterize. Additionally, since serotonin can affect the spike probability directly through action at the axon initial segment, (Ko et al., 2016), short-term facilitation of this serotonergic transmission may play complex roles in dictating spiking dynamics.

Neural coding strategies in the dorsal raphe nucleus

The results presented in this thesis suggest several new ideas about how serotonergic circuits can encode information and broadcast it to downstream targets. Broadly, the findings demonstrate that neural activity DRN can be thought of as a dynamical system that evolves over time, whose coding properties depend on synaptic and intracellular properties as well as input source and input statistics. Multiple streams of information can likely be represented at the population level, through mechanisms of spike synchrony and functionally segregated submodules. Manuscript III describes novel local interactions within DRN microcircuits, where serotonergic neurons form recurrent networks characterized by connections that are slow, inhibitory, and highly facilitating. The emergent dynamics have the potential to fundamentally shape population-level computations occurring within DRN, and complement recent work

highlighting substantial anatomical and functional diversity present in this system. Manuscript IV demonstrates how temporal coding strategies (spike synchrony at the population level) can coexist with rate coding strategies (asynchronous spiking), and can be activated by distinct inputs to multiplex information and communicate it to downstream targets. In this section I discuss implications for how serotonergic circuits represent information.

Our results indicate that serotonin neurons form sparse, local recurrent connections, and the resulting dynamics generate heterogeneous responses to particular inputs at a single-cell level. When photostimulating long-range input from lateral habenula, we observed evidence of at three distinct subpopulations of serotonin neurons. One subpopulation of serotonin neurons was excited by lateral habenula input, a second subpopulation was both directly excited and also received delayed recurrent inhibitory conductances, and a third subpopulation received only recurrent inhibition with no direct excitation. Our network model in Manuscript III, consisting of unstructured recurrent connections, largely reproduced this diversity of serotonergic responses. While further work will be required to determine the precise nature of submodule interactions and structure of recurrent connectivity, such interactions could explain the diversity of single-cell responses during behaviors (Bromberg-Martin et al., 2010; Cohen et al., 2015; Matias et al., 2017; Zhong et al., 2017; Grossman et al., 2022), which reflect not only anatomical segregations within DRN, but also trial-by-trial variability in recurrent neural dynamics. In particular, it will be important to determine whether other long-range inputs to DRN (such as from prefrontal cortex, lateral hypothalamus, or are capable of generating similar recurrent inhibitory effects, and to compare these with the results presented in Manuscript III.

Several possibilities emerge for the computational function of such recurrent inhibition and subsequent single-cell response diversity. Taken from the perspective of anatomically distinct submodules within DRN (Ren et al., 2018), these results support a model where recurrent connections underlie both within-submodule interactions, enacting nonlinear gain modulation of activity, as well as between-submodule interactions, enacting a form of winner-

take-all dynamics between functionally distinct submodules (both explored within Manuscript III). Of these two types of interactions, the potential computational role of winner-take-all dynamics is clearer: to provide a mechanism for contrastive representation between functionally segregated circuits. This idea, building on work proposed several decades ago (Coultrip et al., 1992) and consistent with recent results describing largely non-overlapping functional and behavioral roles for major output pathways of the DRN (Ren et al, 2018; Paquelet et al., 2022), suggests that once a barrage of defined inputs to one subnetwork crosses some activity threshold, other subnetworks are inhibited through facilitating recurrent inhibitory connections. These dynamics are reminiscent of influential work on central pattern generators in spinal circuits, which act as reciprocally inhibitory oscillators to generate alternating motor movements (Marder & Bucher, 2001; Grillner, 2006). However, the fast short-term facilitation present at DRN recurrent inhibitory connections is novel, to our knowledge, and enhances this winner-take-all effect. In addition, the protracted short-term depression provides a practical mechanism for restricting the temporal window with which activity can summate within a submodule and influence the activity of other submodules. As DRN provides widespread serotonergic output to many cortical and subcortical areas, we hypothesize that such winner-take-all dynamics within DRN allow it to act in a switch-like fashion to trigger brainwide state transitions. Such a role for the serotonergic system was shown during transitions between exploration and exploitation in freely moving zebrafish (Marques et al, 2020).

The computational role of within-submodule interactions is less clear at present. We show that such interactions can effectively invert the polarity of high-frequency long-range excitatory input contacting DRN serotonergic neurons, converting the response to net inhibition over protracted timescales. This inhibition is sufficient to reduce spiking of serotonergic neurons, showing it is functionally relevant. In some ways, this is surprising in the context of extensive literature showing that a fundamental balance is maintained between excitatory and inhibitory conductances at a single-cell level in many networks (reviewed in Denève &

Machens, 2016). It is possible that temporary alterations to the excitatory/inhibitory balance within submodules can transiently modulate network properties in a manner that favors certain coding schemes (e.g. rate coding), although further work will be required to elucidate such principles.

One final hypothesis for the computational role of recurrent inhibition and heterogeneous cell responses profiles is to generate population dynamics that perform a suite of computations in low-dimensional subspaces. These dynamics and their computational role would be difficult to comprehend or visualize at a single-cell level, as they would be emergent properties of the population. Recurrent dynamics have been shown to be important in generating flexible task representations, and performing dimensionality reduction on the activity patterns of neurons has been shown to be a viable method for understanding computations through such dynamics (covered in Introduction section 2.6; Churchland et al., 2012; Mante et al., 2013; Vyas et al., 2020; Sylwestrak et al., 2022; Nair et al., 2023). While such studies have predominantly focused on dynamics that emerge from recurrent excitatory networks, notably that can generate persistent internal activity, it is not clear how recurrent connections that are both inhibitory and facilitating would generate population trajectories that perform defined computations. Typically, strong, *excitatory* recurrent connections are capable of generating attractor dynamics that perform a variety of computations, including memory storage and perceptual integration (Fiete et al., 2022). It is tempting to speculate that recurrent *inhibitory* connections would prevent the formation of stable attractor states within DRN, although this would necessarily depend on the precise balance between recurrence within each submodule and recurrence between submodules, as well as the connectivity profile.

Anatomical tracing experiments could be utilized to resolve the spatial distribution of recurrent inhibitory connections. These experiments could be used to distinguish between connections within submodules and between submodules, and could take the form of either retrograde tracer injections or rabies tracing studies. However, it is important to note that while some studies demonstrate anatomically segregated afferents and efferents to DRN circuitry

(Ren et al., 2018), other work has determined that stimulus-selective subpopulations of serotonin neurons are in fact spatially intermingled (Paquelet et al., 2022). Thus, it may be difficult to map recurrent connectivity between and within functional subpopulations without the use of advanced genetic tools or other methods.

It is likely that large-scale recordings from identified serotonergic neurons would be needed to test these hypotheses about recurrent computation, either with an implanted miniendoscope for calcium imaging or with Neuropixels electrodes for electrical recordings. First, the presence of recurrent inhibition could be validated by performing cross-correlation analyses of spiketrains from individual neurons. A negative correlation would be consistent with inhibitory connections, with the time of peak negative cross-correlation consistent with the protracted timescale of GIRK activation as demonstrated in Manuscript III. The strength of these correlations would necessarily be dependent on the activity levels of neurons, with higher activity leading to greater negative correlation between connected neurons. These cross-correlation analyses could additionally be used to supplement experiments investigating the anatomical distribution of recurrent connectivity. Second, one could examine the activity of hundreds of serotonergic neurons in an appropriate low-dimensional subspace to parse out the representation of, for example, context (auditory tone) and choice (left or right lickport) in a two-port dynamic foraging task, similar to seminal work in prefrontal cortex of monkeys (Mante et al., 2013). By pharmacologically blocking recurrent inhibition in DRN through infusion of 5-HT_{1A} receptor antagonists, the role of this population structure in generating dynamic computations could be assessed. Further work would be required to assess how these low-dimensional population dynamics relate to serotonergic release in cortical and subcortical areas.

Manuscript IV suggests that DRN serotonergic neurons encode distinct streams of information through synchronous and asynchronous population spiking activity. This result adds to a body of work describing the diverse coding features of serotonergic neurons over multiple timescales (Ranade et al., 2009; Cohen et al., 2015; Matias et al., 2017; Paquelet et al.,

2022). It has been clear from this body of work, spanning decades, that populations of serotonergic neurons can represent multiple distinct behavioral or environmental variables, often over multiple timescales. However, Manuscript IV provides a putative intracellular mechanism: certain synaptic inputs can differentially recruit a calcium conductance that causes stochastic, delayed spiking at a single-neuron level, and desynchronizes spiking at a population level. This suggests several further questions, as well as links with the recurrent inhibition presented in Manuscript III.

In this Manuscript, we focused on two prominent glutamatergic inputs to serotonergic neurons, from prefrontal cortex and lateral habenula. Further work is required to determine whether other cortical and subcortical inputs, such as from orbitofrontal cortex, amygdala, or hypothalamus (Weissbourd et al., 2014), generate synchronous or asynchronous spike patterns. One hypothesis is that several cortical inputs jointly contribute to asynchronous spiking, while multiple subcortical inputs contribute to synchronous spiking. In this model, cortical input could communicate higher-order features extracted from the environment, such as mean reward rate or expected reward, to control background spiking rates of serotonergic neurons and tonic serotonin output to other brain regions. Conversely, subcortical input could communicate more immediate environmental and internal features, such as threats or satiety levels, through the production of synchronized spiking in serotonergic ensembles; these could transiently increase serotonin levels in recipient brain regions, acting as a separate channel of information transmission compared to background spiking. Further work will be needed to determine how the broad profile of synaptic inputs to serotonergic neurons recruits either synchronous or asynchronous spiking, and furthermore, whether subcortical and cortical areas target similar populations of serotonergic neurons or instead contact more spatially segregated populations. While work has shown that prefrontal cortex and lateral habenula both innervate the ventromedial DRN (Zhou et al., 2017), making these two inputs an ideal target for investigating theories of population multiplexing, future work will be needed to determine whether other long-range inputs jointly target the same individual serotonergic neurons.

One outstanding question arising from Manuscript IV is whether multiplexing through spike synchrony in DRN is a viable coding scheme *in vivo*. While other work has obtained strong support for the idea of synchrony as a multiplexing scheme in cortical areas (Lankarany et al., 2019; Kamaledin et al., 2022), it is unclear whether neuromodulatory systems such as the serotonergic DRN would employ this scheme, and indeed whether defined projection targets would be able to successfully decode multiplexed signals through variable serotonin release. One could test this with two relatively straightforward experiments, to be conducted in mice performing a simple task in which cued rewards and punishments are delivered in an alternating block structure, similar to that previously employed in studies of DRN (Cohen et al., 2015). First, using Neuropixels probes, one could record spiking activity from large populations of genetically identified serotonergic neurons in DRN. Manuscript IV predicts that the spiking cross-correlation between neurons would be higher during periods of punishment delivery (synchronous activity, lateral habenula input), and conversely that mean firing rate would mainly be altered without affecting synchrony in the inter-trial intervals during blocks of punishment versus reward delivery (asynchronous activity, prefrontal cortex input). Second, one could express G-protein coupled receptor activation-based fluorescent serotonin sensors (GRAB_{5-HT}, Wan et al., 2021) in cortical regions (*e.g.* orbitofrontal cortex, a major target of serotonin release in the forebrain; Ren et al., 2018). One could image serotonin release to determine whether transient, synchronized fluorescent signals are associated with punishment delivery, and whether the rate of desynchronized background fluorescent signals across cortical areas changes with average reward rate or reward delivery. The findings in Manuscript IV makes a further prediction that stimulus identity should be decodable from the timing of firing patterns across serotonin neurons, or from temporal patterns of serotonin release, but not from time-averaged (low-pass-filtered) rate of firing or serotonin release.

What is the mechanism by which distinct long-range inputs can differentially generate spike timing? Our results in Manuscript IV implicate an intracellular calcium conductance mechanism, but further work is required to determine how each long-range input differentially

couples to these calcium channels. One possibility is that long-range axons from prefrontal cortex and lateral habenula could preferentially contact either the proximal or distal dendritic compartments of serotonergic neurons. Assuming a dendritic expression gradient of T-type calcium channels, this could provide a putative mechanism for binding certain sources of long-range input with the generation of calcium channel-induced noise and delayed spikes. To test this hypothesis, it is likely that tracing experiments will be required to determine the subcellular localization of each input. Two-photon glutamate uncaging could be used to causally verify the hypothesis that subcellular targeting of excitatory inputs (*e.g.* at proximal or distal dendritic compartments) is sufficient to produce differential spike timing. One potential complication is the fact that the amplitude and decay of excitatory postsynaptic potentials evoked by prefrontal and habenular inputs were approximately matched (Manuscript IV); in contrast, differential dendritic targeting would predict a decreased amplitude and increased event decay for the distal input. Therefore, while the experiments detailed here are a useful starting point, further work will be required to fully determine the mechanistic basis of the results presented in Manuscript IV.

It is possible that the interplay between frequency-dependent recurrent inhibition (Manuscript III) and coding through spike synchrony (Manuscript IV) generates unique rules for information representation in serotonergic circuits. At face value, our metrics for synchrony deal with single spikes produced with variable timing across a population of serotonin neurons. This seems incompatible with the multiple, high-frequency spiking discharges required in single serotonergic neurons to elicit strong recurrent inhibitory drive to other serotonin neurons. However, two simple thought experiments reveal that there may be overlap between these two concepts. First, consider a train of high-frequency excitatory input, from either lateral habenula or prefrontal cortex, contacting a single serotonin neuron. Does the spiking output of this neuron follow the input train with sufficiently high precision to generate the reliable high-frequency spiking that is required for strong recurrent inhibition? The results from Manuscript IV indicate that lateral habenula input would generate more temporally precise postsynaptic

spikes in serotonin neurons, which would be faithfully transformed through short-term facilitation into reliably strong recurrent inhibition onto other serotonin neurons. This is consistent with the results presented in Manuscript III. On the other hand, the same high-frequency train of prefrontal cortex input would likely generate less temporally precise spikes in serotonin neurons, some with large time lags from the synaptic input, and others with shorter time lags. This decreased temporal precision in coupling between presynaptic input train and postsynaptic spikes would seem to decrease the reliability by which a given high-frequency burst of prefrontal input would be capable of generating strong recurrent inhibition through facilitation. This would be especially true for shorter input trains (<200ms), where many postsynaptic spikes elicited would occur after the train has ended. Of course, these inferences are subject to validation by both experiments and theoretical analysis. Therefore, Manuscript IV suggests that the reliability of generating strong recurrent inhibition between serotonin neurons may not be equal for all long-range inputs: some inputs, like lateral habenula, may more reliably trigger recurrent inhibition through precise serotonergic spiking, while other inputs, such as prefrontal cortex, may be less reliable in generating this inhibition.

Consider a separate thought experiment in which strong recurrent inhibition is activated onto a population of cells that are receiving multiplexed inputs from lateral habenula and prefrontal cortex. Can recurrent inhibition change the coding regime of this subpopulation? Recall that Manuscript IV suggested that prefrontal input generates delayed spikes in serotonergic neurons through a stochastic random walk process of membrane voltage. 5-HT_{1A}R activation would lead to protracted inhibition as well as a slight decrease in membrane resistance. If the onset of recurrent inhibition was coincident with either lateral habenula or prefrontal input, it would likely not impact the fast synchronous habenular-evoked spiking (due to the protracted nature of the GIRK conductance), but may reduce the number of temporally delayed prefrontal-evoked spikes due to slow inhibition. On the other hand, if the onset of recurrent inhibition preceded inputs from lateral habenula or prefrontal cortex slightly, it would likely abolish most synchronous habenular-evoked spiking due to strong inhibition that

coincides with habenular input, but would allow some temporally delayed prefrontal-evoked spikes to remain. Thus, while speculative, it appears logical that recurrent inhibition could have an all-or-nothing effect in suppressing synchronous lateral habenula-evoked spiking, depending on the relative timing of the two events, but would have a more graded effect on asynchronous prefrontal-evoked spiking due to the relative lack of precision of spike timing from prefrontal cortex. Further work is required to investigate these possibilities.

More generally, Manuscripts III and IV suggest a slightly revised view of how information is represented and processed within DRN. Evidence has accumulated for several decades indicating that populations of serotonergic neurons can represent multiple environmental and internal variables through changes in their firing rate, including, on short timescales, cues predicting rewards and punishments, punishments and rewards themselves, and motor variables; and on long timescales, the current phase of the sleep/wake cycle, recent reward history, and reward uncertainty (covered in more detail in Introduction, Section 2.3). How is such information communicated and decoded by downstream circuits? The work in this thesis suggests that viewing the activity of DRN serotonergic circuitry as a dynamical system, with multiple functionally segregated submodules, may help reconcile current theories of the variables represented by serotonin neurons. Manuscript III suggests a model where the coding properties of serotonergic circuits can vary dynamically through time, depending on whether certain long-range inputs are active at sufficiently high frequencies to generate strong recurrent inhibition. This view, emphasizing the neural computations being performed by a population, rather than the coding of environmental or behavioral variables by individual neurons, may represent a mechanism for rapid state transitions in changing environments. These state transitions could be read out by neurons receiving serotonergic output in functionally segregated circuitry, leading to behavioral changes. Manuscript IV suggests that temporal coding may be a viable strategy in serotonergic circuits, alongside the more commonly investigated rate coding (changes in firing rate of single serotonin neurons). This view emphasizes spike timing correlations between populations of serotonergic neurons as a

mechanism for representing and communicating information, a view that has obtained recent support from miniendoscopic imaging studies in DRN (Paquelet et al., 2022).

Silent synapses as a substrate for learning in cortical networks

Manuscripts I and II focus on methods to accurately quantify synaptic ensembles (containing silent and nonsilent synapses) in the developing and adult brain. A body of recent work has indicated that the fraction of silent synapses can be regulated well into adulthood, through mechanisms such as administration of drugs of addiction and stress (Huang et al., 2009; Brown et al., 2011; Ma et al., 2014; Ma et al., 2016; Graziane et al., 2016). Manuscript I strongly suggests that to quantify these silent synapse fractions with high confidence, researchers should employ either our synthetic likelihood-based estimator with electrical stimulation approaches, using directed hypothesis testing, or should employ two-photon glutamate uncaging techniques for higher-throughput inference along with subcellular localization data. The latter technique could allow these data to be integrated with current theories on how learning rules and network architecture enable circuits to function, in order to obtain a mechanistic description of how silent synapses interact with learning processes to shape behavior (covered in Introduction, Section 1.4). While beyond the scope of this thesis, it is tremendously exciting to imagine how dynamically regulated structural features of neural circuits (such as silent synapse composition) could interact with learning rules to change the nature of how and what these circuits learn well into adulthood.

This work additionally demonstrates that when interrogating synaptic populations through electrical stimulation or other methods, subsampling problems need to be carefully considered for accurate inference of features of the underlying population. At the core of Manuscript I is the idea that when electrically stimulating a population of synapses and targeting a particular failure rate (e.g. 50% failures), the number of synapses sampled from is intimately linked to their respective release probabilities, creating severe subsampling problems in some cases. These ideas may need to be carefully considered as a general principle when

designing electrical stimulation experiments. In the case of electrical stimulation of serotonergic axons to elucidate short-term facilitation properties (Manuscript III), this consideration likely not a factor since a particular failure rate was not targeted; however, it is possible that the profile of release probabilities of serotonergic axons influenced some experiments (such as the low-frequency stimulation experiments demonstrating successes and failures). Further work should carefully consider release probability constraints at central synapses when trying to infer the makeup of synapse populations.

Synthetic likelihood functions as a general, scalable approach for simulation-based inference

To date, likelihood-based inference techniques are popular for accurately inferring complex features of brains and behavior. For example, maximum likelihood estimation techniques are commonly employed in reward learning paradigms to estimate learning parameters. These techniques rely on finding an analytical expression for a likelihood function (expressing the probability of some data given model parameters), and then maximizing the likelihood for the precise trial-by-trial sequence of observed data. However, such approaches can be intractable if a likelihood function cannot be written in closed form (for example, for some iterative simulations), and can experience difficulties with time-varying processes exhibiting strong stochasticity, since they attempt to maximize the likelihood for the exact sequence of data observed, which may not be observed again even for subsequent runs with identical parameters. Manuscript I, along with earlier work (Wood, 2010), suggests that summary statistics extracted from iterative simulations can be used to build a synthetic likelihood function (not requiring a precise analytical solution), providing a potential generalizable solution to these problems.

While the idea of simulation-based inference is not novel (e.g. Sunnåker et al., 2013; Gonçalves et al., 2020), our results in Manuscripts I and II suggest a simple and scalable framework for relating the output of a time-varying model, with various parameters, to real data. By focusing on an interpretable model, and extracting reasonable summary statistics, a

likelihood function can be rapidly approximated for efficient inference. Similar approaches have been employed recently for fitting parameters in spiking neural networks, where optimization using summary statistics can yield parameters which generate more physiologically realistic simulated activity patterns compared to the more classical approach of maximizing the likelihood of a set of spiketrains (Bellec et al., 2021). The strength of this approach is in its simplicity, interpretability, and high performance. We speculate that application of a similar approach to other domains (for example, more complex reinforcement learning paradigms) could improve inference accuracy.

Concluding remarks

Here, I have presented several major lines of investigation into the synaptic basis of neural computations, in the serotonergic system and in hippocampus. In the dorsal raphe nucleus, I have focused on how synaptic and neural dynamics can imbue this system with the ability to perform contextual computations and represent multiple streams of information in parallel. Slow, serotonergic recurrent dynamics, I showed, can generate contextual operations that rely on synaptic facilitation, allowing threat signals from lateral habenula to control serotonergic output in functionally segregated submodules. These results suggest a new principle of local processing of information within serotonergic circuits through synaptic facilitation mechanisms. I have also documented how temporal coding is a viable mechanism for multiplexing data in serotonergic circuits, and how long-range synaptic input to serotonin neurons can lead to temporally distinct spike patterns to support such a coding scheme. Finally, in hippocampus, I have demonstrated that the composition of synaptic ensembles (silent synapses) can be inferred with high precision using a simulation-based synthetic likelihood approach, allowing the developmental profile of silent synapses to be quantified with high precision. Together, these results demonstrate the power of synaptic dynamics and composition in dictating circuit-level computations.

Bibliography

- Abbott, L.F. & Regehr, W.G. (2004). Synaptic computation. *Nature* 431, 796-803.
- Abbott, L.F., Varela, J.A., Sen, K. & Nelson, S.B. (1997). Synaptic depression and cortical gain control. *Science* 275, 220-224.
- Abrahamsson, T., Gustafsson, B. & Hanse, E. (2007). Reversible synaptic depression in developing rat CA3 CA1 synapses explained by a novel cycle of AMPA silencing-unsilencing. *J Neurophysiol* 98, 2604–2611.
- Aghajanian, G.K., Foote, W.E. & Sheard, M.H. (1968). Lysergic acid diethylamide: sensitive neuronal units in the midbrain raphe. *Science* 161, 706-708.
- Aghajanian, G.K. & Wang, R.Y. (1977). Habenular and other midbrain raphe afferents demonstrated by a modified retrograde tracing technique. *Brain Res* 122, 229-242.
- Aghajanian, G.K., Wang, R.Y. & Baraban, J. (1978). Serotonergic and non-serotonergic neurons of the dorsal raphe: reciprocal changes in firing induced by peripheral nerve stimulation. *Brain Res* 153, 169-175.
- Aitkin, K. & Mihalas, S. (2023). Neural Population Dynamics of Computing with Synaptic Modulations. *eLife* <https://doi.org/10.7554/eLife.83035>.
- Amo, R., Fredes, F., Kinoshita, M., Aoki, R., Aizawa, H., Agetsuma, M., Aoki, T., Shiraki, T., Kakinuma, H., Matsuda, M., Yamazaki, M., Takahoko, M., Tsuboi, T., Higashijima, S.-I., Miyasaka, N., Koide, T., Yabuki, Y., Yoshihara, Y., Fukai, T. & Okamoto, H. (2014) The habenulo-raphé serotonergic circuit encodes an aversive expectation value essential for adaptive active avoidance of danger. *Neuron* 84, 1034-1048.
- Andrade, R. & Haj-Dahmane, S. (2013). Serotonin Neuron Diversity in the Dorsal Raphe. *ACS Chem Neurosci* 4, 22-25.
- Andrade, R., Malenka, R.C. & Nicoll, R.A. (1986). A G protein couples 5-HT and GABAB receptors to the same channels in hippocampus. *Science* 234, 1261-1265.
- Araya, R. (2014). Input transformation by dendritic spines of pyramidal neurons. *Front Neuroanat* 8:141, doi: 10.3389/fnana.2014.00141.
- Asaoka, N., Nishitani, N., Kinoshita, H., Kawai, H., Shibui, N., Nagayasu, K., Shirakawi, H., Nagakawi, T. & Kaneko, S. (2017). Chronic antidepressant potentiates spontaneous activity of dorsal raphe serotonergic neurons by decreasing GABAB receptor-mediated inhibition of L-type calcium channels. *Sci Rep* 7, 13609.
- Ashby, M.C. & Isaac, J.T.R. (2011). Maturation of a recurrent excitatory neocortical circuit by experience-dependent unsilencing of newly formed dendritic spines. *Neuron* 70, 510–521.
- Ayaz, A. & Chance, F.S. (2009). Gain modulation of neuronal responses by subtractive and divisive mechanisms of inhibition. *J Physiol* 101, 958-968.
- Azimi, Z., Barzan, R., Spoida, K., Surdin, T., Wollenweber, P., Mark, M.D., Herlitze, S. & Jancke, D. (2020). Separable gain control of ongoing and evoked activity in the visual cortex by serotonergic input. *eLife* 9:e53552.

- Azouz, R. & Gray C.M. (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proc Natl Acad Sci USA* 97, 8110–8115.
- Ballestro, J., de San Martín, J.Z., Goutman, J., Elgoyhen, A.B., Fuchs, P.A. & Katz, Eleonora. (2011). Short-Term Synaptic Plasticity Regulates the Level of Olivocochlear Inhibition to Auditory Hair Cells. *J Neurosci* 31, 14763–14774.
- Bang, S.J. & Commons, K.G. (2012). Forebrain GABAergic Projections From the Dorsal Raphe Nucleus Identified by Using GAD67–GFP Knock-In Mice. *J Comp Neuro* 520, 4157–4167.
- Barak, O., Tsodyks, M. (2007). Persistent Activity in Neural Networks with Dynamic Synapses. *PLOS Comp Biol* 3, e35.
- Barak, O., Tsodyks, M. (2014). Working models of working memory. *Curr Opinion Neurobio* 25, 20-25.
- Bari, B.A., Grossman, C.D., Lubin, E.E., Rajagopalan, A.E., Cressy, J.I. & Cohen, J.Y. (2019). Stable Representations of Decision Variables for Flexible Behavior. *Neuron* 103, 922–933.
- Barnes, N.M. & Sharp, T. (1999). A review of central 5-HT receptors and their function. *Neuropharmacol* 38, 1083–1152.
- Barri, A., Wiechert, M.T., Jazayeri, M. & DiGregorio, D.A. (2022). Synaptic basis of a sub-second representation of time in a neural circuit model. *Nat Comms* 13, 7902.
- Beaumont, M. A., Zhang, W., & Balding, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics*, 162(4), 2025-2035.
- Béïque, J.-C., Imad, M., Mladenovic, L., Gingrich, J.A. & Andrade, R. (2007). Mechanism of the 5-hydroxytryptamine 2A receptor-mediated facilitation of synaptic activity in prefrontal cortex. *Proc Natl Acad Sci USA* 104, 9870–9875.
- Béïque, J.-C., Lin, D.-T., Kang, M.-G., Aizawa, H., Takamiya, K., & Huganir, R.L. (2006). Synapse-specific regulation of AMPA receptor function by PSD- 95. *Proc Natl Acad Sci USA* 103, 19535–19540.
- Bellec, G., Wang, S., Modirshanechi, A., Brea, J. & Gerstner, W. (2021). Fitting summary statistics of neural data with a differentiable spiking network simulator. *arxiv* 2106.10064v2.
- Beroun, A. et al. (2018). Generation of silent synapses in dentate gyrus correlates with development of alcohol addiction. *Neuropsychopharmacol* 43, 1989 – 1999.
- Best, J., Duncan, W., Sadre-Marandi, F., Hashemi, P., Nijhout, F. & Reed, M. (2020). Autoreceptor control of serotonin dynamics. *BMC Neurosci* 21, 40.
- Bi, G. & Poo, M. (1998). Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *J Neurosci* 18, 10464–10472.

- Bliss, T.V.P. & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anesthetized rabbit following stimulation of the perforant path. *J Physiol* 232, 331-356.
- Boureau, Y.-L. & Dayan, P. (2011). Opponency Revisited: Competition and Cooperation Between Dopamine and Serotonin. *Neuropsychopharm Rev* 36, 74-97.
- Branco, T., Staras, K., Darcy, K.J. & Goda, Y. (2008). Local dendritic activity sets release probability at hippocampal synapses. *Neuron* 59, 475-485.
- Broadhurst, P.L. (1957). Emotionality and the Yerkes-Dodson Law. *J Exp Psychol* 54, 345-352.
- Bromberg-Martin, E.S. & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nat Neurosci* 14, 1209-1216.
- Bromberg-Martin, E.S., Hikosaka, O. & Nakamura, K. (2010). Coding of Task Reward Value in the Dorsal Raphe Nucleus. *J Neurosci* 30, 6262-6272.
- Brown, T.E., Lee, B.R., Mu, P., Ferguson, D., Dietz, D., Ohnishi, Y.N., Lin, Y., Suska, A., Ishikawa, M., Huang, Y.H., et al. (2011). A silent synapse-based mechanism for cocaine-induced locomotor sensitization. *J. Neurosci.* 31, 8163-8174.
- Brunel, N. (2016). Is cortical connectivity optimized for storing information? *Nat Neurosci* 19, 749-755.
- Brunel, N., Hakim, V., Isope, P., Nadal, J.-P., & Barbour, B. (2004). Optimal information storage and the distribution of synaptic weights: perceptron versus Purkinje cell. *Neuron* 43, 745-757.
- Bunin, M.A. & Wightman, R.M. (1998). Quantitative Evaluation of 5-Hydroxytryptamine (Serotonin) Neuronal Release and Uptake: An Investigation of Extrasynaptic Transmission. *J Neurosci* 18, 4854-4860.
- Buonomano, D.V. (2000). Decoding Temporal Information: A Model Based on Short-Term Synaptic Plasticity. *J Neurosci* 20, 1129-1141.
- Buonomano, D.V., Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat Rev Neurosci* 10, 113-125.
- Burhis, T.M. & Aghajanian, G.K. (1987). Pacemaker Potentials of Serotonergic Dorsal Raphe Neurons: Contribution of a Low-Threshold Ca²⁺ Conductance. *Synapse* 1, 582-588.
- Busetto, G., Higley, M.J. & Sabatini, B.L. (2008). Developmental presence and disappearance of postsynaptically silent synapses on dendritic spines of rat layer 2/3 pyramidal neurons. *J Physiol* 586, 1519-1527.
- Cabezas, C. & Buno, W. (2006). Distinct transmitter release properties determine differences in short-term plasticity at functional and silent synapses. *J Neurophysiol* 95, 3024-3034.
- Cain, S.M. & Snutch, T.P. (2013). T-type calcium channels in burst-firing, network synchrony, and epilepsy. *Biochimica et Biophysica Acta* 1828, 1572-1578.

Calizo, L.H., Akanwa, A., Ma, X., Pan, Y.-Z., Lemos, J.C., Craige, C., Heemstra, L.A. & Beck, S.G. (2011). Raphe serotonin neurons are not homogenous: Electrophysiological, morphological and neurochemical evidence. *Neuropharmacol* 61, 524-543.

Campagnola, L., Seeman, S.C., Chartrand, T., Kim, L., Hoggarth, A., Gamlin, C., Ito, S., Trinh, J., Davoudian, P., Radaelli, C., et al., (2022). Local connectivity and synaptic dynamics in mouse and human neocortex. *Science* 375, eabj5861. DOI: 10.1126/science.abj5861.

Campolongo, P., Roozendaal, B., Trezza, V., Hauer, D., Schelling, G., McGaugh, J.L. & Cuomo, V. (2009). Endocannabinoids in the rat basolateral amygdala enhance memory consolidation and enable glucocorticoid modulation of memory. *Proc Natl Acad Sci USA* 106, 4888-4893.

Carta, M. et al. (2014). Membrane lipids tune synaptic transmission by direct modulation of presynaptic potassium channels. *Neuron* 81, 787-799.

Carter, A.G. & Regehr, W.G. (2002). Quantal events shape cerebellar interneuron firing. *Nat Neurosci* 5, 1309-1318.

Caruso, V.C., Mohl, J.T., Glynn, C., Lee, J., Willett, S.M., Zaman, A., Ebihara, A.F., Estrada, R., Freiwald, W.A., Tokdar, S.T. & Groh, J.M. (2018). Single neurons may encode simultaneous stimuli by switching between activity patterns. *Nat Comms* 9, DOI: 10.1038/s41467-018-05121-8

Cavaccini, A., Gritti, M., Giorgi, A., Locarno, A., Heck, N., Migliarini, S., Bertero, A., Mereu, M., Margiani, G., Trusel, M., Catelani, T., Marotta, R., De Luca, M.A., Caboche, J., Gozzi, A., Pasqualetti, M. & Tsononi, R. (2018). Serotonergic Signaling Controls Input-Specific Synaptic Plasticity at Striatal Circuits. *Neuron* 98, 801-816.

Cazettes, F., Reato, D., Morais, J.P., Renart, A. & Mainen, Z.F. (2021). Phasic Activation of Dorsal Raphe Serotonergic Neurons Increases Pupil Size. *Curr Biol* 31, 192-197.

Celada, P., Puig, M.V. & Artigas, F. (2013). Serotonin modulation of cortical neurons and networks. *Front Integ Neurosci* 7, 25.

Celada, P., Puig, M.V., Casanovas, J.M., Guillazo, G. & Artigas, F. (2001). Control of dorsal raphe serotonergic neurons by the medial prefrontal cortex: Involvement of serotonin-1A, GABAA, and glutamate receptors. *J Neurosci* 21, 9917-9929.

Chancey, J.H., Poulsen, D.J., Wadiche, J.I. & Overstreet-Wadiche, L. (2014). Hilar mossy cells provide the first glutamatergic synapses to adult-born dentate granule cells. *J Neurosci* 34, 2349-2354.

Chazal, G. & Ralston, H.J. (1987). Serotonin-containing structures in the nucleus raphe dorsalis of the cat: An ultrastructural analysis of dendrites, presynaptic dendrites, and axon terminals. *Comp Neurol* 259, 317-329.

Chevalleyre, V., Takahashi, K.A. & Castillo, P.E. (2006). Endocannabinoid-mediated synaptic plasticity in the CNS. *Ann Rev Neurosci* 29, 37-76.

Chizat, L., Oyallon, E. & Bach, F. (2019). On Lazy Training in Differentiable Programming. *arxiv* 1812.07956v5.

- Cho, J.R., Treewek, J.B., Robinson, J.E., Xiao, C., Bremner, L.R., Greenbaum, A. & Gradinaru, V. (2017). Dorsal Raphe Dopamine Neurons Modulate Arousal and Promote Wakefulness by Salient Stimuli. *Neuron* 94, 1-15.
- Christoffel, D.J., Walsh, J.J., Hoerbelt, P., Heifets, B.D., Llorach, P., Lopez, R.C., Ramakrishnan, C., Deisseroth, K. & Malenka, R.C. (2021). Selective filtering of excitatory inputs to nucleus accumbens by dopamine and serotonin. *Proc Natl Acad Sci U.S.A.* 118, e2106648118.
- Chung, S., Li, X., Nelson, S.B. (2002). Short-Term Depression at Thalamocortical Synapses Contributes to Rapid Adaptation of Cortical Sensory Responses In Vivo. *Neuron* 34, 437–446.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I. & Shenoy, K.V. (2012). Neural population dynamics during reaching. *Nature* 487, 51-56.
- Churchland, M.M. & Shenoy, K.V. (2007). Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J Neurophysiol* 97, 4235–4257.
- Clopath, C., & Brunel, N. (2013). Optimal properties of analog perceptrons with excitatory weights. *PLoS Comput Biol* 9, e1002919.
- Coffey, K.R., Marx, R.G., Vo, E.K., Nair, S.G. & Neumaier, J.F. (2020). Chemogenetic inhibition of lateral habenula projections to the dorsal raphe nucleus reduces passive coping and perseverative reward seeking in rats. *Neuropsychopharm* 45, 1115-1124.
- Cohen, J.Y., Amoroso, M.W. & Uchida, N. (2015). Serotonergic neurons signal reward and punishment on multiple timescales. *eLife* 4, e06346.
- Cohen, J.Y., Haesler, S., Vong, L., Lowell, B.B. & Uchida, N. (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature* 482, 85-88.
- Colgan, L.A., Cavolo, S.L., Commons, K.G. & Levitan, E.S. (2012). Action potential-independent and pharmacologically unique vesicular serotonin release from dendrites. *J Neurosci* 32, 15737-15746.
- Commons, K.G. (2015). Two Major Network Domains in the Dorsal Raphe Nucleus. *J Comp Neuro* 523, 1488–1504.
- Commons, K.G. (2016). Ascending serotonin neuron diversity under two umbrellas. *Brain Struct Funct* 221, 3347–3360.
- Constantinidis, C., Funahashi, S., Lee, D., Murray, J.D., Qi, X.-L., Wang, M. & Arnsten, A.F.T. (2018). Persistent Spiking Activity Underlies Working Memory. *J Neurosci* 38, 7020 –7028.
- Cools, R., Roberts, A. & Robbins, T.W. (2008). Serotonergic regulation of emotional and behavioural control processes. *TRENDS in Cog Sci* 12, 31-40.
- Coultrip, R., Granger, R. & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks* 5, 47-54.
- Courtney, N.A. & Ford, C.P. (2016). Mechanisms of 5-HT1A receptor-mediated transmission in dorsal raphe serotonin neurons. *J Physiol* 594, 953–965.

Cragg, S. (2003). Variable dopamine release probability and short-term plasticity between functional domains of the primate striatum. *J Neurosci* 23, 4378-4385.

Cramér, H. (1945). *Mathematical Methods of Statistics* (Uppsala: Almqvist & Wiksells).

Dahlstroem, A. & Fuxe, K. (1964). Evidence for the existence of monoamine-containing neurons in the central nervous system. I. Demonstration of monoamines in the cell bodies of brain stem neurons. *Acta Physiol Scand Suppl* 232, 1-55.

Dankoski, E.C. & Wightman, R.M. (2013). Monitoring serotonin signaling on a subsecond time scale. *Front Integ Neurosci* 7, 44.

Daszuta, A. & Portalier, P. (1985). Distribution and Quantification of 5-HT Nerve Cell Bodies in the Nucleus Raphe Dorsalis Area of C57BL and BALBc Mice. Relationship between Anatomy and Biochemistry. *Brain Res* 360, 58-64.

Davis, M., Sheard, M.H. (1974). Habituation and sensitization of the rat startle response: effects of raphe lesions. *Physiology & Behavior* 12, 425-431.

Daw, N.D., Kakade, S. & Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Networks* 15, 603-616.

Deakin, J.F.W. (1983). Roles of serotonergic systems in escape, avoidance and other behaviours. In: Cooper SJ, (ed.). *Theory in Psychopharmacology* Vol 2, 2nd edn. Academic Press: New York, pp 149-193.

Deakin, J.F.W. & Graeff, F.G. (1991). 5-HT and mechanisms of defence. *J Psychopharm* 5, 305-315.

Debanne, D., Boudkkaszi, S., Campanac, E., Cudmore, R.H., Giraud, P., Fronzaroli-Molinieres, L., Carlier, E. & Caillard, O. (2008). Paired-recordings from synaptically coupled cortical and hippocampal neurons in acute and cultured brain slices. *Nat Protocols* 3, 1559-1568.

de Hoz, L. & Nelken, I. (2014). Frequency tuning in the behaving mouse: Different bandwidths for discrimination and generalization. *PLOS One* 9, e91676.

Denève, S. & Machens, C.K. (2016). Efficient codes and balanced networks. *Nat Neurosci* 19, 375-382.

Descarries, L., Watkins, K.C., Garcia, S. & Beaudet, A. (1982). The serotonin neurons in nucleus raphe dorsalis of adult rat: A light and electron microscope radioautographic study. *J Comp Neurol* 207, 239-254.

Destexhe, A., Rudolph, M., Fellous, J.-M. & Sejnowski, T.J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience* 107, 13-24.

Di, S., Itoga, C.A., Fisher, M.O., Solomonow, J., Roltsch, E.A., Gilpin, N.W. & Tasker, J.G. (2016). Acute stress suppresses synaptic inhibition and increases anxiety via endocannabinoid release in the basolateral amygdala. *J Neurosci* 36, 8461-8470.

- Di, S., Malcher-Lopes, R., Halmos, K.C & Tasker, J.G. (2003) Nongenomic glucocorticoid inhibition via endocannabinoid release in the hypothalamus: a fast feedback mechanism. *J Neurosci* 23, 4850–4857.
- Diggle, P.J. & Gratton, R.J. (1984). Monte Carlo methods of inference for implicit statistical models. *J Royal Stat Soc Ser B (Methodological)* 46, 193–212.
- Dobrunz, L.E. & Stevens, C.S. (1997). Heterogeneity of release probability, facilitation and depression at central synapses. *Neuron* 18, 995-1008.
- Dolzani, S.D., Baratta, M.V., Amat, J., Agster, K.L., Sadoris, M.P., Watkins, L.R. & Maier, S.F. (2016). Activation of a habenulo-raphé circuit is critical for the behavioral and neurochemical consequences of uncontrollable stress in the male rat. *eNeuro* 3:e0229-16.2016.
- Dorocic, I.P., Fürth, D., Xuan, Y., Johansson, Y., Pozzi, L., Silberberg, G., Carlén, M. & Meletis, K. (2014). A Whole-Brain Atlas of Inputs to Serotonergic Neurons of the Dorsal and Median Raphe Nuclei. *Neuron* 83, 663–678.
- Dougalis, A.G., Matthews, G.A.C., Bishop, M.W., Brischoux, F, Kobayashi, K. & Ungless, M.A. (2012). Functional properties of dopamine neurons and co-expression of vasoactive intestinal polypeptide in the dorsal raphe nucleus and ventro-lateral periaqueductal grey. *Eur J Neurosci* 36, 3322–3332.
- Douglas, R.J., Koch, C., Mahowald, M., Martin, K.A.C. & Suarez, H.H. (1995). Recurrent excitation in neocortical circuits. *Science* 269, 981-985.
- Driscoll, L., Shenoy, K. & Sussillo, D. (2022). Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *bioRxiv* <https://doi.org/10.1101/2022.08.15.503870>.
- Durand, G.M., Kovalchuk, Y., & Konnerth, A. (1996). Long-term potentiation and functional synapse induction in developing hippocampus. *Nature* 381, 71–75.
- Eccles, J.C., Katz, B. & Kuffler, S.W. (1941). Nature of the “endplate potential” in curarized muscle. *J Neurophysiol* 4, 362–387.
- Echeveste, R., Aitchison, L., Hennequin, G. & Lengyel, M. (2020). Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nat Neurosci* 23, 1138-1149.
- Eriksson, A., Raczkowska, M., Navawongse, R., Choudhury, R., Stewart, J.C., Tang, Y.L., Wang, Z., Claridge-Chang, A. (2017). Neuromodulatory circuit effects on *Drosophila* feeding behaviour and metabolism. *Sci Rep* 7, 8839.
- Ermentrout, G.B., Galan, R.F. & Urban, N.N. (2008). Reliability, synchrony and noise. *Trends in Neurosci* 31, 428-434.
- Evans, D.A., Stempel, A.V., Vale, R., Ruehle, S., Lefler, Y. & Branco, T. (2018). A synaptic threshold mechanism for computing escape decisions. *Nature* 558, 590-594.
- Feng, T.P. (1940). Studies on the neuromuscular junction. XVIII. The local potentials around n-m junctions induced by single and multiple volleys. *Chin J Physiol* 15, 367–404.

Fioravante, D. & Regehr, W.G. (2011). Short-term forms of presynaptic plasticity. *Curr Opin Neurobiol* 21, 269-274.

Fisher, S.A., Fisher, T.M. & Carew, T.J. (1997). Multiple overlapping processes underlying short-term synaptic enhancement. *TINS* 20, 170-177.

Flesch, T., Saxe, A. & Summerfield, C. (2022a). Continual task learning in natural and artificial agents. *TINS* 46, 199-210.

Flesch, T., Juechems, K., Dumbalska, T., Saxe, A. & Summerfield, C. (2022b). Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* 110, 1258-1270.

Fonseca, M.S., Murakami, M. & Mainen, Z.F. (2015). Activation of Dorsal Raphe Serotonergic Neurons Promotes Waiting but Is Not Reinforcing. *Curr Biol* 25, 306-315.

Fortune, E.S., Rose, G.J. (2001). Short-term synaptic plasticity as a temporal filter. *TINS* 24, 381-385.

Fredj, N.B. & Burrone, J. (2009). A resting pool of vesicles is responsible for spontaneous vesicle fusion at the synapse. *Nat Neurosci* 12, 751-758.

Friedrich, R.W., Habermann, C.J. & Laurent, G. (2004). Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7, 862-871.

Fu, W., Le Maître, E., Fabre, V., Bernard, J.-F., Xu, Z.-Q. D. & Hökfelt, T. (2011). Chemical Neuroanatomy of the Dorsal Raphe Nucleus and Adjacent Structures of the Mouse Brain. *J Comp Neurol* 518, 3464-3494.

Funahashi, S., Bruce, C.J. & Goldman-Rakic, P.S. (1989). Mnemonic Coding of Visual Space in the Monkey's Dorsolateral Prefrontal Cortex. *J Neurophysiol* 61, 331-349.

Fuster, J.M. & Alexander, G.E. (1971). Neuron activity related to short-term memory. *Science* 173, 652-654.

Fuxe, K. & Agnati, L.F. (1991). Two principal modes of electrochemical communication in the brain: volume versus wiring transmission. In Fuxe K, Agnati LF (eds): "Volume Transmission in the Brain: Novel Mechanisms for Neural Transmission." New York: Raven Press, *Adv Neurosci* 1, pp 1-9.

Geddes, S.D., Assadzada, S., Lemelin, D., Sokolovski, A., Bergeron, R., Haj-Dahmane, S. & Béique, J.-C. (2016). Target-specific modulation of the descending prefrontal cortex inputs to the dorsal raphe nucleus by cannabinoids. *PNAS* 113, 5429-5434.

George, J., Soares, C., Montersino, A., Béique, J.-C. & Thomas, G.M. (2015). Palmitoylation of LIM Kinase-1 ensures spine-specific actin polymerization and morphological plasticity. *Elife* 4, e06327.

Gerstein, G.L. & Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophys J* 4, 42-68.

Gillespie, D.T. (1996). Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral. *Phys Rev E* 54(2), 2084-2091.

Gonçalves, P.J., Lueckmann, J.-M., Deistler, M., Nonnenmacher, M., Öcal, K., Bassetto, G., Chintaluri, C., Podlaski, W.F., Haddad, S.A., Vogels, T.P., Greenberg, D.S. & Macke, J.H. (2020). Training deep neural density estimators to identify mechanistic models of neural dynamics. *eLife* 9, e56261.

Gozlan, H., El Mestikawy, S., Pichat, L., Glowinski, J. & Hamon, M. (1983). Identification of presynaptic 5-HT autoreceptors using a new ligand: 3H-PAT. *Nature* 305:140-142.

Graeff, F.G., Guimaraes, F.S., de Andrade, T.G.C.S. & Deakin, J.F.W. (1996). Role of 5-HT in Stress, Anxiety and Depression. *Pharmacol Biochem Behav* 54, 129-141.

Gray, J.M., Vecchiarelli, H.A, Morena, M., Lee, T.T.Y., Hermanson, D.J., Kim, A.B., McLaughlin, R.J., Hassan, K.I., Kühne, C., Wotjak, C.T., Deussing, J.M., Patel, S., Hill, M.N. (2015). Corticotropin-releasing hormone drives anandamide hydrolysis in the amygdala to promote anxiety. *J Neurosci* 35, 3879–3892.

Graziane, N. & Dong, Y. (2016). Measurement of silent synapses. In *Electrophysiological Analysis of Synaptic Transmission, Neuromethods* (New York, NY: Springer New York), pp. 217–224.

Graziane, N.M., Sun, S., Wright, W.J., Jang, D., Liu, Z., Huang, Y.H., Nestler, E.J., Wang, Y.T., Schluter, O.M., & Dong, Y. (2016). Opposing mechanisms mediate morphine- and cocaine-induced generation of silent synapses. *Nat Neurosci* 19, 915–925.

Greenberg, D., Nonnenmacher, M., & Macke, J. (2019). Automatic Posterior Transformation for Likelihood-Free Inference. In *Proceedings of the 36th International Conference on Machine Learning*.

Grienberger, C., Milstein, A.D., Bittner, K.C., Romani, S. & Magee, J.C. (2017). Inhibitory suppression of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat Neurosci* 20, 417-426.

Grillner, S. (2006). Biological Pattern Generation: The Cellular and Computational Logic of Networks in Motion. *Neuron* 52, 751–766.

Grossman, C.D., Bari, B.A. & Cohen, J.Y. (2022). Serotonin neurons modulate learning rate through uncertainty. *Curr Biol* 32, 1-14.

Hansel, D. & Mato, G. (2013). Short-Term Plasticity Explains Irregular Persistent Activity in Working Memory Tasks. *J Neurosci* 33, 133–149.

Harkin, E.F., Lynn, M.B., Payeur, A., Boucher, J.-F., Caya-Bissonnette, L., Cyr, D., Stewart, C., Longtin, A., Naud, R. & Béique, J.-C. (2023). Temporal derivative computation in the dorsal raphe network revealed by an experimentally driven augmented integrate-and-fire modeling framework. *eLife* 12, e72951.

Harvey, C. & Svoboda, K. (2007). Locally dynamic synaptic learning rules in pyramidal neuron dendrites. *Nature* 450, 1195-1200.

Harvey, C., Yasuda, R., Zhong, H. & Svoboda, K. (2008). The spread of ras activity triggered by activation of a single dendritic spine. *Science* 321, 136-140.

Hayashi, K., Nakao, K. & Nakamura, K. (2015). Appetitive and Aversive Information Coding in the Primate Dorsal Raphe Nucleus. *J Neurosci* 35, 6195– 6208.

Hernández-Vázquez, F., Garduño, J. & Hernández-López, S. (2019). GABAergic modulation of serotonergic neurons in the dorsal raphe nucleus. *Rev Neurosci* 30, 289–303.

Higley, M.J. & Sabatini, B.L. (2008). Calcium Signaling in Dendrites and Spines: Practical and Functional Considerations. *Neuron* 59, 902-913.

Hill, M.N., Patel, S., Campolongo, P., Tasker, J.G., Wotjak, C.T. & Bains, J.S. (2010). Functional interactions between stress and the endocannabinoid system: from synaptic signaling to behavioral output. *J Neurosci* 30, 14980–14986.

Hioki, H., Nakamura, H., Ma, Y.-F., Konno, M., Hayakawa, T., Nakamura, K.C., Fujiyama, F. & Kaneko, T. (2010). Vesicular Glutamate Transporter 3-Expressing Nonserotonergic Projection Neurons Constitute a Subregion in the Rat Midbrain Raphe Nuclei. *J Comp Neurol* 518, 668–686.

Howe, M.W. & Dombeck, D.A. (2016). Rapid signalling in distinct dopaminergic axons during locomotion and reward. *Nature* 535, 505-510.

Howe, M.W., Tierney, P.L., Sandberg, S.G., Phillips, P.E.M. & Graybiel, A.M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature* 500, 575-579.

Howlett, A.C. et al. (2004). Cannabinoid physiology and pharmacology: 30 years of progress. *Neuropharmacology* 47, 345-358.

Hoyer, D., Hannon, J.P. & Martin, G.R. (2002). Molecular, pharmacological and functional diversity of 5-HT receptors. *Pharmacol Biochem Behav* 71, 533-554.

Huang, K.W., Ochandarena, N.E., Philson, A.C., Hyun, M., Birnbaum, J.E., Cicconet, M. & Sabatini, B.L. (2019). Molecular and anatomical organization of the dorsal raphe nucleus. *eLife* 8, e46464.

Huang, X. et al. (2015). Progressive maturation of silent synapses governs the duration of a critical period. *Proc Natl Acad Sci USA* 112, E3131–40.

Huang, Y.H., Lin, Y., Mu, P., Lee, B.R., Brown, T.E., Wayman, G., Marie, H., Liu, W., Yan, Z., Sorg, B.A., et al. (2009). In vivo cocaine experience generates silent synapses. *Neuron* 63, 40–47.

Isaac, J.T., Crair, M.C., Nicoll, R.A., & Malenka, R.C. (1997). Silent synapses during development of thalamocortical inputs. *Neuron* 18, 269–280.

Isaac, J.T.R., Nicoll, R.A. & Malenka, R.C. (1995). Evidence for silent synapses: implications for the expression of LTP. *Neuron* 15, 427–434.

Izhikevich, E.I., Desai, N.S., Walcott, E.C. & Hoppensteadt, F.C. (2003). Bursts as a unit of neural information: selective communication via resonance. *TINS* 26, 161-167.

Jackman, S.L. & Regehr, W.G. (2017). The Mechanisms and Functions of Synaptic Facilitation. *Neuron* 94, 447-464.

Jacobs, B.L. & Azmitia, E.C. (1992). Structure and function of the brain serotonin system. (1992). *Phys Rev* 72, 165-229.

Jakab, R.L. & Goldman-Rakic, P.S. (1998). 5-Hydroxytryptamine_{2A} serotonin receptors in the primate cerebral cortex: Possible site of action of hallucinogenic and antipsychotic drugs in pyramidal cell apical dendrites. *Proc Natl Acad Sci USA* 95, 735-740.

Jones, E., Oliphant, E., Peterson, P., et al. (2001-). Scipy: Open Source Scientific Tools for Python. <http://www.scipy.org/> [Accessed 2019-09-24].

Kandel, E.R., Dudai, Y. & Mayford, M.R. (2014). The molecular and systems biology of memory. *Cell* 157, 163-86.

Kamaleddin, M.A., Shifman, A., Abdollahi, N., Sigal, D., Ratté, S. & Prescott, S.A. (2022). Physiological noise facilitates multiplexed coding of vibrotactile-like signals in somatosensory cortex. *PNAS* 119, e2118163119.

Kaneko, K., Tamamaki, N., Owada, H., Kakizaki, T., Kume, N., Totsuka, M., Yamamoto, T., Yawo, H., Yagi, T., Obata, K., Yanagawa, Y. (2008). Noradrenergic excitation of a subpopulation of GABAergic cells in the basolateral amygdala via both activation of nonselective cationic conductance and suppression of resting K⁺ conductance: a study using glutamate decarboxylase 67-green fluorescent protein knock-in mice. *Neuroscience* 157, 781-797.

Karst, H., Berger, S., Erdmann, G., Schütz, G., Joëls, M. (2010). Metaplasticity of amygdalar responses to the stress hormone corticosterone. *Proc Natl Acad Sci USA* 107, 14449-14454.

Kass, R.E., Eden, U.T. & Brown, E.N. (2014). *Analysis of Neural Data* (New York: Springer).

Kato, T., Mitsukura, Y., Yoshida, K., Mimura, M., Takata, N. & Takana, K.F. (2022). Oscillatory Population-Level Activity of Dorsal Raphe Serotonergic Neurons Is Inscribed in Sleep Structure. *J Neurosci* 42, 7244-7255.

Khanbabaie, R., Nesse, W.H., Longtin, A. & Maler, L. (2010). Kinetics of fast short-term depression are matched to spike train statistics to reduce noise. *J Neurophysiol* 103, 3337-3348.

Klapoetke, N.C. *et al.* (2014). Independent optical excitation of distinct neural populations. *Nat Methods* 11, 338-346.

Ko, K.W., Rasband, M.N., Meseguer, V., Kramer, R.H. & Golding, N.L. (2017). Serotonin modulates spike probability in the axon initial segment through HCN channels. *Nat Neurosci* 19, 826-834.

Koya, E. *et al.* (2012). Silent synapses in selectively activated nucleus accumbens neurons following cocaine sensitization. *Nat Neurosci* 15, 1556-1562.

- Kuhn, A., Aertsen, A. & Rotter, S. (2004). Neuronal integration of synaptic input in the fluctuation-driven regime. *J Neurosci* 24, 2345–2356.
- Lammel, S., Lim, B.K., Ran, C., Huang, K.W., Betley, M.J., Tye, K.M., Deisseroth, K. & Malenka, R.C. (2012). Input-specific control of reward and aversion in the ventral tegmental area. *Nature* 491, 212–217.
- Lankarany, M., Al-Basha, D., Ratté, S. & Prescott, S.A. (2019). Differentially synchronized spiking enables multiplexed neural coding. *PNAS* 116, 10097–10102.
- Lecca, S., Namboodiri, V.M.K., Restivo, L., Gervasi, N., Pillola, G., Stuber, G.D. & Mameli, M. (2020). Heterogeneous habenular neuronal ensembles during selection of defensive behaviors. *Cell Rep* 31, 107752.
- Lee, B.R., et al. (2013). Maturation of silent synapses in amygdala-accumbens projection contributes to incubation of cocaine craving. *Nat Neurosci* 16, 1644–1651.
- Lee, H.S., Kim, M.-A., Valentino, R.J. & Waterhouse, B.D. (2003). Glutamatergic afferent projections to the dorsal raphe nucleus of the rat. *Brain Res* 963, 57–71.
- Lee, J.-H., Gomora, J.C., Cribbs, L.L. & Perez-Reyes, E. (1999). Nickel block of three cloned T-type calcium channels: low concentrations selectively block 1H. *Biophys J* 77, 3034–3042.
- Lee, K.F.H., Soares, C., Thivierge, J.-P., & Béique, J.-C. (2016). Correlated synaptic inputs drive dendritic calcium amplification and cooperative plasticity during clustered synapse development. *Neuron* 89, 784–799.
- Lewis, J.E., Lindner, B., Laliberté, B. & Groothuis, S. (2007). Control of neuronal firing by dynamic parallel fiber feedback: Implications for electrosensory reafference suppression. *J Exp Biol* 210, 4437–4447.
- Li, Y., Zhong, W., Wang, D., Feng, Q., Liu, Z., Zhou, J., Jia, C., Hu, F., Zeng, J., Guo, Q., Fu, L. & Lu, M. (2016). Serotonin neurons in the dorsal raphe nucleus encode reward signals. *Nat Comms* 7, 10503.
- Liao, D., Hessler, N.A., & Malinow, R. (1995). Activation of postsynaptically silent synapses during pairing-induced LTP in CA1 region of hippocampal slice. *Nature* 375, 400–404.
- Lisman, J.E. (1997). Bursts as a unit of neural information: making unreliable synapses reliable. *TINS* 20, 38–43.
- Liu, C., Goel, P. & Kaeser, P.S. (2021). Spatial and temporal scales of dopamine transmission. *Nat Rev Neurosci*, doi:10.1038/s41583-021-00455-7.
- Liu, C., Kershberg, L., Wang, J., Schneeberger, S. & Kaeser, P.S. (2018). Dopamine Secretion Is Mediated by Sparse Active Zone-like Release Sites. *Cell* 172, 706–718.
- Liu, Z., Lin, R. & Luo, Minmin. (2020). Reward Contributions to Serotonergic Functions. *Annu Rev Neurosci* 43, 141–162.

Liu, Z., Zhou, J., Li, Y., Hu, F., Lu, Y., Ma, M., Feng, Q., Zhang, J., Wang, D., Zeng, J., Bao, J., Kim, J.-Y., Chen, Z.-F., Mestikawy, S.E. & Luo, M. (2014). Dorsal Raphe Neurons Signal Reward through 5-HT and Glutamate. *Neuron* 81, 1360–1374.

Lottem, E., Banerjee, D., Vertechi, P., Sarra, D., oude Lohuis, M. & Mainen, Z.F. (2018). Activation of serotonin neurons promotes active persistence in a probabilistic foraging task. *Nat Comms* 9, 1000.

Lueckmann, J. M., Bassetto, G., Karaletsos, T., & Macke, J. H. (2019, January). Likelihood-free inference with emulator networks. In *Symposium on Advances in Approximate Bayesian Inference* (pp. 32-53).

Luo, M., Li, Y. & Zhong, W. (2016). Do dorsal raphe 5-HT neurons encode “beneficialness”? *Neurobiol Learn Mem* 135, 40-49.

Ma, Y.-Y., Lee, B.R., Wang, X., Guo, C., Liu, L., Cui, R., Lan, Y., Balcita-Pedicino, J.J., Wolf, M.E., Sesack, S.R., et al. (2014). Bidirectional modulation of incubation of cocaine craving by silent synapse-based remodeling of prefrontal cortex to accumbens projections. *Neuron* 83, 1453–1467.

Ma, Y.-Y., Wang, X., Huang, Y., Marie, H., Nestler, E.J., Schluter, O.M., and Dong, Y. (2016). Re-silencing of silent synapses unmasks anti-relapse effects of environmental enrichment. *Proc. Natl. Acad. Sci. U S A* 113, 5089–5094.

MacAskill, A.F., Little, J.P., Cassel, J.M. & Carter, A.G. (2012). Subcellular connectivity underlies pathway-specific signaling in the nucleus accumbens. *Nat Neurosci* 15, 1624-1626.

Magee, J.C. & Grienberger, C. (2020). Synaptic Plasticity Forms and Functions. *Ann Rev Neurosci* 43, 95-117.

Mainen, Z.F. & Sejnowski, T.J. (1995). Reliability of spike timing in neocortical neurons. *Science* 268, 1503-1506.

Mante, V., Sussillo, D., Shenoy, K.V. & Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78-84.

Marcoux, C.M., Clarke, S.E., Nesse, W.H., Longtin, A. & Maler, L. (2016). Balanced ionotropic receptor dynamics support signal estimation via voltage-dependent membrane noise. *J Neurophysiol* 115, 530-545.

Marder, E. & Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Curr Biol* 11, R986–R996.

Markram, H., Lubke, J., Frotscher, M. & Sakmann, B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213-215.

Marowsky, A., Yanagawa, Y., Obata, K. & Vogt, K.E. (2005). A specialized subclass of interneurons mediates dopaminergic facilitation of amygdala function. *Neuron* 48, 1025–1037.

Marques, J.C., Li, M., Schaak, D., Robson, D.N. & Li, J.M. (2020). Internal state dynamics shape brainwide activity and foraging behavior. *Nature* 577, 239-243.

- Martín-Ruiz, R., Puig, M.V., Celada, P., Shapiro, D.A., Roth, B.L., Mengod, G. & Artigas, F. (2001) Control of serotonergic function in medial prefrontal cortex by serotonin-2A receptors through a glutamate-dependent mechanism. *J Neurosci* 21, 9856–9866
- Mason, A., Nicoll, A. & Stratford, K. (1991). Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. *J Neurosci* 11, 72:84.
- Masse, N.Y., Yang, G.R., Song, H.F., Wang, X.-J. & Freedman, D.J. (2019). Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nat Neurosci* 22, 1159-1167.
- Mathew, S.S, Pozzo-Miller, L. & Hablitz, J.J (2008). Kainate modulates presynaptic GABA release from two vesicle pools. *J Neurosci* 28, 725–731.
- Matias, S., Lottem, E., Dugué, G.P. & Mainen, Z.F. (2017). Activity patterns of serotonin neurons underlying cognitive flexibility. *eLife* 6, e20552.
- Matsumoto, M. & Hikosaka, O. (2009). Representation of negative motivational value in the primate lateral habenula. *Nat Neurosci* 12, 77-84.
- Matsuzaki, M., Honkura, N., Ellis-Davies, G.C. & Kasai, H. (2004). Structural basis of long-term potentiation in single dendritic spines. *Nature* 429, 761-766.
- Matthews, G.A., Nieh, E.H., Vander Weele, C.M., Halbert, S.A., Pradhan, R.V., Yosafat, A.S., Glober, G.F., Isadmehr, E.M., Thomas, R.E., Lacy, G.D., Wildes, C.P., Ungless, M.A. & Tye, K.M. (2016). Dorsal Raphe Dopamine Neurons Represent the Experience of Social Isolation. *Cell* 164, 617–631.
- Maxwell, D.J., Leranath, C. & Verofstad, A.A.J. (1983). Fine structure of 5-HT-containing axons in the marginal zone of the rat spinal cord. *Brain Res* 266:253-259.
- McGinty, D.J. & Harper, R.M. (1976). Dorsal raphe neurons: depression of firing during sleep in cats. *Brain Res* 101, 569-575.
- Mejias, J.F., Payeur, A., Selin, E., Maler, L. & Longtin, A. (2014). Subtractive, divisive and non-monotonic gain control in feedforward nets linearized by noise and delays. *Front Comp Neurosci*, doi: <https://doi.org/10.3389/fncom.2014.00019>.
- Miles, R. & Wong, R.K.S. (1986). Excitatory synaptic interactions between CA3 neurones in the guinea-pig hippocampus. *J Physiol* 373, 397-418.
- Milshtein-Parush, H., Frere, S., Regev, L., Lahav, C., Benbenishty, A., Ben-Eliyahu, S., Goshen, I. & Slutsky, I. (2017). Sensory Deprivation Triggers Synaptic and Intrinsic Plasticity in the Hippocampus. *Cereb Cortex* 27, 3457–3470.
- Miyazaki, K., Miyazaki, K.W. & Doya, K. (2011). Activation of Dorsal Raphe Serotonin Neurons Underlies Waiting for Delayed Rewards. *J Neurosci* 31, 469-479.
- Miyazaki, K., Miyazaki, K.W., Sivori, G., Yamanaka, A., Tanaka, K.F. & Doya, K. (2020). Serotonergic projections to the orbitofrontal and medial prefrontal cortices differentially modulate waiting for future rewards. *Sci Adv* 6, eabc7246.

Miyazaki, K.W., Miyazaki, K., Tanaka, K.F., Yamanaka, A., Takahashi, A., Tabuchi, S. & Doya, K. (2014). Optogenetic Activation of Dorsal Raphe Serotonin Neurons Enhances Patience for Future Rewards. *Curr Biol* 24, 2033-2040.

Miyazaki, K., Miyazaki, K.W., Yamanaka, A., Tokuda, T., Tanaka, K.F. & Doya, K. (2018). Reward probability and timing uncertainty alter the effect of dorsal raphe 5-HT neurons on patience. *Nat Comms* 9, 2048.

Mlinar, B., Montalbano, A., Piszczek, L., Gross, C. & Corradetti, R. (2016). Firing Properties of Genetically Identified Dorsal Raphe Serotonergic Neurons in Brain Slices. *Front Cell Neurosci* 10, 195.

Mongillo, G., Barak, O. & Tsodyks, M. (2008). Synaptic Theory of Working Memory. *Science* 319, 1543-1546.

Montage, P.R., Dayan, P. & Sejnowski, T. (1996). A Framework for Mesencephalic Dopamine Systems Based on Predictive Hebbian Learning. *J Neurosci* 16, 1936-1947.

Moore, F.L. & Orchinik, M. (1994) Membrane receptors for corticosterone: a mechanism for rapid behavioral responses in an amphibian. *Horm Behav* 28, 512-519.

Morozov, A., Sukato, D. & Ito, W. (2011). Selective suppression of plasticity in amygdala inputs from temporal association cortex by the external capsule. *J Neurosci* 31, 339-345.

Moukhles, H., Bosler, O., Bolam, J.P., Vallée, A., Umbriaco, D., Geffard, M., Doucet, G. (1997). Quantitative and morphometric data indicate precise cellular interactions between serotonin terminals and postsynaptic targets in rat substantia nigra. *Neuroscience* 76, 1159-1171.

Muller, M., Goutman, J.D., Kochubey, O., Schneggenburger, R. (2010). Interaction between Facilitation and Depression at a Large CNS Synapse Reveals Mechanisms of Short-Term Plasticity. *J Neurosci* 30, 2007-2016.

Muzerelle, A., Scotto-Lomassese, S., Bernard, J.F., Soiza-Reilly, M. & Gaspar, P. (2016). Conditional anterograde tracing reveals distinct targeting of individual serotonin cell groups (B5-B9) to the forebrain and brainstem. *Brain Struct Funct* 221, 535-561.

Nair, A., Karigo, T., Yang, B., Ganguli, S., Schnitzer, M.J., Lindermann, S.W., Anderson, D.J. & Kennedy, A. (2023). An approximate line attractor in the hypothalamus encodes an aggressive state. *Cell* 178-193.

Nakamura, K., Mastumoto, M. & Hikosaka, O. (2008). Reward-Dependent Modulation of Neuronal Activity in the Primate Dorsal Raphe Nucleus. *J Neurosci* 28, 5331-5343.

Naud, R. & Sprekeler, H. (2018). Sparse bursts optimize information transmission in a multiplexed neural code. *PNAS* 115, E6329-E6338.

Neumann, P.A. et al. (2016). Cocaine-induced synaptic alterations in thalamus to nucleus accumbens projection. *Neuropsychopharmacol* 41, 2399-2410.

Newberry, N.R. & Nicoll, R.A.A. (1984). A bicuculline-resistant inhibitory post-synaptic potential in rat hippocampal pyramidal cells in vitro. *J Physiol* 348, 239-254.

- Ngo-Anh, T.J., Bloodgood, B.L., Lin, M., Sabatini, B.L., Maylie, J. & Adelman, J.P. (2005). SK channels and NMDA receptors form a Ca²⁺-mediated feedback loop in dendritic spines. *Nat Neurosci* 5, 642-649.
- Nichols, D.E. & Nichols, C.D. (2008). Serotonin receptors. *Chem Rev* 108, 1614–1641.
- Nogueira, R., Rogers, C.C., Bruno, R.M. & Fusi, S. (2023). The geometry of cortical representations of touch in rodents. *Nat Neurosci* 26, 239-250.
- O'Donnell, C., Nolan, M.F. & van Rossum, M. (2011). Dendritic spine dynamics regulate the long-term stability of synaptic plasticity. *J Neurosci* 31:16142–16156.
- Ogawa, S.K., Cohen, J.Y., Hwang, D., Uchida, N. & Watabe-Uchida, M. (2014). Organization of Monosynaptic Inputs to the Serotonin and Dopamine Neuromodulatory Systems. *Cell Rep* 8, 1105–1118.
- Oh, W.C., Hill, T.C & Zito, K. (2013). Synapse-specific and size-dependent mechanisms of spine structural plasticity accompanying synaptic weakening. *Proc Natl Acad Sci USA* 110, E305-E312.
- Ohmura, Y., Tanaka, K.F., Tsunematsu, T., Yamanaka, A. & Yoshioka, M. (2014). Optogenetic activation of serotonergic neurons enhances anxiety-like behaviour in mice. *Int J Neuropsychopharmacol* 17, 1777–1783.
- Okaty, B.W., Commons, K.G. & Dymecki, S.M. (2019). Embracing diversity in the 5-HT neuronal system. *Nat Rev Neurosci* 20, 397-424.
- Okaty, B.W., Freret, M.E., Rood, B.D., Brust, R.D., Hennessy, M.L., deBairos, D., Kim, J.C., Cook, M.N. & Dymecki, S.M. (2015). Multi-Scale Molecular Deconstruction of the Serotonin Neuron System. *Neuron* 88, 774-791.
- Okaty, B.W., Sturrock, N., Lozoya, Y.E., Chang, Y.J., Senft, R.A., Lyon, K.A., Alekseyenko, O.V. & Dymecki, S.M. (2020). A single-cell transcriptomic and anatomic atlas of mouse dorsal raphe Pet1 neurons. *eLife* 9:e55523.
- Paquelet, G.E., Carrion, K., Lacefield, C.O., Zhou, P., Hen, R. & Miller, B.R. (2022). Single-cell activity and network properties of dorsal raphe nucleus serotonin neurons during emotionally salient behaviors. *Neuron* 110, 1-16.
- Payeur, A., Guergiev, J., Zenke, F., Richards, B.A. & Naud, R. (2021). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat Neurosci* 24, 1010-1019.
- Pazos, A. & Palacios, J.M. (1985). Quantitative autoradiographic mapping of serotonin receptors in the rat brain. I. Serotonin-1 receptors. *Brain Res* 346, 205-230.
- Perez-Reyes, E. (2003). Molecular physiology of low-voltage-activated T-type calcium channels. *Physiol Rev* 83, 117-161.
- Peroutka, S.J. & Snyder, S.H. (1979). Multiple Serotonin Receptors: Differential Binding of [3H]5- Hydroxytryptamine, [3H]Lysergic Acid Diethylamide and [3H]Spiroperidol. *Mol Pharmacol* 16, 687-699.

Peyron, C., Petit, J.-M., Rampon, C., Jouvét, M. & Luppi, P.-H. (1998). Forebrain afferents to the rat dorsal raphe nucleus demonstrated by retrograde and anterograde tracing methods. *Neuroscience* 82, 443–468.

Pineyro, G. & Blier, P. (1999). Autoregulation of Serotonin Neurons: Role in Antidepressant Drug Action. *Pharmacol Rev* 51, 533-591.

Pinto, D.F.C., Yang, H., Dorocic, I.P., de Jong, J.W., Han, V.J., Peck, J.R., Zhu, Y., Liu, C., Beier, K.T., Smidt, M.P. & Stephan Lammel. (2019) Characterization of transgenic mouse models targeting neuromodulatory systems reveals organizational principles of the dorsal raphe. *Nat Comms* 10:4633.

Pozzorini, C., Mensi, S., Hagens, O., Naud, R., Koch, C. & Gerstner, W. (2015). Automated high-throughput characterization of single neurons by means of simplified spiking neurons. *PLAS Comp Biol* 11, e1004275.

Proulx C.D., Hikosaka, O. & Malinow R. (2014). Reward processing by the lateral habenula in normal and depressive behaviors. *Nat Rev Neurosci* 17, 1146-1152.

Prouty, E.W., Chandler, D.J., Waterhouse, B.D. (2017). Neurochemical differences between target-specific populations of rat dorsal raphe projection neurons. *Brain Res* 1675, 28–40.

Puig, M.V. & Gullledge, A.T. (2011). Serotonin and Prefrontal Cortex Function: Neurons, Networks, and Circuits. *Mol Neurobiol* 44, 449-464.

Puig, M.V., Santana, N., Celada, P., Mengod, G. & Artigas, F. (2004). In vivo excitation of GABA interneurons in the medial prefrontal cortex through 5-HT₃ receptors. *Cereb Cortex* 14, 1365–1375.

Puig, M.V., Watakabe, A., Ushimaru, M., Yamamori, T. & Kawaguchi, Y. (2010). Serotonin modulates fast-spiking interneuron and synchronous activity in the rat prefrontal cortex through 5-HT_{1A} and 5-HT_{2A} receptors. *J Neurosci* 30, 2211–2222.

Ranade, S.P. & Mainen, Z.F. (2009). Transient Firing of Dorsal Raphe Neurons Encodes Diverse and Specific Sensory, Motor, and Reward Events. *J Neurophysiol* 102, 3026–3037.

Ratté, S., Hong, S., De Schutter, E. & Prescott, S.A. (2013). Impact of Neuronal Properties on Network Coding: Roles of Spike Initiation Dynamics and Robust Synchrony Transfer. *Neuron* 78, 758-772.

Ren, J., Friedmann, D., Xiong, J., Liu, C.D., Ferguson, B.R., Weerakkody, T., DeLoach, K.E., Ran, C., Pun, A., Sun, Y., Weissbourd, B., Neve, R.L., Huguenard, J., Horowitz, M.A. & Luo, L. (2018). Anatomically Defined and Functionally Distinct Dorsal Raphe Serotonin Sub-systems. *Cell* 175, 472–487.

Ren, J., Isakova, A., Friedmann, D., Zeng, J., Grutzner, S.M., Pun, A., Zhao, G.Q., Kolluru, S.S., Wang, R., Lin, R., Li, P., Li, A., Raymond, J.L., Luo, Q., Luo, M., Quake, S.R. & Luo, L. (2019). Single-cell transcriptomes and whole-brain projections of serotonin neurons in the mouse dorsal and median raphe nuclei. *eLife* 8, e49424.

- Richards, B., Lillicrap, T.P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R.P. et al. (2019). A deep learning framework for neuroscience. *Nat Neurosci* 22, 1761–1770.
- Ridet, J.-L., Tamir, H. & Privat, A. (1994). Direct Immunocytochemical Localization of 5-Hydroxytryptamine Receptors in the Adult Rat Spinal Cord: A Light and Electron Microscopic Study Using an Anti-Idiotypic Antiserum. *J Neurosci Res* 38, 109-121.
- Rooszendaal, B., Okuda, S., de Quervain, D.J.F. & McGaugh, J.L. (2006). Glucocorticoids interact with emotion-induced noradrenergic activation in influencing different memory functions. *Neuroscience* 138, 901–910.
- Rosenbaum, R., Rubin, J., Doiron, B. (2012). Short Term Synaptic Depression Imposes a Frequency Dependent Filter on Synaptic Information Transfer. *PLOS Comp Bio* 8, e1002557.
- Rossbroich, J., Trotter, D., Beninger, J., Toth, K. & Naud, R. (2021). Linear-nonlinear cascades capture synaptic dynamics. *PLOS Comp Bio* 17, e1008013.
- Rothman, J.S., Cathala, L., Steuber, V. & Silver, R.A. (2009). Synaptic depression enables neuronal gain control. *Nature* 457, 1015-1018.
- Rotman, Z., Deng, P.-Y. & Klyachko, V.A. (2011). Short-Term Plasticity Optimizes Synaptic Information Transmission. *J Neurosci* 31, 14800 –14809.
- Rumelhart, D.E., Hinton, G.E. & Williams, R.J. (1986) Learning representations by back-propagating errors. *Nature* 323, 533-536.
- Rumpel, S., Hatt, H., & Gottmann, K. (1998). Silent synapses in the developing rat visual cortex: evidence for postsynaptic expression of synaptic plasticity. *J Neurosci* 18, 8863–8874.
- Salehi, B., Cordero, M.I. & Sandi, C. (2010). Learning under stress: the inverted-U-shape function revisited. *Learn Memory* 17, 522-530.
- Schreiber, S., Fellous, J.-M., Tiesinga, P. & Sejnowski, T.J. (2004). Influence of Ionic Conductances on Spike Timing Reliability of Cortical Neurons for Suprathreshold Rhythmic Inputs. *J Neurophysiol* 91, 194-205.
- Schultz, W., Dayan, P. & Montague, P.R. (1997). A Neural Substrate of Prediction and Reward. *Science* 275, 1593-1599.
- Schwartzkroin, P.A. (1975). Characteristics of CA1 neurons recorded intracellularly in the hippocampal in vitro slice preparation. *Brain Res* 85, 423-436.
- Schwartzkroin, P.A. (1981). To slice or not to slice. In *Electrophysiology of Isolated Mammalian CNS Preparations*. (New York: Academic), pp. 15-49.
- Sengupta, A., Bocchio, M., Bannerman, D.M., Sharp, T. & Capogna, M. (2017). Control of Amygdala Circuits by 5-HT Neurons via 5-HT and Glutamate Cotransmission. *J Neurosci* 37, 1785–1796.

Seo, C., Guru, A., Jin, M., Ito, B., Sleezer, B.J., Ho, Y.-Y., Wang, E., Boada, C., Krupa, N.A., Kullakanda, D.S., Shen, X.C., Warden, M.R. (2019). Intense threat switches dorsal raphe 5-HT neurons to a paradoxical operation mode. *Science* 363, 538-542.

Shenoy, K.V., Sahani, M. & Churchland, M.M. (2013). Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Ann Rev Neurosci* 36, 337-359.

Sigworth, F.J. (1980). The variance of sodium current fluctuations at the node of Ranvier. *J Physiol* 307, 97-129.

Soares, C., Lee, K.F.H., & Béïque, J.-C. (2017). Metaplasticity at CA1 synapses by homeostatic control of presynaptic release dynamics. *Cell Rep.* 21, 1293–1303.

Soares, C., Lee, K.F.H., Cook, D. and Béïque, J.-C. (2014). A cost-effective method for preparing, maintaining, and transfecting neurons in organotypic slices. In Martina M., Taverna S. (eds) *Patch-Clamp Methods and Protocols. Methods in Molecular Biology (Methods and Protocols), vol 1183.* (New York: Humana Press),

Soares, C., Trotter, D., Longtin, A., Béïque, J.-C. & Naud, R. (2019). Parsing Out the Variability of Transmission at Central Synapses Using Optical Quantal Analysis. *Front Synaptic Neurosci* 11, 10.3389/fnsyn.2019.00022.

Soletto, C., Cholley, B., Mestikawy, S.E., Gozlan, H. & Hamon, M. (1990). Direct Immunohistochemical Evidence of the Existence of 5-HT_{1A} Autoreceptors on Serotonergic Neurons in the Midbrain Raphe Nuclei. *Eur J Neurosci* 2, 1144-1154.

Soubrie, P. (1986). Reconciling the role of central serotonin neurons in human and animal behavior. *Behav Brain Sci* 9, 319–364.

Sprouse, J.S. & Aghajanian, G.K. (1987). Electrophysiological Responses of Serotonergic Dorsal Raphe Neurons to 5-HT_{1A} and 5-HT_{1B} agonists. *Synapse* 1, 3-9.

Stephens, D.W. & Krebs, J.R. (1986). *Foraging Theory.* Princeton: Princeton University Press.

Stimberg, M., Brette, R. & Goodman, D.F.M. (2019). Brian 2, an intuitive and efficient neural simulator. *eLife* 8, e47314.

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M. & Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364, eaav7893.

Sunnåker, M., Busetto, A.G., Numminen, E., Corander, J., Foll, M., and Dessimoz, C. (2013). Approximate Bayesian computation. *PLoS Comput Biol* 9, e1002803–e1002810.

Suvrathan, A., Bennur, S., Ghosh, S., Tomar, A., Anilkumar, S. & Chattarji, S. (2013). Stress enhances fear by forming new synapses with greater capacity for long-term potentiation in the amygdala. *Phil Trans R Soc B* 369, 20130151.

Sylwestrak, E.L., Jo, Y., Vesuna, S., Wang, X., Holcomb, B., Tien, R.H., Kim, D.Y., Fenno, L., Ramakrishnan, C., Allen, W.E., Chen, R., Shenoy, K.V., Sussillo, D. & Deisseroth, K. (2022). Cell-type-specific population dynamics of diverse reward computations. *Cell* 185, 3568–3587.

- Takahashi, A., Lee, R.X., Iwasato, T., Itohara, S., Arima, H., Bettler, B., Miczek, K.A. & Koide, T. (2015). Glutamate input in the dorsal raphe nucleus as a determinant of escalated aggression in male mice. *J Neurosci* 35, 6452-6436.
- Tazerart, S., Blanchard, M.G, Miranda-Rottman, S., Mitchell, D.E, Pina, B.N., Thomas, C.I., Kamasawi, N. & Araya, R. (2022). Selective activation of BK channels in small-headed dendritic spines suppresses excitatory postsynaptic potentials. *J Physiol* 600, 2165-2187.
- Tejero-Cantero, A., Boelts, J., Deistler, M., Lueckmann, J. M., Durkan, C., Gonçalves, P. J., ... & Macke, J. H. (2020). sbi--a toolkit for simulation-based inference. *arXiv preprint arXiv, 2007.09114*.
- Templin, J.S., Bang, S.J., Soiza-Reilly, M., Berde, C.B. & Commons, K.G. (2012). Patterned expression of ion channel genes in mouse dorsal raphe nucleus determined with the Allen Mouse Brain Atlas. *Brain Res* 1457, 1-12.
- Ting, J.T., Lee, B.R., Chong, P., Soler-Llavina, G., Cobbs, C., Koch, C., Zeng, H. & Lein, E. (2018). Preparation of Acute Brain Slices Using an Optimized N-Methyl-D-glucamine Protective Recovery Method. *J. Vis. Exp.* (132), e53825, doi:10.3791/53825.
- Trinh, A.-T., Clarke, S.E., Harvey-Girard, E. & Maler, L. (2019). Cellular and network mechanisms may generate sparse coding of sequential object encounters in hippocampal-like circuits. *eNeuro* 6, ENEURO.0108-19.2019.
- Trinh, A.-T., Harvey-Girard, E., Teixeira, F. & Maler, L. (2016). Cryptic laminar and columnar organization in the dorsolateral pallium of a weakly electric fish. *J Comp Neurol* 524, 408-428.
- Urbain, N., Creamer, K. & Debonnel, Guy. (2006). Electrophysiological diversity of the dorsal raphe cells across the sleep-wake cycle of the rat. *J Physiol* 573, 679-695.
- van der Walt, S., Colbert, S.C. & Varoquaux, G. (2011). The NumPy array: A structure for efficient numerical computation. *Comput Sci Eng* 13, 22-30.
- van Rossum, M.C.W., O'Brien, B.J. & Smith, R.G. (2003). Effects of noise on the spike timing precision of retinal ganglion cells. *J Neurophysiol* 89, 2406-2419.
- Vandermaelen, C.P. & Aghajanian, G.K. (1983). Electrophysiological and Pharmacological Characterization of Serotonergic Dorsal Raphe Neurons Recorded Extracellularly and Intracellularly in Rat Brain Slices. *Brain Res* 289, 109-119.
- Vardalaki, D., Chung, K. & Harnett, M.T. (2022). Filopodia are a structural substrate for silent synapses in adult neocortex. *Nature* 612, 323-327.
- Vereczki, V.K., Veres, J.M., Müller, K., Nagy, G.A., Rácz, B., Barsy, B. & Hájos, N. (2016) Synaptic organization of perisomatic GABAergic inputs onto the principal cells of the mouse basolateral amygdala. *Front Neuroanat* 10, 20.
- Vyas, S., Golub, M.D., Sussillo, D. & Shenoy, K.V. (2020). Computation through neural population dynamics. *Annu Rev Neurosci* 43, 249-275.

- Wan, J., Peng, W., Li, X., Qian, T., Song, K., Zeng, J., Deng, F., Hao, S., Feng, J., Zhang, P., Zhang, Y., Zou, J., Pan, S., Shin, M., Venton, B.J., Zhu, J.J., Jing, M., Xu, M. & Li, Y. (2021). A genetically encoded sensor for measuring serotonin dynamics. *Nat Neurosci* 24, 746–752.
- Wang, H.-L., Zhang, S., Qi, J., Wang, H., Cachope, R., Mejjas-Aponte, C.A. *et al.* (2019). Dorsal Raphe Dual Serotonin-Glutamate Neurons Drive Reward by Establishing Excitatory Synapses on VTA Mesoaccumbens Dopamine Neurons. *Cell Rep* 26, 1128–1142.
- Warden, M.R., Selimbeyoglu, A., Mirzabekov, J.J., Lo, M., Thompson, K.R., Kim, S.-Y., Adhikari, A., Tye, K.M., Frank, L.M. & Deisseroth, K. (2012). A prefrontal cortex–brainstem neuronal projection that controls response to behavioural challenge. *Nature* 492, 428–432.
- Weissbourd, B., Ren, J., DeLoach, K.E., Guenther, C.J., Miyamichi, K. & Luo, L. (2014). Presynaptic Partners of Dorsal Raphe Serotonergic and GABAergic Neurons. *Neuron* 83, 645–662.
- Weissmann-Nanopoulos, D., Mach, E., Magre, J., Demasse, Y. & Pujol, J. F. (1985). Evidence for the localization of 5HT_{1A} binding sites on 5-HT containing neurons in the raphe dorsalis and raphe centralis nuclei of the rat brain. *Neurochem Int* 7, 1061–1072.
- Wert-Carvajal, C., Reneaux, M., Tchumatchenko, T. & Clopath, C. (2022). Dopamine and serotonin interplay for valence-based spatial learning. *Cell Rep* 39, 110645.
- Whitaker, L.R., Carneiro de Oliveira, P.E., McPherson, K.B., Fallon, R.V., Planeta, C.S., Bonci, A. & Hope, B.T. (2016). Associative Learning Drives the Formation of Silent Synapses in Neuronal Ensembles of the Nucleus Accumbens. *Biol Psych* 80, 246–256.
- Womelsdorf, T., Valiente, T.A., Sahin, N.T., Miller, K.J. & Tiesinga, P. (2014). Dynamic circuit motifs underlying rhythmic gain control, gating and integration. *Nat Neurosci* 17, 1031–1039.
- Wood, S.N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature* 466, 1102–1104.
- Wright, W.J. *et al.* (2020). Silent synapses dictate cocaine memory destabilization and reconsolidation. *Nat Neurosci* 23, 32–46.
- Xu, S., Das, G., Hueske, E. & Tonegawa, S. (2017). Dorsal Raphe Serotonergic Neurons Control Intertemporal Choice under Trade-off. *Curr Biol* 27, 1–9.
- Yang, Y., Cui, Y., Sang, K., Dong, Y., Ni, Z., Ma, S. & Hu, H. (2018). Ketamine blocks bursting in the lateral habenula to rapidly relieve depression. *Nature* 554, 317–322.
- Zhang, Y.-P. & Oertner, T.G. (2007). Optical induction of synaptic plasticity using a light-sensitive channel. *Nat Methods* 4, 139–141.
- Zhong, W., Li, Y., Feng, Q. & Luo, M. (2017). Learning and Stress Shape the Reward Response Patterns of Serotonin Neurons. *J Neurosci* 37, 8863–8875.
- Zhou, L., Liu, M.-Z., Li, Q., Deng, J., Mu, D. & Sun, Y.G. (2017). Organization of Functional Long-Range Circuits Controlling the Activity of Serotonergic Neurons in the Dorsal Raphe Nucleus. *Cell Rep* 18, 3018–3032.

Zito, K., Scheuss, V., Knott, G., Hill, T. & Svoboda, K. (2009). Rapid functional maturation of nascent dendritic spines. *Neuron* 61, 247–258.

Zucker, R.S., Regehr, W.G. (2002). Short-term synaptic plasticity. *Annu Rev Physiol* 64, 355–405.

Appendices

Appendix A: Manuscript VII: Reconciling current theories of consciousness**Reconciling current theories of consciousness**

Sébastien Maillé* & Michael B. Lynn*

*: **Contributed equally to the work. Co-corresponding authors.**

Review of Noel et al. (2019).

This manuscript was published in the *Journal of Neuroscience* in 2020:

Maillé S & Lynn MB. Reconciling current theories of consciousness. *J Neurosci*. 2020 March 4;40(10):1994-1996. doi: 10.1523/JNEUROSCI.2740-19.2020.

Statement of contribution

This was a collaborative effort between myself and Sébastien Maillé, another graduate student in the Béïque Laboratory. We jointly led a review of the literature, guided the project, and wrote and edited all drafts.

We are grateful to Dr. Leonard Maler and Dr. Vincent Bergeron for useful discussions and comments, and to Dr. Jean-Claude Béïque for providing an excellent training environment that allowed us the time and support to write such a review.

Understanding the neural basis of consciousness is one of the fundamental challenges in modern neuroscience. A number of sophisticated models and theories have attempted to formalize how the brain implements consciousness using insights from philosophy, psychology, computer science and neuroscience. These include two major and perhaps competing theories, the integrated information theory (IIT) and the global neuronal workspace (GNW) theory, which differ mainly in their level of conceptual abstraction and anatomical specificity.

The IIT, first proposed by Giulio Tononi (2004), focuses on defining what a conscious system should look like with respect to information processing and architecture without considering particular brain areas or temporal structures. One prediction of IIT is that neural networks supporting consciousness must be highly interconnected, effectively integrating different components of a state into a unified experience. A crucial advantage of the IIT is that it provides a mathematical metric of irreducibility (or integration), Φ , that can be related to the level of consciousness. Proponents of IIT point to its explanatory power: for instance, it can explain why the cortex is capable of producing conscious experience while the cerebellum is not (Lemon & Edgley, 2010; Yu et al., 2015), even though the cerebellum possesses up to 4 times more neurons. While the IIT has not received unambiguous validation, (possibly due to the abstract nature of its description of consciousness; see Tononi et al. (2016) for a comprehensive review), it provides one of the most detailed accounts for the emergence of conscious experience from an information processing network.

The GNW theory (Dehaene & Changeux, 2011), in contrast to the IIT, was empirically derived from EEG and imaging studies in humans and primates. These studies have shown that when a stimulus is presented but not consciously perceived, activation can be seen mainly in the associated primary sensory cortices. When the stimulus is consciously perceived, however, activation in primary cortical areas is followed by a delayed “neural ignition” in which a sustained wave of activity propagates across prefrontal and parietal association cortices. According to the GNW model, this allows relevant information to be broadcast across the brain to other subsystems for use in decision-making, reporting, memory consolidation and other

processes. Thus, while IIT focuses on abstract connectivity and information processing structure, GNW proposes a concrete spatiotemporal locus for conscious processes.

Unfortunately, while both IIT and GNW have obtained experimental support, testable predictions from both theories are seldom compared within the same dataset. In a recent issue of *The Journal of Neuroscience*, Noel et al. (2019) leveraged a previously published experimental dataset to directly compare IIT and GNW at the single-unit level. In the experiments published by Ishizawa et al. (2016) (and later re-analyzed by Noel et al.), monkeys were subjected to non-aversive stimuli while extracellular microelectrode arrays recorded single-unit activity in S1 (primary somatosensory) and vPM (ventral premotor; involved in multisensory integration) cortices. The stimuli consisted of an auditory stimulus, a tactile stimulus, or concurrent auditory and tactile stimuli. Crucially, partway through the task monkeys were anesthetized with propofol, permitting a sophisticated comparison of single-neuron activity across states of consciousness.

The authors exploited the multisensory nature of the stimuli to categorize neurons based on information processing rules, which they could relate to key mathematical predictions from IIT. Neurons were classified as integrative (AND gate; exhibiting a multisensory response greater than the largest unisensory response), or convergent (XOR gate; exhibiting a multisensory response smaller than or equal to the largest unisensory response). According to IIT, an AND gate possesses a value of Φ which is more than three-fold higher than that of an XOR gate ($\Phi=0.78$ vs $\Phi=0.25$, respectively). Noel et al. did not find large enough pools of integrative and convergent neurons in vPM to generate sufficient statistical power, so this analysis was restricted to S1 neurons. The authors reasoned that if integrative neurons underlie conscious perception, then their multisensory representations should closely track the state of consciousness, while the representations of convergent neurons should not track conscious state. Contrary to this, 69% of convergent neurons but only 37.1% of integrative neurons changed their multisensory response classification after propofol administration. The authors additionally considered single-neuron physiological properties including Lempel-Ziv complexity

(a measure of the statistical complexity of stimulus-driven responses) and noise correlations (the amount of shared response variability between neurons). They found that both of these measures were less correlated with consciousness state in integrative neurons than in convergent neurons. Together, these findings represent strong evidence against the IIT theory of consciousness.

Noel et al. (2019) also considered a prediction of GNW: that consciously perceived stimuli will generate a wave of activation which simultaneously spreads across multiple brain areas (neural ignition). Supporting the GNW, the authors found more coactivation of S1 and vPM during conscious awareness (beyond what would be expected from increased activation of each individual area). Based on these results, Noel et al. concluded that the data support the GNW model over IIT at the level of single neurons.

Should GNW and IIT be viewed as strictly non-overlapping hypotheses about how neural circuits implement consciousness? Placing the findings by Noel et al. in a broader context, it is possible that the two theories are not inherently incompatible, and that GNW and IIT could provide a powerful, overlapping explanation of conscious experience at the level of anatomy and connectivity structure, respectively. Indeed, the current findings, taken at face value, simply rule out an anatomical locus (S1) and spatial scale (single neurons) at which IIT might be expressed. Proponents of IIT could point out that “integrative” neurons in this study might represent local, but not absolute maxima of integrative power; the real locus of consciousness, then, would be located in an area or at a spatiotemporal grain not recorded here. Thus, testing the IIT model and identifying the substrate of consciousness proposed by Tononi will likely require advanced techniques, including simultaneous recording of thousands of neurons in behaving animals (Marshall et al., 2019; Jun et al., 2017). In the future, these approaches could potentially help reconcile the GNW and IIT by defining distinct levels of abstraction for each.

One caveat of this reconciliatory perspective is that it must contend with evidence showing primary sensory cortices can, indeed, reflect conscious processes. Kulics (1982)

recorded from monkey S1 in a go/no-go task, and found that a late component (~50ms) of the evoked potential correlated with perception. Similar perception-related delayed responses have been observed in human primary auditory cortex (Wiegand & Gutschalk, 2012) and visual cortex (Maier et al., 2008). Intriguingly, anatomical data suggested that this later component likely resulted from synaptic input to cortical layers I/II, suggesting it was generated by cortico-cortical loops rather than bottom-up sensory stimulation (Cauller & Kulics, 1991). These loops have been suggested to be a general feature of conscious perception (Meyer et al., 2011). If S1 indeed supports consciousness, and Noel et al. failed to find support for IIT in this area, one might argue that this invalidates IIT itself. However, the findings of Noel et al. are also compatible with a perspective in which primary sensory cortices simply track conscious experiences, due to delayed feedback projections from cortico-cortical loops, but cannot actually generate consciousness. This would explain why integrative single neurons are not as affected by anaesthesia.

The work by Noel et al. also contributes to an important debate about the spatial scale over which conscious processes occur. Earlier work reported that single-unit activity in monkey somatosensory cortex did not covary with perceptual reports, suggesting that consciousness may only be present in network readouts (de Lafuente & Romo, 2005). Further evidence came from work in monkey V1 (Maier et al, 2008) showing that spiking behavior of neurons did not track perception of visual stimuli, yet blood oxygen level-dependent (BOLD) responses and low-frequency local field potentials did. However, recent work in medial frontal cortex of humans has highlighted that the responses of individual cells can, in fact, be correlated with conscious reporting of stimuli (Reber et al, 2017; Gelbard-Sagiv et al., 2018). Noel et al. (2019) helps reconcile these findings by demonstrating that single-neuron responses are heterogeneous and dependent on information processing type: while some neurons (convergent) mostly track consciousness state, other neuron types (integrative) are seemingly not altered to the same extent. However, this also demonstrates the difficulty of inferring network properties from mixed single-unit activity, implying that perhaps a direct comparison of

integrative and convergent neurons is not the ideal spatial scale for testing the IIT framework. Further work is needed to define the input-output architecture of such consciousness-tracking neurons, as well as to link the heterogeneous activity of single neurons with more abstract metrics of brain activity shown to be reflective of conscious processes. Relevant metrics might include criticality, which quantifies the extent to which a system is near the inflection point between stability and disorder (Alonso et al., 2019), and dynamic signal coordination, which quantifies the patterns of phase correlations between brain areas (Demertzi et al., 2019).

Consciousness remains one of the great unsolved mysteries in systems neuroscience. The work of Noel et al. provides an incisive direct comparison of two major theories of consciousness within the same single-unit dataset, ruling out a spatial scale and anatomical locus over which IIT formalism could potentially explain consciousness. This exciting field will, we suspect, continue to work towards a more detailed understanding of how the brain's connectivity structure interacts with spatiotemporal activity patterns to generate human consciousness.

References

Alonso LM, Solovey G, Yanagawa T, Proekt A, Cecchi GA, Magnasco MO (2019) Single-trial classification of awareness state during anesthesia by measuring critical dynamics of global brain activity. *Sci Rep* 9:4927.

Cauler LJ, Kulics AT (1991) The neural basis of the behaviorally relevant N1 component of the somatosensory-evoked potential in S1 cortex of awake monkeys: evidence that backward cortical projections signal conscious touch sensation. *Exp Brain Res* 84:607-619.

Dehaene S, Changeux J-P (2011) Experimental and theoretical approaches to conscious processing. *Neuron* 70, 200–227

de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. *Nature Neuroscience* 12:1698-1703

Demertzi A, Tagliazucchi E, Dehaene S, Deco G, Barttfeld P, Raimondo F, Martial C, Fernández-Espejo D, Rohaut B, Voss HU, Schiff ND, Owen AM, Laureys S, Naccache L, Sitt JD (2019) Human consciousness is supported by dynamic complex patterns of brain signal coordination. *Sci Adv* 5:eaat7603

Gelbard-Sagiv H, Mudrik L, Hill MR, Koch C, Fried I (2018) Human single neuron activity precedes emergence of conscious perception. *Nat Comm* 9:2057.

Ishizawa Y, Ahmed OJ, Patel SR, Gale JT, Sierra-Mercado D, Brown EN, Eskandar EN (2016) Dynamics of Propofol-Induced Loss of Consciousness Across Primate Neocortex. *J Neurosci* 36:7718–7726.

Jun JJ et al. (2017) Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551:232–236.

Kulics AT (1982) Cortical neural evoked correlates of somatosensory stimulus detection in the rhesus monkey. *Electroencephalog Clin Neurophysiol* 53:78-93.

Lemon RN, Edgley SA (2010) Life without a cerebellum. *Brain* 133:652–654.

Maier A, Wilke M, Aura C, Zhu C, Ye F, Leopold D (2008) Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nat Neurosci* 11:1193-1200.

Marshall JH, Kim YS, Machado TA, Quirin S, Benson B, Kadmon J, Raja C, Chibukhchyan A, Ramakrishnan C, Inoue M, Shane JC, McKnight DJ, Yoshizawa S, Kato HE, Ganguli S, Deisseroth K (2019) Cortical layer-specific critical dynamics triggering perception. *Science* 365:558.

Meyer K (2011) Primary sensory cortices, top-down projections and conscious experience. *Progress in Neurobiology* 94:408-417.

Noel J-P, Ishizawa Y, Patel SR, Eskandar EN, Wallace MT (2019) Leveraging Nonhuman Primate Multisensory Neurons and Circuits in Assessing Consciousness Theory. *J Neurosci* 39:7485–7500.

Reber TP, Faber J, Niediek J, Bostrom J, Elger CE, Mormann F (2017) Single-neuron correlates of conscious perception in the human medial temporal lobe. *Curr Biol* 27:2991-2998.

Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD (2019) Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364:eaav7893

Tononi G (2004) An information integration theory of consciousness. *BMC Neurosci* 5:42.

Tononi G, Boly M, Massimini M, Koch C (2016) Integrated information theory: from consciousness to its physical substrate. *Nat Rev Neurosci* 17:450–461.

Wiegand K, Gutschalk A (2012) Correlates of perceptual awareness in human primary auditory cortex revealed by an informational masking experiment. *NeuroImage* 61:62-69.

Yu F, Jiang QJ, Sun XY, Zhang RW (2015) A new case of complete primary cerebellar agenesis: clinical and imaging findings in a living patient. *Brain* 138:e353.

Appendix B:

Python Package I: SynapPy: Automated synaptic waveform analysis in Python

This Python package was created near the start of my time as a graduate student in the Beique Lab. The objective was to offer fast, automated and repeatable analysis of synaptic waveforms in a Pythonic (object-oriented) framework.

Much of synaptic physiology is still effectively analyzed manually, with cursors designating areas of baseline or peak conductances, using software such as ClampFit. This is not only time-consuming - it also makes it difficult for researchers to reproducibly analyze data, and encourages a kind of homogeneity in the metrics used to analyze data. For example, in voltage clamp, integrating currents over time can provide accurate readouts of conductances, which are useful in some situations of high stimulation frequency where post-synaptic potentials overlap, and thus peak current amplitude metrics may be difficult to interpret due to summation. (This exact situation is encountered in Manuscript III, Figure 3, when electrically stimulating slow 5-HT_{1A}Rs that couple to GIRK conductances). The following package aims to fill the gap by extracting statistics of stimulus-evoked synaptic events across neurons, in a flexible manner.

The full package is available at <https://github.com/miclynn/synappy>. Below, I have provided the github documentation associated with this package.

Statement of contribution

I wrote all code and performed all formal analysis.

Table of Contents

- [Introduction](#)
- [Installation](#)
- [Getting started](#)
 - [Previewing data](#)
 - [Detecting and quantifying events](#)
 - [Visualizing event statistics](#)
 - [Retrieving values](#)
- [Event statistics](#)
 - [Amplitude, baseline and latency](#)
 - [Normalized amplitude](#)
 - [Decay](#)
 - [Integral](#)

Introduction

SynapPy is a Pythonic data visualization and analysis tool for patch-clamp synaptic physiologists who work with synaptic events stored in Axon Binary Format files (.abf).

First, SynapPy provides preview and visualization capabilities for .abf files, including sweep highlighting and averaging. Second, SynapPy provides a set of Pythonic tools for automated detection and quantification of post-synaptic events (either electrically or optically evoked)

based on a stimulus trigger channel. SynapPy can auto-magically quantify amplitude, baseline, integral, decay kinetics and rise-time, latency, and release probability. Moreover, these event attributes can be easily overlaid with the raw signal for visual inspection afterwards. Finally, SynapPy includes basic support for detecting spontaneous (non-stimulus-triggered) events and, of course, quantifying their statistics.

Installation

The main dependencies are: python3, numpy, scipy, matplotlib and neo (ver 0.4+ recommended). Neo is used to parse raw data from the .abf files.

A basic installation using the included setup.py file:

```
git clone https://github.com/miclynn/synappy/  
cd synappy  
python3 setup.py install  
(Make sure you substitute your binary of python here!)
```


Getting started

Previewing data

Let's start by previewing a file:

```
import synappy  
  
p = syn.preview('ex_file.abf')
```

This plots all channels and sweeps from the file. Sweeps can be advanced by the indicated buttons, and an average across all sweeps can be taken.



Detecting and quantifying events

To work with synaptic events, we first need to import some files.

```
import synappy
```

```
files = ['ex_file_1.abf', 'ex_file_2.abf']
d = syn.load(files, input_channel=0, stim_channel=2)
```

Here, files can be thought of as a dataset (for example, a related group of files with the same parameters for stimulation or drug infusion.) They're all quantified together. Also note that if needed, we can specify the signal channel number and stimulus trigger channel number. (If no arguments are provided, these default to first and last channels, respectively.)

`syn.load()` stores the signal, time and stimulus information in the following class attributes:

- `d.sig[neuron][trial, t_ind]`
 - The raw signal for a given neuron, trial and time index.
- `d.sig_stim[neuron][trial, t_ind]`
 - The stimulus trigger signal for a given neuron, trial and time index.
- `d.t[neuron]`
 - Time, in seconds, for a given neuron

Next, we add stimulus-triggered events.

```
d.add_events(stim_thresh=5)
```

This searches the stimulus channel provided for any pulses, and adds these locations as event onsets.

Note that the stimulus threshold can be specified (`stim_thresh`), and may need to be altered if the stimulus channel is incorrectly scaled.

Most event statistics can then be added using specific class methods, described in detail below. Here, we first add the simplest event statistic, amplitude:

```
d.add_ampli(event_sign='pos')
```

This quantifies amplitude from baseline, in the direction specified by `event_sign` ('pos' or 'neg' to deal with excitatory or inhibitory events), after the stimulus trigger.

A convenient library of general event statistics, including event decay and integral, can be added with the following command:

```
d.add_all(kwarg_add_ampli={'event_sign': 'neg'})
```

Note that a dictionary of keywords for each method can be provided, which are passed to the appropriate method (`kwarg_add_ampli`, `kwarg_add_decay`, `kwarg_add_integral`).

Visualizing event statistics

Most event statistics can be overlaid with the recorded signal in a single plot for convenient quality control.

```
d.preview(neur=0, attrs=['ampli', 'baseline'])
```

Here, we've specified a neuron to preview, as well as a list of attributes (event statistics) to annotate.

Retrieving values

The values for each statistic (eg amplitude, etc.) is stored as attributes within the class instance. Taking `.ampli` as an example:

- `d.ampli.data[neuron][trial, event]`
 - Baseline-subtracted maximum amplitude data (in pA or mV) for a given neuron, trial and event index.
- `d.ampli.inds[neuron][trial, event]`
 - Indices in `.sig` of maximum amplitudes.
- `d.ampli.params`
 - Parameters related to the `kwarg`s specified for the associated class method. (For example, `d.ampli.params.t_event_upper` specifies the max post-stimulus time to search for the maximum event amplitude.

Event statistics

Here, we detail all the class methods available for measuring and quantifying synaptic events.

All class methods are fully documented (`help(d.example_method)`).

Amplitude, baseline and latency

- `.add_ampli()`.
- Computes pre-event baseline values, event amplitudes, and event latencies (computed in a number of ways).
- This requires a `self.events` attribute, created by calling the method `self.add_events()`. For each stimulus, a baseline signal is calculated between `t_baseline_lower` and `t_baseline_upper` before the stimulus onset. Next, a maximum event amplitude is calculated. Finally, the event latency (time to the peak amplitude, or alternately other latency metrics) is computed.
- These values are stored as the following attributes in the `EphysObject` instance: `.ampli` `.baseline` `.latency`

Parameters

- `event_sign` : str The sign of the events. Can either be 'pos', reflecting EPSPs/IPSCs in IC/VC, or 'neg', reflecting IPSPs/EPSCs in IC/VC.
- `t_baseline_lower` : float The time before stimuli, in ms, from which to start computing a pre-event baseline.
- `t_baseline_upper` : float The time before stimuli, in ms, from which to stop computing a pre-event baseline.
- `t_event_lower` : float The time after stimuli, in ms, from which to start searching for events.
- `t_event_upper` : float The time after stimuli, in ms, from which to stop searching for events.
- `t_savgol_filt` : int Width of savgol filter applied to data, in ms, before computing maximum amplitude.
- `latency_method` : str The method used to calculate latency from stimulus to event. Possible methods are: - 'max_ampli': Time from stimulus to maximum amplitude. -

'max_slope': Time from stimulus to maximum first deriv. - 'baseline_plus_4sd': Time from stimulus to time where signal exceeds baseline levels + 4 standard deviations. - '80_20_line': Computes the times where the signal reaches 20% and 80% of the maximum amplitude, then draws a straight line between these and determines where this line first intersects the signal. The time from the stimulus to this point gives the latency. (Provides a best guess of when the signal starts to rise.)

Attributes added to class instance

- `.ampli`
 - `.ampli.data[neuron][trial, event]`: Baseline-subtracted maximum amplitude data (in pA or mV).
 - `.ampli.inds[neuron][trial, event]`: Indices in `.sig` of maximum amplitudes.
 - `.ampli.params`: SimpleNamespace storing key params from `.add_ampli()` method related to amplitude.
- `.baseline`
 - `.baseline.mean[neuron][trial, event]`: Mean baseline values (in pA or mV).
 - `.baseline.std[neuron][trial, event]`: Standard deviation of baseline values (in pA or mV).
 - `.baseline.inds_start[neuron][trial, event]`: Indices of the start of baseline period in `.sig`
 - `.baseline.inds_stop[neuron][trial, event]`: Indices of the end of baseline period in `.sig`
 - `.baseline.params`: SimpleNamespace storing key params from `.add_ampli()` method related to baseline.
- `.latency`
 - `.latency.data[neuron][trial, event]`: Event latency from stimulus onset (sec).
 - `.latency.inds[neuron][trial, event]`: Indices in `.sig` of event latency.
 - `.latency.params`: SimpleNamespace storing key params from `.add_ampli()` method related to latency.

Normalized amplitude

- `.add_ampli_norm()`.
- Adds normalized amplitude measurement to the class instance as `.ampli_norm`. Amplitudes are normalized to the mean `ampli` for each stimulus delivered to each neuron.
(`ampli` must be an existing attribute, through the `.add_ampli()` method.)

Attributes added to class instance

- `self.ampli`
 - `.ampli_norm.data[neuron][trial, event]` Baseline-subtracted normalized max amplitude data (in pA or mV).
 - `.ampli_norm.inds[neuron][trial, event]` Indices in `.sig` of normalized max amplitudes.

Decay

- `.add_decay()`.
- Fits each post-synaptic event with an exponential decay function and stores the fitted parameters in `self.decay`.
- Decay equation variables correspond to the fitted variables for the equation used (see the kwarg `fn` for more info).
- monoexponential decay: `lambda1`, `b`.
- biexponential decay: `lambda1`, `lambda2`, `vstart2`, `b`.

Parameters

- `t_prestim`: Time before stimulus, in ms, to include in signal used to compute decay.
- `t_poststim`: Time after stimulus, in ms, to include in signal used to compute decay.
- `plotting`: Whether to plot examples of decay fits (True) or not (False).
- `fn`: Exponential decay function to use.
 - `'monoexp'`: $y = e^{-(t * \text{lambda1})} + b$
 - `'biexp_normalized_plusb'`: $y = e^{-(t * \text{lambda1})} + vstart * e^{-(t / \text{lambda2})} + b$
- (In all cases, decay tau can be computed as $\tau = 1/\text{lambda}$).

Attributes added to class instance

- `.decay` : SimpleNamespace
 - `.decay.vars[neuron][trial, stim, decay_var]` Fitted variables for monoexponential decay.
 - If `fn='monoexp'`, then `decay_var=0` : `lambda1` `decay_var=1` : `b`
 - If `fn='biexp_normalized_plusb'`, then `decay_var=0` : `lambda1` `decay_var=1` : `lambda2` `decay_var=2` : `vstart2` `decay_var=3` : `b`
 - `.decay.covari[neuron][trial, stim, decay_param]` Covariance matrices for fitted variables, as documented in `.decay.vars`
 - `.decay.params` Parameters related to the decay fitting.

Integral

- `.add_integral()`.

- Computes the integral for each post-synaptic event.

Parameters

- `t_integral`: The total post-stimulus time to integrate, in milliseconds.
- `cdf_bins`: Number of bins for the cumulative integral

Attributes added

- `.integral`
 - `.integral.data[neuron][trial, event]`: Integral values for each event (in pAsec or mVsec).
 - `.integral.inds_start[neuron][trial, event]`: Indices of the start of integral period in `.sig`
 - `.integral.inds_stop[neuron][trial, event]`: Indices of the end of integral period in `.sig`
 - `.integral.cdf[neuron][trial, event]`: Cumulative distribution of the integral over time for each event.
 - `.integral.params`: SimpleNamespace storing key params from `.add_integral()` method related to integral.

Appendix C:

Python Package II: MouseBerry: Raspberry Pi-based interface for controlling behavioral experiments

This Python package was created when I started running mouse behavior experiments in the Beique Lab. The objective was to run behavioral experiments using a low-cost, flexible Raspberry Pi-based framework. This Python package provides an object-oriented, composable framework for designing behavior experiments from building blocks of events, trial-types and task structures, provides automated video monitoring and measurement logging. It greatly simplifies the mental overhead of designing and running mouse behavior experiments, and also makes it possible to design tasks for mice that are complex, probabilistic and employ multiple stimuli.

The full package is available at <https://github.com/micclynn/mouseberry>. Below, I have provided the github documentation associated with this package.

Statement of contribution

I wrote all code and performed all analysis. I am indebted to Dr. Kirk Mulatz, the lab technician in Beique lab, for electrical expertise with connecting and configuring the Raspberry Pi microcomputers.

Table of Contents

- [Introduction](#)
- [Installation](#)
- [Getting started](#)
 - [The basics](#)
 - [Built-in event types](#)
 - [Rewards](#)
 - [GPIO stimuli](#)
 - [Audio](#)
 - [Looming visual stimulus](#)
 - [Built-in measurement types](#)
 - [Continuous polling from GPIO:](#)
 - [Notes on acquiring measurements](#)
 - [Video streaming](#)
- [Advanced usage](#)
 - [Defining stochastic event start times](#)
 - [Defining stochastic ITIs](#)
 - [Constructing more complex experiments](#)
- [Stored data format: HDF5](#)
 - [Experiment attributes](#)
 - [Trial attributes](#)
 - [Measurements](#)
 - [Events](#)
- [Creating custom classes](#)
 - [Events](#)
 - [Measurements](#)

Introduction

Mouseberry is a hackable, Python-based rodent behavior controller designed around the Raspberry Pi and operating at a moderate-to-high level of abstraction.

It provides a full scheduler and threading for stochastic behavioral events, background (threaded) acquisition of measurements, background PiCam video streaming to a monitor, and automated data storage in the HDF5 format.

With mouseberry, Events are first defined based on a set of steps which occur when they are triggered. Then, TrialTypes are constructed as a collection of events. Finally, the Experiment is built as a collection of TrialTypes (each with some characteristic occurrence probability), background Measurements, and a background Video stream. The core routines of Mouseberry then procedurally construct and run trials based on dynamic scheduling and threading of events that are occurring in those trials. There is full logging and console output, and all relevant parameters of each event, even user-defined parameters, are auto-magically stored in an .hdf5 file.

These successive layers of abstraction make it easier to write and edit complex behavioral tasks. Since the storage of Measurement and Event attributes (eg logged start time and any attributes created on-the-fly) is fully automated and baked into the hdf5 data storage routines, events can be added, subtracted and altered with pure abandon.

Finally, the package is fully compatible with running the Raspberry Pi in headless mode through SSH. Console outputs can be viewed on the controlling computer, and the video stream natively displays on an attached monitor. This conveniently allows one master computer to control many Raspberry Pi's, each running their own behavior task.

Installation

Mouseberry requires a Python3 installation.

Basic installation:

```
git clone https://github.com/micllynn/mouseberry/  
cd mouseberry  
python3 setup.py install  
Additional system dependencies:
```

- A system installation of the UNIX binary `sox` is required to play tones.
 - MacOS: `brew install sox`
 - Linux: `sudo apt-get install sox`

Getting started

The basics

Let's start by importing the package. We will set up a simple experiment consisting of one trialtype with one tone event that occurs. This is the most basic setup for creating and running a behavior experiment.

```
import mouseberry as mb
```

```
tone = mb.Tone(name='tone_low', t_start=1, t_dur=5, freq=2000)
trial = mb.TrialType(name='trial_low', p=1, events=[tone])
```

```
exp = mb.Experiment(n_trials=10, iti=2)
exp.run(trial)
```

Note that for `mb.TrialType`, `p` refers to the probability of a trialtype occurring. It's useful if an experiment consists of multiple trialtypes.

Next, we will set up a more complex experiment with two trialtypes. We will additionally initialize a measurement object which polls from a lickometer, and a video object which acquires from a PiCam and routes the feed to a monitor attached to the RPi.

```
licks = mb.Lickometer(name='lickometer', pin=5, sampling_rate=200)
vid = mb.Video( framerate=30)
```

```
tone_low = mb.Tone(name='tone_low', t_start=1, t_dur=5, freq=2000)
tone_high = mb.Tone(name='tone_high', t_start=1, t_dur=5, freq=10000)
```

```
trial_low = mb.TrialType(name='trial_low', p=0.5, events=[tone_low])
trial_high = mb.TrialType(name='trial_high', p=0.5, events=[tone_high])
```

```
exp = mb.Experiment(n_trials=10, iti=2)
exp.run(trial_low, trial_high, licks, vid)
```

That's the extent of code needed to run this experiment, acquire a measurement in the background, stream a video, and save all relevant trial parameters to an `.hdf5` file! By default, the console prints basic statistics about the current trialtype and event onsets/offsets.

What is going on in the background? `exp.run` takes as input any trials, measurements or videos associated with the experiment. The `Experiment` class then procedurally generates a valid experimental plan and runs it, dynamically scheduling and triggering events as needed. Meanwhile, all measurements, data storage and video streaming are organized in the background in an automated manner. This can remove much of the tedium needed to organize a behavioral experiment, and allows rapid modification of the task.

How is the data stored? Data is stored as an `.hdf5` file in the `data` folder. The filename depends on both the input given at the start of the experiment (Mouse ID) and on the timestamp of when the experiment was started.

What data is included in the .hdf5 file? See the section entitled *Stored HDF5 data format* for more information. By default, basic information about the `Experiment` and about each `Trial` (eg which `TrialType`, start times and end times, etc.) are stored. Additionally, a unique feature of the package is that for any `Events` specified, all attributes, from user-defined ones to temporary trial-specific variables, are auto-magically converted to strings and intelligently stored in the `.hdf5` file. This automatic storage can make it much easier to add and subtract `Events`, or

create new custom Event classes, without having to deal with the low-level data storage routines.

Note that weak-internal-use attributes (e.g. starting with a single underscore, `_`) in any Event class instance are, by default, not included in the .hdf5 file.

Where are the logs? All logs are both printed to the console (STDOUT) and are stored in the log folder with a filename matching that in the data folder.

What are class instance names used for? Events, as well as other base classes, require a name. Names help provide an intelligible format for the dataset and groups in the hdf5 file.

Built-in event types

Other event types are built-in to mouseberry. Full docstrings are available via the `help` command.

Rewards

Rewards can be delivered via two built-in Event classes:

- `mb.RewardSolenoid(name, pin, rate, volume, t_start)`:
 - Delivers liquid rewards through a solenoid connected to a single GPIO pin.
 - The rate of water flow, and the total volume desired, are both specified. A total opening time is then calculated.
- `mb.RewardSteppename, pin_motor_off, pin_step, pin_dir, pin_not_at_lim, rate, volume, t_start)`:
 - Delivers liquid reward through a stepper motor with four associated GPIO pins.
 - Reward direction is tuneable, and there is a limit input pin.

GPIO stimuli

A generic GPIO stimulus can be activated:

- `mb.GenericStim(name, pin, duration, t_start)`:
 - The GPIO pin is triggered for a set amount of time.

Audio

- `mb.Tone(name, t_start, t_dur, freq, db=-10)` :
 - Plays a pure tone at a given frequency. Uses the `sox` command.

Video


```
t_min=2, t_max=10)
```

```
tone_high = mb.Tone(name='tone_high', t_start=t_tone, t_dur=1, freq=10000)
trial = mb.TrialType(name='trial', p=1, events=[tone_high])
exp = mb.Experiment(n_trials=10, iti=1)
exp.run(trial)
```

Here, we have defined a normal distribution with mean (loc) of 4 and standard deviation (scale) of 2. See `help(scipy.stats.norm)` for detailed information on each distribution and the arguments that it takes.

Note that `t_min` and `t_max` define the minimum and maximum time that the draws from the distribution are permitted to take. For convenience, they are by default set to `-math.inf` and `math.inf`, respectively.

Defining stochastic ITIs

Instances of the `TimeDist` class can also be passed to `Experiment` to generate variable ITIs:

```
import mouseberry as mb
import scipy.stats
```

```
tone_high = mb.Tone(name='tone_high', t_start=t_tone, t_dur=1, freq=10000)
trial = mb.TrialType(name='trial', p=1, events=[tone_high])
```

```
iti_dist = mb.TimeDist(t_dist=scipy.stats.expon,
                      t_args={'scale': 1/5}, # corresponds to lambda=5
                      t_max=15)
```

```
exp = mb.Experiment(n_trials=10, iti=iti_dist)
exp.run(trial)
```

Constructing more complex experiments

Complex experiments consisting of many Events per `TrialType`, each with stochastic onset times, can be easily created. Since each event is by default threaded, events can overlap in time with minimal issues:

```
import mouseberry as mb
import scipy.stats as sp_st
```

```
events = []
t_start = mb.TimeDist(t_dist=sp_st.norm,
                     t_args={'loc': 5, 'scale': 3},
                     t_min=0, t_max=10)
```

```
for ind_event in range(40):
    event = mb.Tone(name='tone'+str(ind_event),
```

```

        t_start=t_start,
        t_dur=0.1,
        freq=ind_event*500)
    events.append(event)

```

```

trial = mb.TrialType('main_trial', p=1, events=events)
exp = mb.Experiment(n_trials=10, iti=2)
exp.run(trial)

```

The script above procedurally generates an experiment where 40 tones from 0 to 20KHz stochastically occur with start times following a normal distribution around 5s.

Stored data format: HDF5

By default, mouseberry stores all data in a logical, hierarchical data structure which is dynamically adjusted based on the contents of the trial-types and events.

All data is stored in the /data folder, and can be opened by the h5py package in Python, or by any similar package in other languages to read hdf5 files:

```

>> import h5py
>> f = h5py.File('test.hdf5', 'r')

```

The .hdf5 data format consists of groups, a named file-system-like hierarchy, and datasets, which store array data. Both are employed by mouseberry.

Experiment attributes

The experimental attributes are stored as .attrs, a dict of attributes, in the root group (/)

```

>> f.attrs.keys() # prints all attribute keys available for the experiment
>> f.attrs['mouse_id'] # prints the mouse id
>> f.attrs['sysinfo'] # prints a full POSIX system information summary (machine, etc.)

```

Trial attributes

The first group of the .hdf5 hierarchy is trials. This contains groups for the events and measurements classes. It also contains datasets that hold general information about each trial:

```

>> f['trials'].keys() # prints all keys within trials
>> f['trials/name'][:] # names of each trialtype
>> f['trials/t_start'][:] # start times of each trial, in seconds
>> f['trials/t_end'][:] # end times of each trial, in seconds

```

Measurements

Measurements are stored in `trials/measurements`. Each measurement has its own named group, and each group has both `t` and `data` datasets.

```
>> f['trials/measurements'].keys() # prints the keys for all acquired measurements
>> f['trials/measurements/licks/t'][0] # All measurement times (sec) in trial 0
>> f['trials/measurements/licks/data'][3] # All measurement values in trial 3
```

Events

Basic event data is stored in `trials/events` in an array-like datastructure (`[trial, event]`).

```
>> f['trials/events/name'][0, 1] # Prints name of trial 0's event 1
>> f['trials/events/t_start'][3, 0] # Prints start time (sec) of trial 3's event 0.
>> f['trials/events/t_end'][0, 0] # Prints end time (sec) of trial 0's event 0.
```

More complete event data, including any class instance attributes present that trial, are stored in a group directory structure in the root directory:

```
>> f['trial0/tone_low'].attrs.keys()
    # Prints all attribute keys for the tone_low event in trial 0
>> f['trial0/tone_low'].attrs['freq']
    # Prints the frequency attribute for tone_low in trial 0
>> f['trial0/rew'].attrs['t_start']
    # Prints the logged start-time attribute for reward event in trial 0
```

Creating custom classes

Mouseberry is designed to be hackable. Custom classes can be easily defined as follows, and they will be automatically processed and stored correctly by the rest of the package.

Events

Event subclasses must inherit from the Event base class in `mouseberry.groups.core`, and have the following methods defined, some which are optional and some which are required:

- `on_init`: optional
 - Method can define a set of steps to occur right when the trial starts, to initialize conditions for the event before they occur
- `on_assign_tstart`: required
 - Method must return a start time for the event this trial.
 - Method is called at the start of the trial, but after `.on_init()`.
- `on_trigger` : required
 - Method must define a set of steps to occur at the precise time when the event is triggered.
- `on_cleanup` : optional
 - Method can define a set of steps to occur when the experiment ends, to clean up variables, etc.

A simple example:

```

from mouseberry.groups.core import Event
import random
import time

```

```

class MockEvent(Event):
    """An event which prints both a string and a random number
    from 0 to 100, with a 2 second delay between each print.

    Parameters
    -----
    name : str
        Name of event
    t_start : float
        Start time of the event
    string_to_say : str
        String to print out
    """

    def __init__(self, name, t_start, string_to_say):
        super().__init__(name=name) # Super call for Event init
        self.t_start = t_start
        self.str = string_to_say

    def on_init(self):
        # pick random number from 1 to 100
        self.num = int(random.random()*100)

    def on_assign_tstart(self):
        return self.t_start

    def on_trigger(self):
        print(f'A string: {self.str}')
        time.sleep(2)
        print(f'A number: {self.num}')

    def on_cleanup(self):
        print('Cleaned up.')

```

Note that the dynamic attributes `str` and `num` will be automatically stored as `.attrs` in the hdf5 file without any extra work.

Measurements

Custom measurement subclasses can also be defined from the Measurement base class in `mouseberry.groups.core`. The following methods must be defined:

- `on_start` : required
 - Method must define a set of steps to occur in order to start a measurement

- Method must include threading to run in background, and must include on-the-fly logging of measurements to the following attributes:
 - `.data` (storing actual measurement data)
 - `.t` (storing measurement times)
- `on_stop` : required
 - Method must define a set of steps to occur in order to stop a measurement.
 - Method must include a way to stop the measurement thread.

Let's make a simple mock measurement:

```
from mouseberry.groups.core import Measurement
```

```
import time
```

```
import random
```

```
import threading
```

```
class MockMeasurement(Measurement):
```

```
    """A mock measurement.
```

```
    Parameters
```

```
    -----
```

```
    name : str
```

```
        Name of the measurement
```

```
    rate : float
```

```
        Sampling rate (Hz)
```

```
    """
```

```
    def __init__(self, name, rate):
```

```
        super().__init__(name=name)
```

```
        self.rate = rate
```

```
    def on_start(self):
```

```
        self.data = []
```

```
        self.t = []
```

```
        self.thread = SimpleNamespace()
```

```
        self.thread.stop_signal = threading.Event()
```

```
        self.thread.measure = threading.Thread(target=self.measure_loop)
```

```
        self.thread.measure.start()
```

```
    def measure_loop(self):
```

```
        while not self.thread.stop_signal.is_set():
```

```
            self.data.append(int(random.random()<0.1)) # random binary data
```

```
            self.t.append(time.time())
```

```
            time.sleep(1 / self.rate)
```

```
    def on_stop(self):
```

```
        self.thread.stop_signal.set()
```

```
        self.thread.measure.join()
```

Appendix D

Python Package III: SilentMLE Python package for silent synapse estimation

This Python package is associated with Manuscripts I and II. It performs two main tasks.

First, it allows us to disseminate the synthetic likelihood estimator for silent synapse estimation in an open-source, easy-to-use package, which will be of use for experimenters wanting to use this method.

Second, this package is concerned with code reproducibility of Manuscripts I and II. It contains all code needed for simulations and analysis for these two manuscripts, and will regenerate *any* main figure, in its entirety, simply by running a function (e.g. `sil.figures.plot_fig1()`). These figures are fully formatted exactly as they appear in the paper, and will exhibit slight variability from run to run due to the stochastic nature of the simulations. This allows complete reproducibility of the findings of these manuscripts.

The full package is available at <https://github.com/miclynn/SilentMLE>. Below, I have provided the github documentation associated with this package.

Statement of contribution

I wrote all code for this package.

SilentMLE

SilentMLE allows for optimal estimation of the silent fraction of some defined synaptic population. It works on electrophysiology data collected using the failure-rate analysis method (FRA; failure rates collected from neurons voltage-clamped at -70mV and +40mV while synapses are electrically stimulated).

SilentMLE works by running a set of detailed, constrained experimental simulations which provide a mapping between ground truth silent values, and biased subsamples of synapses recorded by the traditional failure-rate methods in most experimental setups. Thus, a numerically simulated likelihood function can be constructed which allows for optimal estimation, with significant reductions in both bias and variance of the estimator.

The estimator can work on either raw failure rate data, or on previously estimated fractions using the failure-rate analysis method. It initializes an estimator object where electrophysiologists can specify the parameters of their particular experimental setup, and can provide a release probability distribution if they so choose. The estimator then simulates a likelihood function which should match their particular experiments. Given a set of observations, the estimator returns a joint likelihood function, which allows for maximum-likelihood estimation as well as the specification of confidence intervals and comparison between experimental groups, if so desired.

Getting Started

Prerequisites

The main dependencies are: python3.6, numpy, scipy, matplotlib.

Installation on UNIX-like systems (MacOS/Linux)

A setup.py file allows for simple package installation:

```
$ cd {path-to-package}
$ python3 setup.py install
```

Performing maximum-likelihood estimation

Initializing the estimator

First, the package is imported and an instance of the `sil.Estimator` class is initialized. This runs a set of constrained simulations to determine the mapping between the ground-truth fraction silent in some population, and the biased measurement returned by electrical stimulation experiments. The resulting likelihood function is then stored in the estimator (`estimator.likelihood`), which is then ready to estimate.

```
import silentmle as sil
```

```
estimator = sil.Estimator()
```

Alternately, the `sil.Estimator` class can be initialized with a number of experimental constraints and/or simulation parameters. A full list of possible parameters can be found in the full class documentation of `sil.Estimator`, accessible through `help(sil.Estimator)`. As an example:

```
estimator = sil.Estimator(n_likelihood_points=500,
    pr_dist_sil=sil.PrDist(sp_stats.uniform),
    pr_dist_nonsil=sil.PrDist(sp_stats.uniform),
    num_trials=100, failrate_low=0.4,
    failrate_high=0.6)
```

This simulates using a fine-grained 500 points in the hypothesis-space (fraction silent); using a uniform release probability distribution for all synapses; and run the experimental simulations with 100 trials per Vm and an accepted failure rate range at hyperpolarized values of $0.4 < F < 0.6$.

All simulation and experimental parameters are stored in the class instance as `estimator.params` for easy access. The observation space and hypothesis space employed are stored as `estimator.obs` and `estimator.hyp` respectively. The numerically simulated likelihood function is stored as `estimator.likelihood`.

Advanced estimator usage

More advanced control of the Pr distributions is available with the `PrDist` class, which is passed as an argument to the `Estimator()` class:

```
import scipy.stats as sp_stats
```

```
pr_dist = sil.PrDist(dist=sp_stats.gamma,
    args={'a':1, 'scale':1/5.8})
```

```
estimator = sil.Estimator(pr_dist_sil=pr_dist, pr_dist_nonsil=pr_dist)
```

This creates a distribution object for release probability which is sampled during the experimental simulations. Any distribution from the `scipy.stats` library can be specified, and any keyword arguments should be specified as a dictionary as above.

MLE on experimental data

One can perform MLE on either a set of previous FRA estimates, or on pairs of raw failure rates at -70mV and $+40\text{mV}$. Both types are demonstrated below.

- *For data in FRA estimate form:*

```
data = [0.18, 0.27, 0.11, -0.10, -0.02, 0.08]
likelihood = estimator.estimate(data, dtype='est')
```

- *For data in raw failure rate form:*

```
data = [[0.2, 0.18], [0.4, 0.36], [0.38, 0.48], [0.58, 0.6], [0.46, 0.31]]
likelihood = estimator.estimate(data, dtype='failrate')
```

Here, `likelihood` returns a vector of length `h` containing the joint likelihood across all observations, where `h` is the hypothesis space size (i.e. `len(estimator.hyp)`)

Example of full estimation procedure and output:

In this section, we provide an example code block for the full SilentMLE estimation, as well as the resulting figure.


It is important to tailor the parameters to one's unique set of experimental parameters for optimal estimation (see above). However, this can act as a starting point for the use of the SilentMLE package.

```
import silentmle as sil
import scipy.stats as sp_stats

estimator = sil.Estimator(n_likelihood_points=500,
                          pr_dist_sil=sil.PrDist(sp_stats.uniform),
                          pr_dist_nonsil=sil.PrDist(sp_stats.uniform),
                          num_trials=50, failrate_low=0.2,
                          failrate_high=0.8)

# replace data here with fra-calculated silent fraction for each cell
data = [-0.36, 0.37, 0.57, 0.89, -0.43,
        0.49, 0.20, 0.22, 0.16, 0.51,
        0.67, 0.17, -0.48, -0.25, 0.00,
        0.16, 0.16, 0.90, 0.23, -0.15]
likelihood = estimator.estimate(data, dtype='est')
```

The resulting figure is shown below, and corresponds to Figure 1b in the STAR Protocols paper. It can be used to verify correct installation and usage of the package.



Advanced commands

Core simulation/analysis functions

The following are core simulation and analysis functions. Full docstrings are available with `help(desired-function)`.

- `sil.core.draw_subsample()`: Experimental simulation drawing a subsample of synapses from a larger population
- `sil.core.fra(fh, fd)`: Returns FRA estimate of fraction silent, given a hyperpolarized failure rate (`fh`) and a depolarized failure rate (`fd`).
- `sil.core.gen_fra_dist()`: Generates distribution of FRA values given a ground truth silent.
- `sil.core.power_analysis()`: Performs a power analysis (number of samples required for some specified statistical power) on two distributions of estimates for fraction synapses silent, which ostensibly come from two separate populations with distinct true fractions of silent synapses.
 - Utilizes a highly efficient divide-and-conquer algorithm (binary search) which has complexity $O(\log_2(n))$, where n is the maximum sample size required.

- `sil.core.power_analysis_llr()`: As above, but performs a power analysis using log-likelihood ratio testing (hypothesis-based testing) using the maximum-likelihood estimates.

Plotting function

Each of these functions will redo all analysis/simulations and will return a near-fully formatted figure.

- `sil.figures.plot_fig1()`
- `sil.figures.plot_fig2()`
- `sil.figures.plot_fig4()`
- `sil.figures.plot_figS1()`
- `sil.figures.plot_figS2()`
- `sil.figures.plot_figS3()`
- `sil.figures.plot_figS4()`

Note that some figures (eg FigS2-3) are constructed from base figures with different kwargs.

APPENDIX E: STAR Methods and Supplemental Figures associated with Manuscript I

EXPERIMENTAL MODEL AND SUBJECT DETAILS

MICE

This study employed male mice (C57BL/6; P14-P15; 10-15g). Mice were group housed before experimental procedures, and were provided with ad libitum access to food and water. All experiments and procedures were performed in accordance with approved procedures and guidelines set forth by the University of Ottawa Animal Care and Veterinary Services.

METHOD DETAILS

DERIVATION OF FRA EQUATION

Let n_s represent the number of silent synapses in some ensemble, n_a the number of active synapses, F_h the failure rate at hyperpolarized potentials (n_a synapses active) and F_d the failure rate at depolarized potentials ($n_a + n_s$ synapses active since NMDA-receptor-containing silent synapses can only conduct under depolarized settings). Then:

$$F_h = (1 - \text{Pr})^{n_a} \rightarrow n_a = \log_{1-\text{Pr}}(F_h)$$

$$F_d = (1 - \text{Pr})^{n_a+n_s} \rightarrow n_a + n_s = \log_{1-\text{Pr}}(F_d)$$

The silent fraction is $s = \frac{n_s}{n_a + n_s}$, so:

$$s = 1 - \frac{n_a}{n_s + n_a}$$

$$s = 1 - \frac{\log_{1-\text{Pr}}(F_h)}{\log_{1-\text{Pr}}(F_d)}$$

$$s = 1 - \frac{\log F_h}{\log F_d}$$

This can be written as a percentage:

$$s_{pc} = \left(1 - \frac{\log F_h}{\log F_d} \right) \cdot 100\%$$

(1)

SIMULATIONS WITH ACTIVE SYNAPSES ONLY

We simulated single active (nonsilent) synapses or groups of active synapses. For each synapse i we assigned a release probability by drawing from a uniform distribution $\text{Pr}_i \sim U(0,1)$.

We numerically stimulated stochastic neurotransmitter release at hyperpolarized voltages (V_h ; only active synapses) and depolarized voltages (V_d ; active and silent synapses). 50 trials were simulated at

each of V_h and V_d . Trials were modelled as Bernoulli trials with failure rate $F = \prod_{i=1}^m (1 - Pr_i)$ over m synapses considered, where $m = \begin{cases} n_a & \text{if } V = V_h \\ n_a + n_s & \text{if } V = V_d \end{cases}$. This yielded \widehat{F}_h and \widehat{F}_d , which are estimates of the true failure rate over 50 trials at V_h and V_d respectively. The FRA calculation was performed as reported in equation (1), with \widehat{F}_h and \widehat{F}_d as the equation inputs. Simulated currents (for purposes of illustration only) were generated by modeling the total current as a sum of all successfully releasing synapses, where each success is treated as generating a stochastic unitary current with an (arbitrary) mean amplitude of 10pA. However, all simulations proceeded using idealized Bernoulli trials to compute the stochastic failure rates (i.e. no mis-classification of successes and failures)

SIMULATIONS WITH ACTIVE AND SILENT SYNAPSES

In this section, we describe our simulations of minimum-simulation electrophysiological experiments, where a larger set of synapses with defined release probabilities and some fraction silent are iteratively reduced into some experimentally feasible subset on which to perform experiments. This iterative reduction models a gradual reduction in electrical stimulation intensity that an electrophysiologist would perform.

Assume the true fraction of silent synapses in a population is s_t . For each replicate, an initial set of silent and active synapse numbers $\{n_s, n_a\}$ were each chosen from a binomial distribution with an initial population size of $n_t = 100$ synapses and s_t silent probability (e.g. the distribution of n_s silent synapses was $p(n_s = k) = \binom{n_t}{k} s_t^k (1 - s_t)^{n_t - k}$). Release probabilities for each synapse were I.I.D and were drawn from either a uniform or gamma distribution as noted in the text, producing a vector of release probabilities $\mathbf{Pr}_a = [Pr_1, Pr_2, \dots, Pr_{n_a}]$ for active synapses and $\mathbf{Pr}_s = [Pr_1, Pr_2, \dots, Pr_{n_s}]$ for silent synapses. The synapse set $\{n_s, n_a\}$ and associated release probability set $\{\mathbf{Pr}_s, \mathbf{Pr}_a\}$ were then subjected to iterative rounds of reduction, which yielded an experimentally feasible subset.

Iterative reduction of the synapse sets proceeded until the failure rate criterion was met: $0.2 < F_{-70} < 0.8$. Here, $F_{-70} = \prod_{i=1}^{n_a} (1 - Pr_i)$ where $Pr_i \in \mathbf{Pr}_a$. F is analytically calculated at each iteration, rather than estimated using Monte Carlo simulations, for purposes of computational efficiency. Iterative reduction consisted of the following sequential steps: 1) Using $p_{elim} = 0.1$ as the elimination probability, stochastically eliminate $p_{elim}(n_s + n_a)$ synapses along with their associated release probabilities in $\{\mathbf{Pr}_s, \mathbf{Pr}_a\}$. 2) After each elimination, calculate F_{-70} . 3) If F_{-70} falls within criterion range, then store $\{n_s, n_a\}$ and $\{\mathbf{Pr}_s, \mathbf{Pr}_a\}$. If F_{-70} does not fall within criterion range, loop back to 1) and repeat.

Together, this iterative decrease yielded Monte Carlo-simulated subsamples (replicates) of active and silent synapses along with associated release probabilities $\{\mathbf{Pr}_s, \mathbf{Pr}_a\}$. For each replicate, we used $\{\mathbf{Pr}_s, \mathbf{Pr}_a\}$ to simulate 50 Bernoulli trials at each of $V_m = -70mV$ (n_a synapses) and $V_m = +40mV$ ($n_a + n_s$ synapses) similar to above, yielding stochastic failure rate estimates, \widehat{F}_{-70} and \widehat{F}_{+40} . These estimates were used as input to the FRA estimator, to return one silent fraction estimate, \widehat{s}_{FRA} , per Monte Carlo simulation.

For each ground truth silent value, s_t , we obtained ~20 000 replicates which each yielded an estimate, \hat{s}_{FRA} . These numerically simulated distributions formed the basis of our bias and variance analysis the FRA estimator. For each s_t , we can then compute:

$$\text{Bias}(\hat{s}_{FRA, s_t}) = E(\hat{s}_{FRA}) - s_t$$

$$\text{Var}(\hat{s}_{FRA}) = E[(\hat{s}_{FRA} - E(\hat{s}_{FRA}))^2].$$

EXPERIMENTAL MEASUREMENTS OF FRA

Hippocampal slices (300um thick) were prepared from P14-15 wild-type (C67/Blk) mice using a previously described technique (36). In brief, mice were anesthetized using isoflurane, decapitated, and the brain was extracted and placed in ice-cold choline-based cutting solution for the duration of the cut. Cutting solution was comprised of 119 mM choline-Cl, 2.5 mM KCl, 1 mM CaCl₂, 4.3 mM MgSO₄-7H₂O, 1 mM NaH₂PO₄, 1.3 mM sodium L-ascorbate, 26.2 mM NaHCO₃, and 11 mM glucose, and was equilibrated with 95% O₂ and 5% CO₂. 300uM coronal sections were made of the hippocampus. Slices were recovered in a chamber containing standard Ringer's solution of the following composition: 119 mM NaCl, 2.5 mM CaCl₂, 1.3 mM MgSO₄- 7H₂O, 1 mM NaH₂PO₄, 26.2 mM NaHCO₃, and 11 mM glucose, at a temperature of 37 °C, continuously bubbled with a mixture of 95% O₂ and 5% CO₂. Slices were recovered for 1 h in the recovery chamber and equilibrated to a temperature of ~25 °C until recordings were performed.

Neurons were visualized in the CA1 subfield using an upright microscope (Examiner D1 equipped with differential interference contrast (DIC) (Zeiss; 40×/0.75 N.A. objective). Using 4-6mOhm fire-polished glass electrodes (Sutter Corporation) pulled on a on a Narishige PC-10 pipette puller, CA1 subfield neurons were recorded in the whole-cell patch-clamp configuration at room temperature (20-22C) with a cesium-based internal comprised of 115 mM cesium methane-sulfonate, 5 mM tetraethylammonium-Cl, 10 mM sodium phosphocreatine, 20 mM Hepes, 2.8 mM NaCl, 5 mM QX-314, 0.4 mM EGTA, 3 mM ATP (Mg²⁺), and 0.5 mM GTP (pH 7.25) (adjusted with CsOH; osmolarity, 280–290 mOsmol/L). As an external solution, Ringer's solution was continuously circulated. Whole-cell recordings were carried out using an Axon Multiclamp 700B amplifier, filtered at 2 kHz and sampled at 10 kHz with an Axon Digidata 1440A (or 1550) digitizer. An electrical stimulation electrode was placed apposed to the apical dendritic arbor to activate Schaffer collaterals. Stimulation intensity was tuned such that ~50% synaptic failures and successes each were recorded.

The FRA protocol consisted of 100 sweeps of electrical stimulation (0.1ms pulse width; 10-15 seconds sweep-to-sweep time). 50 sweeps were performed at $V_m = -70mV$ and 50 sweeps were performed at $V_m = +40mV$. This FRA protocol was performed in triplicate on single neurons, with FRA iterations immediately following one another. Recordings were discarded when the access resistance changed by >30%. Liquid junction potential was not compensated for. The data was analyzed offline with the aid of a custom Python module written by the author and available online at: <https://github.com/miclynn/synappy>. Synaptic failures and successes were determined in an automated manner by setting a threshold for the baseline-subtracted electrically evoked current to exceed to be considered a success. The FRA calculation was performed as indicated in equation (1), by comparing the failure rate at -70mV and +40mV. Mean estimate spread was calculated by subtracting the minimum FRA-estimated silent synapse fraction from the maximum FRA-estimated silent synapse fraction within the same cell.

THE IMPROVED MAXIMUM-LIKELIHOOD ESTIMATOR

For a jointly observed pair of failure rates at hyperpolarized and depolarized potentials, \hat{F}_h and \hat{F}_d , we wish to estimate the most likely s_t . If we had an analytically tractable likelihood function, we could use likelihood-based inference techniques (e.g. maximum likelihood estimation) to find our most likely s_t . However, two factors make it difficult to analytically derive a likelihood function: first, the number of active and silent synapses in each sampled subset are only probabilistically related to s_t ; and second, each synapse has release probabilities Pr which are a complex function of the starting population's Pr and a bias imposed by the number of other (active) synapses present in that particular subset.

To circumvent these obstacles, we use a statistical inference approach inspired by recent treatments of noisy dynamical systems (Wood et al., 2010). While each individual simulation of synaptic release trains contains a large amount of information, we can extract summary statistics for each simulation and construct a synthetic likelihood function yielding the probability of obtaining these summary statistics at each underlying value of s_t . We choose the biased estimator \hat{s}_{FRA} as the summary statistic as it contains information about both \hat{F}_h and \hat{F}_d , (see below for a detailed discussion about this choice of summary statistic). Our synthetic likelihood function can therefore be written in terms of the summary statistic: $p(\hat{s}_{FRA} = d | s_t, Pr) = \frac{n_d^{s_t}}{n_{tot}^{s_t}}$, where d is a given value of observed data, $n_d^{s_t}$ is the count of Monte Carlo-simulated replicates using the parameters s_t, Pr which satisfy $\hat{s}_{FRA} = d$ (i.e. the observation is obtained), and $n_{tot}^{s_t}$ is the total number of Monte Carlo-simulated replicates using the parameters s_t, Pr . As above, s_t refers to the true silent fraction and Pr denotes a synaptic release probability distribution.

When more than one piece of data is acquired, is convenient to write a joint log-likelihood function over a given vector of observations $\mathbf{D} = (d_1, d_2, \dots, d_n)$. We can express this as

$\ell(s_t | \mathbf{D}) = \sum_{i=1}^n \log[p(d_i | s_t)]$. We can then construct an MLE estimator for s_t given some vector of measured data \mathbf{D} , expressed as $\hat{s}_{FRA-MLE}(\mathbf{D}) = \underset{s_t \in [0,1]}{\operatorname{argmax}}(\ell(s_t | \mathbf{D}))$. This estimator's bias and variance was

evaluated in an identical way as the FRA estimator (above), and the MLE estimator was found to exhibit no bias and significantly lower variance.

Due to the stochastic and finite nature of the simulations needed to generate $\ell(s_t | \mathbf{D})$, and to save computational resources, we have smoothed the numerical estimate for $\ell(s_t | \mathbf{D})$ using a Savitsky-Golay filter (window length = 0.1 silent fraction; polynomial order = 3) before computing $\underset{s_t \in [0,1]}{\operatorname{argmax}}(\ell(s_t | \mathbf{D}))$.

Two caveats to this approach: First, note that instead of constructing $p(\hat{s}_{FRA} | s_t)$ point-by-point with some step size $\delta s_t \rightarrow 0$, we could conceivably model the likelihood function as a continuous probability density over all values of s_t , by estimating parameters of a normal distribution similar to (33). Here, one could compute $\mu_{s_t} = \mathbf{E}[\hat{s}_{FRA}]$ and $\sigma_{s_t}^2 = \mathbf{E}[(\hat{s}_{FRA} - \mathbf{E}[\hat{s}_{FRA}])^2]$ for each s_t . However, we have established that the distribution of \hat{s}_{FRA} for each s_t has nonzero skewness and kurtosis, making it ineligible for normal distribution fitting.

Second, while we chose one summary statistic \hat{s}_{FRA} to construct a synthetic likelihood function, it is also possible to directly use the failure rates \hat{F}_h and \hat{F}_d as two summary statistics. This would appear to

be advantageous because we are using two summary statistics instead of one. However, this does not provide better estimation capabilities because the distributions for \hat{F}_h show no relationship with s_t in our simulations. Therefore, we chose a single summary statistic \hat{s}_{FRA} .

It is also formally possible that some summary statistic extracted from the raw event traces themselves (such as the variance of amplitudes of successful trials) could provide more information about the fraction of silent synapses in some ensemble than the failure rates alone. However, the unitary event amplitudes of AMPA-R and NMDA-R responses are typically not equivalent, nor are their standard deviations. Furthermore, there is typically more noise in recordings when a neuron is clamped at a depolarized potential. Since all of these factors can vary between experimental setups and brain regions, it is not trivial to provide a generalizable solution based on summary statistics of raw event traces. For these reasons, we took a failure-rate-based approach for the summary statistic.

POWER ANALYSIS OF ESTIMATORS FOR s_t

We employed a power analysis approach to compare the statistical power of different techniques for silent synapse detection. Parameters for power analyses were $\alpha = 0.05$ (false positive rate) and $\beta = 0.2$ (false negative rate) unless otherwise noted.

Given two populations of synapses, P_1 and P_2 , having distinct silent fractions $s_t^{(1)}$ and $s_t^{(2)}$, respectively. We want to determine n_{min} , the minimum number of samples required to discriminate these populations with some estimation technique. To do this, we wish to express statistical power, p , as a function of the number of samples: $p(n)$. We can then find an n_{min} such that $p(n > n_{min}) < \beta$. Our approach to explore $p(n)$ was to rely on Monte Carlo sampling methods paired with an efficient binary search algorithm to search the space of possible $n \in [0, 2^{11}]$ (an arbitrary maximum limit on n to search, due to computational constraints). Here, at each iteration, a sample size guess n_{guess} is made. After Monte Carlo simulations and an evaluation of statistical power (whether $p(n_{guess}) < \beta$ or $p(n_{guess}) > \beta$), n_{guess} is adjusted up (if insufficient statistical power), or down (if sufficient statistical power). The amount that n_{guess} is adjusted by is halved at each time point in order to converge on the true value for n_{min} in precisely $t = 11$ iterations (i.e. the algorithm has order $O(\log_2(n))$ where n is the initial guess.) The algorithm is described in the Python source code. Briefly:

0) Using the Monte Carlo methodology described above, generate numerical probability distributions for estimates of $s_t^{(1)}$ and $s_t^{(2)}$. These distributions are denoted $\mathbf{p}^{(1)}(\hat{s}_{FRA})$ and $\mathbf{p}^{(2)}(\hat{s}_{FRA})$.

1) Set $n_{guess} = 2^{11} = 2048$. Draw $k = 10000$ independent replicates, each composed of n_{guess} paired experimental samples drawn from $\mathbf{p}^{(1)}(\hat{s}_{FRA})$ and $\mathbf{p}^{(2)}(\hat{s}_{FRA})$. This yields corresponding replicate matrices \mathbf{D}_1 and \mathbf{D}_2 , which have dimension (k, n_{guess}) .

2) *Evaluation of statistical power.* For each one of the k columns in \mathbf{D}_1 and \mathbf{D}_2 , use a Wilcoxon Rank-Sum statistical test and record whether the replicate is classified as a positive with $p < \alpha$. Since the populations have distinct silent fractions, all positives are true positives so we can record the number of true positive n_{TP} . Compute the false-negative rate as $FNR = 1 - n_{TP}/k$.

3a) If $FNR < \beta$: Update $n_{guess}(t+1) \leftarrow n_{guess}(t) - \frac{|n_{guess}(t) - n_{guess}(t-1)|}{2}$. Loop back to (1).

- 3b) If $FNR > \beta$: Update $n_{guess}(t+1) \leftarrow n_{guess}(t) + \frac{|n_{guess}(t) - n_{guess}(t-1)|}{2}$. Loop back to (1).
- 4) Repeat until $|n_{guess}(t) - n_{guess}(t-1)| = 1$. Then $n_{min} = \min\{n_{guess}(t), n_{guess}(t-1)\}$.

In Fig. 4C, we test each of the four estimators (*fra*, *binary*, *fra-mle*, *binary(llr)*) by testing the null hypothesis that P_1 is silent. So we fix $s_t^{(1)} = 0$ while $s_t^{(2)} \in [0,1)$ and compute $n_{min}(s_t^{(2)})$ for each estimator:

- i) For \hat{s}_{FRA} (*fra*) the computation of $n_{min}(s_t^{(2)})$ proceeds numerically as described in the algorithm above.
- ii) For the binary sampling case (*binary*), the algorithm was identical to above, except instead of drawing estimates from a Monte Carlo-simulated distribution, we simply considered each draw of n_{guess} samples as constituting Bernoulli trials with $p = s_t^{(1)}$ or $s_t^{(2)}$. In step 2, a Chi-Squared test was used instead of a Rank-Sum test as the data is categorical (silent or active binarization).
- iii) For $\hat{s}_{FRA-MLE}$ (*fra-mle*), we employ a log-likelihood (LLR) approach. We first calculate likelihood functions as described above. In all cases the null hypothesis is $H_0: s_t = 0$ and the alternate hypothesis is $H_A: s_t = \hat{s}_{FRA-MLE}$. We apply Wilks' theorem, calculating a test statistic $D = -2[\ell(s_t = \hat{s}_{FRA-MLE}) - \ell(s_t = 0)]$ where $D \sim \chi^2(df = 1)$. This allows us to employ a Chi-Squared test on D in place of a Wilcoxon Rank-Sum test (step 2 of the algorithm above). All other components of the algorithm are identical.
- iv) For a binary sampler incorporating a model (*binary(llr)*), we derive an analytical expression for minimum sample size, n_{min} , with a LLR testing approach (starting from the binomial distribution) :

$$n_{min}(s_t) = \frac{\log(\beta)}{1 - s_t}.$$

QUANTIFICATION AND STATISTICAL ANALYSIS

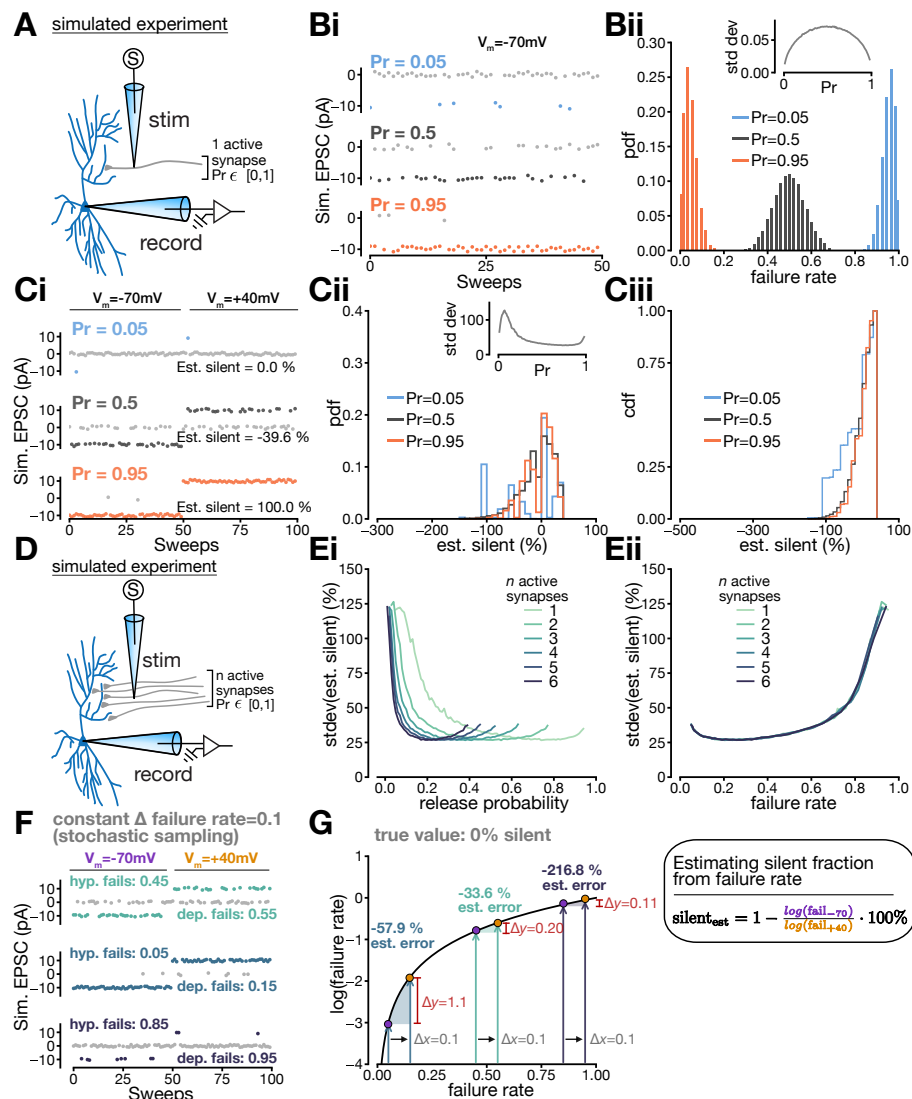
All numerical simulations and statistical analysis were performed using custom scripts employing standard functions of technical computing libraries NumPy (39) and SciPy (40) in a Python 3.6 environment. Further details about statistical tests are described in the Figure captions, Results and Methods of each respective section. In all cases, significance was defined as $p < 0.05$.

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Mouse: C57BL/6	Charles River Labs	C57BL/6NCrl
Software and Algorithms		
Python	Python Software Foundation	Python 3.6

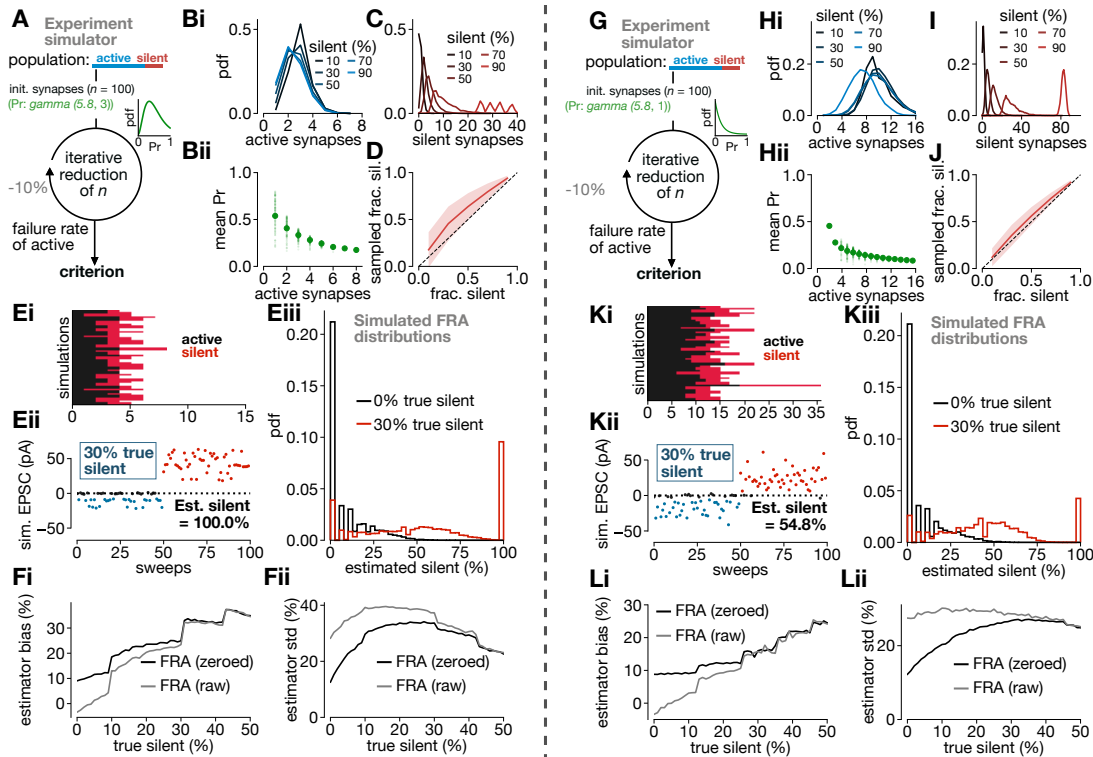
NumPy	NumPy	NumPy 1.18.1
SciPy	SciPy	SciPy 1.4.1
SilentMLE Python Package	This paper	https://github.com/miclynn/silentmle

Supplemental Figures

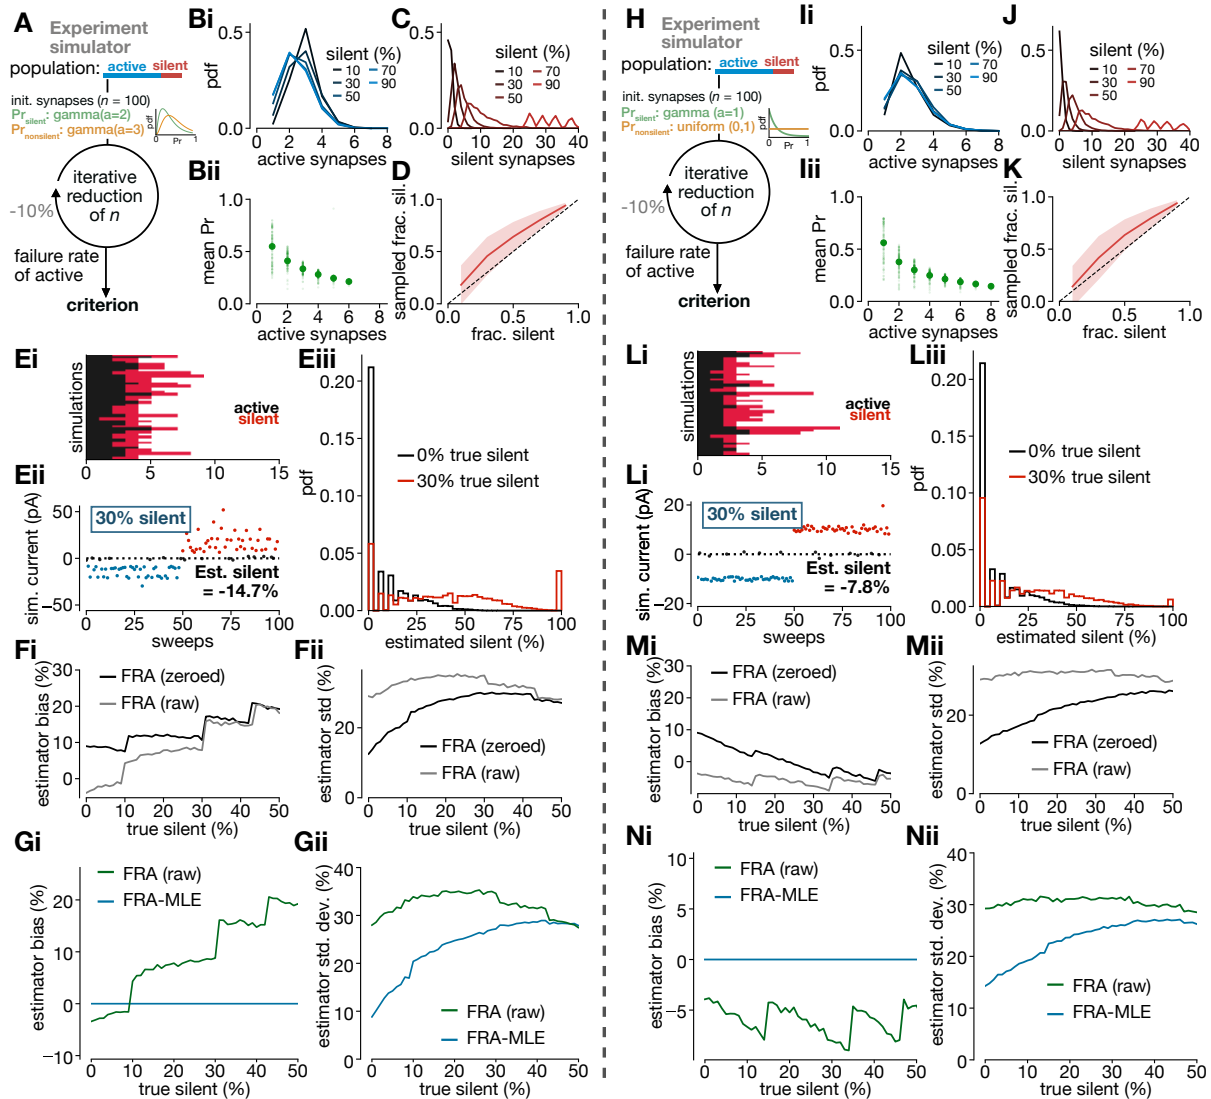


Supplemental Figure 1: The poor estimation characteristics of the FRA equation do not vary significantly over the parameter space of release probabilities and synapse numbers. *Related to Figure 1. A:* Depiction of simulation with one active synapse. **Bi:** Representative simulations for failures and successes at $V_m = -70\text{mV}$ for $Pr = 0.05$, $Pr = 0.5$ and $Pr = 0.95$. **Bii:** Representative distributions for failure rate at $Pr = 0.05$, $Pr = 0.5$ and $Pr = 0.95$. *Inset;* failure rate distributions have standard deviations which are parabolically related to release probability. **Ci:** Representative simulations and silent fraction estimates for full FRA experiments with $Pr = 0.05$, $Pr = 0.5$ and $Pr = 0.95$. **Cii:** Representative distributions for estimated silent fraction at $Pr = 0.05$, $Pr = 0.5$ and $Pr = 0.95$. *Inset;* in contrast to failure rate, the distributions for estimated silent fraction exhibit similar standard deviations across the range of possible release probability values. **Ciii:** Same as **Cii**, but the distributions are plotted as cumulative probability densities. **D:** Depiction of simulation with

multiple active synapses. **Ei**: Standard deviation of the estimated silent fraction as a function of release probability and number of active synapses. **Eii**: Standard deviation of the estimated silent fraction as a function of failure rate and number of active synapses. **F-G** provide mathematical intuition for the poor performance of the silent synapse estimation equation across the entire failure rate range. **F**: Simulations depicting 3 distinct experiments wherein probing consecutively (at -70 mV, then at +40 mV) the same all-active synaptic population leads to a slightly different failure rate of $\Delta 0.1$, well within the expected variability of successive probing of a stochastic process. **G**: These failure rate differences between successive experimental probing of the same synaptic population (plotted on abscissa) are transformed into $\log(\text{failure rate})$ changes (plotted on ordinate) which become smaller as absolute failure rate approaches 1. Since silent estimate is a function of $\log(\text{failure rate})$ (*box*), one might expect that estimate precision should be related to the absolute failure rate of the experiment (perhaps representing more deterministic or stochastic synapses). Despite this, the estimate error actually increases as failure rate approaches 1 or 0 (*overhead text*). This counterintuitive result occurs for several reasons. First, under high-Pr conditions, as failure rates approach 0, the natural logarithm of failure rates decreases nonlinearly, which amplifies failure rate variability under high-Pr conditions (*left side*). Second, under low-Pr conditions, as failure rates approach 1, the natural logarithm approaches 0 (*right side*). Taking the fraction of two stochastic numbers approaching 0, as in the FRA equation, amplifies the variability generated by small differences in failure rates. Therefore, the fundamental structure of the FRA equation yields substantial variability in silent synapse estimates, across the continuum of release probabilities, due to amplification of the inevitable fluctuations in failure rates when a given population of stochastic synapses are repeatedly probed.

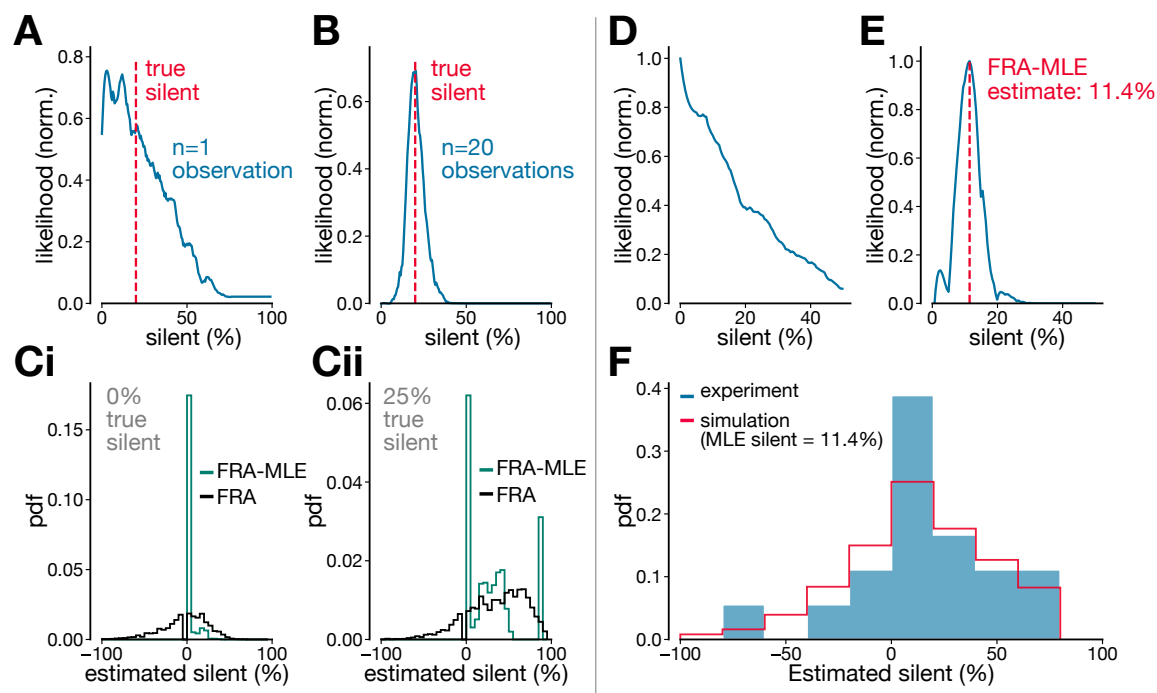


Supplemental Figure 2: Simulations exploring low synaptic release probabilities. *Related to Figure 2.* **A-F** depict simulations utilizing the gamma release probability distribution described in Branco et al., 2008. **A:** Schematic depicting experimental simulator approach. **Bi:** Probability distribution of number of active synapses obtained during each simulation. **Bii:** Release probability of active synapses. Small dots denote release probability for each simulation, while large dots denote the mean across simulations. **C:** Probability distribution of number of silent synapses obtained during each simulation. **D:** Sampled fraction silent is shown against the true fraction silent. **Ei:** Depiction of returned samples of silent/active synapses. **Eii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction. **Eiii:** Distribution of FRA estimates for two true silent values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii:** FRA estimator standard deviation for non-zeroed and zeroed distributions. **G-L** depict simulations utilizing a gamma release probability distribution biased towards very low values. **G:** Schematic depicting experimental simulator approach. **Hi:** Probability distribution of number of active synapses obtained during each simulation. **Hii:** Release probability of active synapses. Small dots denote release probability for each simulation, while large dots denote the mean across simulations. **I:** Probability distribution of number of silent synapses obtained during each simulation. **J:** Sampled fraction silent is shown against the true fraction silent. **Ki:** Depiction of returned samples of silent/active synapses. **Kii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction. **Kiii:** Distribution of FRA estimates for two true silent values. **Li:** FRA estimator bias for non-zeroed and zeroed distributions. **Lii:** FRA estimator standard deviation for non-zeroed and zeroed distributions.



Supplemental Figure 3: Simulations exploring an FRA assumption violation with unequal release probabilities between silent and active synapses. *Related to Figure 2.* **A-G** depict simulations with mild unequal release probabilities; *silent*: $\text{gamma}(a=2, \text{lambda}=5.8)$; *active*: $\text{gamma}(a=3, \text{lambda}=5.8)$. **A:** Schematic depicting experimental simulator approach. **Bi:** Probability distribution of number of active synapses obtained during each simulation. **Bii:** Release probability of active synapses. Small dots denote release probability for each simulation, while large dots denote the mean across simulations. **C:** Probability distribution of number of silent synapses obtained during each simulation. **D:** Sampled fraction silent is shown against the true fraction silent. **Ei:** Depiction of returned samples of silent/active synapses. **Eii:** Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction. **Eiii:** Distribution of FRA estimates for two true silent values. **Fi:** FRA estimator bias for non-zeroed and zeroed distributions. **Fii:** FRA estimator standard deviation for non-zeroed and zeroed distributions. **Gi:** Estimator bias for FRA (green) and corrected FRA-MLE (blue) estimators. **Gii:** Estimator standard deviation for FRA (green) and corrected FRA-MLE (blue) estimators. **H-N**

depict simulations with extreme unequal release probabilities; *silent*: $\text{gamma}(a=1, \text{lambda}=5.8)$; *active*: $\text{uniform}(0, 1)$. **H**: Schematic depicting experimental simulator approach. **Ii**: Probability distribution of number of active synapses obtained during each simulation. **Iii**: Release probability of active synapses. Small dots denote release probability for each simulation, while large dots denote the mean across simulations. **J**: Probability distribution of number of silent synapses obtained during each simulation. **K**: Sampled fraction silent is shown against the true fraction silent. **Li**: Depiction of returned samples of silent/active synapses. **Lii**: Depiction of a simulated FRA experiment performed on a sample of synapses returned from the iterative reduction. **Liii**: Distribution of FRA estimates for two true silent values. **Mi**: FRA estimator bias for non-zeroed and zeroed distributions. **Mii**: FRA estimator standard deviation for non-zeroed and zeroed distributions. **Ni**: Estimator bias for FRA (green) and corrected FRA-MLE (blue) estimators. **Nii**: Estimator standard deviation for FRA (green) and corrected FRA-MLE (blue) estimators.



Supplemental Figure 4: A maximum likelihood estimator for silent synapse fractions. *Related to Figure 4.* **A-C** depict simulations verifying the performance of the FRA-MLE estimator. **A:** Given a single observation from an FRA distribution, FRA-MLE returns likelihood function which shows that multiple silent fractions are consistent with this one observation. **B:** When 20 observations from the same FRA distribution are considered together by FRA-MLE as a joint likelihood function, a highly precise estimate of the true silent fraction is possible. **C:** Given a true silent fraction of 0% (**Ci**) or 25% (**Cii**), FRA-MLE on single observations returns low-variance, precise estimates, while FRA returns high-variance, imprecise estimates. **D-F** depict an application of the FRA-MLE estimator to experimental data in hippocampus at the CA3-CA1 synapse (*Figure 3*). **D:** Likelihood function returned by FRA-MLE from a single experimental observation. **E:** The joint likelihood function, across all experimental observations, returned by FRA-MLE. The FRA-MLE estimate from this joint likelihood function is 11.4% silent. **F:** A comparison of the distribution of FRA-estimated silent values (blue) overlaid with a simulated distribution of estimated silent values, calculated at the maximum likelihood estimate from FRA-MLE (11.4% silent). The distributions are not statistically different ($p = 0.99$; KS-test, two-tailed).

Appendix F: Supplemental Information for Manuscript III




Figure S1: Validation of SERT-CRE-tdTomato mouse line and Lhb injection site. Confocal image of SERT^{cre} (A) and TPH2-expressing (B) 5-HT neurons in the DRN. (C) Merge confocal image of SERT^{cre} and TPH2 channels. (D) Bar plot of overlap between SERT-cre and TPH2-expressing 5-HT neurons in the DRN (98.38 ± 0.8; n = 3, mean ± SE). (E) Representative image of Lhb injection site (AAV-ChR2-mCherry).

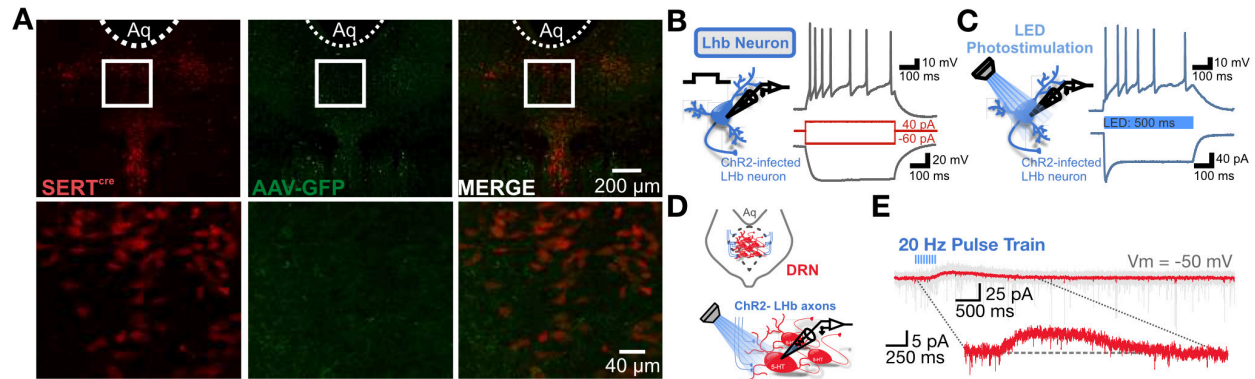


Figure S2: Lhb sends axonal projections to the DRN. **(A)** Confocal image of GFP-expressing Lhb axons in the DRN with 5-HT neurons identified with SERT-TdT. **(B)** ChR2-infected Lhb neuron can be induced to spike with current injection. **(C)** LED photostimulation produces qualitatively similar spiking patterns to those induced by current injection. **(D)** Experimental schematic for long-range optogenetic interrogation. **(E)** Example serotonin neuron which does not receive direct Lhb input, but displays a delayed heterosynaptic inhibitory current.

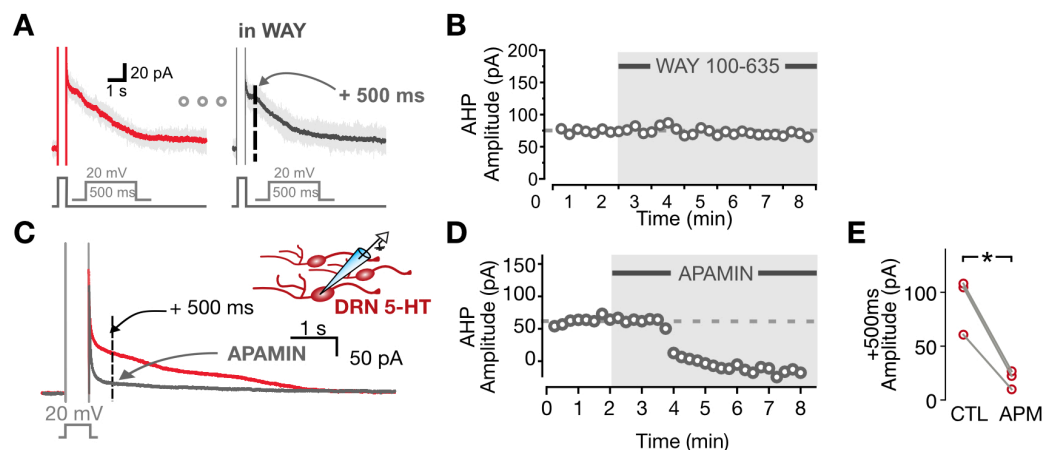


Figure S3: Voltage-clamp experiments provide evidence that the afterhyperpolarization of serotonin neurons is mediated by SK channels, not 5-HT_{1A}Rs. **(A)** Representative traces of AHP current (IAHP) driven by a 500 ms depolarization step to 20 mV in voltage clamp. **(B)** Representative scatter plot showing no differences in IAHP amplitude following WAY 100-635 application. **(C)** Representative IAHP traces before and after APAMIN application (300 nM). **(D)** Representative scatter plot showing the IAHP +500 ms amplitudes in response to APAMIN. **(E)** +500 ms IAHP amplitudes before and after APAMIN application (CTL, 91.1 ± 15.3; APM, 19.6 ± 5.0; n = 3; p = 0.02, paired two-sided t-test).




Figure S4: Fluorescent dextran injections in DRN reveal local 5-HT projections. **A)** Experimental schematic illustrating visually guided dextran injections in DRN of acute slices. **B)** Confocal image showing dextran injection site and labeled neurons. **C)** Confocal image showing tdTomato-labeled serotonin neurons. **D)** Overlaid confocal image showing overlap between distant dextran-labeled neurons and tdTomato fluorescence.

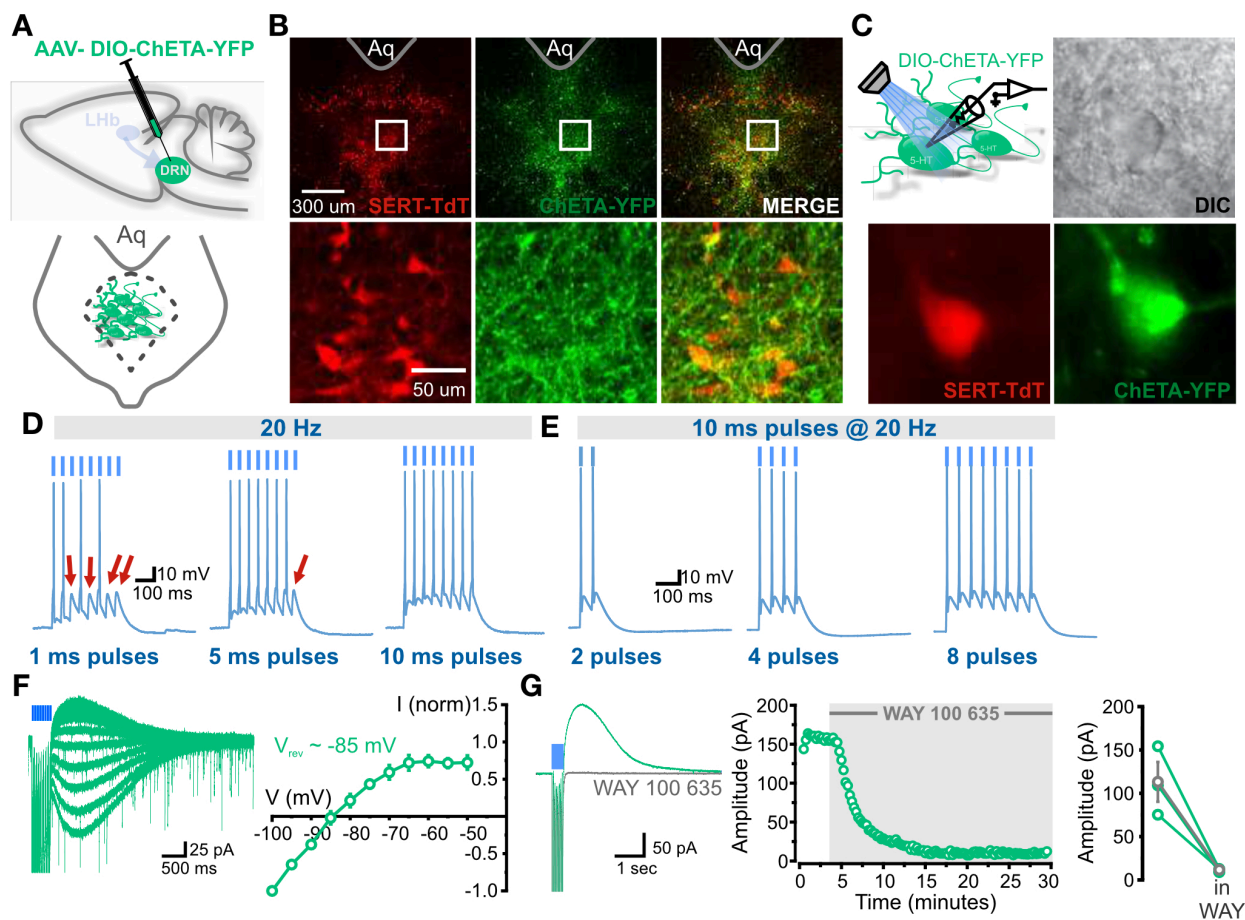


Figure S5: Validation of local CHETA optogenetic strategy in DRN. **(A)** Schematic depicting AAV-DIO-ChETA-YFP injection into the DRN of SERT^{cre}-TdT mice. **(B)** Confocal image of ChETA-YFP expression in tdTomato-positive neurons in the DRN. Above, low magnification (scale, 300 μ m); below, high magnification (scale, 50 μ m). **(C)** Schematic and EPI-fluorescent image of whole-cell recordings from ChETA-YFP-expressing SERT-TdT neurons. **(D)** LED stimulation of 5-HT neuron spiking with different pulse durations at frequency of 20 Hz. **(E)** LED stimulation of 5-HT neuron spiking in response to varying number of 10 ms pulses at 20 Hz. **(F)** Representative IV relationship traces and IV relationship ($n = 5$, mean \pm SE). **(G)** Representative traces and scatter plot showing complete blockade of outward current by WAY 100-635. Right, population data is displayed (Pre-WAY, 112.9 \pm 22.9; Post-WAY, 11.5 \pm 1.1; $n = 3$; $p = 0.02$, paired two-sided t-test).

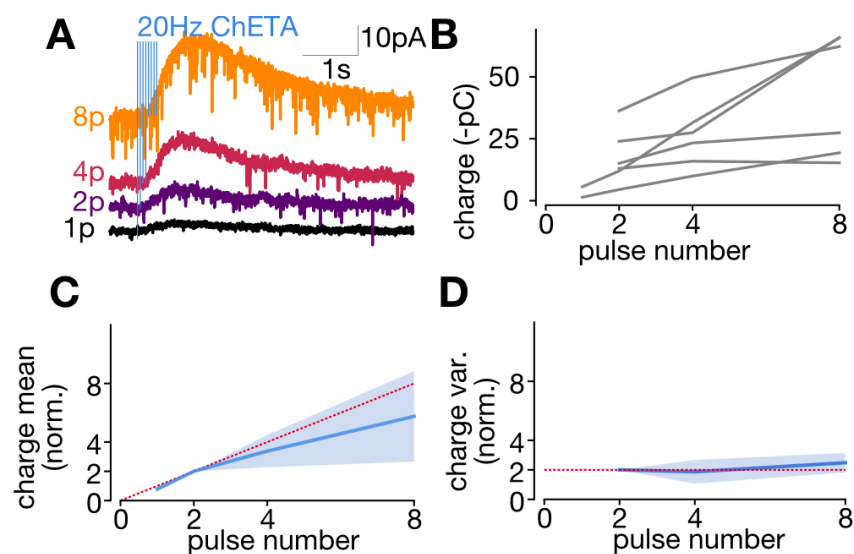


Figure S6: Postsynaptic dynamics of recurrent 5-HT1A connections indicate non-desensitizing receptors. **(A)** Representative averaged traces from a single non-CHETA-expressing 5-HT neuron depicting responses to 20Hz optogenetic stimulation of DRN for various pulse numbers. **(B)** Event charge for each neuron as a function of pulse number at 20Hz (n=6). **(C)** Charge mean as a function of pulse number, normalized to the mean at two pulses (n=6, mean \pm SD). Dotted red line represents a linear relationship between pulse number and evoked charge. **(D)** Charge variance as a function of pulse number, normalized to the variance at two pulses (n=3, mean \pm SD). Dotted red line represents predicted variance for release probability $p=1$.

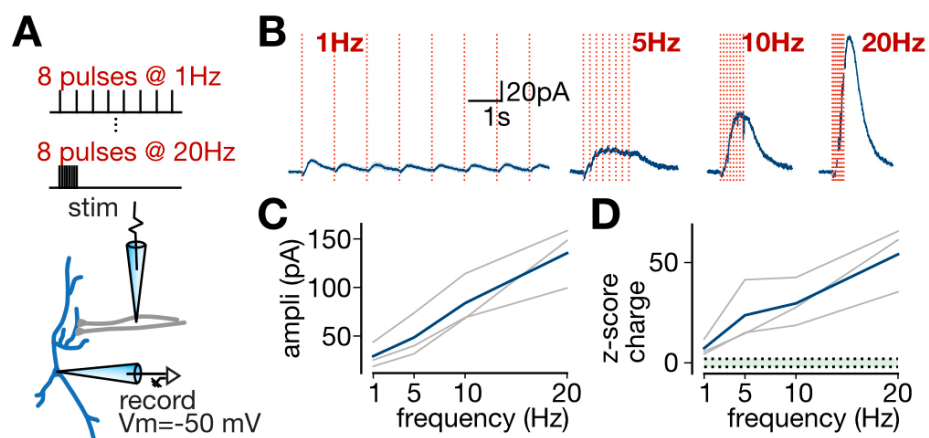


Figure S7. Frequency dependence of 5-HT_{1A} connections at suprathreshold stimulation frequencies. **(A)** Experimental setup depicting electrical stimulation of 5-HT_{1A} synapses while recording from an identified 5-HT neuron in the whole-cell configuration. **(B)** Representative averaged traces from a single neuron depicting responses to electrical stimulation at various frequencies, with pulse number held constant. **(C)** IPSC amplitude as a function of frequency, for individual neurons (grey) and the mean of all neurons (blue, $n=3$). **(D)** Event charge, expressed as a z-score relative to a noise distribution, as a function of frequency. Depicted are individual neurons (grey) and the mean of all neurons (blue, $n=3$). Green shaded area represents 95% confidence interval for noise distribution.

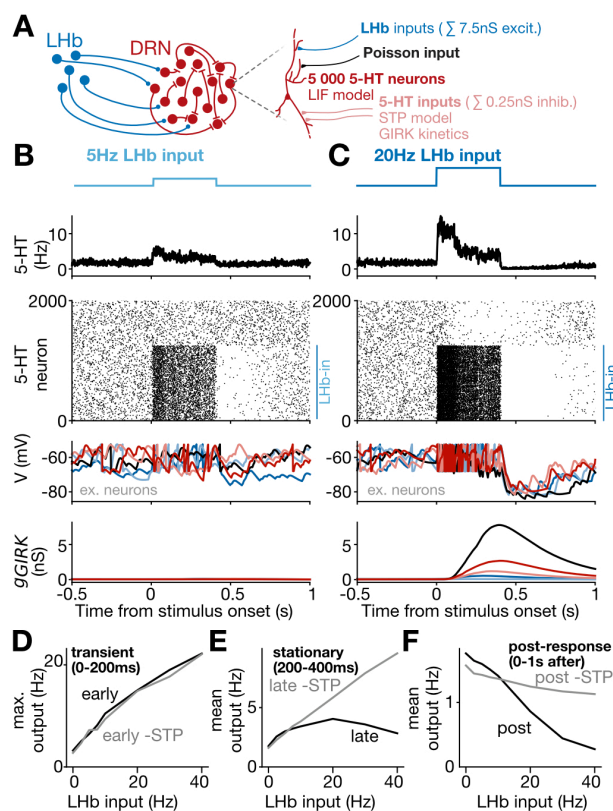
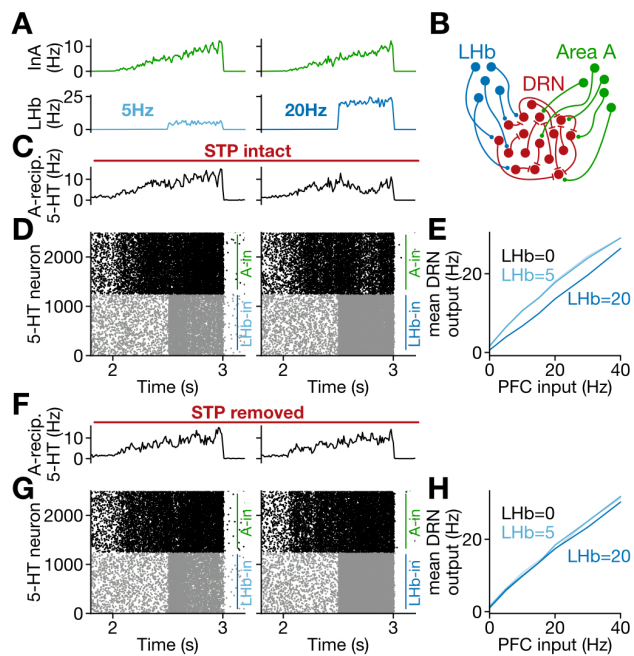


Figure S8: Nonlinear transformation of habenulo-raphé afferents generated by recurrent plasticity rules, with tenfold decrease in recurrent inhibitory connection probability (sum 0.25nS conductance received by each serotonin neuron). **(A)** Schematic of the LIF network model. **(B)** Results of a simulation where a 5Hz step input from Lhb was delivered to a subset of serotonin neurons. **(C)** Results of a simulation where a 20Hz step input from Lhb was delivered to a subset of serotonin neurons. **(D)** Input-output transformation for the transient (0-200ms) response phase, with (black) and without (grey) STP rules at recurrent serotonergic connections. **(E)** Input-output transformation for the stationary (200-400ms) response phase, with (black) and without (grey) STP rules at recurrent serotonergic connections. **(F)** Input-output transformation for the post-response (0-1s after stimulus) phase, with (black) and without (grey) STP rules at recurrent serotonergic connections.



removed. **(G)** Input-output function, for input A, with STP rules removed and coincident LHB input at 0, 5 or 20Hz.

Figure S9: Plasticity-dependent winner-take-all dynamics in raphe, with tenfold decrease in recurrent inhibitory connection probability (sum 0.25nS conductance received by each serotonin neuron). **(A)** Simulations with Lhb inputs of either 5Hz (left) and 20Hz (middle) were delivered to serotonin neurons during ramp input from area A. Schematic of LIF network model with two inputs (right). **(B)** Mean responses of serotonin neurons to 5Hz (left) and 20Hz (right) input with STP rules intact. **(C)** Raster plot showing serotonin neuron responses to 5Hz (left) and 20Hz (right) input with STP rules intact. **(D)** Input-output function, for input A, with STP rules intact and coincident Lhb input at 0, 5 or 20Hz. **(E)** Mean responses of serotonin neurons to 5Hz (left) and 20Hz (right) input with STP rules removed. **(F)** Raster plot showing serotonin neuron responses to 5Hz (left) and 20Hz (right) input with STP rules

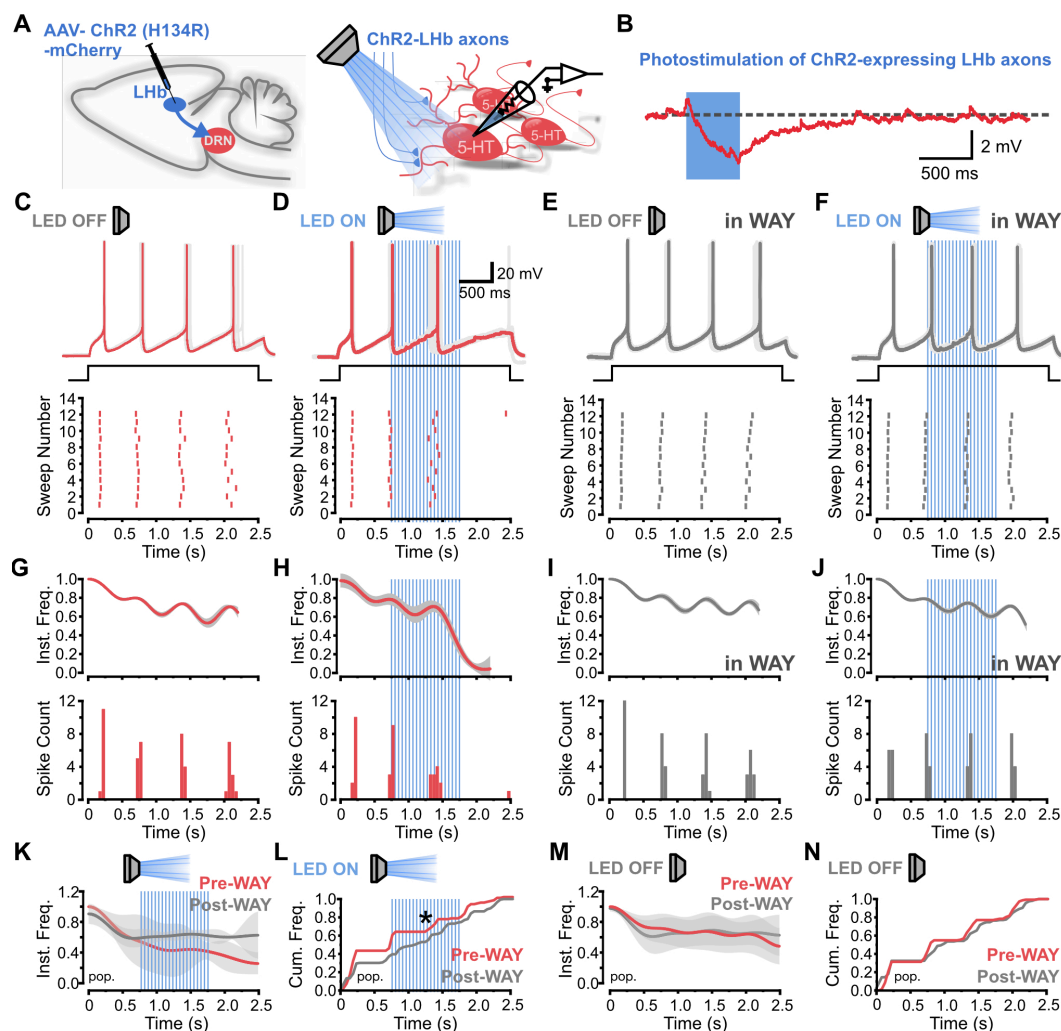


Figure S10: Lhb-driven, 5HT1AR mediated, recurrent feedforward modulation of 5-HT firing activity in the DRN. **(A)** Schematic of experimental design. Left, AAV-ChR2(H134R)-mCherry is injected into the Lhb. Right, whole-cell recordings from TdT-positive 5-HT neurons. **(B)** Lhb-driven heterosynaptic 5HT1AR-mediated potential in SERT-TdT 5-HT neuron. **(C)** Spiking of a 5-HT neuron in response to a 2.5 s current step (scale, 20 mV, 500ms). Below, raster plot of the same 5-HT neuron spiking. **(D)** Spiking of a 5-HT neuron in response to a 2.5 s current step with LED activation of ChR2-expressing Lhb terminals in the slice for 1 sec in the middle of the current step. Below, raster plot of the same 5-HT neuron spiking. **(E)** Spiking of a 5-HT neuron in response to a 2.5 s current step in the presence of WAY 100-635. Below, raster plot of the same 5-HT neuron spiking. **(F)** In the presence of WAY 100-635, spiking of a 5-HT neuron in response to a 2.5 s current step with LED activation of ChR2-expressing Lhb terminals in the slice for 1 sec in the middle of the current step. Below, raster plot of the same 5-HT neuron spiking. **(G)** Instantaneous frequency line (top) and spike count histogram plots (below) from recording in **C**. **(H)** Instantaneous frequency line (top) and spike count histogram plots (below) from recording in **D**. **(I)** Instantaneous frequency line (top) and spike count histogram plots (below) from recording in **E**. **(J)** Instantaneous frequency line (top) and spike count histogram plots (below) from recording in **F**. **(K)** Instantaneous frequency population plot in pre- and post-

WAY recordings with LED stimulation (n=3 neurons, mean \pm SE). **(L)** Cumulative frequency population plot in pre- and post-WAY recordings with LED stimulation (3 cells, $p = 0.0017$, 2-sample Kolmogorov-Smirnov test). **(M)** Instantaneous frequency population plot in pre- and post-WAY recordings in the absence of LED stimulation (n=3 neurons, mean \pm SE). **(N)** Cumulative frequency population plot in pre- and post-WAY recordings in the absence of LED stimulation (3 cells, $p = 0.07$, 2-sample Kolmogorov-Smirnov test).

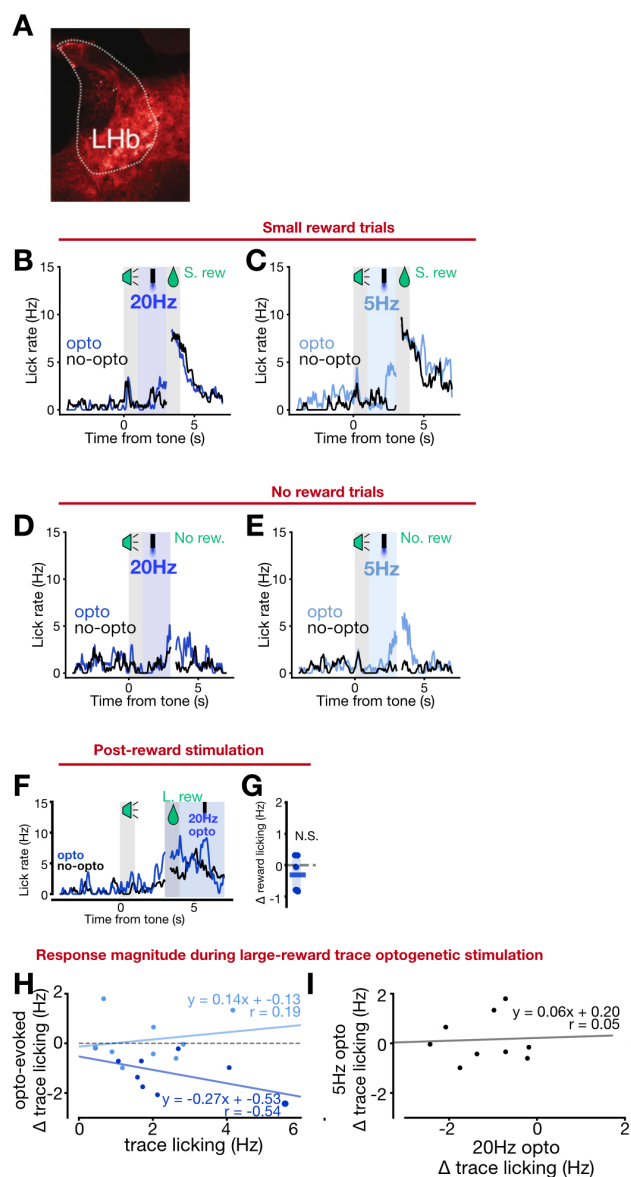
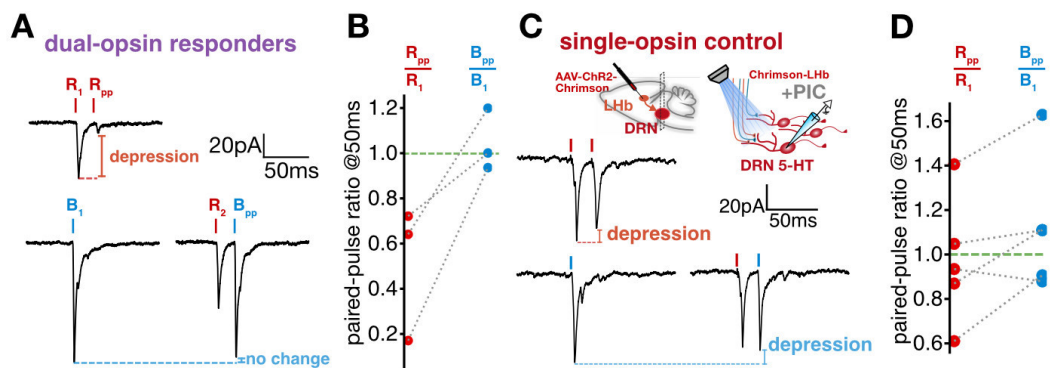
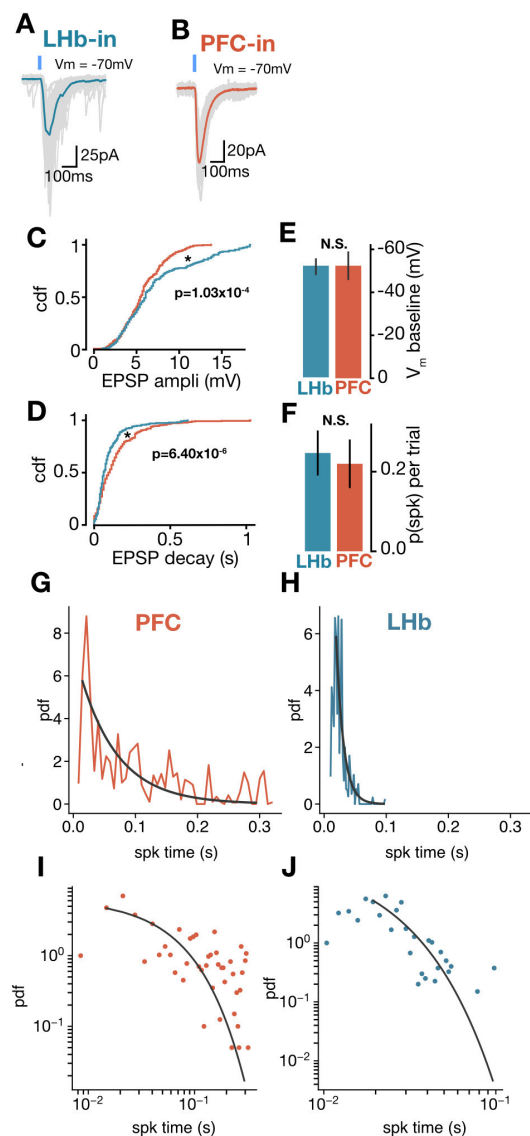


Figure S11: The habenulo-raphe pathway disrupts reward-conditioned responses in a value-dependent manner and does not affect unconditioned response. **(A)** Representative example of LHb injection site. **(B)** Mean licking across all small reward trials for a single mouse depicting control trials (black) or 20Hz optogenetic stimulation trials (blue). **(C)** Mean licking across all small reward trials for a single mouse depicting control trials (black) or 5Hz optogenetic stimulation trials (blue). **(D)** Mean licking across all no reward trials for a single mouse depicting control trials (black) or 20Hz optogenetic stimulation trials (blue). **(E)** Mean licking across all no reward trials for a single mouse depicting control trials (black) or 5Hz optogenetic stimulation trials (blue). **(F-G)** depict post-reward optogenetic stimulation experiments. **(F)** Mean licking across all large reward trials for a single mouse depicting control trials (black) or trials where 20Hz optogenetic stimulation was delivered during the post-reward period (blue). **(G)** Change in post-reward licking during optogenetic stimulation trials, compared to pre-stimulated baseline period ($n=6$ mice, paired two-sided t-test, unstimulated vs stimulated reward licking). **(H)** Correlation between mean trace licking (large reward trials) and optogenetic effect size for 20Hz stimulation (dark blue) and 5Hz stimulation (light blue) for trace stimulation during large reward trials. **(I)** Correlation between effect size for 20Hz stimulation and 5Hz stimulation for trace stimulation during large reward trials.


Appendix G: Supplemental for Manuscript IV




Supplemental Figure 1: Dual-color opsin dissection of long-range inputs to 5-HT neurons. **A:** For 5-HT neurons that respond to both blue and red light, red paired-pulse stimuli produce depression (20Hz, top) while a red prepulse does not significantly affect the amplitude of a blue light-evoked current (bottom). **B:** Quantification of paired pulse ratio for red paired-pulse stimuli (red) and blue paired-pulse stimuli with red prepulse (blue). **C:** Single-opsin control with Chrimson injected in LHb. Red paired-pulse stimuli produce depression (20Hz, top) and a red prepulse significantly affect the amplitude of a blue light-evoked current (bottom). **D:** Quantification of paired pulse ratio for red paired-pulse stimuli (red) and blue paired-pulse stimuli with red prepulse (blue).



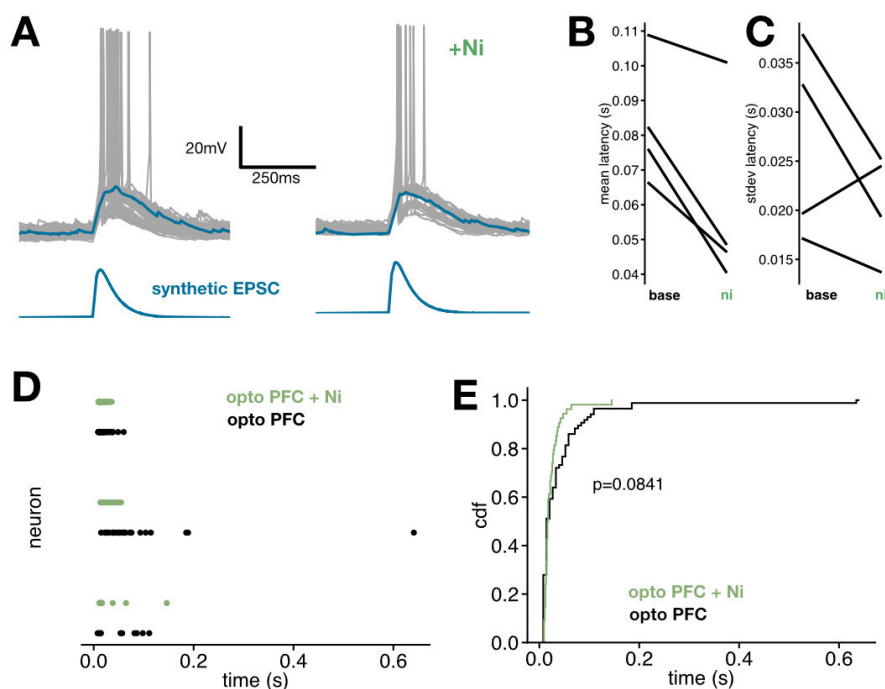
Supplemental Figure 2: Further quantification of subthreshold and suprathreshold statistics of light-evoked events in 5-HT neurons. **A-B:** Example traces in voltage clamp of light-evoked EPSCs in 5-HT neurons. **C:** Cumulative distribution of EPSP amplitudes of all Lhb-evoked events (blue) and PFC-evoked events (red) (equal sampling for each neuron, K-S test). **D:** Cumulative distribution of EPSP decay of all Lhb-evoked events (blue) and PFC-evoked events (red) (equal sampling for each neuron, K-S test). **E:** Mean baseline membrane potential before light-evoked events (n=9 neurons for Lhb and PFC groups). **F:** Mean probability for photostimulation to trigger a spike (n=9 neurons for Lhb and PFC groups). **G-H:** Probability distribution of spike times triggered by PFC (red) and Lhb (blue). Black lines indicate exponential fit. **I-J:** Log-log plot of probability distribution of spike times triggered by PFC (red) and Lhb (blue). Black lines indicate exponential fit.



Supplemental Figure 3: Low-threshold calcium channels modulate the kinetics of EPSPs in 5-HT neurons. **A:** Example traces in current-clamp showing that injection of fast ($\tau=50\text{ms}$) synthetic EPSCs produce EPSP responses that are modulated by bath application of NiCl. **B:** Example traces in current-clamp showing that injection of fast ($\tau=100\text{ms}$) synthetic EPSCs produce EPSP responses that are modulated by bath application of NiCl. **C:** Example traces in current-clamp showing that injection of fast ($\tau=500\text{ms}$) synthetic EPSCs produce EPSP responses that are modulated by bath application of NiCl.



Supplemental Figure 4: Subthreshold noise properties are not affected by calcium channels in 5-HT neurons. **A:** Example traces of a single neuron recorded in current-clamp during injection of constant subthreshold depolarizing current (current step immediately before spiking), before (top) and after (bottom) application of NiCl. **B:** Probability distribution of membrane potentials for an example neuron before (top) and after (bottom) application of NiCl. **C:** Membrane potential distribution skewness is not significantly affected by bath application of NiCl (n=7, paired T-test). **D:** Membrane potential distribution kurtosis not is significantly affected by bath application of NiCl (n=8, paired T-test).



Supplemental Figure 5: Calcium channels affect spike timing in 5-HT neurons. **A:** Example trace in current clamp showing injection of synthetic EPSCs with long decay triggers delayed spike timing (left), and bath application of NiCl₂ reduces occurrence of delayed spikes (right). **B:** Mean spike latency before and after NiCl₂ application (n=4). **C:** Standard deviation of spike latency before and after NiCl₂ application (n=4). **D:** Timing of spikes triggered by photostimulation of long-range PFC axons before (black) and after (green) bath application of NiCl₂. Each dot is a spike and each horizontal row is one neuron. **E:** Cumulative distribution of spike times triggered by photostimulation of PFC axons before (black) and after (green) bath application of NiCl₂ (K-S test).