Okinawa Institute of Science and Technology
Graduate University

Thesis submitted for the degree

Doctor of Philosophy

Characterisation of Serotonin, Noradrenaline and Dopamine Release in the Motor Cortex

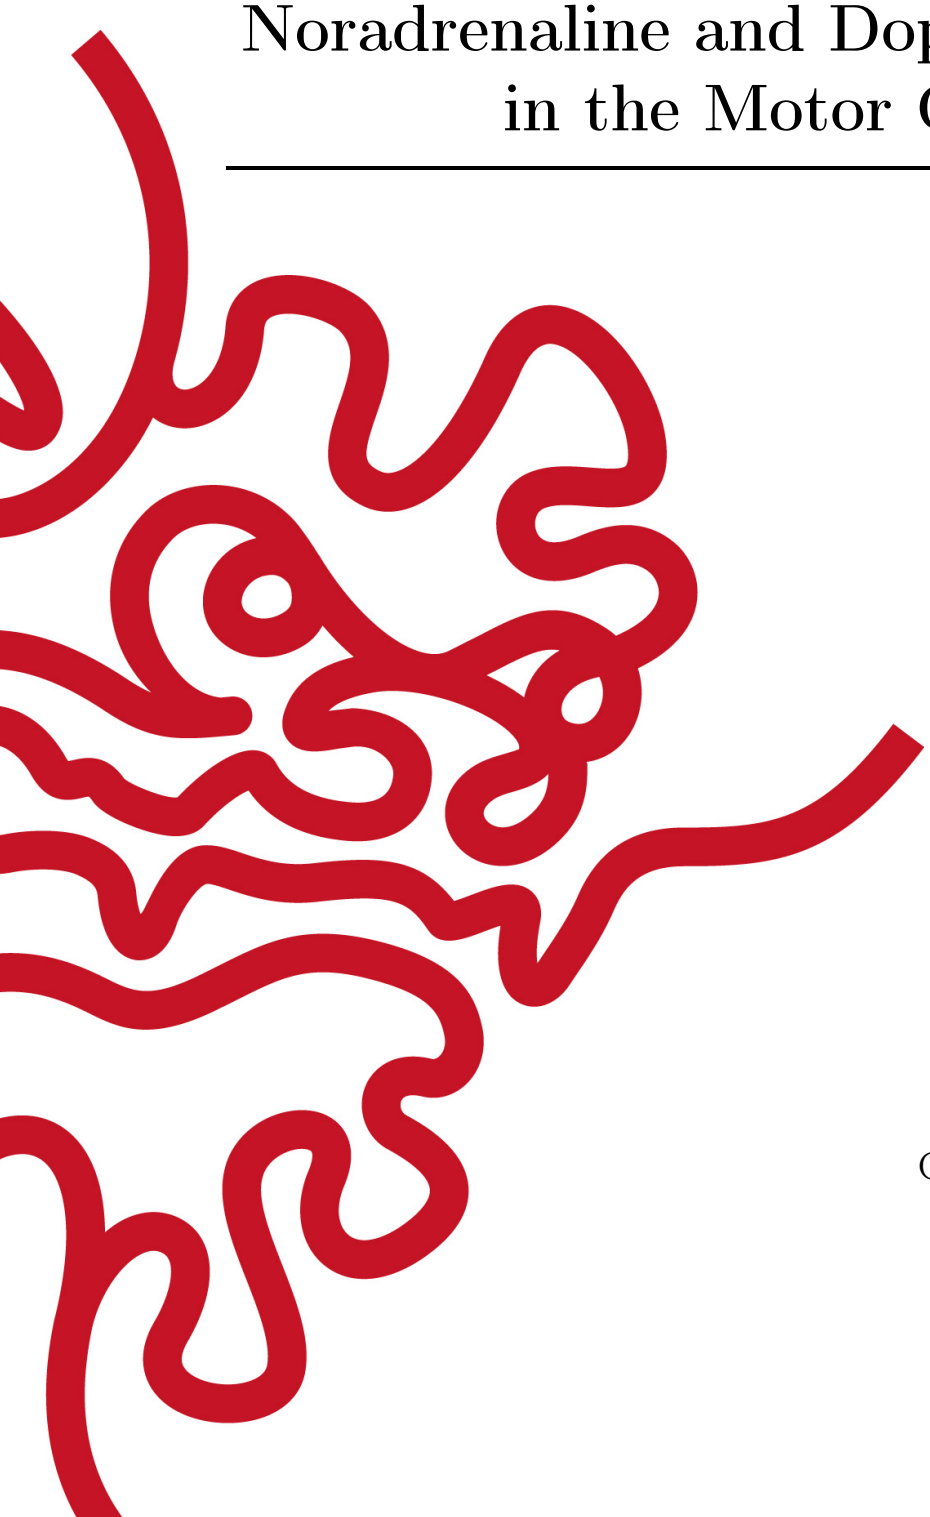

by

Miles Desforges

Supervisor: **Kenji Doya**
Co-Supervisor: **Bernd Kuhn**

April 2024

Declaration of Original and Sole Authorship

I, Miles Desforges, declare that this thesis entitled *Characterisation of Serotonin, Noradrenaline and Dopamine Release in the Motor Cortex* and the data presented in it are original and my own work.

I confirm that:

- No part of this work has previously been submitted for a degree at this or any other university.
- References to the work of others have been clearly acknowledged. Quotations from the work of others have been clearly indicated, and attributed to them.
- In cases where others have contributed to part of this work, such contribution has been clearly acknowledged and distinguished from my own work.
- None of this work has been previously published elsewhere.

Date: April 2024

Signature:

Abstract

The precise spatiotemporal dynamics of neuromodulator release in the cortex remain elusive. These release patterns determine receptor activation, as receptors with different affinities are activated by different concentrations, and thereby distances, from the release site. While neuromodulator nuclei have been extensively characterised, technological challenges have hindered detailed cortical investigation. This thesis clarifies the release patterns of dopamine, noradrenaline, and serotonin in the secondary motor cortex (M2) using a novel orthogonalised Go/NoGo task developed for head-fixed mice. We dissect neuromodulator activity at the intersection of locomotion and unconditioned stimuli (US). Employing two-photon microscopy and novel genetically encoded sensors, we achieved high-resolution imaging unattainable with previous techniques. Our findings indicate that serotonin, noradrenaline, and dopamine are robustly correlated with locomotion transitions, exhibiting significant increases from rest to locomotion and decreases during the reverse. Cross-correlation analysis suggests neuromodulator activity precedes locomotion onset, potentially facilitating motor behaviours. Unexpectedly, both appetitive (sucrose) and aversive (air-puff) stimuli elicited fluorescence decreases in all neuromodulators. Serotonin and noradrenaline showed stronger, more consistent responses during movement than rest, while dopamine responses were more consistent across locomotion states. Detailed analysis of activity maxima and minima within individual trials revealed significant variability and complexity. Contrary to expectations, no neuromodulator responses to US-predicting cues were observed, despite anticipatory behavioural changes. This highlights the need for further investigation into sensory cue processing in M2. Our findings suggest a complex multiplexing of information in M2, whereby neuromodulator activity is more influenced by locomotive state than appetitive or aversive stimuli delivery.

Acknowledgment

- Kazuo for his patience teaching me various techniques and keeping the lab in tip top condition.
- Philipp for all his help customising Flow Registration for my data.
- Leo for all his time saving tips and tricks.
- Shinobu for help with the pharmacology.
- All the lab members from Doya and Kuhn units for their help, inspiration, and patience over the years.
- Sumire for her incredibly persistent feedback in formatting.
- ChatGPT for help making the text clearer and more concise.

Abbreviations

5HT	Serotonin
AUC	Area under the Curve
DA	Dopamine
GENI	Genetically Encoded Neuromodulator Indicators
GPCR	G-protein coupled receptor
ITI	Inter-Trial Interval
ISI	Inter-Stimulus Interval
LC	Locus Coeruleus
M2	Secondary motor cortex
NA	Noradrenaline / Epinephrine
Neg-Go	Negative Reinforcement of No-Go action
Neg-NoGo	Negative Reinforcement of Go Action
NRI	Noradrenaline re-uptake inhibitor
OGNG	Orthogonalised Go/No-Go
Pos-Go	Positive Reinforcement of Go Action
Pos-NoGo	Positive Reinforcement of No-Go Action
PLC	Protein Lipase C
ROI	Region of Interest
SEM	Standard Error of the Mean
SNc	Substantia Nigra pars compacta
SVD	Single Value Decomposition
VTA	Ventral Tegmental Area

Contents

Declaration of Original and Sole Authorship ii

Abstract ii

Acknowledgment iv

Abbreviations v

Contents vi

List of Figures ix

List of Tables xi

1 Introduction 1

1.1 Motivation . . . 1
1.2 Neuromodulators . . . 2
1.2.1 What are neuromodulators? . . . 2
1.3 Noradrenaline . . . 7
1.3.1 Anatomy & Physiology . . . 8
1.3.2 Receptors . . . 8
1.3.3 Modulation of Behaviour . . . 10
1.3.4 Summary . . . 12
1.4 Dopamine . . . 12
1.4.1 Anatomy & Physiology . . . 13
1.4.2 Receptors . . . 13
1.4.3 Modulation of behaviour . . . 17
1.4.4 Summary . . . 20
1.5 Serotonin . . . 20
1.5.1 Anatomy & Physiology . . . 21
1.5.2 Receptors . . . 22
1.5.3 Modulation of Behaviour . . . 25
1.5.4 Summary . . . 29
1.6 Genetically Encoded Neuromodulator Indicators . . . 30
1.6.1 Noradrenaline biosensor . . . 30
1.6.2 Dopamine biosensor . . . 31
1.6.3 Serotonin biosensor . . . 32

1.7 Orthogonalised Go/NoGo Task . . . 34
1.7.1 Previous Orthogonalised Go/NoGo Tasks . . . 35
1.8 Open questions . . . 36
1.8.1 What is the spatial extent of volume transmission? . . . 36
1.8.2 Does *in vivo* neuromodulator release cause an increase or decrease in neuronal activity? . . . 37
1.8.3 How do neurones response to the mixture *in vivo* neuromodulator signals? . . . 37
1.8.4 What is the overall role of each neuromodulator? . . . 37

2 Methods **39**
2.1 Animals . . . 39
2.1.1 Animal ethics statement . . . 39
2.2 Animal habituation . . . 40
2.3 Cranial window surgery . . . 40
2.4 AAV viruses . . . 42
2.5 Two-photon imaging . . . 42
2.6 Behavioural setup . . . 43
2.7 Behavioural paradigm . . . 44
2.7.1 Main experiment . . . 44
2.8 Perfusion and Histology . . . 45
2.9 Image Processing . . . 45
2.9.1 Motion correction . . . 45
2.9.2 Neuromodulator hotspots . . . 46
2.9.3 Statistical analysis . . . 47

3 Results **50**
3.1 Behavioural results . . . 50
3.2 Imaging results . . . 55
3.2.1 Rest versus locomotion . . . 56
3.2.2 Locomotion transitions . . . 56
3.2.3 Unconditioned and Conditioned Stimuli . . . 62
3.2.4 Spatial release patterns . . . 75
3.3 Pharmacology . . . 77

4 Discussion **81**
4.1 Neuromodulator release strongly correlated with locomotion transitions 81
4.2 Distinct and nuanced neuromodulator responses to appetitive and aversive stimuli . . . 82
4.2.1 Responses to cues . . . 86
4.3 Spatiotemporal release dynamics . . . 87
4.4 Orthogonalised Go/NoGo task . . . 89
4.5 Pharmacology . . . 91
4.6 Methodological limitations . . . 91
4.7 Concluding thoughts on the open questions . . . 93

Conclusion **97**

Bibliography **99**

A Appendices **111**
A.1 Preliminary experiments 111

List of Figures

1.1 Figure from Doya, 2002 on neuromodulator nuclei 5
1.2 Figure from Avery et al, 2017 on neuromodulators disease 7
1.3 Figure from Arnsten et al, 2011 on performance and arousal for noradrenaline and dopamine 9
1.4 Figure from Radnikow et al, 2018 on layer-wise distribution of dopamine receptors 17
1.5 Figure from Floresco et al, 2013 on dopamine and working memory and set-shifting tasks 18
1.6 Figure from Cohen et al, 2020 on the role of serotonin 21
1.7 Figure from Pehrson et al, 2016 on serotonergic receptors 26
1.8 Figure from Patriarchi et al, 2018 on dLight 33

2.1 Injection site map 41
2.2 Behavioural setup graphic 44
2.3 Paradigm schedule 45
2.4 Hot-spot methodology 49

3.1 Task performance: logistic regression 51
3.2 Task performance: linear regression of mean locomotion velocity 52
3.3 Linear regression of velocity for each animal 53
3.4 Mixed-effects model of velocity changes over sessions 54
3.5 Dual GENI and GECI expression patterns 55
3.6 Neuromodulator activity time courses during sustained locomotion and rest 57
3.7 Analysis of GENI fluorescence during rest and locomotion 58
3.8 Neuromodulator activity time courses during locomotion transitions . . 58
3.9 Analysis of GENI fluorescence during locomotor transitions 59
3.10 Correlation and cross-correlation between neuromodulator, calcium and velocity during onset 60
3.11 Correlation and cross-correlation between neuromodulator, calcium and velocity during offset 61
3.12 Neuromodulator activity during cued and uncued unconditioned stimuli 63
3.13 Analysis of GENI activity for cued and uncued unconditioned stimuli . 64
3.14 Correlation and cross-correlation between neuromodulator, calcium and velocity during cued sucrose 65

3.15 Correlation and cross-correlation between neuromodulator, calcium and velocity during uncued sucrose . . . 66
3.16 Correlation and cross-correlation between neuromodulator, calcium and velocity during cued air-puff . . . 67
3.17 Correlation and cross-correlation between neuromodulator, calcium and velocity during uncued air-puff . . . 68
3.18 Neuromodulator activity time courses for unconditioned stimuli during locomotion or rest . . . 70
3.19 Analysis of GENI activity for unconditioned stimuli during locomotion or rest . . . 71
3.20 Correlation and cross-correlation between neuromodulator, calcium and velocity during sucrose while moving . . . 72
3.21 Correlation and cross-correlation between neuromodulator, calcium and velocity during air-puff while moving . . . 73
3.22 Analysis of peristimulus fluorescence in response to auditory cues . . . 74
3.23 Correlated neuromodulator activity as a function of distance . . . 77
3.24 GENI-specific pharmacological manipulations . . . 79
3.25 All pharmacological manipulations . . . 80

A.1 Pilot experiments paradigm schedule . . . 113

List of Tables

1.1 Table of noradrenaline, dopamine and serotonin receptors and effects . 4
1.2 Table from Subramaniam et al, 2016 of dopaminergic projections 14

2.1 Table of sound cues 43

A.1 Summary of logistic regression performance data 111

Chapter 1

Introduction

1.1 Motivation

The aim of this research was to investigate the release of neuromodulators in the cortex. It is well established that the neuromodulators serotonin, noradrenaline and dopamine are involved in almost every important aspect of behaviour. Moreover, most psychiatric diseases are linked to dysfunction in these systems [6]. However, much of what is known about neuromodulation during behaviour comes from recordings of deep neuromodulator nuclei rather than at their efferent targets. Yet, neuromodulator nuclei are heterogeneous in terms of cell type, response profile, and projection targets [15, 73, 96]. Different projections can bias distinct behaviours [21, 95]. Moreover, from somatic activity recordings we cannot be certain of the actual neuromodulator release amount, if any is released at all [69], nor what effect this has on cortical activity. For example, in some regions the majority of dopamine axonal varicosities do not appear active in response to depolarisation; cholinergic heteroreceptors can directly trigger dopamine release [69]; dopaminergic neurones can co-release glutamate [88]; noradrenergic neurones can co-release dopamine [93]; and the serotonergic dorsal raphe sends significant GABAergic projections [104]. The amount and constitution of neuromodulator release is important as target cells can express a variety of neuromodulator receptors, each with their own affinity and down-stream effects. At different extracellular concentrations specific receptors are activated which can produce distinct behaviours [4]. Ultimately, to understand the role of neuromodulation in cortex we will need direct measurements of released neuromodulator concentrations, in combination with the cortical neuronal responses.

One of the motivations for this project was to utilise the recently developed genetically-encoded neuromodulator indicators (GENI(s)) to directly measure released neuromodulator concentrations at the release site. With the advent of GENIs nanomolar fluctuations in extracellular neuromodulator concentrations have become detectable *in vivo* [123], and if combined with two-photon imaging, micrometer and millisecond precision. GENIs greatly improve on the spatio-temporal precision that was possible with the prevailing microdialysis and Fast Scan Cyclic Voltammetry (FSCV) methods; which are limited in spatiotemporal precision and unable to simultaneously record intracellular calcium and extracellular neuromodulator. The specifics of the project were decided

to build on the specialities of both of my co-supervisors, Bernd Kuhn and Kenji Doya, who are experts in two-photon microscopy and neuromodulation, respectively. I was particularly interested in the combinatorial effects of the neuromodulators serotonin, dopamine and noradrenaline in reinforcement learning proposed by Doya [31, 32]. I chose to investigate these neuromodulators in the secondary motor cortex (M2), an area directly imageable by two-photon microscopy and also implicated in reinforcement learning [109]. I hoped to shed light on some major open questions about neuromodulator dynamics in the cortex: firstly, the scale of neuromodulator release, from cellular to whole cortical area; and secondly, how spontaneous *in vivo* release affects the activity of nearby neurones, as indicated by fluctuations intracellular calcium activity. Surprisingly, these two fundamental questions are still relatively unknown.

1.2 Neuromodulators

Interest in neuromodulators has spread beyond the research sphere into the public one. In particular, dopamine and serotonin are known for their roles in diseases, such as Parkinson's and depression, as well as recreational drugs such as Ecstasy. Phrases such as 'dopamine rush' and 'serotonin sucker' have entered the common parlance. However, much remains unknown of the multifarious ways in which these molecules influence and regulate our behaviour.

1.2.1 What are neuromodulators?

Neuromodulators are signalling molecules which, when detected by receptors on neurones, enact diverse effects including altering neuronal membrane potential, presynaptic release probability, postsynaptic detection, and synaptic plasticity [100]. Neuromodulation works across multiple spatial and temporal timescales, from the subcellular to the network, from milliseconds to days. These effects often occur without directly exciting or inhibiting the neurones, a key distinction with neurotransmitters, but through changing the balance of active ion channels by adjusting channel kinetics and permeability. The three neuromodulators investigated in this thesis are some of the most well known of the roughly one hundred molecules which fit the neuromodulator definition [47]. Each are involved in most, if not all, cognitive processes, including attention, emotion, decision-making, learning, social-interactions and goal-directed behaviour. Each system shows increasing activity from quiescence during slow-wave sleep to dynamic responses in task-engaged alertness.

Neuromodulators versus Neurotransmitters

It will be useful to review the role and characteristics of neuromodulators through comparison with neurotransmitters. Neurotransmission is the standard route of conveying information from one neurone to another, through fast, direct, synaptic communication. Sensory and motor information is generally conveyed directly by fast neurotransmitters, such as glutamate or GABA, within a millisecond. Fast transmission is attributable to opening of ionotropic receptors, ligand-gated ion channels which directly alter the

ionic gradient at the post-synaptic membrane. In contrast, the definition of neuromodulators, as with their effects, is much broader and more complex. Neuromodulators, which can be biogenic amines, peptides or amino acids, predominantly act through slow transmission via metabotropic G-protein coupled receptors (GPCR) receptors.

While fast neurotransmission permits rapid, stable and consistent information transfer, connections are fairly rigid and plastic changes are relatively slow. In contrast, neuromodulation allows for dynamic shifts between different network ensembles. Fast transmission may be akin to roads in a city, robustly allowing information to flow, while neuromodulators subtlety change the network dynamics by shifting the timings of traffic lights.

Synaptic versus volume transmission

To allow rapid communication, neurotransmitter receptors are clustered at active zones on synapses. These, predominantly ionotropic, receptors can be trafficked away from the active zones to reduce the strength of the connection. Due to rapid neurotransmitter clearance only axon boutons and dendritic spines that have a direct synaptic connection are considered active. However, there are also extrasynaptic neurotransmitter receptors, predominantly GPCR, which can monitor synaptic spillover after excessive activity, or release from other sources such as neuromodulator axons [88]. Whereas, neuromodulator axons rarely form direct connections. The large number of asynaptic neuromodulator varicosities, and extrasynaptic receptors on neurones and glia lead to the idea of volume transmission, as an accompanying mechanism to synaptic transmission [44]. Volume transmission is the extracellular diffusion chemical signals across distances much greater than a synapse. Through volume transmission neuromodulators can influence a large number of nearby neuronal and astrocyte networks. However, neuromodulator release events have also been shown at a neuronal and even compartment scale [60], leading to a debate on the spatial extent of action. Volume transmission adds a non-linearity to networks based on spatial proximity rather than direct synaptic contact. This may create micro-environments depending on different neuromodulator innervation and receptor patterns and could itself lead to the intrinsic development of to cortical modular specialisations. What is unknown, and what this thesis hopes to shed light on, is the spatial scope of neuromodulator release.

Receptors

Neuromodulators in general have multiple families of different receptors, with each receptor setting off a variety of downstream cascades resulting in a myriad of distinct effects. Neuromodulator receptors are predominantly GPCR. These receptors do not possess an ion pore but instead influence intracellular activity by releasing α and β/γcomplex subunits, which act as second messengers that act across the whole cell or within small cellular compartments. There are a variety of GPCR families, including G_s, $G_{i/o}$ or $G_{q/11}$, each with distinct downstream effects. This indirect mechanism greatly increases the time-delay of action onset, lasting from milliseconds to minutes, but also increases the variety of effects. A single neuromodulator system may have receptors for each family. Table 1.1 summarises the receptors, their downstream ef-

Receptor class	Cortical Receptors	Coupled	Downstream	Excitability Effect	Localisation
NA α1-family	α1A, α1B, α1D Low affinity	G_q	↑PLC, ↑IP_3, ↑Ca^{2+},↑K_{Ca}, ↑PKC, ↓AMPA, ↓NMDA, ↑LTD	Pre & Post: Suppress activity	Pre- & Post-synaptic, α1A: GABAergic > pyramidal (L6 > L5 > L2/3) α1D: GABAergic >> pyramidal (L2/3 > L5)
NA α2-family	α2A, α2B, α2C High affinity	G_i	↓AC, ↓cAMP, ↓PKA, ↓HCN channels, ↓AMPA, ↓NMDA	Pre: ↓glut release Post: ↓weaker synapses, ↓spontaneous, ↑persistent firing	Pre-synaptic autoreceptor Post-synaptic, spines, dendrites, soma pyramidal (L2/3, L5/6)
NA β-family	β1, β2, β3 Lowest affinity	G_s	↑AC, ↑cAMP, ↑PKA, ↑HCN, ↓K- channels (sAHP, TREK-2), ↑AMPA, ↑NMDA, ↑LTP	Pre: ↑glut release Post: ↑spontaneous activity	Pre-synaptic Post-synaptic (L2/3, L4, L5/6)
DA D1-like	D1, D5 Low affinity	G_s / G_q	↑AC, ↑cAMP, ↑PKA, ↓K-channels (sAHP, K_{ir}), ↑Na channels, ↓N- type Ca-channels, / ↑PLC, ↑IP_3, ↑Ca^{2+},↑K_{Ca}	↓excitability L2/3 & L5 ↑max firing L2/3	Post-synaptic, pyramidal, D1: dendritic spines, D5: soma. Pyramidal > GABAergic (L6 > L5 > L2/3) Pyramidal L4
DA D2-like	D2, D3, D4 High affinity	G_i	↓AC, ↓cAMP, ↓PKA, ↑K-channels (sAHP, K_{ir}), ↓N- P/Q- R- type Ca-channels, ↑β-arrestin	↓excitability L2/3 & L5	Pre- & Post-synaptic D2: pyramidal > GABAergic (L5 > L6 > L2/3) D3: L5 pyramidal subtype D4: GABAergic > pyramidal (L2/3 > L4 > L5/6)
$5HT_1$	5HT1A, 5HT1B High affinity	G_i	↑AC, ↑cAMP, ↑PKA, ↓NMDA,↑Kir, ↓VGCC, ↓NaV1.2	↓excitability L2/3 & L5	$5HT_{1A}$: Somatodendritic. L1 pyramidal apical dendrite, L2/3, L5/6, pyramidal soma, L2/3, L5/6 FS soma $5HT_{1B}$: presynaptic auto/hetero receptor. L5/6 pyramidal soma
$5HT_2$	5HT2A, 5HT2C Moderate affinity	G_q	↑PLC, ↑IP_3, ↑Ca^{2+}, ↑PKC, ↓K-leak, ↓Kir, ↑NC, ↑Kca, ↓AMPA, ↑LTD	↑Excitability L2/3 & L5	$5HT_{2C}$: L1 GABAergic, L5/6 pyramidal soma, L5/6 FS soma $5HT_{2A}$: L2/3, L5, pyramidal soma, L2/3, L5/6 FS soma
$5HT_3$	5HT3 Low affinity	Iono-tropic	↑NC	Depolarise L1,2,3. Rapidly activating & desensitising.	L1, L2, L3 SS soma L5/6 non-PV-FS soma
$5HT_5$	5HT5A,B Low affinity	G_i	↑Kir	Inhibitory	L5 pyramidal
$5HT_{4,6,7}$	5HT6,7 Moderate affinity	G_s	↑AC, ↑cAMP, ↑PKA, ↓K-leak, ↓Kca,↑NC	Excitatory	$5HT_4$: L5/6 pyramidal soma, $5HT_6$: L1 pyramidal apical dendrite, L1 SS soma, L2/3 pyramidal soma, L5/6 pyramidal apical dendrite $5HT_7$: L1,L2/3, L5/6 pyramidal soma

Table 1.1: Noradrenaline, dopamine and serotonin receptors and effects. *Kir: inwardly rectifying potassium channel, Kca: calcium-activated potassium channel, HCN: hyperpolarisation-activated cyclic nucleotide–gated channel, NaV1.2: voltage gated sodium channel, NC: non-selective cation channel, sAHP: slow afterhyperpolarisation, TREK-2: two-pore potassium channel, VGCC: voltage-gated calcium channel*

fects and localisation for noradrenaline, dopamine and serotonin. It should be noted that the three systems often influence overlapping intracellular cascades, wherein the overall effect is determined by the aggregation of cascades, and the somatic location of receptors. For example, activation of noradrenergic $\alpha 1$ and serotonergic 2A receptors both increase PLC activity, yet their effects are inhibitory and excitatory, respectively [4, 7]. It is these nuances that make neuromodulators so intriguing, but also difficult to study.

Global or Modular

Neuromodulators predominantly originate from small nuclei in the brainstem and basal forebrain and are released throughout the central nervous system, Figure 1.1. However, the density of axonal projections differs depending on the cortical region and species, which has functional implications. For example, Chandler and Waterhouse [20] traced cholinergic and noradrenergic projections to the medial-prefrontal, anterior cingulate and orbitofrontal cortices. They found more than half of the cholinergic projections went to two or more regions, whereas $< 5\%$ of noradrenergic projections went to two regions and none went to three. Furthermore, cholinergic afferents enter the cortex through a superficial layer 1 route and an ascending route from the deep layers, whereas noradrenergic projections only have the ascending route. The anatomy implies the noradrenergic system can differentially modulate specific regions, as found

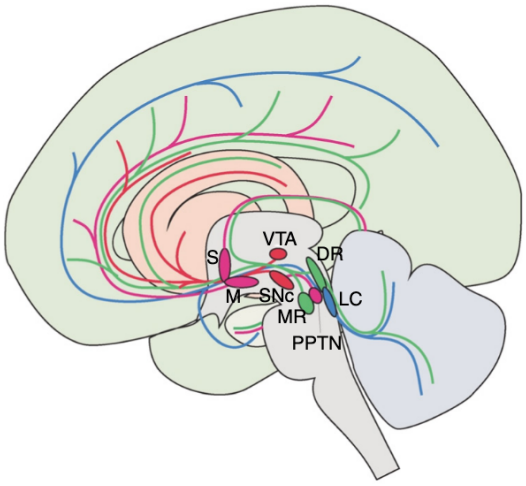

Figure 1.1: Figure from [31] depicting the basal nuclei and efferents of the 'prototypical' neuromodulator. *Magenta: Cholinergic (M: Meynert nucleus, S: medial septum, PPTN: pedunculopontine tegmental nucleus); Blue: Noradrenergic (LC: locus coeruleus); Red: Dopaminergic(VTA: ventral tegmental area, SNc: substantia nigra pars compacta); Green: Serotonergic (DR: dorsal raphe, MR: median raphe)*

by Breton-Provencher et al. [15], while the cholinergic system is more regionally homogeneous, as found by Teles-Grilo Ruivo et al. [113], but may exert different layer-wise control. The widespread, yet heterogeneous, projection sites of neuromodulator may also accounts for the multitude of different diseases associated with atypical neuromodulator functioning, see Fig 1.2.

Tonic and phasic activity

Neuromodulatory systems often display both tonic and phasic activity [72]. Tonic activity changes have been found in noradrenaline during stress [38], and dopamine and serotonin in value processing [24] and acetylcholine in arousal [113]. Tonic activity is also often associated with slow time-course changes such as sleep-wake states. Noradrenaline, serotonin and histamine agonists have been shown to enhance wakefulness and antagonists increase sleep [65]. Interestingly, while the monoaminergic neuromodulators dopamine, serotonin and noradrenaline cease firing during REM sleep [4, 30, 123], cholinergic neurones are active during both wakefulness and REM sleep [113], suggesting acetylcholine produces distinct type of arousal. Neuromodulator tonic activity may also be important for attention through desynchronisation of cortical oscillatory activity [54].

Phasic neuromodulatory activity is generally linked to more dynamic task-related variables. Dopaminergic and serotonergic systems have been shown to respond to appetitive and aversive stimuli and prediction errors [131]. While noradrenergic phasic activity occurs during unexpected cue or state changes and for self-initiated movements [127]. Phasic responses relating to locomotion, conditioned and unconditioned stimuli are investigated in this thesis.

Interestingly, tonic and phasic activity patterns may also alter the composition neuromodulator vesicular release. Neuromodulatory neurones have been found to contain both fast neurotransmitters and neuromodulators. For example, cholinergic neurones can co-release GABA [103] and dopaminergic neurones can release glutamate [88]. Fast neurotransmitters are contained in smaller vesicles, which are more readily released by single action-potentials, whereas, neuromodulators are contained in larger vesicles, which often need volleys of activity such as phasic bursts to be released.

Signal-to-noise

In the cortex, neuromodulators are believed to influence the signal-to-noise ratio of neural activity, enhancing behaviourally relevant sensory or motor signals by suppressing less currently relevant ongoing activity [72, 114]. However, which signal is being accentuated will vary by system, state of that system and stimuli being attended. For example, increasing the gain of a sound may alter its salience and draw attention to it. In motor areas silencing competing ensembles may allow one action to stand out and be initiated, a proposed role for noradrenaline by Aston-Jones and Cohen [5]. In working memory tasks noradrenaline acts to boost signal gain by increasing the firing of neurones which prefer the correct choice, while dopamine suppresses activity of incorrect choices [4]. And in dysfunction changes in sensory signal strength may produce hallucinations [75] or produce distractibility [102]. This diverse modulation of signal-to-noise ratios, affecting both bottom-up and top-down processes, could underlie the involvement of neuromodulator dysfunctions in a wide range of neurological disorders (Figure 1.2).

In the cortex neuromodulators are thought to also influence the signal-to-noise ratio of activity [72, 114]. Altering the signal-to-noise allows certain signals to stand out above the other ongoing activity. Which signal is being accentuated will vary by system, state of that system and stimuli being attended. For example, increasing the gain of a sound may alter its salience and draw attention to it. In motor areas silencing competing ensembles may allow one action to stand out and be initiated, a proposed role for noradrenaline by Aston-Jones and Cohen [5]. In working memory tasks noradrenaline acts to boost signal gain by increasing the firing of neurones which prefer the correct choice, while dopamine suppresses activity of incorrect choices [4]. And in dysfunction changes in sensory signal strength may produce hallucinations [75] or produce distractibility [102]. The myriad of roles modulating the signal-to-noise ratio of both bottom-up and top-down activity may explain why dysfunction of neuromodulators is implicated in the majority of neurological diseases (Figure 1.2).

Neuromodulators influence learning

Neuromodulatory responses to the salient stimuli, the cues which predict them and the actions performed to obtain them play an indispensable role in learning. Importantly, in the cortex both appetitive and aversive unexpected salient stimuli are associated with phasic bursts of activity in cholinergic [67, 113], noradrenergic [5], [121], serotonergic and dopaminergic [24] systems; although each system has different activity patterns for each type of stimuli. For example, in well trained mice the

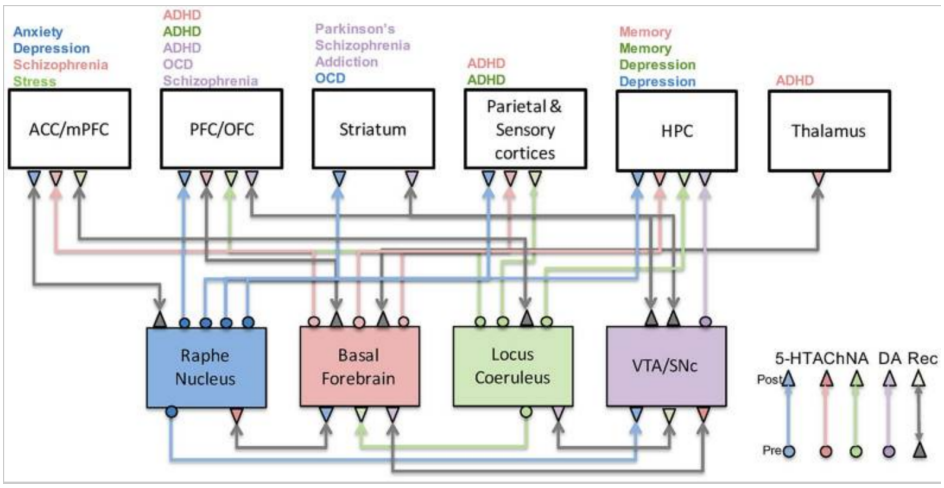

Figure 1.2: Figure from Avery and Krichmar [6] depicting neuromodulators known to be involved in various diseases. This diagram also highlights the complexity of the interactions between the neuromodulator nuclei and their targets. *5-HT: serotonin, ACh: acetylcholine; NA: noradrenaline; DA: dopamine; Rec: recurrent connections; ACC: anterior cingulate cortex;(m)PFC: (medial) prefrontal cortex; OFC: orbitofrontal cortex; HPC: hippocampus; VTA: ventral tegmental area; SNc: substantia nigra pars compacta*

cholinergic system shows phasic responses at key points in a maze relating to decision points and rewards [113], whereas the noradrenergic system signals unexpected cues and self-initiated movements [127], but not for reward itself [38]. These fluctuations in release can significantly influence behaviour. For example, phasic dopaminergic responses increases the signal-to-noise of a subset of cortical pyramidal neurones in response to aversive stimuli, and bias avoidance over approach responses [119]. These studies suggest cortical neuromodulator dynamics play a key role in both Pavlovian and instrumental learning.

1.3 Noradrenaline

Noradrenaline plays an important role in alertness, arousal, working memory, and stress responses. Drugs targeting the noradrenergic system are often used to treat stress, ADHD and depression [6]. The overall role of cortical noradrenaline has received various hypothesis including to: increase neural gain [5]; reset cortical networks [14]; signify unexpected uncertainty [129]; or state-value prediction errors [98]; and switch between exploration and exploitation behaviours [5], akin to the inverse temperature component of reinforcement learning models [31]. A key theme running throughout is the role in initiating task related actions and switching action strategies. Despite this, the majority of the studies have focused on the prefrontal cortex rather than M2.

1.3.1 Anatomy & Physiology

Cortical noradrenaline originates almost exclusively from the locus coeruleus (LC) in the brainstem, from which projections extend to the virtually the entire central nervous system including all of the neocortex, although interestingly not to the basal ganglia. Tracing studies show that noradrenergic innervation of the cortex is largely diffuse and homogeneous. Nomura et al. [82] found lower density of noradrenergic fibres in frontal association area than the parietal or occipital cortices, however, also a higher density in the cingulate cortex. Expression was fairly consistent across all layers of the cortex, although with a peak in the superficial layers of the cingulate, somato-sensory and visual areas. Cerpa et al. [19] note a subtle rostro-caudal gradient in the medial and orbitofrontal areas of the frontal cortex. The orbitofrontal and medial areas may be involved in evaluating rewards and costs, and both areas provide direct and indirect inputs to the LC, creating a positive feedback loop. Prefrontal regulation of LC is thought to drive transitions from tonic to phasic activity.

The LC displays distinct firing modes: virtually silent in REM sleep; low firing in sleep and drowsiness; phasic mode with a moderate baseline and task-relevant bursts of firing; and a tonic mode with elevated baseline but absent of bursting. Phasic activity, in particular in the PFC, is associated with external or novel salient stimuli and task performance. Phasic responses to cues rapidly habituate in the absence of reinforcement and in well-predicted situations. However, when the state changes, such as a new contingency rule or unexpected reward or omission, the phasic activity returns. Bouret and Sara [14] found these phasic noradrenergic responses occur before behavioural adaption in the PFC and suggested a role in facilitating behavioural flexibility. In contrast, increased tonic activity is associated with distractibility and stress. Arnsten [4] describes how recruitment of different noradrenergic receptors produces the inverted-U of task performance. Too little noradrenaline does not engage any noradrenergic receptors, the optimal recruits $\alpha 2$ receptors, and too much activates $\alpha 1$ receptors, Figure 1.3.

1.3.2 Receptors

In the cortex noradrenaline has three receptor types: high affinity G_i coupled $\alpha 2$-family; and lower affinity G_q coupled $\alpha 1$-family and even lower affinity G_s β-family.

$\alpha 2$

$\alpha 2$ receptors are synaptic and localised more to dendrites. $\alpha 2$ activation can suppress the strength of individual inputs while increasing responsiveness to coordinated activity. $\alpha 2$ activation can also promote persistent firing by inhibiting cAMP production, reducing cAMP-mediated opening of hyperpolarisation-activated cyclic nucleotide-gated (HCN) channels on dendritic spines. $\alpha 2$ also has pre-synaptic effects. In cell cultures of sensorimotor neurones it was shown $\alpha 2$ activation also inhibits glutamate release from cortical nerve terminals through inhibition of N-type and P/Q-type calcium channels. $\alpha 2_A$ agonist enhances firing of neurones in the preferred stimuli orientation, while antagonists reduce or silenced responses [4]. In terms of behaviour, depletion

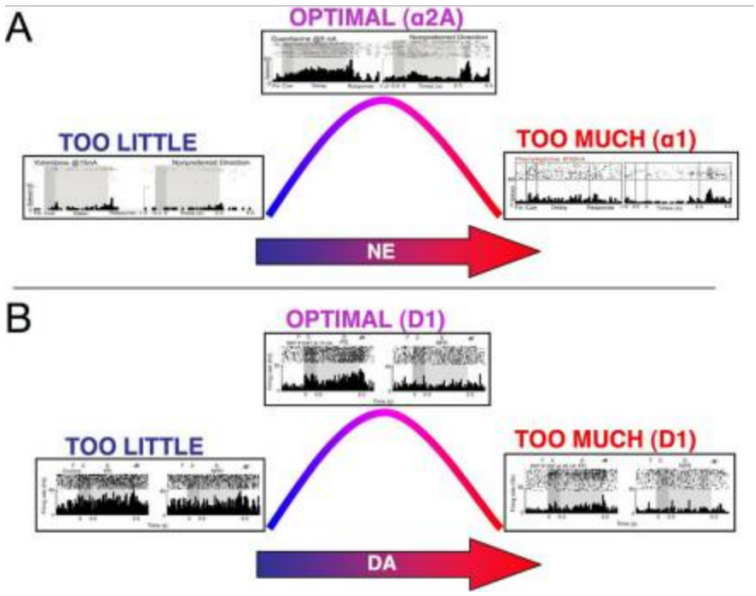

Figure 1.3: Figure from Arnsten [4] illustrating the Yerkes-Dodson relationship of between performance and arousal for noradrenaline (NE) and dopamine (DA)

of noradrenaline, or $\alpha 2$ receptor antagonism, impairs working memory while agonism improves performance. Arnsten [4] list improvements associated with $\alpha 2$ activity including working memory, attention, response flexibility, reversal learning and action withholding (NoGo) performance; and these effects were not seen in $\alpha 1$ or β receptor agonism.

$\alpha 1$

$\alpha 1$ receptors are present on cortical spines, but are often perisynaptic or extrasynaptic [27]. $\alpha 1$ receptors have a low affinity, activation requires much higher concentrations of noradrenaline in the μM range. Activation suppresses activity, most likely as a safeguard mechanism when excessive noradrenaline has been released to the point of spilling over in high concentrations to the extrasynaptic sites. $\alpha 1$ agonists can both enhance LTD and suppress LTP in pyramidal cells, as well as suppress AMPA mediated transmission. $\alpha 1$ receptors are also present on a variety of GABAergic interneurones and suppress activity. Datta et al. [27] found $\alpha 1$ agonism decreased activity in relation to neurones preferred stimuli direction, but not for unpreferred direction. Doze et al. [33] found $\alpha 1$ receptor knockout mice displayed poor cognitive function.

β

β receptors are found at axo-axonic junctions with noradrenergic neurones [2]. $\beta 1$ and $\beta 2$ receptors are more strongly expressed in L4, but also in other layers. β receptors have the lowest affinity. β receptors potentiate excitatory propagation through enhancing increasing AMPA and NMDA currents, increasing LTP and suppressing LTD. β receptors reverse the effects of $\alpha 2$ activation by increasing cAMP production. Presynaptic activation of β receptors increases GABA and glutamate release. Grzelka

et al. [48] found L5 mPFC pyramidal cells were directly depolarised by application of a high concentration of noradrenaline , an effect mediated by β receptor activation of HCN channels. β receptors are also present on GABAergic interneurones and enhance activity.

In the field it does not appear there is an overall model of noradrenergic receptor activation. Speculatively, $\alpha 2$ and β receptors may work together to enhance task-relevant ensembles. In the phasic mode there is a moderate baseline of activity, which Arnsten [4] shows activates $\alpha 2$ receptors and increases the neural gain. In the process AMPA and NMDA currents and trafficking are reduced, promoting LTD in weaker strength inputs. In contrast, it can be expected that burst firing elicits brief localised high concentrations of noradrenaline, enough to recruit β receptor activation, promoting LTP in task-relevant ensembles. In this way, $\alpha 2$ receptors may weaken task-irrelevant, noisy activity, while phasic β receptor activation may strengthen task relevant activity. To keep the system in check, $\alpha 1$ receptors would be activated by excessive release and spillover, likely during stress.

1.3.3 Modulation of Behaviour

The role of noradrenaline in behaviour in behaviour often comes under the broad categories of attention and task performance [5], novelty [52], working memory [4], enhanced sensory cues [121], and descending analgesia [21]. Most of what we know of the role of noradrenaline in behaviour comes from measurements at the LC or PFC, although there is emerging evidence for nuances in the signals between the PFC and M2 [15, 21]. Here let us explore a few important papers which highlight the role of noradrenaline in motor actions and reinforcement learning.

Aston-Jones and Cohen [5] describe how LC neurones fire at a moderate rate during accurate task performance, with brief phasic responses after target stimuli and rewards, but not after irrelevant stimuli. They also note that in instrumental tasks LC activity was more tightly locked to behavioural responses, rather than stimuli presentation. LC activity preceded motor responses irrespective of whether the trial was correct or incorrect. However, LC activation did not occur for spurious lever presses in the inter-trial interval when no stimuli were presented. It therefore appears that LC phasic responses are coupled to task-relevant behaviour. Bouret and Sara [14] also mention how, in rats and monkeys, LC activity shows a phasic response to conditioned cues, but that activation was more tightly aligned to the subsequent behavioural response than cue onset. Bouret and Sara [13] showed that self-initiated actions do not elicit phasic responses in un-cued situations, and conclude that in cued situations the phasic response reflects the recognition of an awaited stimulus. Both Bouret and Sara [14] and Aston-Jones and Cohen [5] also found phasic LC responses when contingencies abruptly changed, such as action-outcome reversal.

Varazzani et al. [120] directly compared the firing rates of noradrenergic LC and dopaminergic SNc neurones in monkeys during a positively reinforced grip task, where reward size and physical effort requirements were systematically varied. They found phasic responses from both nuclei were aligned with cue and action periods. However, only LC activity at the action varied with physical effort requirements. While SNc activity at the cue increased with expected reward and decreased with signalled physical

effort requirements. The focus here on the effort of actions actions is relevant to the behavioural task in this research, which utilises an Orthogonalised Go/No-Go (OGNG) task (see 1.7). The OGNG investigates the effect of positive and negative reinforcement on active "Go" actions, such as running (here) or button pressing (original [50]), and inactive "NoGo" actions, resting or withholding. Effortful Go and less effortful NoGo actions are an essential part of Go-NoGo tasks, however, the majority of paradigms reinforce either the Go or NoGo action. For example, while Bouret and Sara [13] describe their paradigm as a Go-NoGo task, the NoGo action is never reinforced, and therefore it could be interpreted as a Go task with distractors. The OGNG task employed in this research seeks to decouple the action from reinforcement, allowing for a more nuanced exploration of noradrenergic influence on locomotor reinforcement.

Xiang et al. [127] recorded from LC neurones while rats performed a self-paced T-maze task with a rule switch. They found that task-related visual cues did not illicit a phasic response in well-trained rats, although in untrained rats LC activity increased for the reward delivery cues. This may be because visual cues were self-activated by rats passing a photodetector, and so would not be unexpected nor elicit novelty responses. In support of this role signifying novelty they did see an increase in firing to cues in the first trial after a state change, but not otherwise. In contrast to the lack of responses to well-expected cues, they did see consistent increase in firing just before physically crossing the photodetector which initiated the trial. The activity onset occurred 35 ms before locomotor acceleration, a latency too short to be initiating motor responses, and was suggested to relate to mobilising resources and systemic arousal.

Uematsu et al. [118] found that projections to the amygdala and PFC convey different information. While both projections were activated by aversive foot-shocks, the amygdala projections also responded to shock-predicting cues, whereas PFC projections showed responses to extinction cues. Interestingly, inhibiting LC output at neuronal terminals had a greater effect on fear-related behaviours than indiscriminate inhibition of across the LC. If the LC can activate circuits with opposing behavioural effects, this suggests a modular organisation. Moreover, this may confound interpreting behavioural results from experiments which activate the LC en masse.

Feng et al. [38] performed a variety of behavioural tests in freely moving mice while measuring noradrenaline release in the lateral hypothalamus through fibre photometry. They found stressful events such as forced swimming and tail suspension induced a sustained ramp up of noradrenaline which returned back to baseline during rest. Interestingly, the authors also found simply presenting a human hand in the cage resulted in transient increases in fluorescence, and effect not seen with interactions with conspecific intruders of same or opposite sex. Whereas, contact with rewarding food stimuli did not result in changes fluorescence. They concluded that noradrenergic activity in the lateral hypothalamus increases in relation to stressful, but not the non-stressful, stimuli tested.

In an approach complementary to ours, Breton-Provencher et al. [15] used two-photon microscopy to record noradrenergic axons and released noradrenaline in the motor or prefrontal cortices, all while mice performed a Go-NoGo task. Their approach differed from ours by using positive reinforcement and punishment in level-press task; reinforcing 'hit' and punishing 'false alarm' Go actions. They found sharp increases in activity before both hit and false-alarm lever-presses, with a greater response in

the motor cortex than PFC. They showed a strong further increase in activity to the air-puff following the false-alarm lever press. Whereas, they seemed to see a decrease in activity after liquid delivery, although this was not discussed. Axon calcium was correlated with $GRAB_{NE}$ signals. Additionally, recording from soma in the LC, they found both global and modular signals. For aversive events, all neurones responded to aversive air-puffs, but for appetitive rewards only a subset responded.

1.3.4 Summary

The noradrenergic system is involved in arousal and attention. Noradrenergic tone shows a Yerkes-Dodson relationship, with sleep and drowsiness at low concentrations, phasic mode at intermediate concentrations, and tonic mode at high concentrations [4]. In phasic mode a low baseline of activity is interrupted by bursts of activity for novel or salient stimuli or salient task-related actions, and is associated with high task performance. When there is a tonic increase in baseline activity there is a decrease in phasic bursts allowing flexible responses, but also stress if persistent. The phasic mode is commonly associated with $\alpha 2$ receptor activation. As hypothesised here, $\alpha 2$ receptors may be recruited by a phasic response during the low baseline component of the phasic mode, while β receptors are activated by phasic bursts. In the tonic mode, $\alpha 1$ receptors are activated by elevated background release, broadening excitability, but reducing bursting activity. Aston-Jones and Cohen [5], Bouret and Sara [14], Varazzani et al. [120] all found increased phasic noradrenergic activity occurred during task-related actions, particularly those requiring effort. Whereas, for cues phasic activity appears related to novelty or surprise such as unexpected stimuli or changes in the state. Noradrenergic responses in M2 are not as well established as in the PFC, however, Breton-Provencher et al. [15] showed that noradrenergic boutons in M2 encoded actions with greater amplitude than the PFC, while regions both showed similar responses aversive stimuli. What is unknown is the scope of noradrenergic responses, across the whole cortical region or more localised. Moreover, whether these responses generally increase or decrease nearby calcium activity *in vivo*, through activation of alpha1 or alpha2 receptors.

1.4 Dopamine

Dopamine is most well known for its role in the basal ganglia conveying RPE [63] and saliency [94], and involvement in diseases such as Parkinson's, schizophrenia, addiction and ADHD [6]. The role of dopamine in the cortex is much less studied, yet it plays an important role in a variety of tasks including working memory [4], attention to cues [119], set-shifting [39] cost-benefit decisions. Importantly, dopamine release in the cortex has been shown in response to appetitive and aversive stimuli [90]. Dopamine has a generally inhibitory effect in the cortex by suppressing subthreshold activity, however can also boost suprathreshold firing [4, 64, 111]. Dopamine has been shown to increase the signal-to-noise of aversive stimuli and bias aversive responses when challenged with appetitive and aversive predicting cues [119]. Dopaminergic neurones typically fire at 4 Hz in tonic mode, and emit bursts of action potentials at 15 Hz in response to

salient stimuli. Interestingly, while dopamine exerts a Yerkes-Dodson relationship on performance in working memory tasks [4], Floresco [39] did not find this in set-shifting tasks (Figure 1.5). The scarcity of studies on the role of cortical dopamine, along with the complexity of available results has so far obfuscated a complete picture.

1.4.1 Anatomy & Physiology

Forebrain dopamine projections originate in the tegmental area, which can be divided into the ventral tegmental area (VTA) or the substantia nigra pars compacta (SNc). Recent tracing studies highlight the diversity of dopaminergic neurones within these classical divisions. Subramaniam and Roeper [108] summarises the differences between mesostriatal, mesolimbic and mesocortical projections in Table 1.2. Dopamine afferents from the VTA differentially project to cortical layers in an area-specific manner. Hosp et al. [57] found VTA projections to M1 and PFC come from topographically distinct neuronal groups, with minimal colateralisation or contralateral projections. While projections to PFC and nucleus accumbens originated from intermingled groups, cells projecting to M1 are located more dorso-rostrally in the VTA. Nomura et al. [82] found dopaminergic afferents terminate in all layers in the medial frontal cortical areas, such as prelimbic and cingulate cortices, but only deeply in L5-6 for in the more lateral sensory and frontal association cortices; M2 lies in between the medial and lateral areas, but was not investigated by Nomura et al. [82]. Functional imaging with a PKA-biosensor showed dopamine or D1 agonists produced robust increases in PKA activity. While noradrenergic β receptor agonism elicited reactions in 97% of cells, D1 agonism only resulted in PKA fluctuations in 75% of cells. Dopamine induced responses were also almost twenty times smaller in amplitude than those induced by noradrenaline.

Dopaminergic nuclei receive substantial input from the deeper brain structures, and only a little from the frontal and sensorimotor cortices [125]. The calculation of RPE signal appears to be calculated in a distributive fashion across various basal areas, with few neurones conveying a pure RPE signal. Whereas, the characteristic pause in activity during unexpected reward omission disappeared with lesions to the ventral striatum or lateral habenula. Verharen et al. [122] found projections from the VTA to the accumbens altered RPE based behaviours leaving animals insensitive to loss and impaired in reversal learning. Manipulating mesocortical projections spared RPE based behaviours, but did increase lose-stay behaviour, although this was seen in both mesoaccumbal and mesocortical manipulations.

1.4.2 Receptors

Dopaminergic receptors can be divided into two groups, D1 and D5 (D1-class), are mainly coupled to G_s, while D2, D3, and D4 (D2-class) are coupled to $G_{i/o}$. D2-class receptors have 10 to 100 times higher affinity for dopamine than D1-class, with D3 and D4 showing the highest affinities. D1 and D2 receptors are most prevalent in the cortex and thus are most commonly studied, although the importance of other dopaminergic receptors is emerging.

Subtype	Cell body position	Gene expression for maintenance	Axonal projection	Synaptic inputs	In vitro physiology	In vivo physiology	Behavioral function
Mesostriatal	SN	PITX3	Dorsal striatum	CX (M1, M2, S1), dorsal striatum (patch compartment), STN, PPN, etc.		Tonic and burst firing	Action selection
Mesostriatal subclass 1	Rostromedial SN		Posterior dorsomedial striatum		Low-frequency HCN channels, D2 autoreceptor response	K-ATP channel-mediated bursting	Exploration
Mesostriatal subclass 2	Rostrolateral SN	ALDH1A1	Dorsolateral striatum				Stimulus-response learning
Mesostriatal subclass 3	SN dorsal tier (calbindin-positive)		Dorsal striatum		Low-frequency HCN channels, D2 autoreceptor response	Oscillatory bursting	
Mesolimbic	VTA		Ventral striatum	CX (LO), ventral striatum (patch), LH, dorsal raphe, etc.		Tonic and burst firing	Motivation
Mesolimbic subclass 1	Rostrolateral VTA		Lateral shell of NAc		Low-frequency, HCN channels, D2 autoreceptor response		Reward signaling
Mesolimbic subclass 2	Posteromedial VTA	OTX2	Medial shell and core of NAc, and basolateral amygdala		High-frequency, absent HCN channels, D2 autoreceptor response		
Mesocortical	VTA		Prefrontal cortex				Cognition
	Posteromedial VTA	OTX2	Pre- and infralimbic cortex		High-frequency, absent HCN channels, absent D2 autoreceptor response		Aversion signaling

The table summarizes the anatomical and physiological properties of various DA neuron subtypes within the midbrain.
CX, cortex; LH, lateral hypothalamus; LO, lateral orbital; M1,M2, motor cortex; PPN, pedunculopontine tegmental nucleus; SN, substantia nigra; STN, subthalamic nucleus; S1, somatosensory cortex.

Table 1.2: Reproduced from Subramaniam and Roeper [108] summarising the dopaminergic projections. The mesocortical projections are investigated in this thesis.

D1

D1 receptors are more prevalent on dendritic spines on glutamatergic and GABAergic neurones. D1-class, through G_s, generally increase activity through activation of AC, PKA, suppression of slow inactivating and inwardly rectifying K-channels and enhancement of slow inactivating Na-channels. D1 appears suppresses high-threshold calcium currents through N- P/Q type channels on apical dendrites, thus decreasing the effectiveness of dendritic inputs in depolarising the soma [128]. D1-class may also couple to G_q, which activate PLC, IP3 and intracellular Ca release. Swanson et al. [111] found that D1 activation decreases subthreshold activity by increasing the input resistance and HCN-channel mediated hyperpolarisation. They also found increasing suprathreshold activity through intrinsic mechanisms in L2/3 and extrinsic mechanisms (glutamate and GABA receptor dependent) in L5. Interestingly, Kisilevsky and Zamponi [64] showed that both D1 and D2 activation decreased N-type calcium currents and surface expression in the PFC, an effect not seen with D1 activation in the striatum.

D2

D2 receptors are less strictly localised to specific cellular regions and are found on glutamatergic and GABAergic neurones. D2-class decrease AC activity and PKA activity, activate K channels and deactivate N, P/Q and R type Ca channels. D2 receptors also signal through a Akt (protein kinase B) - protein phosphatase-2A (PP2A) - β-arrestin 2 pathway which depress NMDA and GABAa receptors. D2 activation decreased N-type calcium currents and surface expression in the PFC. Swanson et al. [111] also showed the D2 antagonism increased the input resistance of L2/3 neurones, but this was eliminated by glutamate and GABA channel blockade, implying this effect is mediated by AMPA, NMDA and GABAa channels and relies on synaptic transmission. D2 antagonism also hyperpolarised the action potential threshold across layers, but unlike D1, this effect was eliminated by ionotropic receptor blockade. This implies D2 activation raises the firing threshold, but that this is mediated by synaptic inputs. For L5 there was a slight increase in firing in the linear portion of the response curve in an ionotropic input independent mechanism. Overall, in both L2/3 and L5 dopamine on D2 receptors decreases excitability, but mediated by different mechanisms. L2/3 suppression is dependent on external inputs from ionotropic receptors, while L5 there are intrinsic mechanisms involved.

D1 and D2

In contrast to the basal ganglia, it appears D1 and D2 receptor activation in the cortex both generally decrease excitability. This may be surprising as D1 and D2 activation have opposing effects on PKA activation, however, there appear to be congruous downstream effects as well. Both receptors inhibit of N-type calcium receptors, decreasing dendritic calcium spikes, likely through GPC$\beta\gamma$ subunit activity [64]. There are also still some key differences between the receptors. At lower concentrations D2 receptor activation suppresses activity through modulating ionotropic inputs both on

glutamatergic and GABAergic neurones across layers. While, at higher concentrations D1 receptor activation decreases excitability but increases the maximum firing rate.

Comparing D1 and D2 expression across layers, Santana and Artigas [99] found expression of both receptors on both glutamatergic and GABAergic neurones, with higher expression in GABAergic neurones in all layers. They found no expression in L1; some expression in L2/3, slightly more for D1 than D2; robust expression in L5, with higher expression for D2; and robust expression in L6, with D1 highest expression of all, in particular in GABAergic neurones. However, colocalisation of D1 and D2 receptors in the same cell is relatively low, a trait that appears common throughout the brain.

D3

D3 receptors are found on a subtype of L5 IT pyramidal neurones which predominantly lack D1 or D2 receptors. D3 receptor activation was shown to regulate low-voltage-activated calcium channels on the axon initial segment, suppressing action potential bursts and reducing transmission reliability, especially of low probability of release synapses [22]. This would shunt activity away from certain synapses in a dopamine dependent manner. Additionally, reducing burst firing would reduce NMDA receptor activation, which is important for long term plasticity. D3 expression is concentrated in the superficial layers in the PFC, with less expression from the L5 boundary down.

D4

D4 receptors are predominantly expressed in the PFC. D4 receptors are spread throughout cortical layers, including on apical dendrites in L1, although densest in L2/3 [97]. They are expressed on both pyramidal and GABAergic cells, and in particular on striatal projecting neurones. D4 receptors are part of the D2-class and decrease PKA activity. Activation of D4 receptors decreased postsynaptic GABAa-mediated transmission in pyramidal cells in the PFC in a PKA mediated mechanism [124]. Zhong and Yan [130] found D4 activation decreased spontaneous activity in pyramidal cells, while increasing spontaneous activity in fast-spiking interneurones. A variant of D4 receptor, D4.7, has been linked to ADHD, potentially through impaired oligomerisation between D4-D2 receptors [12].

D5

D5 receptors are commonly found around the soma, allowing G_q release of intracellular calcium. In macaques Mueller et al. [78] found D5 receptors expressed particularly on long-range projection pyramidal cells, but the proportion of pyramidal neurones expressing D5 was similar across all layers. D5 receptors are also GABAergic interneurones, but to a much lower degree. The lack of pharmacological agents to differentiate D1 from D5 receptors means most studies group them together. However, using D5 receptor knockout mice Carr et al. [17] show deficits in PFC-dependent spatial working memory and recency memory tasks. Although the authors note there appears to be redundancy in D1-class in the PFC as D1 knockout mice show similar deficiencies.

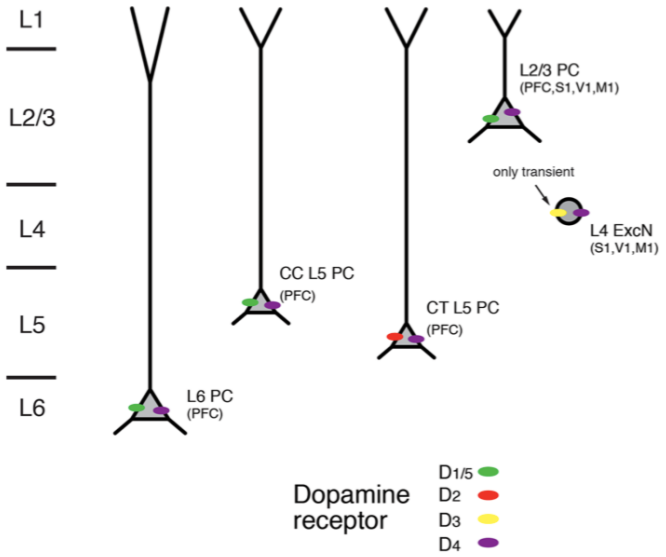

Figure 1.4: Image from Radnikow and Feldmeyer [92]. Layer-wise distribution of dopamine receptors. Regions in which distribution was obtained given in brackets. CC: pyramidal cells with corticocortical projections, CT: pyramidal cells with corticothalamic projections. L4 ExcN: L4 excitatory neurones including spiny stellate, star and pyramidal neurones

1.4.3 Modulation of behaviour

The majority of behavioural experiments investigating dopamine focus on its role in the basal ganglia conferring RPE and saliency signals [63, 94]. However, in the cortex, specifically the PFC, dopamine appears necessary for optimal performance in higher level tasks involving working memory and set-shifting, as well as responses to aversive stimuli. Dopamine also appears necessary for network synchrony between the PFC and M2. Herz et al. [55] found in healthy participants performing motor learning task there is strong coupling between the prefrontal and M2, which is not seen in Parkinson's patients when OFF medication and this is restored by the medication levodopa, which increases dopamine availability.

Looking at the role of dopamine in the PFC, Arnsten [4] describe how D1 activation shows an inverted-U relationship with performance, whereby antagonists or excessive stimulation impairs performance, Figure 1.3. In a cue-orientation working memory task, during optimal performance D1-class receptors enhance spatial tuning by suppressing firing of neuronal ensembles for non-preferred directions during the delay period. Whereas, high doses of D1 agonists silence activity across all ensembles and antagonists increase firing irrespective of preferred direction. Arnsten [4] note how dopamine D1 and noradrenaline $\alpha 2$ have complementary roles, decreasing firing of non-preferred direction ensembles and enhancing firing of preferred direction ensembles, respectively. Dopamine is also important for tasks involving longer delays. In rodents, Floresco [39] found D1, but not D2, modulation altered responses differentially with delays of 30 minutes or 6 hours. D1 agonism altered performance in a dose-dependent inverted-U manner for 30 minute delays. However, the same deleterious concentra-

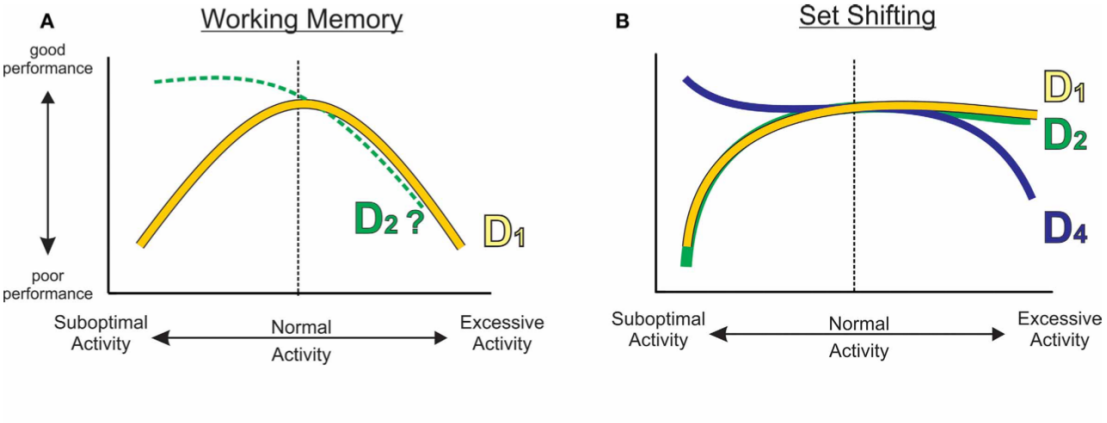

Figure 1.5: Figure from Floresco [39] illustrating the relationship between dopaminergic tone and performance in working memory and set-shifting tasks.

tions improved performance in the 6 hour delay version of the task, during which PFC dopamine concentrations were naturally lower. This highlights the narrow range D1 receptor activation that is needed for optimal working memory and longer-term recall.

Dopaminergic activity in the PFC also appears important for behavioural set-shifting [39]. D1 or D2 receptor antagonism induces preservative errors akin to whole PFC reversible inactivation. Interestingly, D1 or D2 agonism did not impair set-shifting behaviours, arguing for a curvilinear performance relationship which starts with impaired performance with receptor inactivation, increases as dopaminergic tone rises, and plateaus with sufficient activation. In contrast, as seen in Figure 1.5, modulation of D4 receptors bidirectionally altered performance, which followed a negative sinusoid shape, whereby performance decreased with agonism and increased with antagonism. The results of extracellular dopamine recordings through micro-dialysis are also revealing. Dopamine increased by about 100% above baseline during both initial learning and set-shifting, with performance following a curvilinear function with a threshold of minimal dopaminergic tone. Thus dopaminergic tone required for optimal set-shifting without preservative errors is quite distinct from working memory.

In an unconventional way dopamine has been shown to increases the signal-to-noise ratio of aversive signals in the PFC [119]. Inhibition of PFC-projecting VTA dopaminergic neurones attenuated aversive signals in the PFC, an effect not seen in noradrenergic LC inhibition. VTA activation alone did not produce place aversion, but instead increased the salience of aversive cues over appetitive cues when presented together and biased freezing over approach behaviour. The source of the freezing behaviour was found to be a distinct group of PFC neurones projecting to the dorsal periaqueductal grey (dPAG). These projection neurones respond robustly to foot-shock, and optogenetic activation is sufficient to induce place aversion. Interestingly, D1 or D2 receptor expression was not found in identified PFC-dPAG neurones, despite dopamine decreasing calcium event frequency, increasing event amplitude and specifically enhancing air-puff, but not sucrose, responses in these neurones. It therefore appears dopamine increases the output from PFC-dPAG neurones through suppressing neighbouring activity or bottom-up sensory inputs. It should be noted these PFC-dPAG

were predominantly in L5, which as Swanson et al. [111] showed, are heavily modulated by extrinsic inputs.

VTA-PFC projections display other impressive dynamics, showing distinct differences in dopaminergic activity when outcomes were delivered deterministically or with some uncertainty (90% delivery) [106]. In the deterministic state VTA dopaminergic neurones were negatively modulated over time in the inter-stimulus interval (ISI), decreasing between the cue and outcome periods as reward expectation grew. Whereas, in the uncertain state activity increased as a function of time as uncertainty grew. They found inactivating the PFC altered dopamine responses in the non-deterministic case, the increasing-activity pattern in the ISI was abolished. Responses to the cue and activity timing were unaltered by prefrontal inactivation in both deterministic and stochastic cases. The authors conclude that the medial PFC is necessary for computing a belief state when outcomes are uncertain.

Lohani et al. [71] analysed the effect of various types of phasic stimulation of dopaminergic VTA neurones on cortical activity, in freely-moving, well-habituated animals. They found the effect of dopamine on cell firing was weak and heterogeneous in both duration and direction of modulation, although a small number showed persistent excitation or inhibition. Cortical neurone types were only distinguishable by firing properties, and were split into fast-spiking and regular-spiking. Across stimulation protocols 57% percent of fast-spiking but only 6% for regular-spiking neurones were transiently modulated by dopamine, both through activation and inhibition. Whereas, on a prolonged timescale fast-spiking were less responsive, while regular-spiking were more commonly activated than inhibited. It may be speculated that these fast-spiking neurones are fast-spiking interneurones, which is consistent with higher expression of dopamine receptors on GABAergic neurones [99]. Lohani et al. [71] also analysed changes in ensemble activity through a population coding approach, finding that phasic and fast burst activation significantly altered ensemble dynamics with a 5 minute period from stimulation onset to offset. Finally, they found burst firing of dopamine cells exerted greater effects when animals were in active states, implying a behavioural-state dependent effect.

Engelhard et al. [34] were able to classify dopaminergic neurones based on their responses to a variety of variables including reward, reward-predicting cues, reward history, spatial position, kinematics and behavioural choices. The diversity of variables was matched with a diversity of response characteristics during the cue period, including: spatial neurones which exhibited upwards or downward ramps; neurones entrained to certain velocities; and neurones modulated in either direction by previous trial outcome. In contrast, most neurones responded consistently during the outcome period, with stronger responses to rewards than omissions. Activity in the outcome period was consistent with classic RPE during Pavlovian conditioning, in which activity represents discrepancy between actual and predicted value. They found higher reward responses when reward expectation was low due to previous trial outcome or trial difficulty. Dopaminergic neurones were clustered based on specific cue-period responses into five functional groups, and these showed rough topographic divisions. By comparing spatial structure of activity across neurones, they found activity correlations decreased during the cue period but not during the outcome.

1.4.4 Summary

Dopamine appears to play an important role in the cortex, in particular for tasks involving working memory and aversive stimuli. Dopamine generally increases the signal-to-noise ratio by inhibiting responses from surrounding ensembles, thus focusing on decreasing noise. In working memory tasks this includes suppression of ensembles relating to incorrect cues. Whereas, in aversive situations dopamine receptors suppress activity across neurones, but are not present on the subset of neurones which instigate aversive-behaviour. In contrast to the striatum, both D1-class and D2-class receptors appear to suppress activity, despite having opposing effects on PKA activity. Differences in D1/D2 expression between neuronal types and layers may still allow for nuanced responses from each receptor type. Additionally, the convergent effect on calcium channels by both receptors, an effect not seen in the striatum, also appears key. The regional and layer-wise effects of dopaminergic receptor activation is just emerging and will aid in formalising a more complete model. In particular the behavioural-state dependent effects seen in various studies need further investigation.

1.5 Serotonin

The role of the serotonergic system is least well understood of the three neuromodulators studied here. This may be in part because serotonin is closely associated with the complex and nuanced behaviours of emotion and mood, which makes investigating its role through non-verbal animals difficult. In humans, the importance of the serotonergic system is highlighted by the diversity of diseases associated with it, these include depression, anxiety, schizophrenia, addiction, hallucinations and epilepsy [6, 16]. Serotonin has a pervasive influence on behaviour in both aversive and appetitive situations, often in relation to temporality of actions, such as behavioural inhibition, impulsivity and patience [26], [24], [77]. Reduced serotonergic transmission is associated with impulsivity and aggression. The serotonergic system may mediate behaviour when faced with negative expectations [76].

In Figure 1.6, Cohen and Grossman [23] summarise four propositions for the role of serotonin on system dynamics, and note these are not mutually exclusive. First, adjusting the variance of activity, including the membrane potential, firing rate or noise correlation. Increasing variance would decrease synchronous activity, and serotonergic tone has been shown to bidirectionally modulate synchronisation of cortical circuits [18]. Second, change the decay time of neuronal or behavioural processes. Increased integration time alters how much activity history is incorporated into current output, with analogies drawn to the reinforcement learning α component [58] or eligibility trace [112]. Third, modulate the signal-to-noise ratio of activity, akin to the other neuromodulators. Fourth, regulate system robustness to perturbations, whereby serotonin makes certain states more or less stable, effectively changing the 'depth the attractor basin'.

As with other neuromodulators, serotonergic neurones display tonic and phasic activation patterns. Tonic activity is lower than other systems, 0.5 - 2 Hz [28]. Phasic activity occurs within milliseconds of salient stimuli presentation. It has been shown

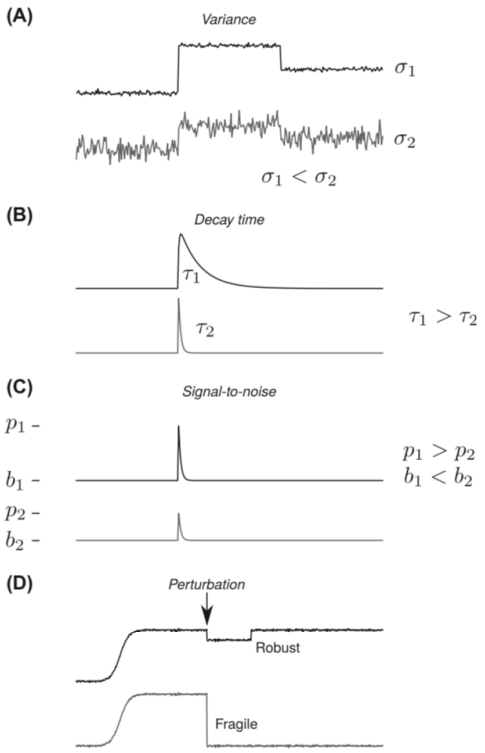

Figure 1.6: Figure from Cohen and Grossman [23] depicting four possible role of serotonin on network dynamics: Variance, Decay time, Signal-to-noise and Robustness

that serotonin activity ramps up during delay periods [131]. This ramping may be inhibiting impulsive actions, resulting in waiting persistence. Another interpretation is that serotonin alters confidence in expectations, promoting patience [76]. Importantly, it appears the serotonergic system is particularly prone to state dependent differences, including basal activity, network and behavioural states.

1.5.1 Anatomy & Physiology

Cortical serotonin predominantly originates from the dorsal raphe and median raphe nuclei. The dorsal raphe is made up of B6 and B7 cell groups, while the median raphe consists of the B5 and B8 cell groups. Soiza-Reilly and Gaspar [105] review how B6 group projects to the hippocampus, the lateral septal nuclei, and the preoptic areas. While B7 group sends projections to the cortex, amygdala and striatum. While projections to the cortex and amygdala don't collateralise, projections appear to innervate multiple cortical areas. The ventral B7 projects heavily to the frontal, somatosensory and motor cortices as well as to subcortical structures like the thalamus, periaqueductal grey, superior colliculus and locus coeruleus; while the dorsal B7 projects to the dopaminergic midbrain and basal ganglia. Up to half dorsal raphe cells are non-serotonergic, with other cell types including, glutamatergic, GABAergic and peptidergic projection neurones, GABAergic interneurones, and some intermixing with A10 dopamine nucleus in the rostral dorsal raphe. Co-transmission appears predominantly glutamatergic, with 80% of ventral B7 serotonergic neurones co-expressing

a glutamatergic transporter. Dorsal raphe projections are also characterised by fine varicosities in comparison to the large spherical varicosities of the median raphe.

The median raphe are more complex neurochemically but with more consistent projection targets. The B5/B8 groups project to all of the septal nuclei and hippocampal formation. Cortical projections are restricted to the mPFC and anterior cingulate cortex. Only 20% of median raphe projection neurones are serotonergic, with a higher proportion being GABAergic. However, there is also a high rate of co-expression of serotonergic and glutamatergic markers in serotonergic neurones. Adding to the complication, a large proportion of non-serotonergic neurones in B8 express Pet1, a gene commonly used to identify serotonergic neurones, which is particularly important for studies using Pet1-Cre. Interestingly, dorsal raphe neurones project to the median raphe, but not vice versa, implying unidirectional modulation [9].

Parent and Descarries [84] review layer-wise serotonergic innervation, noting the highest density of serotonergic afferents is in L1. These varicosities appear largely asynaptic. The synaptic incidence of varicosities in the primary motor cortex of 36% in the superficial layers and 28% in the deep layers. There also appear to be a higher concentration of synaptic varicosities in the medial limbic superficial layers [101]. As noted, this area also receives inputs from the median raphe, known to have axons with larger varicosities. Interestingly, dopamine and serotonin axon terminals are rarely found juxtaposed and are separated by 1 μm on average. Parent and Descarries [84] conclude that a wealth of evidence points to a generally diffuse mode of serotonergic transmission given the common asynaptic character of serotonergic innervation in the cortex.

In summary, the dorsal raphe projects widely across the cortex, with distinct projections to subcortical areas. While the median raphe projects predominantly to the medial areas of the frontal cortex. Importantly, the frontal cortex appears one of the few places overlapping projections. The diversity of serotonergic cell groups, cell types within groups, and often non-overlapping projections, implies serotonin can differentially regulate distinct regions to produce a diverse range of behaviours. However, the prevalence of asynaptic varicosities implies low spatiotemporal precision.

1.5.2 Receptors

There is a great diversity of serotonergic receptors and their downstream effects. However, as the literature stands there is significant insight into the role of only two GPCR families, high affinity inhibitory 5HT1 receptors linked to Gi and lower affinity excitatory 5HT2 receptors linked to Gq, as well as the very low affinity ionotropic 5HT3 receptor. Much less is known about the moderate affinity excitatory 5HT4,6,7 linked to Gs [101]. In the frontal cortex serotonergic receptors 5HT1A, 5HT2A, 5HT2C, and 5HT6 are expressed in both pyramidal neurones and interneurones, 5HT4, 5HT5A, and 5HT7 are predominantly in pyramidal neurones, and 5HT3 expressed are only found on interneurones.

Santana and Artigas [99] found expression of 5HT1A in glutamatergic and GABAergic cells with increasing depth, suggesting 5HT1A exert stronger inhibition in the deeper layers. In contrast 5HT2A receptors on glutamatergic and GABAergic neurones showed higher expression in L2/3 and 5 than L6 suggesting 5HT2A-mediated

excitation would favour top-down processing over reciprocal L6 cortico-thalamic circuits.

Bath or local application of 5HT inhibits most pyramidal neurones via 5HT1A and 5HT5A receptors [46], but excites certain populations. At resting membrane potential, Avesar and Gulledge [7] found application of 5HT resulted in long-lasting 5HT1A receptor-mediated inhibition in 84%, 5HT2A-mediated excitation in 9%, and biphasic responses with 5HT1A-mediated inhibition was followed by 5HT2A-dependent excitation in 5%. The excitatory and biphasic responses were found in L2/3 and L5 callosal and commissural projecting neurones which send axons to the contralateral hemisphere, while corticopontine projecting neurones were invariably inhibited.

5HT1

The action of 5HT1 receptors, which includes subtypes 5HT1A and 5HT1B, is predominantly inhibitory across various cell types. Activating both receptors reduces neurotransmitter release: 5-HT 1A activation causes hyperpolarisation and decreased neural firing, whereas 5-HT 1B inhibits presynaptic voltage-gated calcium channels and release, and 5HT1B receptors enhance the reuptake of serotonin through effects on the serotonin transporter [81]. 5HT1 activation can inhibit both glutamatergic and GABAergic neurones [66], thereby reducing cortical gain.

5HT1A

The 5HT1A receptors are expressed somatodendritically, serving as autoreceptors on serotonergic neurones in the raphe nuclei and as postsynaptic receptors in the cortex. These receptors have the highest affinity among serotonergic receptors, exerting their inhibitory effect by enhancing GIRK and inwardly rectifying potassium channels, decreasing voltage-gated calcium channel currents, and decreasing NMDA-mediated currents through ERK1/2 activity. Notably, 5HT1A receptors are often positioned on axon initial segments, giving them significant control over neuronal output. L2/3 and L5 corticopontine and L6 corticothalamic neurones are highly sensitive to 5HT1A-mediated inhibition. 5HT suppresses the strong excitatory projections from L6 pyramidal neurones onto to fast-spiking and non-fast spiking interneurones in L5 [115].

The distribution of 5HT1A receptors includes a high prevalence of heteroreceptors in the hippocampus and limbic cortex, implicating them in emotional and cognitive processes. Knock-out of 5HT1A heteroreceptors can produce depressive like-symptoms. In contrast, 5HT1A autoreceptors in the raphe nuclei play a crucial role in modulating serotonin release, with activation causing depressive behaviours [81]. Activation of 5HT1A autoreceptors is thought to produce the initial worsening of symptoms following SSRI treatment, which then dramatically improve once these receptors desensitise after a few weeks. The dual role of these receptors, autoreceptors decreasing serotonin release and heteroreceptors mediating antidepressant responses, underscores their profound involvement in anxiety and depression regulation.

5HT1B

5HT1B receptors are primarily expressed presynaptically and exhibit lower affinity compared to 5HT1A receptors. As autoreceptors on serotonergic neurones, they function to reduce serotonin release. 5HT1B are expressed presynaptically on some PV interneurones where they also inhibit release. These receptors modulate adenyl cyclase

and calcium channels, activate potassium channels, and regulate the β-arrestin-Akt pathway.

5HT1B receptors are highly expressed in the basal ganglia, with significant presence also in the cortex, playing a role in movement regulation, impulsivity, addiction and aggression reduction. Both 5HT1B autoreceptors and heteroreceptors are implicated in depression [81].

5HT2

5HT2A are excitatory postsynaptic receptor found throughout the cortex, predominantly in L5 and L2/3, in pyramidal and interneurones. 5HT2A have a lower affinity than 5HT1A, although still generally quite high affinity. Activation increases PLC and PKC, which in turn decreases dendritic excitability by decreasing the maximum of the sodium current and voltage gated calcium channels, while increasing suprathreshold output by inhibiting inwardly rectifying and leaky potassium currents, and increasing non-selective cation and calcium-dependent potassium currents.

Stephens et al. [107] examined the excitatory effect of 5HT application on the contralateral projecting neurones and found 5HT2A-mediated enhanced output given suprathreshold current injection. However, the effect of 5HT2A activation appears state-dependent as 5HT promoted calcium-dependent after-depolarizations, but these responses were also attenuated by intracellular calcium. This suggests under baseline conditions 5HT2A activation enhances excitation, but when activity is high and calcium has built up in the cell this enhancement limits off.

In mice, 5HT2A activation enhances activity in contralateral projecting cortico-cortical neurones, as well as intratelencephalic pyramidal neurones projecting to the striatum and amygdala [101]. Whereas, fascinatingly, Colangelo et al. [25] showed that in rats 5HT2A activation enhances cortico-pontine projecting neurones rather than intratelencephalic neurones, in direct opposition to the mechanism in mice described here. It is unclear how this difference affects behaviour.

5HT3

5HT3 receptors are the only ligand-gated ion channels in the serotonergic family. 5HT3 receptors have a very low affinity, two to three orders of magnitude lower than 5HT1 receptors. Expression is found only in the superficial layers L1-L3 and is limited to certain interneurones. Expression of 5HT3 receptors defines a class of interneurones, which include those expressing cholecystokinin, calretinin, vasoactive intestinal peptide, or neuropeptide Y [117]. 5HT3 activation rapidly excites these interneurones, resulting in an overall inhibitory effect [87]. As these receptors have a very low affinity and rapidly desensitise, activation is likely to be limited to phasic bursts of release when baseline concentrations are low.

5HT5

5HT5A receptors show a broad but moderately low level of expression across the cortex, with greater expression in L2-L4 [87]. 5HT5 have a low affinity similar to 5HT3. 5HT5A receptor activation in L5 pyramidal cells produces an inhibitory response in the rodent

PFC. Genetic deletion of this receptor produces a compensatory increase in 5HT1A currents [46].

5HT4, 5HT6 and 5HT7

5HT4, 5HT6 and 5HT7 receptors all have moderate affinity, and activate the Gs signalling pathway and increase excitatory output. 5HT4 and 5HT6 receptors appear to have low expression in the cortex, while 5HT7 are expressed to a reasonable degree. In the hippocampus activation of 5HT4 and 5HT7 receptors attenuates the calcium-activated potassium current which mediates the slow afterhyperpolarisation and limits repetitive firing, thus causing a facilitation of suprathreshold activity. Co-activation of 5HT1A and 5HT4 receptors causes a strong increase in excitatory gain, in which 5HT1A decreased excitability and 5HT4 facilitated output, and similar interactions might be expected in the cortex. Unfortunately, as Pehrson et al. [87] note, the effects of these receptors in the cortex is so far lacking in conclusive electrophysiological data. Although Fan et al. [36] found 5HT7 agonists delivered locally in the PFC induced generally excitatory responses in pyramidal neurones, but also inhibitory and non-responses. The inhibitory responses were blocked by both GABAa antagonism, implying 5HT7 receptors on both pyramidal and interneurones. 5HT7 receptors likely play an important role in the cortex given their prevalence, although their specific role in behaviour is still unclear [8]

1.5.3 Modulation of Behaviour

The influence of serotonin on behaviour has been seen in experiments on impulsivity [77], learning rate [58], movement suppression [26], and switching between active and passive coping strategies [16, 104]. The dual role promoting either active or passive behaviours depending on the state may explain the difficulties there have been in the field in characterising serotonin. Here 'state' is used broadly to encompass both the environmental and internal state, and mood in humans. Here we will review some papers which explore this dual role of serotonin in both humans and rodents.

Phillips and Robbins [89] describe the role of serotonin in impulsivity, compulsivity and decision making. They note that serotonin is involved with waiting impulsivity, as investigated by Miyazaki et al. [77], whereas noradrenaline improves stop-signal reaction time performance in stopping impulsivity tasks. Miyazaki et al. [77] found serotonin neurone activation in the dorsal raphe prolonged waiting times by about 30%, especially when photoactivation was delivered at the decision point to continue waiting or not. They also found increasing serotonin had a significant effect on persistence when both reward probability and wait-time variability were high, which aligns with the suggested role in impulsivity and risk [76]. They consider these results in a Bayesian framework and propose that serotonin modulates the prior probability, or expectation, that the current trial will be rewarded. They suggest a general role of serotonin in modulating the trade-off between negative sensory evidence and positive subjective belief.

Serotonin appears to have a dual role, not just in passive waiting but also active escape. Seo et al. [104] found serotonergic neurones in the dorsal raphe support con-

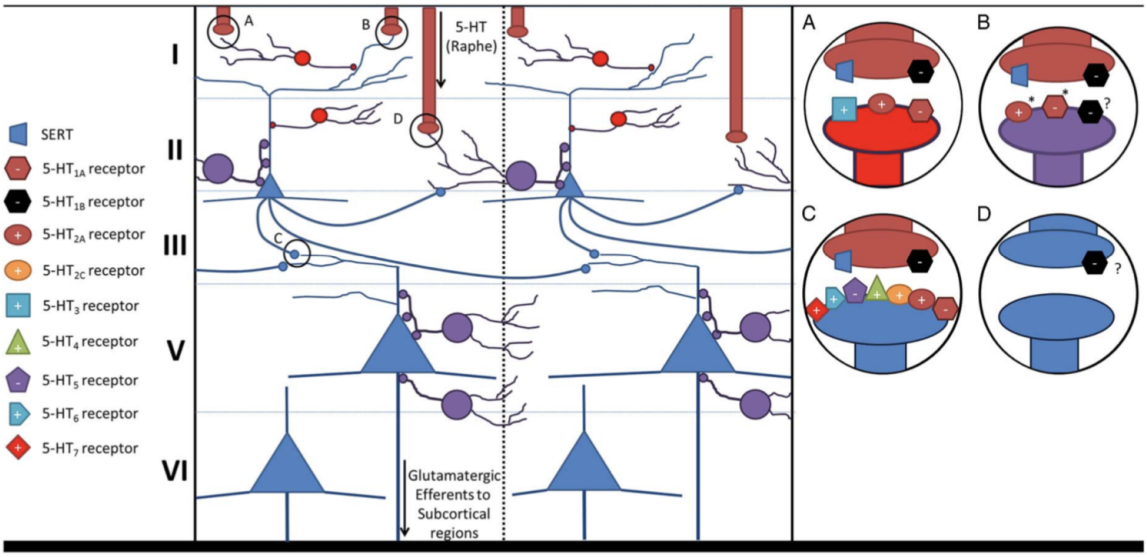

FIGURE 3. Serotonergic influence on a theoretical prefrontal cortex microcircuit. The influence of serotonin on the behavior of cortical pyramidal neurons is mediated by several different factors. Non-parvalbumin-immunoreactive interneurons (depicted in red) in the shallow cortical layers receive serotonergic inputs that have mixed excitatory and inhibitory effects mediated via 5-HT_{1A}, 5-HT_{2A}, and 5-HT_3 receptors (A) and have an inhibitory influence on pyramidal cell dendritic membrane polarization. Additionally, serotonin's effects on parvalbumin-immunoreactive interneurons (depicted in purple) are also probably a mix of excitatory and inhibitory actions that are mediated by 5-HT_{1A}, 5-HT_{2A}, and perhaps 5-HT_{1B} receptors (B). These cells exert a powerful inhibitory effect on pyramidal neuron activity via hyperpolarizing actions at the pyramidal neuron soma or axon hillock. Pyramidal neurons themselves (depicted in blue) also receive a mix of excitatory and inhibitory serotonergic signals, which are mediated by 5-HT_{1A}, 5-$HT_{2A/2C}$, 5-HT_4, 5-HT_5, 5-HT_6, and 5-HT_7 receptors (C). Finally, glutamate release from pyramidal neuron terminals in the prefrontal cortex may be inhibited by 5-HT_{1B} receptor activation (D).

Figure 1.7: Figure from Pehrson et al. [87]. Layer-wise distribution of serotonergic receptors in the PFC. (A) Non-PV interneurones (red) in superficial layers receive direct serotonergic connections, express 5HT1A, 5HT2A and 5HT3 receptors which produce mixed excitatory and inhibitory responses, and inhibit pyramidal dendrites. (B PV-interneurones (purple) express 5HT1A, 5HT2A and 5HT1B receptors producing mixed responses and strongly inhibit pyramidal cells via connections to soma and axon hillock. (C) Pyramidal cells (blue) receive mixed inputs from 5HT1A, 5HT2A, 5HT4, 5HT5, 5HT6 and 5HT7 receptors. (D) Glutamate release may be inhibited via 5HT1B heteroreceptors

trasting behaviours in lower and higher stress environments. In an open field test neurones showed a decreased activity at movement onset. However, in a high-stress tail suspension test activity increased upon movement onset. Photoactivation of serotonergic neurones decreased movement speed in open field, and decreased speed and latency of cued movements in approach and avoidance environments. Whereas, in the tail suspicion, stimulation increased movements.

Carhart-Harris and Nutt [16] corroborate the proposal that 5HT1A and 5HT2A receptors work in compliment to facilitate passive coping via 5HT1A and active coping via 5HT2A. They note the large number of studies which show that increased serotonergic tone or activation of 5HT1A receptors decrease impulsivity, aggression and anxiety. This effect appears to be due to postsynaptic receptor activation in cortical or hippocampal targets, rather than 5HT1A-mediated autoinhibition in the raphe. They argue that 5HT2A signalling opens a window of plasticity which allows for behavioural flexibility. Among their cited examples, Furr et al. [43] found chronic stress induced impairments in reversal learning due to decreased serotonin in the orbitofrontal cortex, an effect replicated by 5HT2A antagonism and alleviated by a selective serotonin reuptake inhibitors. Carhart-Harris et al also discuss how psychedelic 5HT2A agonsits, such as LSD and psilocybin, have been used as novel treatments for depression and mood disorders. They describe a 5HT2A-mediated destabilisation of the Default Mode Network (DMN) producing hyper-connectivity between brain regions, allowing for rewiring. However, they warn that while the increased 5HT2A-mediated plasticity can be useful in therapy, it is sensitive to environmental conditions and can produce acute psychosis and anxiety when administered in less amiable settings. They also report that electro-convulsive shock therapy, which is still used for drug-resistant depression, has a significant effect on 5HT2A receptor densities. Carhart-Harris and Nutt [16] conclude that both receptors can enhance mood and tolerance of stress, and reduce depression. 5HT1A receptors do so through directly decreasing limbic activity and emotional responses, which permits passive coping. Whereas, 5HT2A receptors enhance behavioural plasticity and adaptability and increase cortical entropy, which permits active coping.

Correia et al. [26] found optogenetic activation of the dorsal raphe induced strong suppression of spontaneous movement in an open field test, which could not be attributed to anxiety or motor impairments. As photoactivation was unable to condition a place preference, they suggest previous place preference findings were the result of motor inhibition. It should be noted that while other studies have found conditioned place preference when optogenetically activating cre-Pet1+ cells and this effect was shown to be mainly due to activation Pet1+ glutamatergic projections to the VTA [23]. Correia et al. [26] also found 15 minute daily photostimulation over weeks lead to a persistent increase in locomotive behaviour. At first the locomotor inhibition could cast doubt on the patience hypothesis. However, a key difference between [77] and [26] is the stimulation protocol. Miyazaki et al. [77] increased firing to about 6 Hz, while Correia et al. [26] tested a range of frequencies and found almost no locomotor inhibition at 5 Hz, but dose-dependent inhibition at 15 Hz and above. Given the discussion of cortical receptor affinities, it might be expected 5HT1A activity might dominate at lower levels of activation, while 5HT2A, 5HT7 and 5HT3 would be activated at by stronger stimulation. 5HT2A receptors decrease corticopontine output, which could

decrease speed, alternatively, this could be attributed to 5HT3 activation of superficial interneurones. The long-term effects of strong stimulation seem to concur with overall long-term improvement in mood Carhart-Harris and Nutt [16] describe with 5HT2A receptor activation.

Recording in the dorsal raphe in a block-wise appetitive and aversive Pavlovian conditioning task, Cohen et al. [24] found heterogeneous but consistent serotonergic neurone responses. Most neurones showed phasic excitatory responses to aversive stimuli, although some were inhibited. There were also smaller phasic responses to punishment-predicting cues. Meanwhile, about half were strongly excited by reward-predicting cues, and less so for rewards. Thus, large phasic responses appear to be predominantly for actual aversive stimuli and also for reward predicting cues. These phasic responses of 5 - 10 Hz lasted less than 500 ms. Additionally slower changes in tonic activity were seen. In about 40% of neurones tonic firing rate increased or decreased. Changes depended on the valence of the stimuli, with roughly equal numbers being excited by during appetitive blocks as those excited during aversive. These tonic signals built up slowly over trials and persisted for minutes. Tonic and phasic responses were correlated, with phasic reward-responding neurones more commonly showing tonic increases in during appetitive blocks. Cohen et al. [24] also performed this experiment recording from dopaminergic neurones in the VTA. As expected, dopaminergic neurones showed large, phasic increases to unexpected rewards, and smaller phasic responses to expected rewards. Both serotonergic and dopaminergic neurones showed larger responses unexpected stimuli, but dopaminergic were faster and stronger responses. However, despite dopaminergic responses previously being purported to track long-term value-related changes, valence related ramping was only seen for serotonergic neurones. The heterogeneity of serotonergic responses may suggest that appetitive responding and aversive responding neurones project to different regions, especially as recordings likely included both B6 and B7 cell-groups. Alternatively, the average activity across neurones may provide a more global population code.

Matias et al. [74] performed a similar Pavlovian association task, but with a reversal. Mice learned one of four associations, large or small reward, neutral, or air-puff, for multiple days before the reversal in order to create strong prediction errors. Photometric measurements in the dorsal raphe were compared to dopaminergic neurone activity in the VTA. Before reversal, serotonergic and dopaminergic neurones increased activity for reward predicting cues. Dopaminergic neurones showed mild responses to well-predicted liquid rewards, whereas serotonergic neurones showed mild decrease in activity. For air-puff delivery serotonergic responses increased activity, while dopaminergic showed little change. However, aversive-stimuli predicting cues elicited no response from serotonergic, but a pause in activity for dopaminergic neurones. These results largely fit with [24]. Both serotonergic and dopaminergic neurones showed similar reward prediction error signals for positive reversals, better-than-expected outcomes, with strong excitatory phasic responses post-reversal. For negative reversals, serotonergic neurones showed robust activity to worse-than-expected outcomes, whereas dopaminergic neurones showed reduced responses. Phasic responses were also seen in omission trials, but smaller than reversals. The authors suggest serotonin signal resembles an unsigned prediction error or surprise signal, although, as they note this interpretation was not supported by their finding well-expected air-puffs elicited mild

responses even after extensive training, in line with [24].

Zhong et al. [131] saw a similar lack of response when bitter quinine was unexpectedly added to the sucrose reward. These lack of responses is at odds with the unsigned prediction error and worse-than-expected outcome theories. However, while Matias et al. [74] found the averaged responses are similar to baseline, individual neurones showed either increases or decreases in response, similar to the heterogeneity found by Cohen et al. [24]. It is unclear why Matias et al. [74] found larger serotonergic responses for appetitive outcomes than aversive outcomes, while Cohen et al. [24] found the opposite. However, importantly, both reported increased serotonergic responses to cues predicting liquid-rewards and air-puffs. Finally, Matias et al. [74] also found the difference in adaption rates. Serotonergic post-reversal adaption was markedly slower and more persistent than the dopaminergic system, leading to an interesting hypothesis. In positive reversals, as dopaminergic adaptation is faster it will have a greater influence invigorating behaviour, whereas, in negative reversals, as serotonergic responses persist longer it will lead to long term inhibition of non-adaptive behaviours. They also directly manipulated the serotonergic system with transfected synthetic receptor DREADD-mediated inhibition. For worse-than-expected reversals licking rate adaptation was slower when serotonergic neurones were inhibited. In contrast, there was no significant difference to controls in positive reversals. This implies dorsal raphe serotonergic activity is sufficient, but not necessary, to facilitate behavioural flexibility when contingencies change negatively. Matias et al. [74] summarise that, in stable environments serotonin and dopamine signals may nullify each other, requiring no adaptive change. But in dynamic environments the adaption rate differences may provide a serotonergic inhibitory surround to dopamine's excitatory centre.

1.5.4 Summary

Under basal conditions the firing rate of the raphe is very low, which would favour activation of the higher affinity of 5HT1A receptor, decreasing both glutamatergic and GABAergic currents. The low affinity, synaptic positioning and rapid desensitisation implies 5HT3 receptors are activated by strong phasic activity when baseline activity is low. 5HT3 receptors induce rapid depolarisation of a class of GABAergic interneurones predominantly located in the superficial layers. However, it is unclear under what physiological conditions 5HT3 are activated and what effect this has on network dynamics. Equally unknown are the behavioural conditions 5HT2A receptors are preferentially activated. 5HT2A have a moderately high affinity, although much lower than 5HT1A. Stephens et al. [107] found despite many neurones co-expressing 5HT1A and 5HT2A receptors, projection targets consistently determined whether 5HT1A-mediated inhibitory or 5HT2A-mediated excitatory responses were seen. Importantly, the excitatory responses were found in a subset of pyramidal neurones projecting contralaterally or to the amygdala. It appears serotonin evokes heterogeneous responses which depend on cell type, projection target and activity-dependent state. This complexity helps explain why a layer-wise cortical model has yet to emerge.

Many studies suggest a dual, context dependent, role for serotonin. In lower stress conditions serotonin appears to promote decreased locomotion [26], which has been interpreted as passive coping [16, 104] or patience [76], and may be implemented by

5HT1 receptors. Whereas, in other conditions active coping [16] or escape [104] behaviours are seen, and these may be produced by 5HT2A receptors. Matias et al. [74] suggest serotonin inhibits learned behaviour that is no longer beneficial and driving the plasticity needed for reconfiguration. Similarly, Cohen et al. [24] suggest serotonin-dopamine opponency, whereby serotonin AND dopamine signals reward-predicting, while serotonin AND NOT dopamine signals punishment.

While serotonin has generally been hypothesised to oppose with dopamine, promoting action inhibition and invigoration respectively, in many ways serotonin appears have a tightly linked role to noradrenaline, at least in the cortex. Both systems appear to regulate the breadth or sharpness of activity, permitting or inhibiting action initiation. Serotonin may promote action persistence at lower concentrations but also facilitate flexibility in certain contexts. Similarly noradrenaline allows action initiation but also action flexibility depending on the firing mode. In support of this Seo et al. [104] found a robust decrease in dorsal raphe activity at movement onset, while Xiang et al. [127] noted a noradrenergic phasic peak at movement initiation.

1.6 Genetically Encoded Neuromodulator Indicators

GENIs are recently developed sensors that allow for measuring the extracellular concentrations of neuromodulators. GENIs combine an endogenous GPCR with a conformationally sensitive fluorescent protein such as circularly-permutated GFP. In the unbound state the protein complex exhibits low basal fluorescence, however when bound to a ligand a conformation shape change occurs and produces a highly fluorescent form. There are GPCRs for most if not all neurotransmitters and neuromodulators, making them ideal scaffolds for building this kind of biosensor. As the GPCR base is an endogenous receptor, it has similar temporal dynamics, in the tens of milliseconds, and naturally high specificity. If combined with cranial window microscopy, neuromodulator dynamics can be imaged at wide-field (millimetres) or sub-cellular resolution, and without the insertion of invasive probes.

1.6.1 Noradrenaline biosensor

Feng et al. [38] first released two noradrenaline biosensor variants, medium affinity GRAB_{NE1m} and high affinity GRAB_{NE1h}, both based on an α2AR scaffold. In cultured neurones affinity of GRAB_{NE1m} was 1900 nM, as shown by the dose-dependent measurement of half maximal effective concentration EC50. Whereas a single point mutation increased the affinity GRAB_{NE1h} by 10x to 93 nM. These results are lower than the data published in the table which were obtained from HEK293T cells, however, the affinity of GRAB_{NE1h} is still very impressive. In terms of kinetics, GRAB_{NE1m} has rapid onset and a more gradual offset ($\tau_{on/off} \approx 72$ / 680 ms), whereas, GRAB_{NE1h} has very rapid onset but a delayed offset ($\tau_{on/off} \approx 36$ /1890 ms). The maximal fluorescence change, $\Delta F/F_0$, in response to a 100mM saturating concentration of NE was good for both biosensors, 230% for GRAB_{NE1m} and 150% for GRAB_{NE1h}. Comparing both biosensors, there is tradeoff between sensitivity and speed of offset, with high affinity receptors having high sensitivity but slow offset and low affinity showing lower

sensitivity but fast offset. $GRAB_{NE1m}$ have a much higher change in fluorescence and faster offset, making it more appropriate for behavioural research.

1.6.2 Dopamine biosensor

There are two commonly used dopamine biosensors dLight [86] and $GRAB_{DA}$ [110]. dLight is based on dopamine D1R, whereas, $GRAB_{DA}$ variants, $GRAB_{DA1m}$ and $GRAB_{DA1h}$ are based on D2R. This difference would allow those wishing to target the D1R expressing direct pathway should use $GRAB_{DA}$ and those targeting the D2R indirect pathway to use dLight. The dLight biosensors have faster dynamics and greater signal strength, whereas the GRAB sensors have a higher affinity and so greater sensitivity for small changes. dLight was used in this thesis and the specifics will be presented here. In acute striatal slices, in response to a single electrical stimulus dLight1.2 showed a 220 ± 50% increase in fluorescence and very rapid dynamics ($\tau_{on/off} \approx 9.5$ ms / 90 ms) with a peak signal plateau of ∼150 ms which decayed to baseline within ∼400 ms. In situ dopamine titration on HEK cells showed high affinities, dLight 1.1: $K_d = 330$ nM and dLight1.2: $K_d = 770$ nM. For specificity, dLight1 is more sensitive to dopamine than to noradrenaline, ∼70x, and adrenaline, ∼40x; responses to all other neuromodulators tested were negligible. dLight was shown to have negligible interactions with downstream G-protein pathways and a lack of dopamine-induced internalisation. dLight biosensor responses were abolished by D1R antagonist SCH-23390 but not by D2R antagonist haloperidol.

Sun et al. [110] performed a Pavlovian association paradigm in head fixed mice while dopamine using GRAB recording $GRAB_{DA}$ in the nucleus accumbens using fibre photometry. In naive animals they found responses to the appetitive stimuli were initially high but slowly decreased during training, whereas aversive stimuli consistently produced a decreased dopamine response. This contradicts the standard RPE theory as predictable aversive stimuli continued to elicit the same response. Furthermore, responses to the appetitive cues increased during training, while responses to aversive stimuli remained at baseline throughout training. They also found large dopaminergic responses in male mice during sexual intercourse.

In a pertinent experiment for the proposed research, Patriarchi et al. [86] recorded dLight signals from the motor cortex using two-photon imaging while mice performed a NoGo-to-Go task, Figure 1.8. Mice were required to initially stand still for 10 s in response to a cue, then following another cue, if mice sustained movement for 1s during the 3s window they received a reward with 80% probability. To increase the ease of the task, initially the duration of the NoGo phase was set to 3 s, the Go-window to 20 s and the reward probability was 100%. As individual mice progressed the NoGo phase was lengthened and the Go window shortened. Mice typically reached proficiency within 7 days, after which two-photon microscopy was performed. Imaging targeted L2/3 neurones in the motor cortex, ∼ 0.2 mm deep. $\Delta F/F_0$ normalisation of the fluorescence time series data used the 8th percentile baseline to smooth the signal and a sliding window of ±15s around each point to correct for slow drift and bleaching. ROIs were sorted by task criteria, revealing a heterogeneous group. 63% of ROIs showed significant activity during reward expectation period, with two peaks: an initial ramp up of activity from cue presentation to reward delivery window, in both rewarded and

omitted trials; and a second peak during reward delivery for rewarded trials which was abolished by reward omission. Miss trials showed no increases in activity, Fig 1.8 D Right. 37% of ROI were significantly active during Go-cue presentation and locomotion but showed no responses to reward delivery or omission, Fig 1.8 D Left. Approximately 5% of ROIs showed significant transients during reward expectation but not spontaneous running, Fig 1.8 E centre. The lack of activity in miss trials during the reward expectation phase indicates mice had learned the task and understood when they would be penalised for mistakenly running prematurely. Approximately 32% of ROIs were active during any locomotion. These results highlight myriad of distinct dopamine signals which appear to be modulating different circuits depending on task-related stimuli. The cortex receives dopaminergic afferents from VTA neurones, which are known to heterogeneously respond to specific aspects of behaviour [34]. Pertinently, ROI were heterogeneously spread in the dorsal-ventral plane as well. Across 2 mice they found ∼200 reward related ROI in L2/3 but, almost no reward expectation ROIs and surprisingly only ∼10 locomotion ROI. In contrast, in L1 they found ∼30 location ROI, ∼10 reward expectation ROI and ∼15 reward ROI. Taken together these results suggest that dopaminergic release sites are heterogeneously spread in the XYZ planes.

1.6.3 Serotonin biosensor

Wan et al. [123] developed a genetically encoded biosensor based on a 5HT2C scaffold, referred to as $\text{GRAB}_{5HT1.0}$ The biosensor produces a robust 280% fluorescence increase $\Delta F/F_0$ in cultured neurones in response to bath serotonin application. $\text{GRAB}_{5HT1.0}$ has reasonable onset but quite slow offset dynamics ($\tau_{on/off} \approx 200$ /3100 ms), which is common among the high affinity biosensors; EC50: 22 nM. The biosensor displayed high specificity, with no responses seen for glutamate, GABA, dopamine, noradrenaline, acetylcholine, histamine, octopamine or glycine. There was also a lack of response in both G-protein and β-arrestin coupling assays, implying a lack of intracellular effects. Electrical stimulation of dorsal raphe cells expressing $\text{GRAB}_{5HT1.0}$ elicited increasing fluorescent responses with increasing pulse number or frequency. Stimulation with 10 high frequency 100 Hz pulses produced about 100% increase in $\Delta F/F_0$, with rapid onset but very slow offset ($\tau_{on/off}$ 150 ms/7220 ms). $\text{GRAB}_{5HT1.0}$ responses appear to last 1–2 times longer than FSCV, but with a double or greater signal-to-noise ratio. Wan et al. [123] also measured endogenous serotonergic activity in the PFC using two-photon imaging. They found robust increases in fluorescence in response to intraperitoneal injection MDMA, which increases serotonin release. They also performed long-term fibre-photometry recordings in freely moving mice while recording in the basal fore-brain. They found robust increase in baseline in activity during wakefulness, showing high temporal variability with bursts and dips; low baseline but rhythmic activity in non-REM sleep and quiescence in REM sleep. Activity was strongly reduced by system application of 5HT antagonist metergoline. They also simultaneously recorded from both the orbitofrontal cortex and bed nucleus of the stria terminalis (BNST), and found highly correlated activity between the two areas during non-REM sleep.

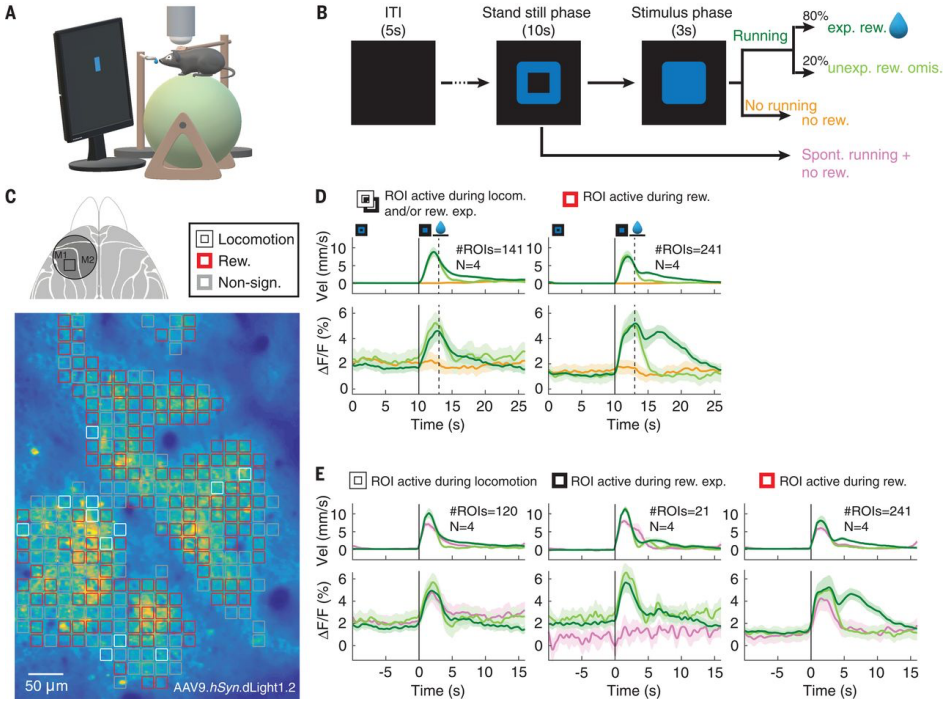

Figure 1.8: Figure from [86] with similar setup to the proposed research. **A:** Spherical treadmills allow movement along the XY plane, however, require longer habituation training than 1D treadmills. **B:** NoGo -> Go task flow diagram: initial NoGo, "Stand still phase"; timeout for premature movement, "Spontaneous running"; Go phase, "Stimulus phase"; probabilistic outcomes. **C:** Representative biosensor fluorescence signal. Regions of interest (ROI) were hand-selected and sorted by task-relevant criteria. **D:**Population data. Dark green: rewarded trials; light green: omitted trials, yellow: miss trials. Top: running speed. Bottom: ROI fluorescence. Left: ROI active for Go-cue (reward expectation) and reward delivery. Right: ROI active in delivery window, rewarded (dark green) and omitted (light green). **E:** Data realigned to locomotion onset. Dark green: rewarded trials; light green: omitted trials; pink: spontaneous run miss trials. **Top:** running speed. **Bottom:** ROI fluorescence. **Left:** locomotion ROI in all trials. No differences in rewarded, unrewarded and miss trials. **Centre:** reward expectation ROI during both rewarded and omitted trials but not miss trials. Activity initially increased for reward, omitted and miss trials; this may represent locomotion as it appears similar to the locomotion signal seen in E Left. Notably, second peak appeared only during reward consumption.

1.7 Orthogonalised Go/NoGo Task

As we have seen the neuromodulators noradrenaline, serotonin and dopamine are involved in various aspects of reinforcement learning. Reinforcement learning is by definition the altering of motor behaviours through appetitive or aversive conditioning, and so is highly likely to involve the secondary motor cortex. However, the exact role of these neuromodulators in M2 is not well known. To investigate this the OGNG task was adapted for use in head-fixed mice with aim of maximising responses in M2. The OGNG task uses a combination of positive or negative reinforcement to modify locomotor behaviour. In positive reinforcement, correct responses elicit an appetitive sucrose solution delivery, while for negative reinforcement incorrect responses elicit an aversive air-puff. Both positive and negative reinforcement aim at increasing the probability of the specified action. For example, doing homework could be done by the student to gain praise (positive reinforcement) or to avoid scolding (negative reinforcement).

Using positive and negative reinforcement to modify locomotor behaviour has been shown to include a Pavlovian and instrumental component. When instrumental actions are reinforced, action-independent Pavlovian cue-associations and anticipatory behaviours can also emerge: appetitively conditioned stimuli come to elicit Pavlovian approach behaviours, while aversive stimuli elicit avoidance behaviours. Conflict between instrumental and Pavlovian systems occurs when avoidance-like behaviours are positively reinforced or approach-like behaviours are negatively reinforced; for example, using aversive stimuli to reinforce locomotion is difficult because fear of punishment can cause freezing. In conflict situations learning is much slower and performance can even be lower than chance in both humans [49] and rodents [62]. Timberlake and Wahl [116] show how in rats innate appetitive behaviours emerge even for a ball-bearing token, and that this impairs their ability to relinquish the token for actual food. They noted that the difficulty arises when instinctive food-related behaviours, such as picking up, holding and gnawing an object, compete with goal-oriented instrumental responses. This suboptimal behaviour emerges later in training as the token comes to signify food itself and elicit Pavlovian responses. Homayoun and Moghaddam [56] demonstrated that in the frontal cortices activity related to Pavlovian and instrumental responses occurred in the same neurones and was correlated with the strength of the influence of Pavlovian cues on instrumental responses. The PFC appears to play a pivotal role in conflict situations, both decision conflict [40], and in overcoming instinctual Pavlovian behaviours [49], and overriding habitual actions [11]. Changes in neuromodulator dynamics may play an important role in regulating these Pavlovian behaviours in the frontal cortex [102].

The OGNG task was developed to tease apart the influence of Pavlovian influences on instrumental learning [49]. In the OGNG task, mobile and immobile behaviours, and positive and negative reinforcement are all counterbalanced. Counterbalancing positive and negative reinforcement, and Go and No-Go actions produces a contingency table with four states, whereby action type and reinforcement valence are on orthogonal axis, hence the name OGNG. This allows the differential effects of valence on instrumental and Pavlovian conditioning to be disentangled.

1.7.1 Previous Orthogonalised Go/NoGo Tasks

The original OGNG task was performed by human subjects in an fMRI scanner who learned state-action contingencies in which cues signalled which action should be taken, button press (Go) or withhold (NoGo), through positive or negative reinforcement [49]. Go and NoGo actions were counterbalanced with positive and negative monetary reinforcement, resulting in four contingencies: “Go to win” and "NoGo to win”, “Go to avoid punishment” and “NoGo to avoid punishment”. When subjects learned the task from trial and error, Go performance was higher when positively reinforced and NoGo performance better when negatively reinforced. However, in the incongruous NoGo to win situation some subjects performed worse than chance.

Guitart-Masip et al. [49] also compared activity in the striatum, the inferior frontal gyrus and the combined dopaminergic nuclei in the midbrain. The striatum and dopaminergic nuclei showed strong increases in activity in the anticipatory phase, before movement initiation. In the striatum they found increased activity for Go and decreased activity for NoGo, independent of outcome valence. This suggests that activity in the striatum is related more to action requirements than outcome valence, and challenges the wealth of studies which show state-value, action-value and temporal-difference reward prediction error signals in the striatum [42, 63]. Results in the VTA for most conditions were identical, except that activity was not significantly greater than baseline in the ‘NoGo to win’. Thus, RPE activity in the dopaminergic nuclei was only seen for positive reinforcement of mobile, but not immobile, actions, and not in negative reinforcement.

The same group repeated the task in participants under pharmacological manipulation of the dopaminergic system with levodopa and the serotonergic system with citalopram [51]. The authors expected to find a clear contrast in the effects of levodopa and citalopram, given the long-hypothesised opponency of dopamine and serotonin in invigorating and inhibiting movements, respectively. In contrast, they found systemic administration of citalopram, a selective serotonin reuptake inhibitor, resulted in high performance in both Go conditions, and low performance in both NoGo conditions. Thus, citalopram had a facilitatory effect on mobile actions, irrespective of valence. Whereas increasing dopamine availability through systemic administration of dopamine precursor levodopa resulted in a decrease in the performance asymmetry seen across states in controls. This indicated levodopa decreased the influence of the Pavlovian system on instrumental responses. They note these unexpected findings may have resulted from systemic drug administration, and different results might have occurred with targeted activation in the basal ganglia or PFC.

In the closest paradigm to ours, Jones et al. [62] found species-specific differences in response between mice and rats in a shuttle box OGNG task. Animals learned to change box (Go) and stay (NoGo) in response to a cue. In the first stage animals were split into groups and reinforced positively, to either Go-Pos or NoGo-Pos. Once animals reached criterion of 70% performance they were moved on to the negative reinforcement stage, whereby the Go group learned NoGo-Neg and NoGo group learned Go-Neg, with incorrect responses receiving an aversive air-puff. Finally, regardless of whether individuals reached criterion within 14 sessions they were moved onto the discrimination test with interspersed trials of positive and negative reinforcement. Rats

performed equally well for both Go conditions. However, performance for NoGo-neg only reached around 70%, and Go-neg was not learned at all. Whereas mice also performed best in Go-Pos, often above 90%, and well in NoGo-Pos at about 80%. For the negative reinforcement situations, the performance for Go-neg was higher, around 70%, while NoGo-neg was lower at about 30%. Interestingly, in positive reinforcement rats performed equally for both Go and NoGo, while mice showed lower performance for positive reinforcement of NoGo responses, akin to humans. Both species performed much worse under negative reinforcement, however in different situations. Rats were completely unable to learn Go-neg, whereas mice could achieve poor performance of NoGo-neg. The authors suggest this may reflect ecological differences as mice are small and herbivorous, so food is static and danger is mobile, while rats are larger and omnivorous, so food may be moving or not and escape is less common. Colangelo et al. [25] also found some key differences in the effects of serotonin on cortical activity between these species, as in rats 5HT2A receptors promotes cortico-pontine projecting neurones, whereas in mice it promotes intratelencephalic projections including cortico-amygdala; and these circuits are known to be involved in fear conditioning. Finally, it is somewhat unclear whether the rats learned either negative reinforcement situation, as despite the 70% performance on NoGo-neg this requires no change in location to succeed and may just reflect a tendency to remain in place. This is supported by the low shuttle rate, only twice per minute. Whereas mice seemed very active, with an average shuttle rate of five times per minute.

In summary, the OGNG task examines the influence of innate Pavlovian responses on instrumental learning. The prefrontal cortex appears pivotal in attenuating Pavlovian interference, which will be particularly important when Pavlovian and instrumental systems are incongruous. The neuromodulatory systems seem to play a key role in regulating this prefrontal control. Manipulating the dopaminergic and serotonergic systems have been shown to influence the Pavlovian and instrumental aspects of the task [50], while Sales et al. [98] suggest that the noradrenergic system is important for Go/NoGo and reversal learning.

1.8 Open questions

1.8.1 What is the spatial extent of volume transmission?

While there is a lot of evidence that neuromodulators must diffuse across much greater distances than synaptic, and extrasynaptic, neurotransmission the actual spatial extend of release events is unknown. There is some recent evidence showing release at the neuronal and even compartment scale [61, 86]. However, neuromodulator axons are sparse and do not extend to all layers equally [82]. The actual scale of release, both quantity and distribution, has significant implications for our theoretical understanding of neuromodulation. At lower molar concentrations, only high-affinity receptors are activated, whereas higher concentrations activate low-affinity receptors as well, potentially triggering opposing downstream effects. Additionally, the pattern of release, whether occurring in heterogeneous pockets or uniformly across the tissue, could influence the spatial scope of functionally connected ensembles. Previous techniques, such

as micro-dialysis or FSCV, have a limited spatial scale. However, using a combination of the novel GENIs and two-photon microscopy we can detect neuromodulator release at the micrometer scale, over hundreds of micrometers. By looking at whether activity is correlated at different distances, we can shed light on to the extent of volume transmission.

1.8.2 Does *in vivo* neuromodulator release cause an increase or decrease in neuronal activity?

Much is known about neuronal responses to pharmacological manipulations or evoked release through optogenetic stimulation. Nevertheless, there are still large gaps in our knowledge of neuronal responses to natural, *in vivo* neuromodulator release. Each neuromodulator has a wide range of receptors that are activated at different concentrations. Moreover, each cortical region and each layer within that region shows a different receptor expression pattern [87, 92]. Therefore, if a low affinity receptor is expressed in a layer far from the axonal release site, it may only be activated in extreme cases of neuromodulator build-up, such as very high baseline activity cause by stress. Conversely, if the same receptor is expressed nearby the axon, it may be often activated. Pharmacological interventions activate receptors ubiquitously and so can lose this nuance. What we would like to know is, how much neuromodulator is released in response to a stimuli and how does this influence the surrounding cortical activity. Recording both intracellular calcium and nearby neuromodulator release will enable us to see how release events are influencing neuronal activity.

1.8.3 How do neurones response to the mixture *in vivo* neuromodulator signals?

The vast majority of neuromodulator experiments investigate each system individually. Often dopamine, noradrenaline and serotonin show overlapping response profiles, in particular to appetitive and aversive stimuli. Yet, for a given stimuli there is also considerable variability in neuromodulator and behavioural response. It is well known from pharmacological studies that increases in one system, but not in another, are associated with certain types of errors or behavioural strategy. It would be interesting to know what degree variance can be explained by different balances of these neuromodulator systems. To answer this we would need to image multiple neuromodulator systems simultaneously. With this aim we performed wide-field recordings of two neuromodulator systems, one in each hemisphere, at the same time.

1.8.4 What is the overall role of each neuromodulator?

Neuromodulators are strongly implicated in cue-response learning, with analogies drawn to hyper-parameters of Reinforcement Learning algorithms [31]. However, there are a variety of models, and each has assigned distinct roles for each neuromodulator. For example, the Reinforcement Learning model proposed by Doya [31], and the Active Inference model proposed by Friston et al. [41], Parr and Friston [85]. Both models agree that both dopamine and noradrenaline make use of previous reinforcement

event history to inform the exploitation-exploration balance of action selection, but with subtle differences. Doya [31] proposes that dopamine provides a RPE signal, as evidenced by [125], while noradrenaline controls the exploitation-exploitation balance [5]. Whereas, in Active Inference dopamine encodes a "precision of state estimation", which incorporates RPE and incentive salience [94, 125]. Whereas, noradrenaline provides a state-value prediction errors [98]. Both of these models aims to provide a high-level overview of the role of each neuromodulator. However, we should not expect that neuromodulator signals within every region reflect this overarching role. As we have seen neuromodulator signals may differ by cortical region. In particular, the well-established dopaminergic RPE signal seen in the basal ganglia is not seen in the cortex [126]. Knowing how noradrenaline, dopamine and serotonin act in the motor region M2 will to inform *how* the above models could be implicated.

Chapter 2

Methods

The OGNG task was selected to examine how the valence of outcomes affects reinforcement learning. In reinforcement, action probability and vigour are increased using the action outcome as feedback. Positive reinforcement enhances actions that increase the probability of appetitive stimuli, whereas negative reinforcement promotes actions that decrease the likelihood of aversive stimuli. Parameters of the task were selected that have been shown to induce phasic release of neuromodulators in the frontal cortex. These include high-probability, but not deterministic, reinforcement; variable intervals to increase uncertainty; and state reversal. The task setup was also selected to require full body responses as the easily accessible, midline M2 cortex relates to trunk and hind limb movement Zingg et al. [132]. Moreover, locomotion and rest are more natural behaviours than licking, and therefore may produce more generalisable results.

2.1 Animals

Female, 1 to 3 month old C57/BL6 mice. Surgery was typically performed on 4 to 6 week old mice, with imaging from 3 weeks post injection, for 5 weeks. Mice were housed under reverse light/dark cycle. Female mice were housed in sororal pairs 2 days after surgery, as a form of environmental enrichment. Group housing improves the quality of life of mice and helps reduce stress. Mice were also given free access to an igloo with a dish-type running wheel, plastic tube, and chew toy.

2.1.1 Animal ethics statement

All procedures involving mice were conducted in strict accordance with the guidelines set by the Animal Care and Use Committee (ACUC) of OIST. This study was approved by the ACUC under protocol number 2022-385. The laboratory mice used in this study were housed in a controlled environment with a 12-hour light/dark cycle. Mice were placed on water restriction and only given ad libitum access to water for one hour a day, after the experiment. Mice were weighed daily to confirm their body weight at 80% of the non-water restricted weight. All experimental procedures, including handling, housing, and euthanasia, were performed with due care to ensure the humane treatment of the animals. Any signs of distress or discomfort observed in the

animals were promptly addressed and additional pain medication administered where appropriate. Efforts were made to minimise the number of animals used.

2.2 Animal habituation

Daily handling began 3 days before the first surgery. In turn, each mouse was ushered into the hollow tube in their home cage. The tube was picked up and mice were allowed to leave the tube to explore gloved hands at their own pace. Picking up mice by their tail was avoided, as this can be stressful [38]. Starting with 3 minutes, handling was gradually increased until mice could happily sit on gloved hands without the tube for 10 minutes.

To increase the amount of time mice had for training, mice were first habituated to head-fixation prior to virus injection. This was achieved through an additional preliminary surgery in which a head-plate was attached to the skull. This preliminary surgery followed the same basic surgery protocol as described in 2.3 with the omission of the cranial window and virus injection steps. The head-plate was removed during the second surgery and reattached following virus injection. The time between the preliminary surgery and virus injection was approximately 2 weeks.

From the second day after the preliminary surgery, mice were placed on water restriction and weighed daily to keep their body weight at 80% of the non-water restricted weight. Each day mice were habituated to moving on treadmill-type wheels (see 2.6). Initially, while the mice explored a training treadmill, the wheel was externally rotated to coincide with the movements of the mouse. Over increasing periods, their head-plate was held between fingers, giving them the resistive force needed to turn the treadmill themselves. Mice were also taken to the behavioural room and allowed to explore the behavioural setup and taught to receive sucrose solution from the spout.

2.3 Cranial window surgery

Cranial window surgery was performed on mice to insert a chronic transparent window over the secondary motor cortex (M2), allowing longitudinal two-photon imaging.

Mice were anaesthetised with 3% isofluorane, weighed and secured in the stereotaxic frame. Carprofen (5 μg/g; i.p), Dexamethasone (2 μg/g; intramuscular) and Buprenorphine (0.1 μg/g) were injected to reduce inflammation, immune response, and pain, respectively. Isofluorane was reduced to 1% and maintained at this level throughout surgery. Breathing and flinch responses were periodically checked to confirm an adequate level of anaesthesia. Mouse eyes were protected with mycochlorin eye ointment (Sato Pharmaceutical Co., Ltd). Hair was trimmed and removed with a hair trimmer and hair removal cream (Veet). The head was cleaned, sterilised with iodine, and numbed with topical analgesic (Lidocaine). The skin was cut with sterile scissors to expose the skull. Connective tissue was removed with a cotton swab, and the skin was held apart with clips. The lambda and bregma points on the skull were used to check that the head was level, and the stereotaxic frame was adjusted if necessary. Using a diamond drill, the area to be removed was measured with respect

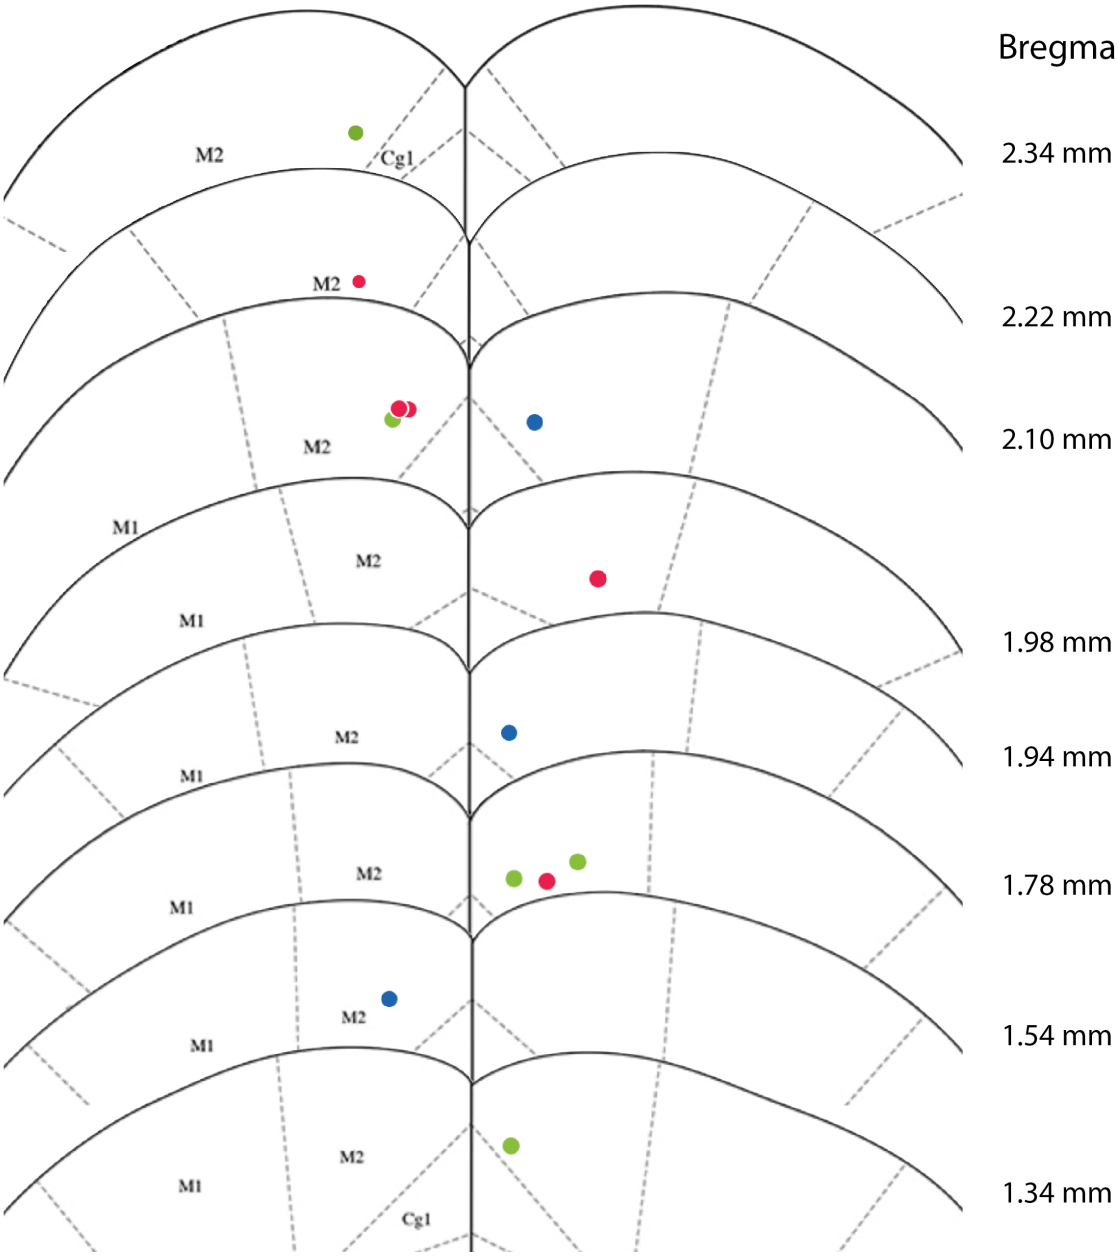

Figure 2.1: Coronal slices with injection sites marked. Colours signify the injected GENI. Green: serotonin, Red: Noradrenaline; Blue: Dopamine

.

to bregma. For the 5 mm diameter glass window to sit securely on a lip of skull, a 4.8 mm hole was excavated. The edge of the hole was thinned by repeated cycling around the rim until the remaining disk of the skull could be depressed with a light touch. A toothpick was glued to the bone disk and gently pulled to expose the dura without damaging it. Any bleeding was carefully stopped and the dura was cleaned with Carprofen-soaked gelfoam (Pfizer). The location of the virus injection site was found in relation to bregma. A glass pipette was used to uptake the virus, pierce the dura, and slowly descend to the required depth. Generally, mixtures of two viruses were used, one for a GECI and one for a GENI (see 2.4). 200 nl of each were mixed using a micropipette, placed on wax film and taken up into the pipette. Once at the correct depth, after a 5 minute wait to allow the tissue to settle around the pipette, approximately 70 nl of virus mixture was injected over the course of 5 minutes. After injection, the pipette was again allowed to rest for 5 minutes and then very slowly ascended. A pre-cut 5 mm glass window was placed directly onto the dura and glued to the skull at the rim of the craniotomy. An aluminium head-plate with central hole was mounted on top of the glass window and fixed with dental adhesive resin cement (Super-Bond). Dental cement was also used to cover any exposed skull, creating a stronger bond. Mice were given additional Carprofen and saline post surgery, placed on a heat mat until awake, and then individually housed for 2 days to recover.

2.4 AAV viruses

To measure local extracellular neuromodulator fluctuations alongside intracellular calcium dynamics, 2 indicators were transfected at each site. 70μl was injected into M2 to a depth of 0.6 mm, for imaging L2/3 above at 0.2 - 0.3 mm. See Figure 2.1 for injection sites.

- Red calcium indicator: AAV1.Syn.NES-jRGECO1a.WPRE.SV40 (AAV1) 2.95×10^{13} GC/ml
- Green Dopamine indicator: pAAV-hSyn-dLight1.2 (AAV5) 8.7×10^{12} GC/ml
- Green Noradrenaline indicator: pAAV-hSyn-GRAB_NE1m (AAV9) 1.8×10^{14} GC/ml
- Green Serotonin indicator: pAVV-hSyn-GRAB_5HT1.0 (AAV9) 6.05×10^{13} GC/ml

2.5 Two-photon imaging

Imaging was performed using a combined wide field and two-photon microscope (MOM, Sutter Instruments) with a 25x NA1.05 Olympus water immersion objective and a femtosecond-pulsed Ti:sapphire laser (Vision II, Coherent). The back aperture of the objective was under filled to create an elongated point spread function and a a spatial resolution 1 μm x 1 μm x 4 μm. A resonant scanner was used to acquire images of 512 x 512 pixels at 30.9 Hz, corresponding to a field of view of 375 μm x 375 μm.

Task variable	Tone part 1	Tone part 2	Total time
Sucrose cue	6kHz for 1000ms		1000ms
Air-puff cue	4kHz for 150ms	8kHz for 150ms	1000ms
Neutral cue	16kHz for 500ms	12kHz for 500ms	1000ms
State cue 1	2kHz for 200ms	0kHz for 200ms	2000ms
State cue 2	12kHz for 200ms	0kHz for 200ms	2000ms
State cue 3	1kHz for 200ms	8kHz for 200ms	2000ms
State cue 4	8kHz for 150ms	16kHz for 150ms	2000ms

Table 2.1: Table of 7 sound cues

Simultaneous excitation of a red calcium indicator (jRGECO1a) and a green neuromodulator indicator (either $\text{GRAB}_{5HT1.0}$, GRAB_{NE1m} or $\text{dLight}_{1.2}$), was performed at 950nm. Fluorescence was separated by a 560 nm dichroic mirror (Chroma) and detected in two channels by GaAsP photomultipliers (Hamamatsu) in spectral windows 490–550 nm (green) and 600–700 nm (red). The microscope was controlled by MScan software (Sutter Instruments).

2.6 Behavioural setup

The treadmill consisted of a metal rod pushed through as foam cylinder. The ends of the rods extend into holes in a metal frame that allows the foam cylinder to rotate freely. Above the wheel, a head fixation attachment was placed, with an adjustable arm to hold the water spout and air-puff delivery nozzles. Bilateral air-puff nozzles were used to mitigate any left-right hemisphere bias during bihemisphere imaging.

Delivered stimuli The behavioural paradigm was inputted to ABET II Software (Lafayette Instruments, 89501) and delivered in combination with ABET-lite Interface Module (Lafayette Instruments, 81427).

- Sound cues: 7 sound cues were delivered using a tone generator (Lafayette Instruments, 81415M).
- Liquid: 10% sucrose solution was delivered through a tube directed at the mouth using a peristaltic pump (Lafayette Instruments, 80204M).
- Air-puff: pressurised air was delivered through two tubes aimed at either side of the lower face and body using a solenoid (Lafayette Instruments, 82601).

Measured behavioural variables

- Locomotion: the velocity of the animal was measured using a rotary encoder (E6A2, Omron) attached to the axle of the treadmill.
- Behaviour video: facial expression and licking behaviour were monitored using a face-directed sCMOS camera and an infrared light source. Data were quantified using SLEAP pose tracking software.

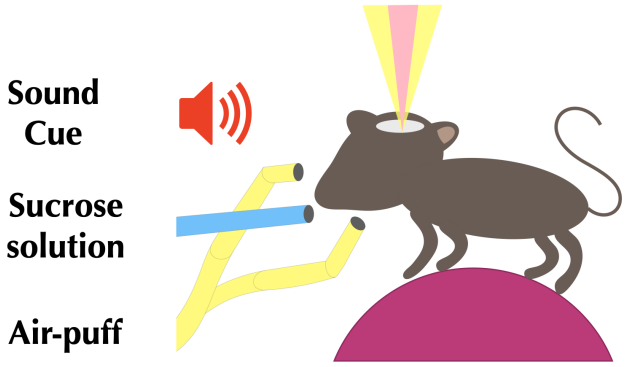

Figure 2.2: Behavioural setup: treadmill, water spout, two air-puff nozzles, and tone generator

2.7 Behavioural paradigm

2.7.1 Main experiment

The behavioural paradigm consisted of three stages, with increasing complexity. While waiting for sufficient virus expression (which takes approximately 3 weeks after injection), mice were habituated daily to the imaging setup. Sessions increased in time, building up to the 20 minutes planned for imaging. During these sessions, mice were presented with uncued appetitive and aversive stimuli, initially only sucrose solution and subsequently with interspersed sucrose and air-puffs.

Once the optimal imaging period began, the mice started a week of Pavlovian conditioning, Figure 2.3. Mice were exposed to three auditory cues followed, 2-4 seconds later, by a paired outcome, appetitive sucrose solution, aversive air-puff and neutral no-stimuli. In each session, 5 trial blocks of each cue-outcome set were rotated pseudo-randomly. Each day, mice performed 2 x 10 minute Pavlovian sessions, with ~30 trials per session.

After 7 days of Pavlovian conditioning, mice progressed to the Instrumental phase Figure 2.3. Prior to the outcome cues, hereinafter referred to as feedback cues, mice were presented with one of 4 state cues, signifying the 4 states: Positive Reinforcement of Go (Pos-Go), Positive Reinforcement of No-Go (Pos-NoGo), Negative Reinforcement of Go (Neg-Go), and Negative Reinforcement of No-Go (Neg-NoGo). Following the state cue was a 5 second action window within which mice were expected to either Go (move the treadmill at an average speed >5 cm/s) or No-Go (keep the average speed below 0.2 cm/s). Correct behaviour during the Positive Reinforcement states elicited the previously conditioned appetitive tone and sucrose delivery; whereas incorrect responses resulted in the neutral tone and no reward. Correct behaviour during the Negative Reinforcement states elicited just the neutral tone, while incorrect responses resulted in the conditioned aversive tone and an air-puff. In 10% of correct Positive Reinforcement trials and 10% of incorrect Negative Reinforcement trials the corresponding appetitive and aversive tones were played but the outcome was omitted,

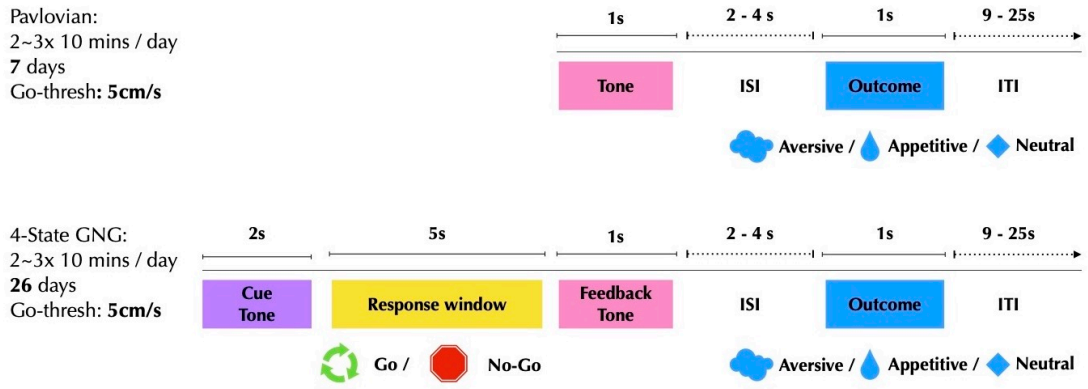

Figure 2.3: Paradigm schedule: 3 feedback cues, Pavlovian and Instrumental stages

thus the reinforcement was delivered with 90% probability. The same cue-outcome pairings were used throughout both Pavlovian and Instrumental stages. In total 4 state cues and 3 outcome cues were presented, each a different auditory cue. Additionally, during the inter-trial interval (ITI) in% 10 of trials mice were presented with an un-cued outcome, either sucrose solution or an air-puff. Each day mice performed 2 x 10 minute OGNG sessions, with ∼20 trials per session; blocks of 8-12 trials each state were pseudorandomly selected, generally resulting in 2 states sampled per session.

2.8 Perfusion and Histology

Within a week of completing all of the imaging mice were perfused. The animals were deeply anaesthetised with the 3-mixture anaesthetic and transcardially perfused with PBS followed by PLP (4% paraformaldehyde, 0.2% periodate and 1.2% lysine in 0.1 M phosphate buffer) fixation. The brains were extracted and stored in PLP at 4 C for a minimum of 48 h. Extracted brains were cut into 100 μm thick coronal sections using a vibratome (VT1000S, Leica). Slices were mounted on glass slides with Mowiol and stored at 4 C. Images were taken using vertical mounted Nixon digital camera.

2.9 Image Processing

2.9.1 Motion correction

After carefully reviewing the imaging data by eye, I found that some videos had some strong Z-motion artefacts. Z-motion is movement of the brain in the dorsal-ventral plane, and is a common problem in 2P microscopy due to the thin imaging plane. My case was more pronounced than others in the lab due to the proximity to the midline. In our experience, the midline cortex shows a greater degree of movement, likely due to increased tension from the spinal cord, as well as, some tissue warping from dilation and constriction of the nearby Superior Saggital Sinus.

Detecting Z-motion through simple observation of ROI time courses can be challenging. Z-motion can mimic calcium activity fluctuations, as brain movement upwards brings certain cells into the imaging plane, enhancing their brightness, and then downwards movement reduces it. Additionally, standard motion correction software may obscure these artefacts by correcting XY motion but inadvertently integrating Z-motion fluctuations into the analysed signal. However, by visually inspecting videos both before and after motion correction, we discovered the persistence of these artefacts.

I worked with a collaborator, Philipp Flotho, on a post hoc method to reduce the impact of Z-motion on the data. They developed a custom version of their Flow Registration motion correction software to target specific difficulties in my data. We noticed that in my recordings the brain tends to vacillate between a higher and lower position depending on the posture of the mouse. These two positions were easily distinguishable by simple K-means clustering. The k-means clusters were used to index each frame and to create 2 reference frames. Flow Registration then uses variational optical flow estimations to align each frame to its reference. The output is effectively a concatenation of interspersed video sections corresponding to one or the other reference frames. For subsequent analysis, the videos were split into peri-stimulus sections and only time courses that occurred entirely within one reference frame were used.

Flow Registration also outputs the mean displacement of each frame from its reference as a 1-dimensional time course. This mean displacement is the mean vector length of the optical flow displacement field, i.e. the mean distance each point has moved in the XY direction from the corresponding point in the reference image. To further reduce the possibility that the motion correction was creating artefacts, I diminished components that were correlated with the mean displacement. This was done by taking the single value decomposition (SVD) of a matrix containing both red and green channels reshaped so that one column represented one time point. The V components of the SVD, which here correspond to the time course, were checked for correlations with the mean displacement of the motion correction. Components that were correlated more than 30% were then reduced by decreasing the eigenvalues of those components. The amount of reduction was inversely proportional to the degree of correlation. The more correlated components received a greater degree of reduction, with a maximum at 20% of the original eigenvalue. Heuristically, it was found that reducing the correlated components to 20% of their initial strength reduced the impact of these components without producing any further artefacts. This SVD method was found to be more effective than the commonly used approach of taking the residuals of a linear regression, as that is subtractive, and thus can overcompensate to produce spurious anticorrelated signals.

2.9.2 Neuromodulator hotspots

To analyse the spatiotemporal dynamics of neuromodulator release, we examined the correlation of neuromodulator signals across different distances. Dividing the frames into 16x16 sections of approximately 25 μm each (Figure 2.4 Top Left) – about the size of a pyramidal cell – we calculated the temporal correlation between each pair of grid sections and plotted these as a function of their distance (Bottom Left). We

encountered high levels of low and anticorrelated activity in areas with minimal GENI expression, likely due to spurious activity at the edges of the imaging window and possible autofluorescence from senescent cells with low dynamic range. To isolate dynamically active areas, we defined 'hotspots' based on two criteria: 1) maximum intensity exceeding 50% of the average intensity across all grid sections, and 2) standard deviation of the time course surpassing the average of all sections. This approach highlighted a distinct set of time courses (Mid Right) with strong correlations (Bottom Right), effectively isolating regions with high signal-to-noise ratios. We found this method could capture both correlated and anticorrelated activity between hotspots.

2.9.3 Statistical analysis

Data was processed using MATLAB version 2023b. All statistical analyses were performed using the GraphPad Prism software version 10.

Behavioural data was first analysed in terms of state-wise performance, defined as the probability of the correct locomotion response in each session (the number of correct trials divided by the total number of trials) for each state. To assess whether performance improved over successive sessions, logistic regression was applied to the raw binary data (correct versus incorrect). Additionally, the average locomotion velocity during the task window was examined through a linear regression across sessions for each state. To compare early and late sessions, the velocities of each mouse in the initial 50 trials and the last 50 trials were compared using a repeated measures ANOVA.

Data was divided into short 4 to 10 second sections for analysis based on mouse locomotion and stimuli timing:

- **Rest** velocity within -0.3 and 0.3 cm/s for 80% of a 10 second section.
- **Movement** velocity exceeding 5 cm/s throughout a 10 second section.
- **Locomotion onset** for 1 second velocity between -0.03 and 0.03 cm/s, followed by 2/3 of the subsequent 3 seconds above 1cm/s.
- **Locomotion offset** for 2/3 of 3 seconds velocity above 1cm/s, followed by 2 seconds velocity between -0.03 and 0.03 cm/s.
- **Peri-stimulus** 1 second pre-stimulus baseline, followed by 3 seconds from stimulus onset.

To increase the sample size, broader velocity definitions of onset and offset were applied compared to rest or locomotion. Locomotion transitions often involved movements in the dorsal-ventral (Z) direction, requiring different reference frames (see Section 2.9.1). Only time courses within a single reference frame were included.

For the analysis of peristimulus activity and locomotion transitions, data was divided into pre- and post-periods. In time course plots, pre-periods are demarcated by the interval from -1 to 0, with the exception of the locomotion offset where the last second is used. The mean activity during the pre-period was employed for both dF/F

calculations and statistical analyses. For the post-period of locomotion onset and offset, the mean activity between 2:3 and -3:-2 seconds was used, respectively. We defined the post-period of unconditioned and conditioned stimuli as the 1 second interval surrounding the maximum absolute change in fluorescence, or peak flux, occurring within 2 seconds of stimulus onset. We also incorporated the delay to peak period in our subsequent analysis. To compare between states we used the fluorescence flux, calculated as the percentage difference between the mean of the pre-period and the mean of the post-period for each time course. Paired data points with percentage changes greater than 200% were excluded as they were commonly found to have remaining motion artefacts. Outliers in the data were identified and removed using the ROUT (Robust regression and Outlier removal) method in Prism, with a false discovery rate of 1%. Analysis was carried out with and without outlier removal, and no significant impact on the findings was observed.

We generated correlation-distance plots by computing the correlations between the time courses of each individual hot spot with each other; see Section 2.9.2 above. Additionally, we recorded the Euclidean distances between these hotspot pairs. The resulting plots illustrate the average correlated activity at varying distances, with distance values converted to micrometres in the plots for clarity.

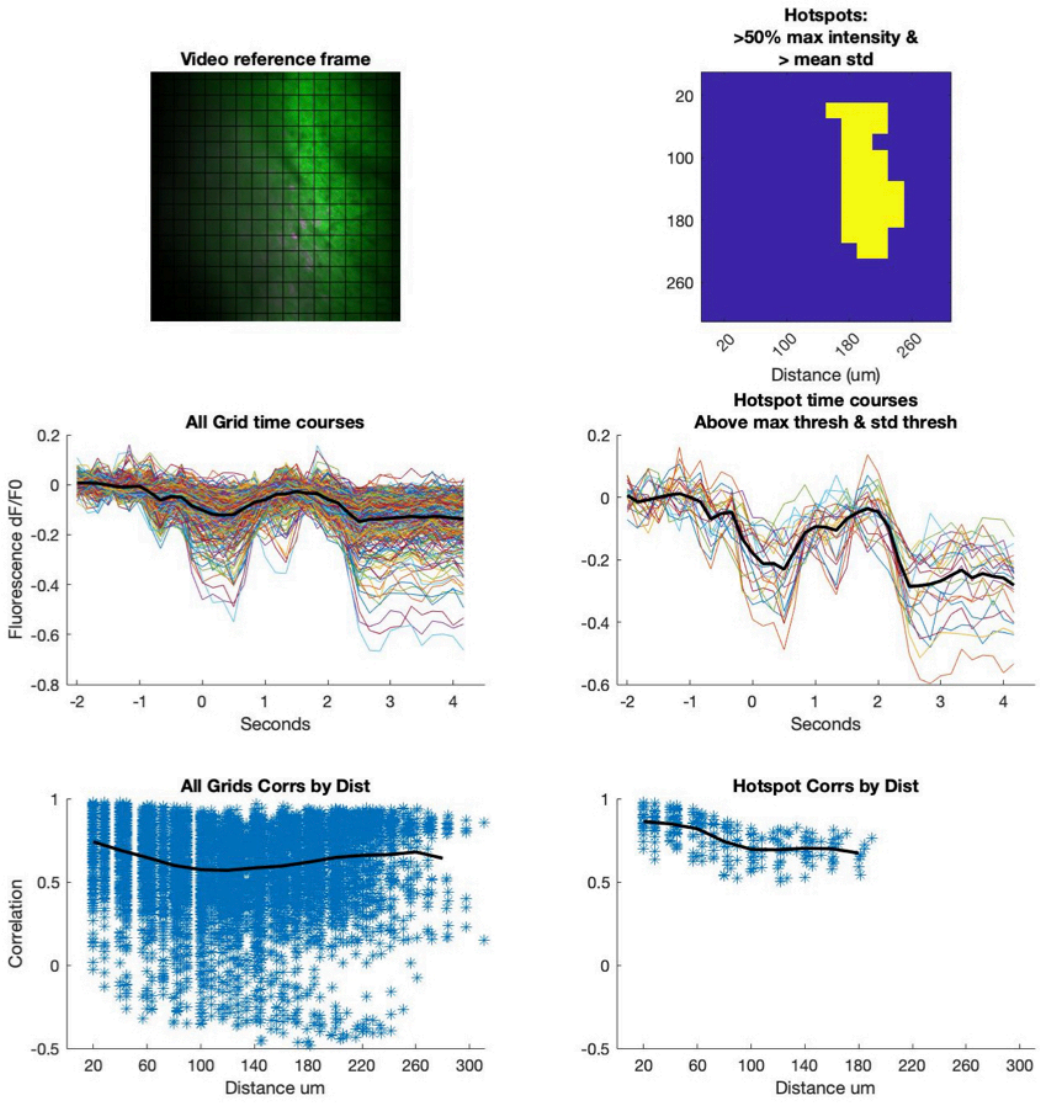

Figure 2.4: Methodology for finding activity hotspots. **Top Left:** Frames divided into 16x16 sections of $\approx 25\mu$m. **Top Right:** Sections are described as hotspots if they 1) have a maximum intensity (at one time point) greater than 50% of the mean intensity (all time), and 2) have a standard deviation (temporal) greater than the mean (all other sections) standard deviation. **Mid Left:** Time courses for all grid sections. **Mid Right:** Time courses for hotspot grid sections. **Bottom Left:** Temporal correlations between each grid section pair. **Bottom Right:** Temporal correlations between each hot spot pair.

Chapter 3

Results

3.1 Behavioural results

Let us begin by analysing the results of the OGNG task, designed to elucidate how the valence of unconditioned stimuli (US) influences instrumental conditioning. We chose this task to explore the motor cortices, where we anticipate dynamic changes in neuromodulator dynamics associated with both actions and outcomes. The OGNG task balances mobile and immobile behaviours, as well as positive and negative reinforcement, resulting in four states: Pos-Go, Pos-NoGo, Neg-Go, and Neg-Go.

As can be seen in Figure 3.1A, on average mice inherently performed well in the Go states of the task and poorly in the NoGo states. Average performance across all mice and in all sessions for the task states that required locomotion was 72% for Pos-Go and 69% for Neg-Go. For the NoGo states, which required stopping locomotion to succeed, the average performance was 21% for Pos-NoGo and 17% for Neg-NoGo. The performance across all mice did not show a significant improvement in any of the states. Logistic regression was used to quantify the relationship between the number of training sessions and the probability of performing the instructed action (Fig 3.1B). The slope of the odds ratio, $\beta 1$, indicates whether performance improves or declines with training. For Pos-Go $\beta 1 = 1.001$, which conveys that for every training session performance improved by only 0. 1%, which as expected was not a significant change $p = 0.9317$. The performance improvement for Neg-Go was only marginally better at $\beta_1 = 1.01$, $p = 0.1194$. The p value reported is for the Likelihood Ratio test, which assesses whether the slope of the performance change is significantly different from zero. Despite the very low starting performance, there were no notable changes in either of the NoGo states: Pos-NoGo $\beta_1 = 1.010$, $p = 0.1842$; Neg-NoGo: $\beta_1 = 0.984$, $p = 0.1194$. Overall, the inherent locomotive tendencies of mice to run while on a treadmill allowed for good performance in the Go states and poor performance in the NoGo states; however, there was no overall change in performance, suggesting that the mice did not learn the task.

Figure 3.1C shows substantial inter-mouse variability in performance. In Go states, some mice consistently performed well (close to 100%), while others showed performance as low as 25%. However, performance in NoGo states was generally very low. This variability explains the Area under the Curve (AUC) values in Figure 3.1A being

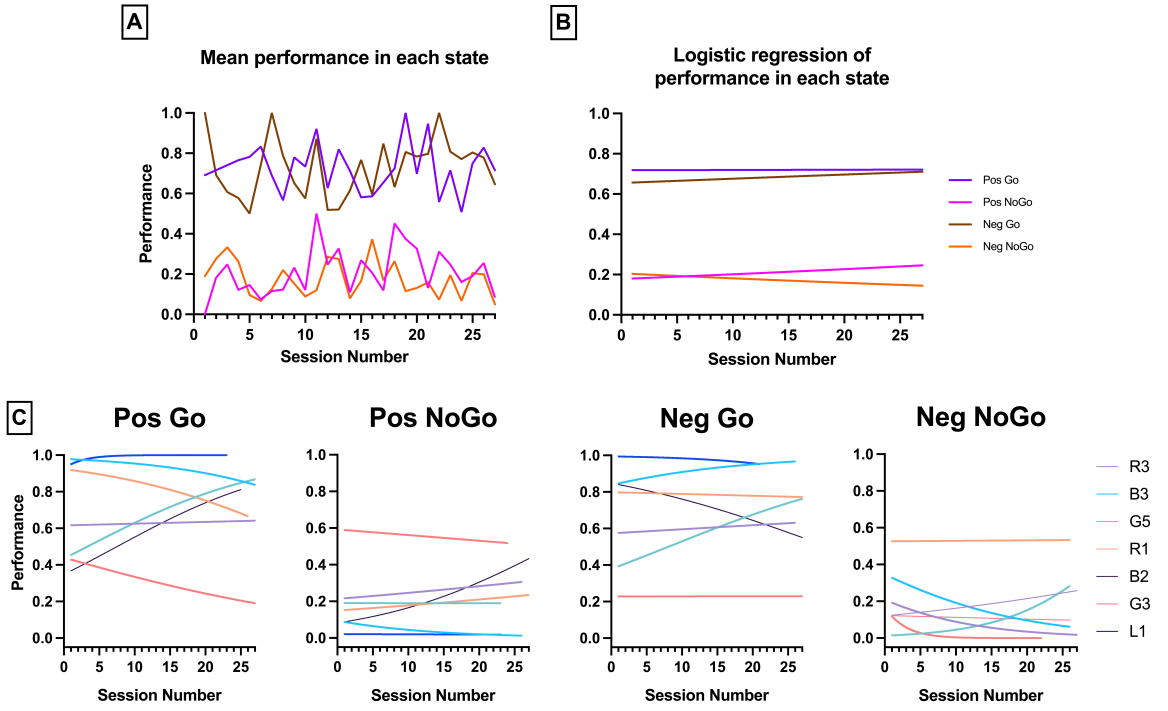

Figure 3.1: **A:** Mean performance per session for all mice. **B:** Logistic regression of performance of all mice. Regression was conducted on binary performance data (1 for success, 0 for failure) for each task state, using grouped data from all animals. The results for each state were: Pos-Go: Odds ratio (β_1) = 1.001, $p = 0.9317$, AUC= 0.5006. Pos-NoGo: Odds ratio (β_1) = 1.015, $p = 0.1194$, AUC= 0.5369. Neg-Go: Odds ratio (β_1) = 1.010, $p = 0.1842$, AUC= 0.5217. Neg-NoGo: Odds ratio (β_1) = 0.984, $p = 0.1194$, AUC= 0.5340. These results indicate that, on average, there was no significant improvement in performance in any state. The p-values reported are for likelihood ratio tests and were adjusted for multiple comparisons using the Holm-Šídák method. **C:** Logistic regression performed on binary performance data for each animal for each task state. Plots highlight variance in performance between mice. Each line represents one animal; colours are consistent between graphs. Statistical data for each regression line is summarised in Table A.1.

close to the chance value of 0.5, as they assess model fit. We conducted individual logistic regression analyses to summarise animal performance during each state Table A.1. We found that the mean performance improved in positively reinforced states and declined in negatively reinforced states, however this was mainly due to one animal. Median performance change per session, showed slight performance decreases across all states: Pos-Go -0.4%, Pos-NoGo -0.1%, Neg-Go 0%, Neg-NoGo -0.43%. Most individual mice experienced significant performance changes in at least one state, but these changes more often reflected decreased performance (6 out of the 10 significant performance changes within a state).

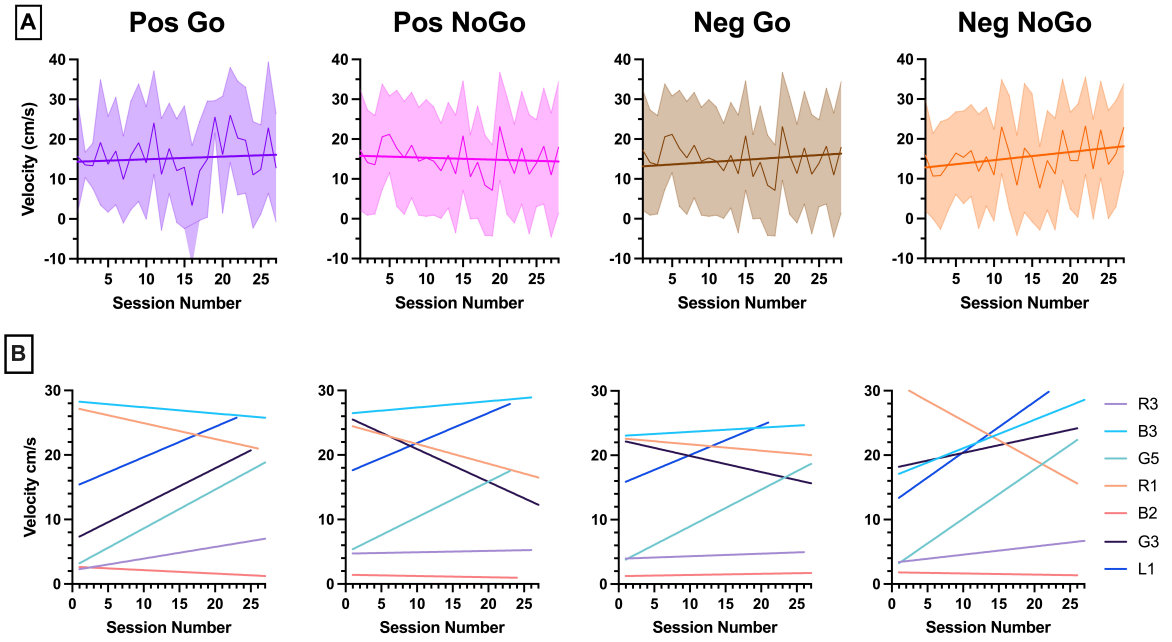

Figure 3.2: A: Mean, standard deviation and linear regression of mean locomotion velocity during the task response window for each task state. Pos-Go: Slope= 0.0671, R^2 = 0.0012, SD_{resid} = 14.20, F= 2.622, p = 0.2. Pos-NoGo: Slope=-0.0528, R^2 = 0.0009, SD_{resid} = 13.86, F= 1.741, p = 0.2. Neg-Go: Slope= 0.1180 R^2 = 0.0047, SD_{resid} = 14.10, F= 10.03, p = 0.0048. Neg-NoGo: Slope= 0.2047, R^2 = 0.0148, SD_{resid} = 13.40, F=, p = 0.0040. **B:** Linear regression performed on mean locomotion velocity for each mouse during the task response window for each task state. Line colours are consistent for each animal between graphs. In line with the task state goal: Positive Reinforcement Go: 4/7 animals *increased* velocity in line with the task requirement; Positive Reinforcement NoGo: 4/7 animals *decreased* velocity in line with the task requirement; Negative Reinforcement Go: 5/7 animals *increased* velocity in line with the task requirement; Negative Reinforcement NoGo: only 2/7 animals *decreased* velocity on in line with the task requirement.

Next, I asked whether changes in performance were related to changes in locomotion. Grouped data were subjected to simple linear regression for each task state, as seen in Figure 3.2A. In positively reinforced states, there appeared to be a velocity trend that might correspond to performance improvement, with slopes of 0.0671 for Pos-Go

and -0.0528 for Pos-NoGo. However, these trends were not statistically significant, with both having Holm-Šídák-adjusted p-values of $p = 0.2$. Moreover, reorganising the individual performance data presented in 3.2B revealed that this trend was primarily driven by mouse G3 and the low starting velocity in the Pos-Go state (Figure 3.3). Conversely, for both negatively reinforced states, the velocity increased significantly: Neg-Go: Slope = 0.1180, $p = 0.0048$; Neg-NoGo: Slope = 0.2047, $p = 0.0040$. This implies that mice generally accelerated in response to air-puffs, regardless of the task requirements. Notably, the high SD_{resid} values in each state (14.20, 13.86, 14.10 and 13.40) are close to the values of the regression line itself, highlighting the considerable variance in velocities measured. Figure 3.2B highlights this variance among individuals, resulting in very low R^2 values, all below 0.0015, indicating a poor model fit due to extensive variance.

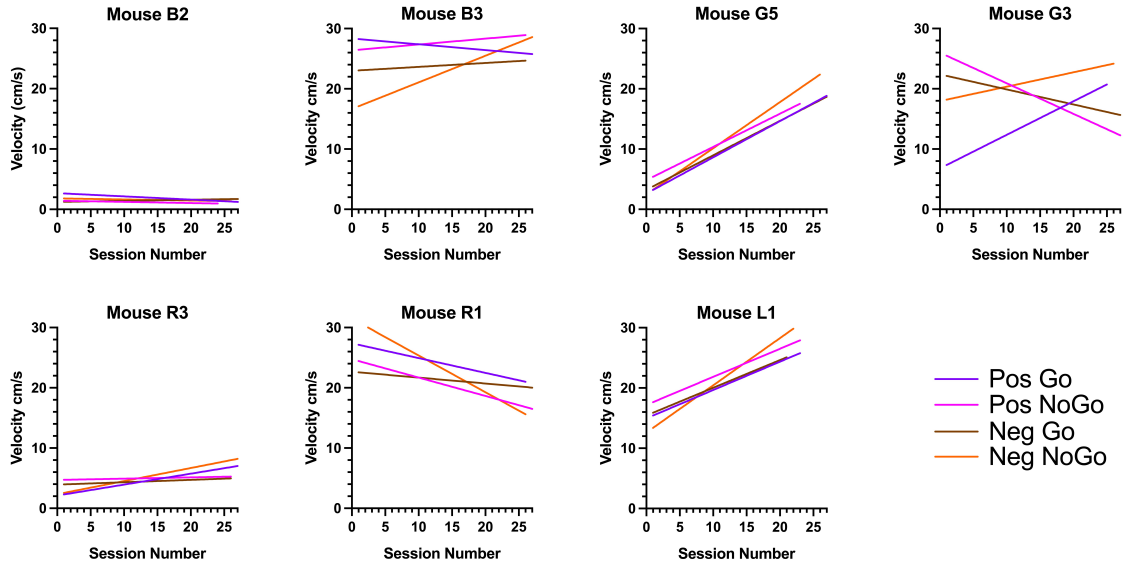

Figure 3.3: Simple linear regression performed on mean locomotion velocity during the task response window for each animal. Data re-plotted from Fig 3.2B to highlight the between animal variability. Animals have been grouped into rows by similarities in their locomotion profile. The four groups can be described as 'non-runners', 'runners', 'increasing speed' and 'mixed'.

To assess the differences between states indicated in Figure 3.2A, we employed a mixed-effects model with REML (Restricted Maximum Likelihood) analysis. For the fixed effects we investigated the influence of two factors, "State" and "Time," as well as their interaction. We compared the 50 first trials against the 50 last trials of each state for each mouse. Paired analysis was used to reduced the impact of individual mouse differences. We found statistically significant effects for all three main factors: "State" ($p = 0.0157$), "Time" ($p < 0.0001$), and the interaction "State x Time" ($p < 0.0001$). The statistical significance was confirmed by the corresponding F-statistics, with "Time" showing the highest effect size ($F(1, 349) = 20.56$). This suggests a strong difference in behaviour between the first and last training sessions.

The analysis also considered random effects, with "Mouse", "Mouse x State",

"Mouse x Time", and components. Notably, Mouse and Mouse x Time showed non-zero standard deviations (SD = 7.667 and SD = 5.401, respectively). This indicates that a significant amount of variability can attributed to differences between mice, and how the velocity of each mouse changed over time. The differences between mice was also captures using a chi-square to confirm effectiveness of matching ($p < 0.0001$), suggesting consistent behaviour of each mouse between time points.

Further post-hoc comparisons, using Holm-Šídák's multiple comparisons test, revealed differences between the first week and last week. Between the early and late trails was a significant increase in velocity in the Pos Go (Mean Diff. = 4.468, $p < 0.0001$), and Neg NoGo (Mean Diff. = 4.324, $p < 0.0001$) states. Moreover, in the last week there were statistically significant differences observed between the different states; no differences were seen between states in the first week. This indicates that mice had adapted the behaviour in relation to task states.

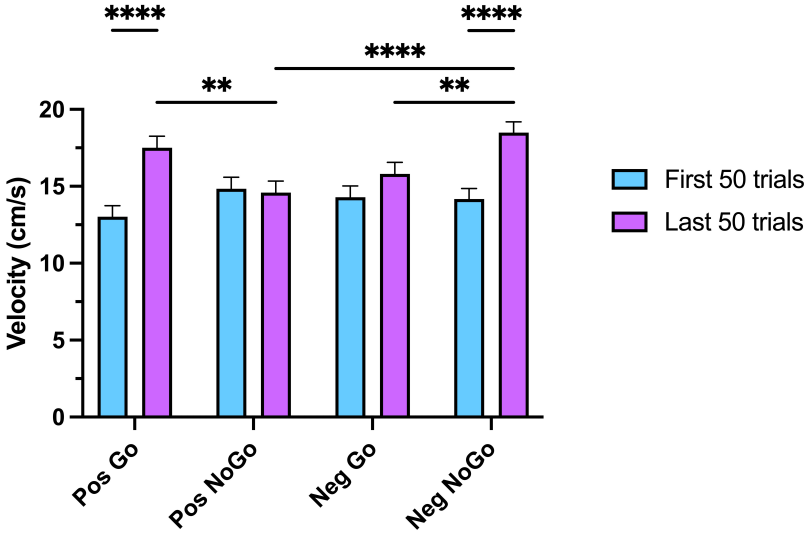

Figure 3.4: Mixed-effects model applied to assess differences in velocity during the response window across task states, using the first and last 50 measurements from each mouse for each state. Results from Holm-Šídák's multiple comparisons test highlight significant increases in velocity in the Pos - Go ($p =< 0.0001$) and Neg NoGo ($p =< 0.0001$) states. This resulted in significant differences emerging between states in the later trials, Pos NoGo - Pos Go ($p = 0.0062$), Neg NoGo - Pos NoGo ($p =< 0.0001$), and Neg NoGo - Neg Go ($p = 0.0097$). Subsequently, the mixed-effects model analysis revealed significant effects of State ($p = 0.0157$), Time ($p < 0.0001$), and their interaction (State x Time, $p < 0.0001$), with the largest effect size attributed to "Time" ($F(1.000,349.0) = 20.56$). Random effects analysis indicated non-zero standard deviations for Mouse (SD = 7.667) and Mouse x Time (SD = 5.401), underscoring individual variations in response. A chi-square test ($p < 0.0001$) confirmed the effectiveness of the matching criteria used. Error bars represent the SEM.

In summary, the behavioural results indicate that while there were minimal changes in task performance, mice did learn to adapt their behaviour in a state dependent manner. Many individual mice exhibited statistically significant performance changes

in certain task states through shifts in locomotion patterns. However, these adaptive behaviours were not optimised to improve task performance, but instead may reflect more Pavlovian approach and avoidance behaviours.

3.2 Imaging results

Our imaging results aim to characterise the temporal and spatial release patterns of serotonin, noradrenaline, and dopamine in the motor M2 region. We initially compared neuromodulator activity during rest and locomotion and transitions between these states. Subsequently, we explored how each neuromodulatory system responded to unconditioned stimuli, including appetitive sucrose and aversive air-puff. For all video sections we ensured the entire video section was stable and used a single reference frame, as determined using Flow Registration. We analysed 10s video sections for rest and locomotion. For locomotion onset and offset, unconditioned and conditioned stimuli were examined in 4s peristimulus sections. Time courses shown here are the mean of activity hotspots, see Section 2.9.2 for details. Lastly, we present the results of our pharmacological tests.

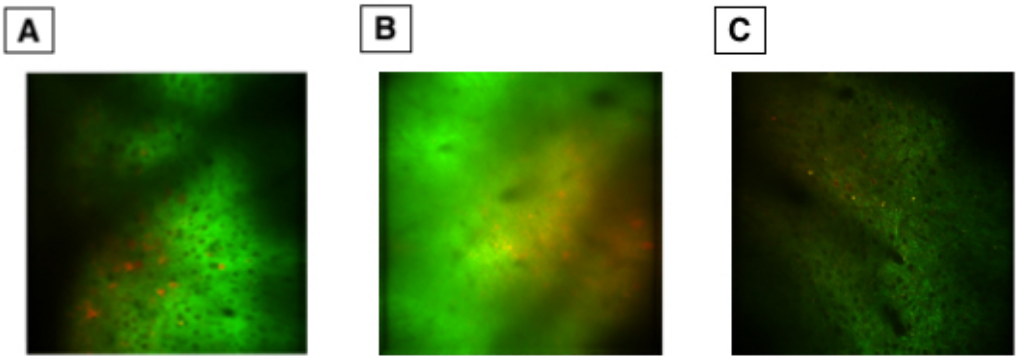

Figure 3.5: Examples of dual GENI and GECI expression for each neuromodulator. jRGECO1a GECI in red, GENIs in green. **A**: $GRAB_{5HT1.0}$, **B**: $GRAB_{NE1m}$, and **C**: $dLight_{1.2}$

Figure 3.5 illustrates the co-expression patterns of GENI and GECI, highlighting an interaction between the viruses used, which resulted in differential expression levels across different areas. This phenomenon, known as 'tropism competition,' is well-documented in epidemiology but seldom discussed in experimental contexts. Despite using a sub-optimal excitation wavelength of 950 nm for jRGECO1a, we believe that the absence of GECI expression in areas of high GENI expression is not due to excitation inefficiency, but the result of tropism competition. Serotonin GENI (Fig 3.5 A) often showed a balanced expression with strong jRGECO1a presence across the imaging window. Noradrenaline GENI (Fig 3.5 B), while bright and strong, produced divided expression due to the apparent repulsion between virus vectors, leading to separation of GENI and GECI. Dopamine GENI (Fig 3.5 C) was broadly expressed but dimmer compared to others; GECI expression was reasonable.

Given these distinct expression patterns and the varied onset/offset dynamics of GENI ($GRAB_{5HT1.0}$:$\tau_{on/off}$ 150ms/7220ms, $GRAB_{NE1h}$:$\tau_{on/off}$ 72 / 680 ms, $dLight_{1.2}$: $\tau_{on/off}$ 9.5ms / 90ms), we refrained from statistical comparisons between neuromodulator systems. The individual differences among mice and the failure of many datasets to pass normality tests led us to use non-parametric paired tests, accommodating individual variability.

3.2.1 Rest versus locomotion

We defined rest as zero movement of the treadmill for 7 seconds; 5 seconds for analysis plus 1 seconds on either side as a buffer from locomotion transitions. Sustained locomotion was defined similarly as sustained movement of the treadmill above 25 cm/s for 7 seconds. Mean time traces are plotted in Fig 3.6. The the F0 value for each 10-minute session video for the dF/F0 was calculated by taking the XYZ mean fluorescence of all of the 7 s rest video sections, and selecting the lowest. The same F0 value was used for all the rest and movement video sections from the same 10-minute session. There was a significant difference between the sustained rest and locomotion states for serotonin, noradrenaline and dopamine ($p = 0.0003$ for all). The cross-correlation between fluorescence and treadmill velocity for each neuromodulator showed high values, with a peak at 0 s and elevated tails in both positive and negative lags (serotonin $p = 0.018$, noradrenaline $p = 0.018$, dopamine $p = 0.027$). These results suggest strong temporal synchronisation or consistent patterns between the two signals.

3.2.2 Locomotion transitions

We investigated neuromodulator responses during locomotion transitions. There was a striking similarity in activity associated with locomotion onset and offset between all three neuromodulatory systems. Here, locomotion onset was defined as a 4-second period with an initial second of treadmill speed between -0.03 and 0.03 cm/s, followed by 2/3 of the subsequent 3 seconds above 1cm/s. Locomotion offset was defined as having a treadmill speed above 1cm/s for 2/3 of the initial 3 seconds, followed by 1 second at a speed between -0.03 and 0.03 cm/s; see Methods Section 2.9.3.

In Figure 3.8a we can see a strong increase in fluorescence occurs with locomotion onset for all three neuromodulators (serotonin $p = 0.009$, noradrenaline $p = 0.009$, dopamine $p = 0.018$). The mean cross-correlations between fluorescence and treadmill velocity show a strong immediate correlation (peak at lag 0), followed by a gradual decrease in correlation as the time lag increases. The raised negative lags indicate that the neuromodulator signals precede changes in velocity. On the contrary, a substantial decrease in fluorescence was observed during locomotion offset, Figure 3.8b. Likewise, the mean cross-correlations show a significantly high positive lags values for serotonin ($p = 0.009$), noradrenaline ($p = 0.009$), and dopamine ($p = 0.018$). This indicates a more rapid decrease in locomotion velocity than neuromodulator activity when stopping locomotion.

We used Wilcoxon paired non-parametric tests to compare the mean pre- and post-period fluorescence (Fig 3.9). Serotonin, noradrenaline and dopamine all showed significant increases during onset and decreases during offset (all $p = 0.0009$). For each

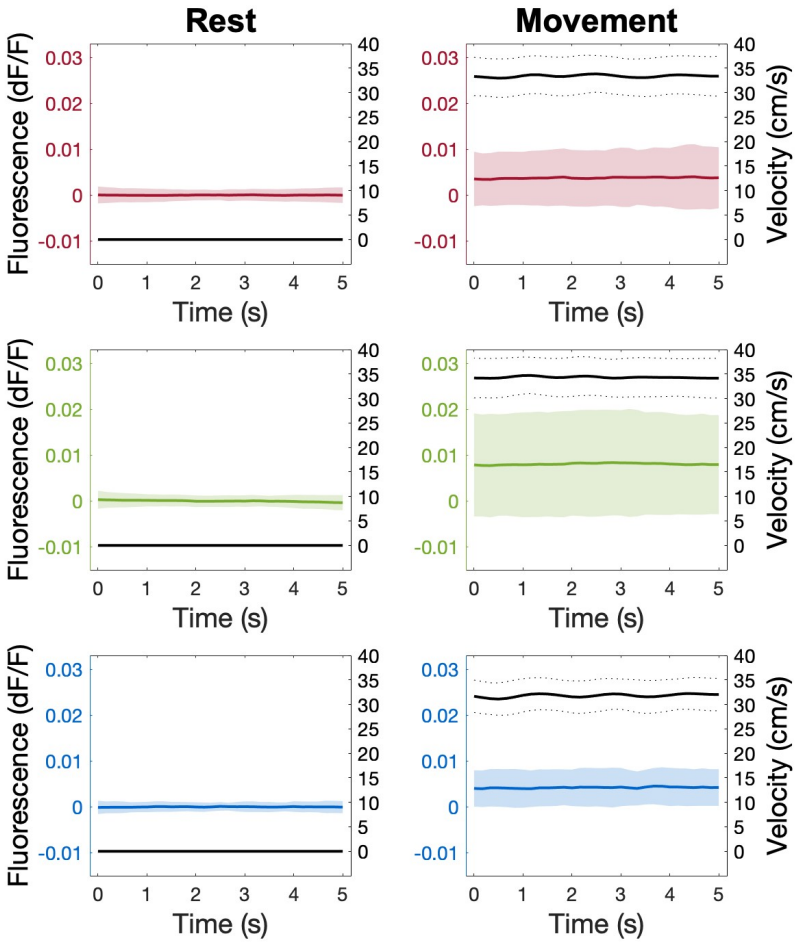

Figure 3.6: The time course of mean fluorescence levels during different locomotion states. Sustained rest corresponds to periods with a sustained treadmill movement speed of 0 cm/s, while sustained locomotion denotes speeds at or above 25 cm/s. The shaded region depicts the standard deviation. The calculation of dF/F involved examining all 1-second time courses from the at-rest video sections of each video and using the lowest XYZ-mean fluorescence value as F0. The legend provides information about the number of video sections used for calculating the mean. Detailed statistical analysis is presented in Fig 3.7.

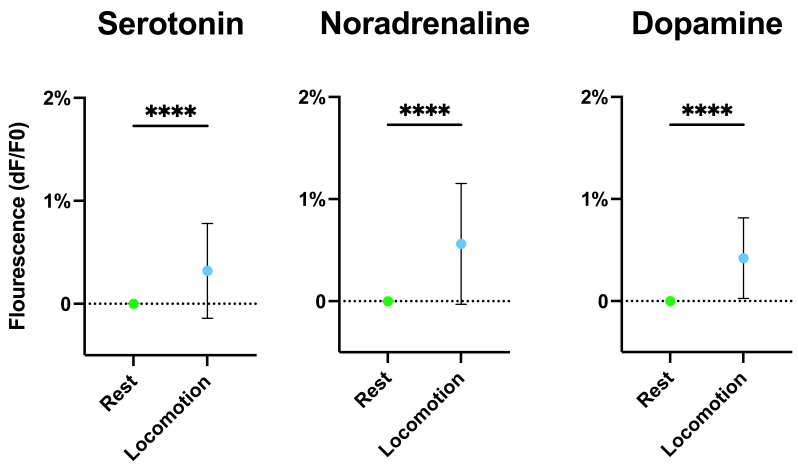

Figure 3.7: Median GENI fluorescence levels during periods of rest and locomotion. Mann-Whitney testing revealed a statistically significant difference between sustained rest and sustained locomotion for all three neuromodulators ($p = 0.0003$). All p-values adjusted for multiple comparisons using the Holm-Šídák method.

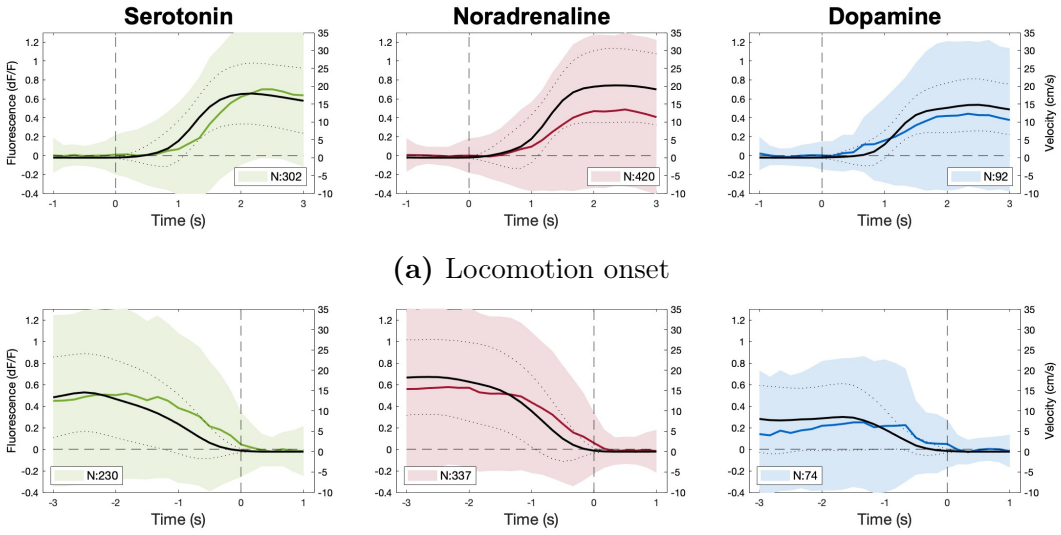

(a) Locomotion onset

(b) Locomotion offset

Figure 3.8: Mean and standard deviation of fluorescence levels during transition between rest and locomotion; velocity mean and standard deviation in black. Locomotion onset defined as 1 second between -0.03 and 0.03 cm/s followed by 2/3 of the subsequent 3 seconds above 1cm/s. Offset follows the same criteria in reverse. The calculation of dF/F used the mean fluorescence of between -1 and 0 for onset and 0 and 1 for offset. The legend states the number of video samples used. Statistical analyses is presented in Fig 3.9.

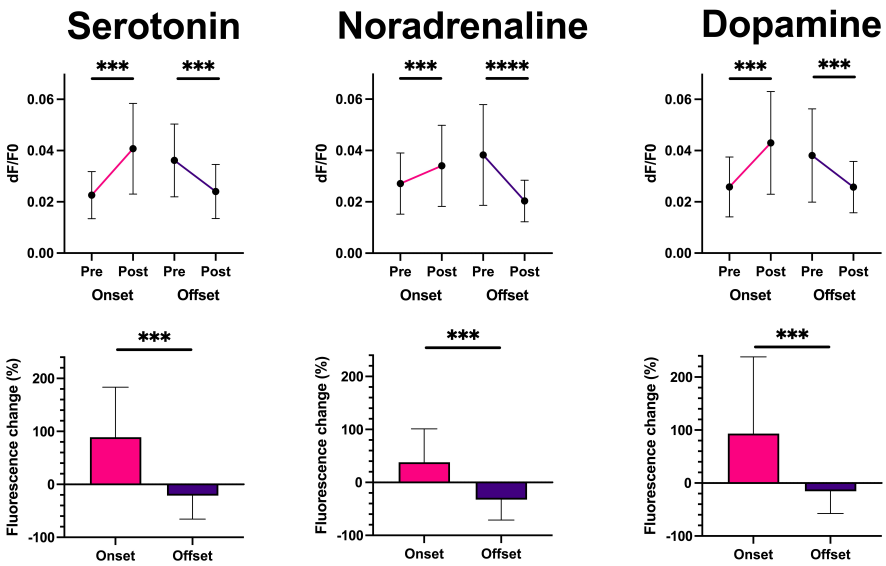

Figure 3.9: Analysis of median GENI fluorescence levels during pre- and post-locomotion transition periods. **Top:** Within each state, significant differences were observed in serotonin, noradrenaline, and dopamine levels during both onset and offset (all p = 0.0009; Wilcoxon test). **Bottom:** Comparing onset and offset states, a significant difference was observed for serotonin, noradrenaline, and dopamine (all p = 0.0009; Mann-Whitney test). All p-values were adjusted using the Holm-Šídák method.

neuromodulator we also confirmed the differences between onset and offset groups using Mann-Whitney tests, all p = 0.0009.

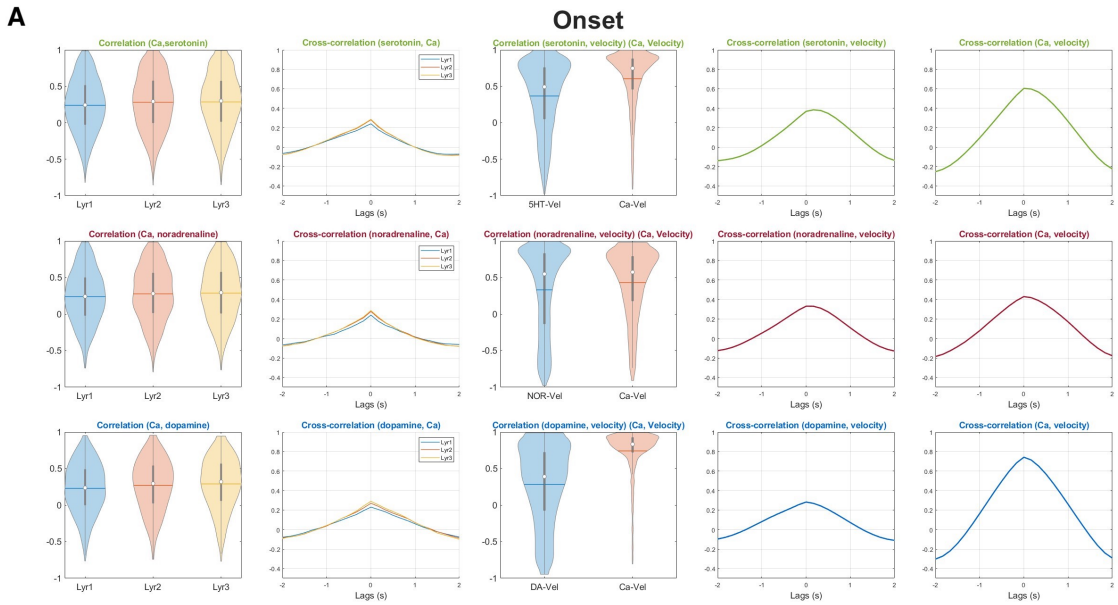

Figure 3.10: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion onset. **Serotonin**: median correlation serotonin-calcium (0.27), cross-correlation peak serotonin-calcium (0.29 at 0s), median correlation serotonin-velocity (0.49), median correlation calcium-velocity (0.75), cross-correlation peak serotonin-velocity (0.38 at 0.17s), cross-correlation peak calcium-velocity (0.6 at 0s). **Noradrenaline**: median correlation noradrenaline-calcium (0.28), cross-correlation peak noradrenaline-calcium (0.29 at 0s), median correlation noradrenaline-velocity (0.55), median correlation calcium-velocity (0.57), cross-correlation peak noradrenaline-velocity (0.33 at 0s), cross-correlation peak calcium-velocity (0.43 at 0s). **Dopamine**: median correlation dopamine-calcium (0.29), cross-correlation peak dopamine-calcium (0.29 at 0s), median correlation dopamine-velocity (0.39), median correlation calcium-velocity (0.83), cross-correlation peak dopamine-velocity (0.28 at 0s), cross-correlation peak calcium-velocity (0.74 at 0s).

We investigated whether changes in neuromodulator release during locomotion transitions influenced neuronal cell activity. By comparing intracellular calcium activity, as an indicator of neurone activity, with extracellular neuromodulator activity within three 5 μm concentric rings around each central soma ROI, we observed distinct patterns during locomotion onset and offset. As depicted in Figure 3.10, during the onset phase, there was a noticeable correlation between each neuromodulator and nearby calcium activity (serotonin 0.27, noradrenaline 0.28, dopamine 0.29), indicating synchronous rises and falls (0 peak cross-correlation). We found only minimal variation between the concentric circles, suggesting homogeneous neuromodulator release within this range. A robust correlation was also seen between neuromodulators and velocity, particularly strong in serotonin (0.49) and noradrenaline (0.55), with dopamine (0.29) showing some anti-correlated activity. Calcium activity correlations with velocity were

even more pronounced, ranging from 0.57 to 0.83.

Shifting to the offset phase, as shown in Figure 3.11, we found that correlations were generally lower. Median correlations with calcium dropped to 0.23 for serotonin, 0.26 for noradrenaline, and 0.19 for dopamine. Similarly, the neuromodulator correlations with velocity were reduced (serotonin 0.35, noradrenaline 0.51, dopamine 0.17), due to a larger proportion of uncorrelated and anti-correlated activity. Variability in calcium activity's correlation with velocity was also evident, ranging between 0.15 and 0.34. Notably, the shift in peak cross-correlation times to positive lags for calcium and velocity implies a precedence of velocity changes over calcium activity alterations during locomotion offset.

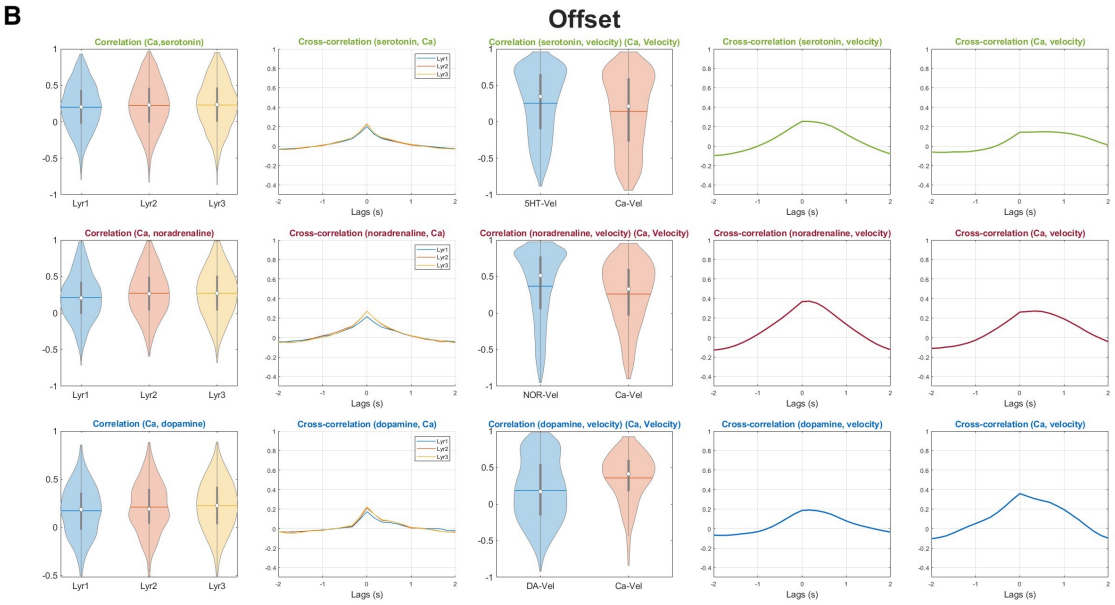

Figure 3.11: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion offset. **Serotonin**: median correlation serotonin-calcium (0.23), cross-correlation peak serotonin-calcium (0.23 at 0s), median correlation serotonin-velocity (0.35), median correlation calcium-velocity (0.21), cross-correlation peak serotonin-velocity (0.25 at 0s), cross-correlation peak calcium-velocity (0.15 at 0.68s). **Noradrenaline**: median correlation noradrenaline-calcium (0.26), cross-correlation peak noradrenaline-calcium (0.27 at 0s), median correlation noradrenaline-velocity (0.51), median correlation calcium-velocity (0.33), cross-correlation peak noradrenaline-velocity (0.37 at 0.17s), cross-correlation peak calcium-velocity (0.27 at 0.34s). **Dopamine**: median correlation dopamine-calcium (0.19), cross-correlation peak dopamine-calcium (0.22 at 0s), median correlation dopamine-velocity (0.17), median correlation calcium-velocity (0.17), cross-correlation peak dopamine-velocity (0.19 at 0.17s), cross-correlation peak calcium-velocity (0.34 at 0s).

.

3.2.3 Unconditioned and Conditioned Stimuli

Upon the presentation of appetitive sucrose solution and the aversive air-puff, each neuromodulator system showed a similar yet subtly heterogeneous response profile (Fig 3.13). We noted that the standard deviation of activity was much greater than in our previous analysis. This diversity of responses occluded the effectiveness in using the greatest absolute peak, as this measure did not capture the intricacies of the profiles observed in the mean activity. By examining both the peak minima and peak maxima in the post-stimulus period, we uncovered a wider array of activity profiles. Each neuromodulator exhibited significant minima and maxima in every state ($p = 0.0024$ for all). However, the magnitude and timing displayed marked variations.

The serotonergic responses to sucrose displayed marked differences between cued and uncued conditions. In the cued scenario, the responses comprised 39% minima-only, 23% maxima-only, and 36% mixed. This led to a biphasic response characterised by an initial drop followed by a subsequent rebound. In contrast, uncued conditions showed a higher frequency of minima-only responses at 55%, alongside 12% maxima-only and 28% mixed. This resulted in a more prolonged dip in mean fluorescence, although some traces suggested a possible rebound. A similar increase in minima-only responses was observed for noradrenaline in uncued compare to cued sucrose conditions, leading to a significantly delayed peak minima ($p < 0.0001$), as shown in Fig 3.12c. Dopaminergic responses, however, maintained a consistent dip in activity for both cued and uncued sucrose. This consistency translated into a smaller standard deviation in activity levels compared to the other neuromodulators.

The mean time courses of neuromodulator responses to air-puffs generally appeared smaller than those to sucrose. For serotonin and dopamine, these reductions in magnitude were not statistically significant. The minima-maxima composition of serotonergic and dopaminergic responses remained largely consistent between cued and uncued air-puff conditions. However, these profiles diverged from the responses to sucrose, displaying a higher proportion of maxima-containing time courses. This elevated ratio of maxima could account for the more subdued mean activity, even though the magnitudes of minima and maxima were statistically indistinguishable for both neuromodulators, as seen in Figure 3.13. For dopamine responses to cued air-puffs, a significant peak-to-lag difference was observed ($p < 0.0001$), with minima preceding maxima. A similar pattern was seen across all US responses and was visible in the time courses. Finally, noradrenergic activity demonstrated a modest yet significant reduction in response to both cued and uncued air-puffs. Interestingly, uncued air-puffs elicited weaker responses compared to cued air-puffs ($p = 0.0106$) and uncued sucrose ($p < 0.0001$). Additionally, a higher proportion of maxima-containing responses was noted for cued air-puffs.

We then examined the correlations and cross-correlations between neuromodulator, calcium, and velocity signals during each US (Figs 3.14, 3.15, 3.16, 3.17). We observed notably weaker correlations compared to those seen during locomotion transitions. Of note, our analysis here differentiated between cued and uncued US but did not factor in velocity, resulting in the inclusion of velocity time courses of 0, corresponding to periods of rest.

For serotonin, we identified very weak correlations with calcium activity during

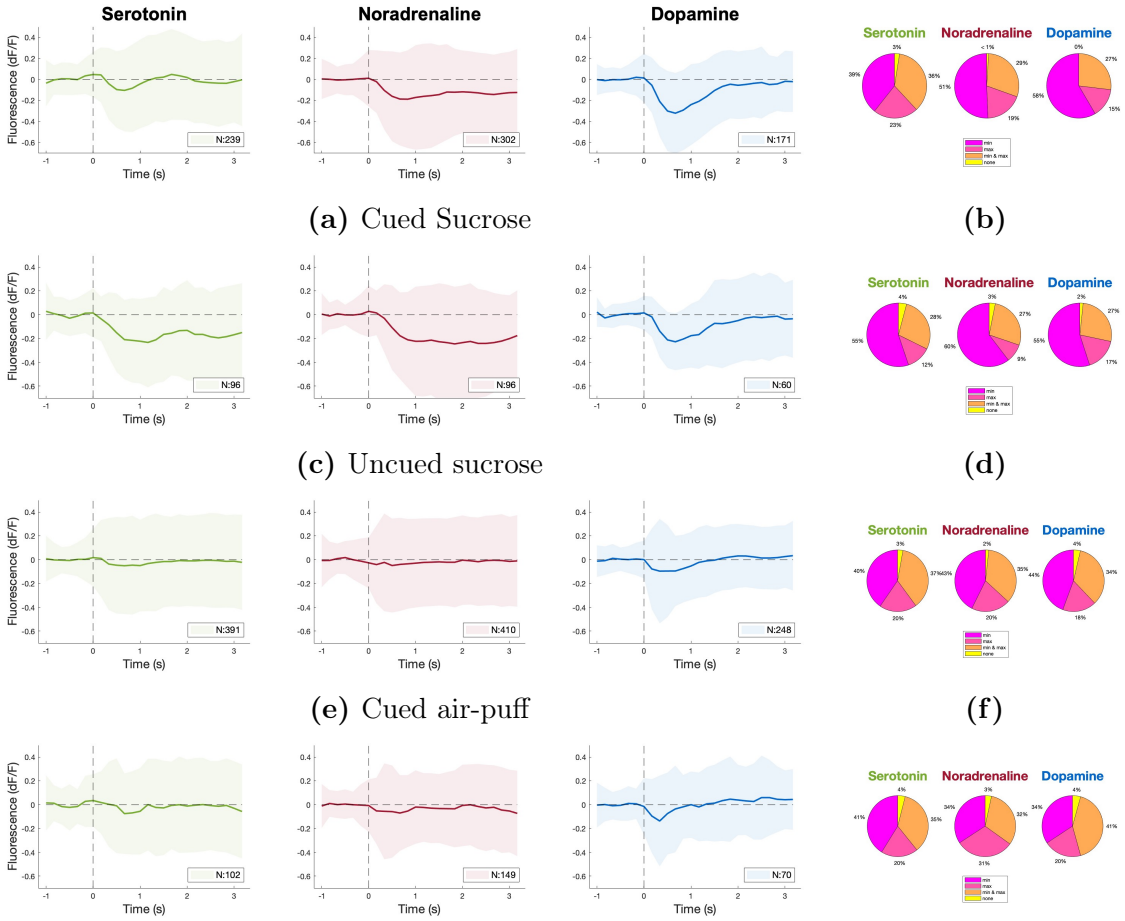

(a) Cued Sucrose **(b)**

(c) Uncued sucrose **(d)**

(e) Cued air-puff **(f)**

(g) Uncued air-puff **(h)**

Figure 3.12: Fluorescence levels around the time of unconditioned stimulus delivery, involving both sucrose solution and air-puff, are depicted. The mean is represented by the line, and the standard deviation is indicated by the shading. To normalise the data using dF/F, the pre-stimulus period, spanning from -1 to 0, was employed. The stimulus onset took place at time point 0. The legend specifies the number of video sections utilised for this analysis. Detailed statistical analysis is available in Figure 3.13.

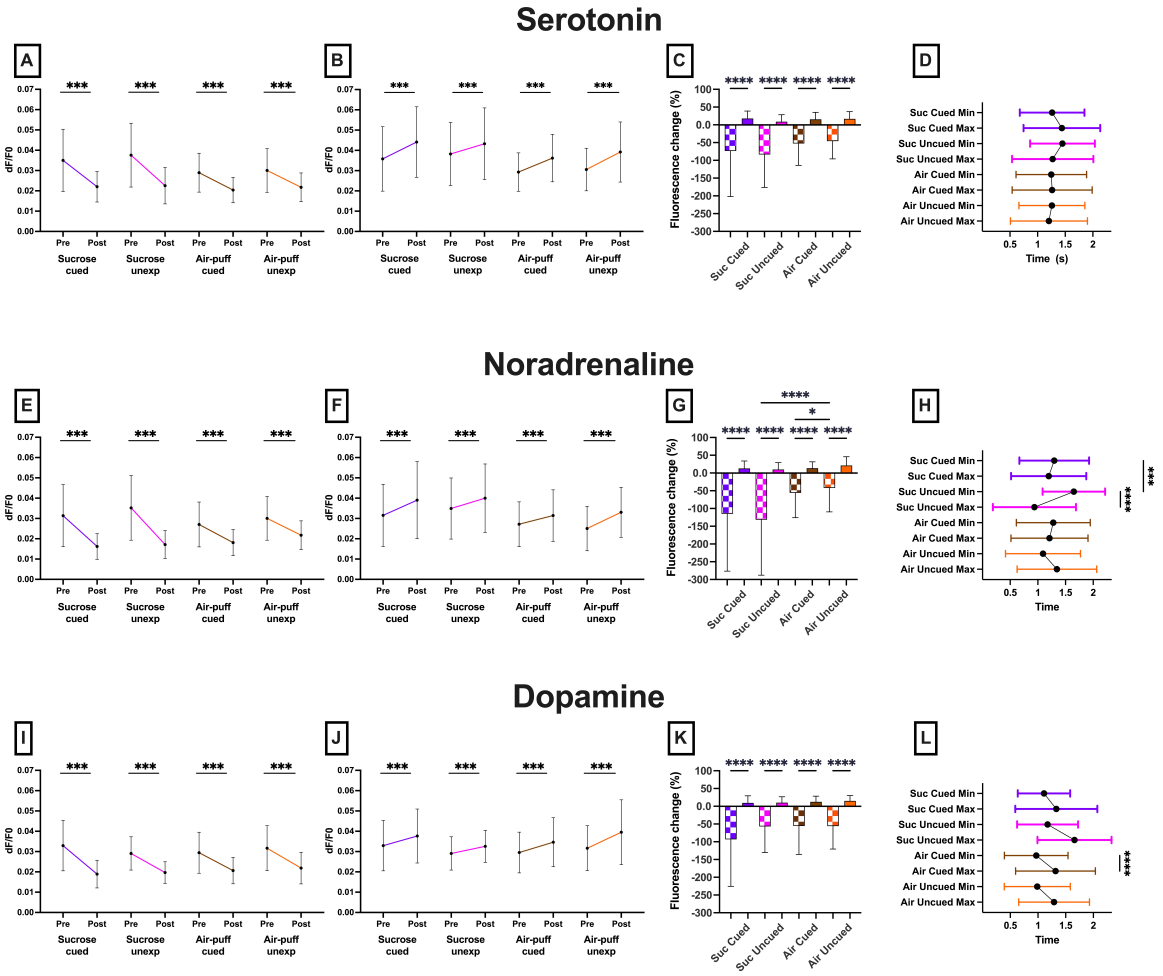

Figure 3.13: Analysis of peristimulus activity around US, as depicted in Fig 3.12. **A, E, I:** Peak post-stimulus minima with respect to pre-stimulus. We observed significant median decreases for serotonin, noradrenaline and dopamine in all states ($p = 0.0024$ for all). **B, F, J:** Peak post-stimulus maxima with respect to pre-stimulus. Significant median increases were observed for serotonin, noradrenaline and dopamine in all states ($p = 0.0024$ for all). **C, G, K:** Comparison between states. Minima chequered, maxima block colours. Significant differences seen between minima and maxima in all states. Furthermore, for noradrenaline in cued vs uncued air-puff minima ($p = 0.0106$), and uncued sucrose vs uncued air-puff ($p < 0.0001$). **D, H, L:** Comparison of delay to peak minima and peak maxima within 2 seconds of stimulus onset compared to the pre-stimulus period. Noradrenaline showed delay between minima and maxima of uncued sucrose ($p < 0.0001$), and between peak minima for uncued to cued sucrose ($p = 0.0003$). Dopamine between minima and maxima of cued air-puff ($p < 0.0001$). Left, centre left used paired Wilcoxon tests, corrected using the Holm-Šídák method. Right, centre right used non-parametric Kruskal-Wallis ANOVA.

sucrose delivery (cued and uncued both at 0.23), but no notable correlations during air-puff trials (cued 0.16, uncued 0.17). Serotonin showed no significant correlation with velocity across these data sets (cued sucrose 0.06, uncued sucrose 0.13, cued air-puffs 0.06, uncued air-puffs 0.17).

Noradrenaline exhibited a weak correlation with calcium activity during cued sucrose (0.23) and uncued air-puffs (0.23), and a mild correlation during uncued sucrose (0.33). Noradrenergic correlation with velocity was weak during sucrose (0.24 for both cued and uncued), but absent during air-puff delivery.

Dopamine activity showed no correlation with calcium activity across all US conditions (cued sucrose 0.14, uncued sucrose 0.13, cued air-puff 0.09, uncued air-puff 0.1). Furthermore, calcium activity during both cued and uncued US generally did not correlate with velocity, underscoring a distinct pattern from locomotion transitions.

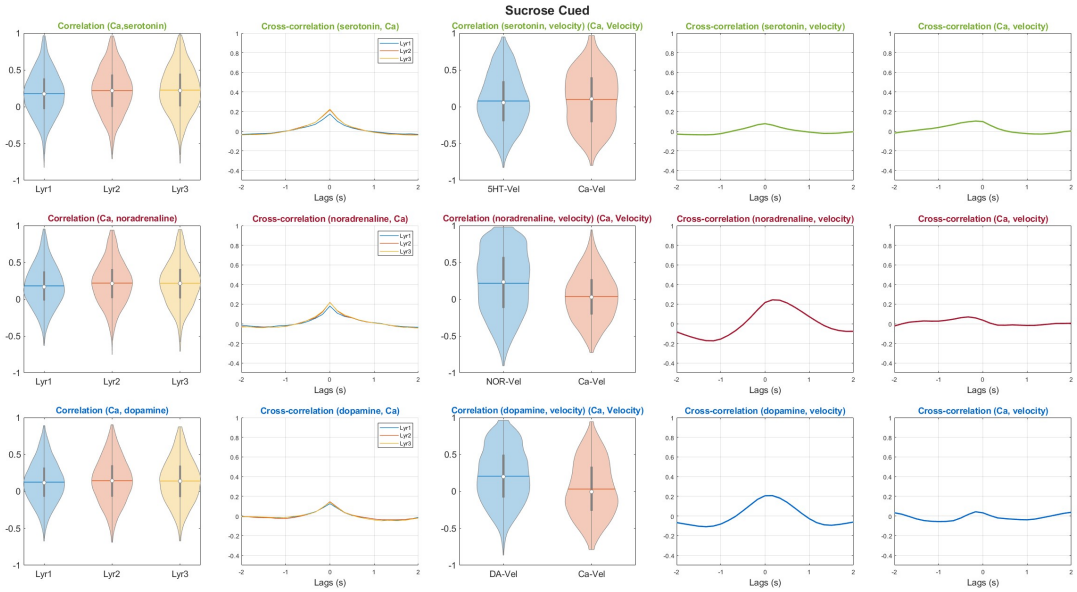

Figure 3.14: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion cued sucrose. **Serotonin**: median correlation serotonin-calcium (0.22), cross-correlation peak serotonin-calcium (0.23 at 0s), median correlation serotonin-velocity (0.06), median correlation calcium-velocity (0.11), cross-correlation peak serotonin-velocity (0.08 at 0s), cross-correlation peak calcium-velocity (0.1 at -0.17s). **Noradrenaline**: median correlation noradrenaline-calcium (0.21), cross-correlation peak noradrenaline-calcium (0.22 at 0s), median correlation noradrenaline-velocity (0.23), median correlation calcium-velocity (0.03), cross-correlation peak noradrenaline-velocity (0.24 at 0s), cross-correlation peak calcium-velocity (0.07 at -0.34s). **Dopamine**: median correlation dopamine-calcium (0.14), cross-correlation peak dopamine-calcium (0.14 at 0s), median correlation dopamine-velocity (0.2), median correlation calcium-velocity (0), cross-correlation peak dopamine-velocity (0.2 at 0s), cross-correlation peak calcium-velocity (0.04 at -0.17s).

In summary, distinct activity profiles were observed for serotonin, dopamine, and noradrenaline in response to appetitive sucrose and aversive air-puffs. These responses

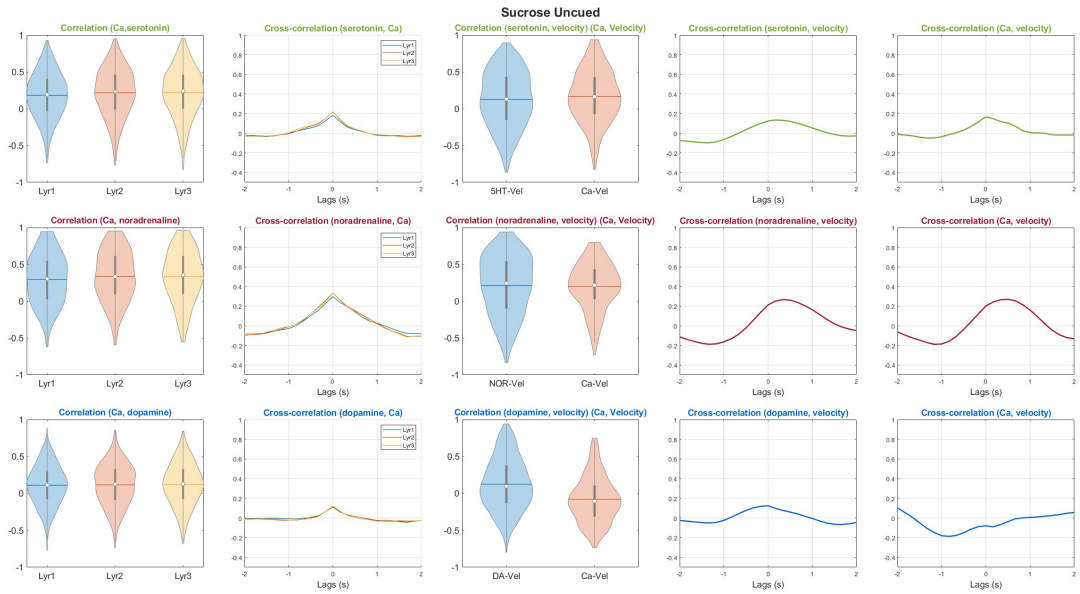

Figure 3.15: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion uncued sucrose. **Serotonin**: median correlation serotonin-calcium (0.23), cross-correlation peak serotonin-calcium (0.22 at 0s), median correlation serotonin-velocity (0.13), median correlation calcium-velocity (0.16), cross-correlation peak serotonin-velocity (0.13 at 0.17s), cross-correlation peak calcium-velocity (0.16 at 0s). **Noradrenaline**: median correlation noradrenaline-calcium (0.33), cross-correlation peak noradrenaline-calcium (0.33 at 0s), median correlation noradrenaline-velocity (0.24), median correlation calcium-velocity (0.22), cross-correlation peak noradrenaline-velocity (0.27 at 0.34s), cross-correlation peak calcium-velocity (0.27 at 0.51s). **Dopamine**: median correlation dopamine-calcium (0.13), cross-correlation peak dopamine-calcium (0.12 at 0s), median correlation dopamine-velocity (0.1), median correlation calcium-velocity (-0.1), cross-correlation peak dopamine-velocity (0.12 at 0s), cross-correlation peak calcium-velocity (-0.19 at -0.83s).

were marked by variations in peak minima and maxima and differed significantly in magnitude and timing across conditions. Notably, serotonergic and noradrenergic responses showed marked differences between cued and uncued sucrose, and between sucrose to air-puffs while dopaminergic activity remained consistent. In relation to calcium activity, serotonin and noradrenaline exhibited weak correlations during sucrose, but not air-puff, trials. Serotonin showed no significant correlation with velocity across all data sets, while noradrenaline demonstrated weak correlations with velocity during sucrose delivery. Dopamine, on the other hand, displayed no correlation with calcium activity in any unconditioned stimulus condition and no significant correlation with velocity. These findings underscore the differences in neuromodulator responses to cued and uncued sucrose compared to air-puffs, in particular for serotonergic and noradrenergic systems. While dopaminergic activity was consistently more heterogeneous across all US conditions.

We then proceeded to explore the data from a different angle, asking whether neu-

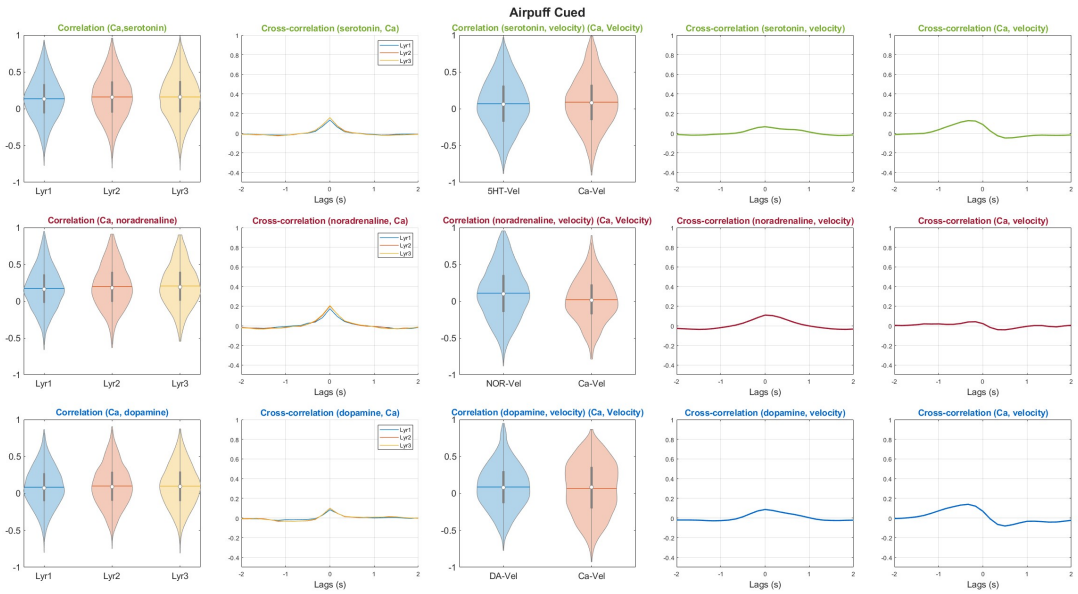

Figure 3.16: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion cued air-puff. **Serotonin**: median correlation serotonin-calcium (0.16), cross-correlation peak serotonin-calcium (0.16 at 0s), median correlation serotonin-velocity (0.06), median correlation calcium-velocity (0.08), cross-correlation peak serotonin-velocity (0.06 at 0s), cross-correlation peak calcium-velocity (0.13 at -0.34s). **Noradrenaline**: median correlation noradrenaline-calcium (0.18), cross-correlation peak noradrenaline-calcium (0.2 at 0s), median correlation noradrenaline-velocity (0.1), median correlation calcium-velocity (0.01), cross-correlation peak noradrenaline-velocity (0.1 at 0s), cross-correlation peak calcium-velocity (0.04 at -0.17s). **Dopamine**: median correlation dopamine-calcium (0.09), cross-correlation peak dopamine-calcium (0.09 at 0s), median correlation dopamine-velocity (0.07), median correlation calcium-velocity (0.08), cross-correlation peak dopamine-velocity (0.09 at 0s), cross-correlation peak calcium-velocity (0.14 at -0.34s).

romodulator responses to US differ between resting and locomoting mice; specifically, how neuromodulator systems react to the interplay of locomotion changes and reinforcing stimuli. Subsets of the data were identified where mice were predominantly at rest (velocity within -0.3 and 0.3 cm/s for 80% of the time) or in motion (velocity exceeding 5 cm/s throughout); cued and uncued data were combined. Figure 3.18a shows the considerable resemblance between mean fluorescence (coloured line) and mean locomotion velocity (black line) during cued sucrose delivery. During the experiment, we observed that the mice tended to slow down or stop to consume the sucrose solution. Interestingly, the plot suggests that decrease in velocity tended to begin during the pre-stimuli period, indicating an anticipatory behavioural response to the US. The mean velocity time course appears to precede the neuromodulator activity changes, as confirmed by cross-correlation analysis (Figures 3.20, 3.21).

The analysis depicted in Figure 3.19 reveals significant decreases in fluorescence levels for serotonin, noradrenaline, and dopamine during locomotion ($p = 0.0024$ for each).

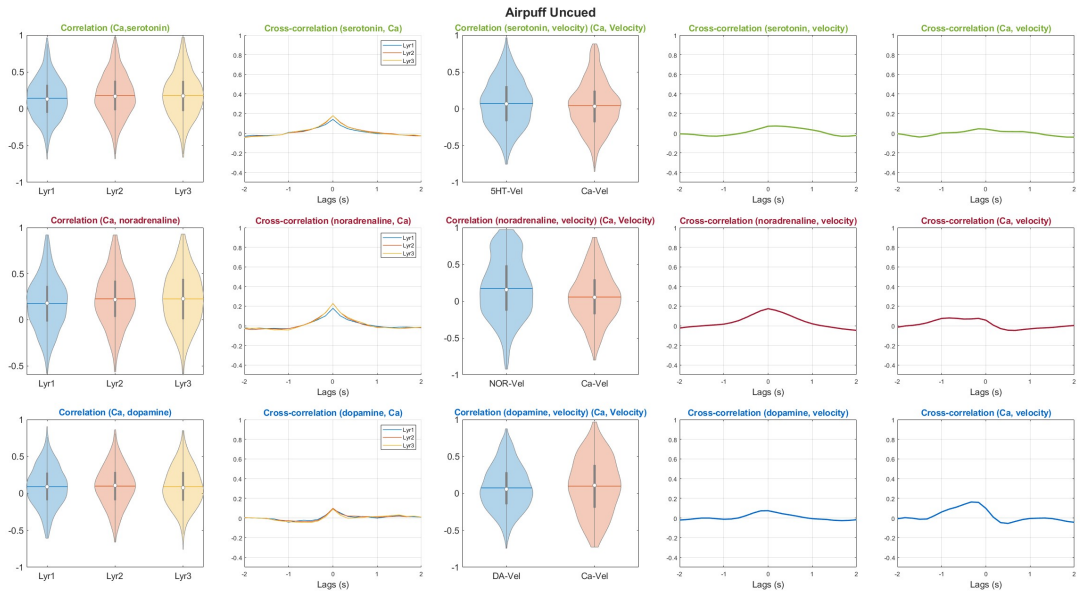

Figure 3.17: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion uncued air-puff. **Serotonin**: median correlation serotonin-calcium (0.17), cross-correlation peak serotonin-calcium (0.18 at 0s), median correlation serotonin-velocity (0.07), median correlation calcium-velocity (0.03), cross-correlation peak serotonin-velocity (0.07 at 0.17s), cross-correlation peak calcium-velocity (0.04 at 0.17s). **Noradrenaline**: median correlation noradrenaline-calcium (0.22), cross-correlation peak noradrenaline-calcium (0.23 at 0s), median correlation noradrenaline-velocity (0.16), median correlation calcium-velocity (0.05), cross-correlation peak noradrenaline-velocity (0.17 at 0s), cross-correlation peak calcium-velocity (0.08 at -0.83s). **Dopamine**: median correlation dopamine-calcium (0.1), cross-correlation peak dopamine-calcium (0.09 at 0s), median correlation dopamine-velocity (0.06), median correlation calcium-velocity (0.1), cross-correlation peak dopamine-velocity (0.07 at 0s), cross-correlation peak calcium-velocity (0.16 at -0.34s).

These profiles are akin to those observed for cued sucrose but with some important distinctions. It's worth noting that we are drawing loose comparisons between these sets of analyses without statistical confirmation, as the data sets overlap and are therefore not independent, making them unsuitable for conventional statistical analyses.

Keeping this in mind, the proportion of minima-only time courses for serotonin and noradrenaline was higher during locomotion (59% and 71%, respectively) compared to cued sucrose conditions (39% and 51%, respectively). Dopamine did not exhibit such variations. This increase in the prevalence of minima was also associated with a noticeable increase in the magnitude of minima. When comparing these findings to those observed during rest, it becomes evident that velocity is a significant contributing factor. Specifically, the effect of sucrose on minima at rest was significantly less pronounced for both serotonin ($p = 0.0094$) and noradrenaline ($p = 0.0033$). This was accompanied by an increase in the proportion of maxima-only time courses when comparing locomotion to rest for serotonin (from 10% to 32%) and noradrenaline (from

11% to 35%). Intriguingly, while dopaminergic responses also showed an increased proportion of maxima-only time courses at rest (from 12% to 33%), but this change was not associated with an increase in response magnitude.

Upon examining the responses to air-puffs, we observed smaller fluctuations in activity for serotonin and noradrenaline compared to sucrose. This difference reached statistical significance for noradrenaline when comparing move vs. rest sucrose minima ($p = 0.0131$). The magnitudes of dopaminergic responses remained consistent across both US and behavioural states. As for the minima-maxima composition, the proportions of serotonergic and dopaminergic responses did not significantly vary with the locomotion state in the case of air-puffs. However, noradrenaline showed an increase in mixed responses — containing both minima and maxima — during locomotion compared to the resting state. Additionally, the composition of responses to air-puffs at rest closely mirrored those to sucrose at rest.

Our analysis indicates that locomotion significantly influences neuromodulator responses to sucrose and air-puffs. Serotonin and noradrenaline show greater response minima during locomotion, while dopamine remains consistent. Noradrenaline also exhibits mixed responses to air-puffs during locomotion compared to rest.

Analysing the correlations and cross-correlations between calcium and neuromodulators during sucrose delivery while in motion, we observed patterns similar to those in the cued and uncued cases for serotonin. Serotonin exhibited a weak median correlation with calcium activity (0.21) and with velocity (0.2). Notably, when velocity was a controlled factor, noradrenaline showed a stronger correlation with calcium (0.3) and a similar correlation with velocity (0.2). Dopamine, consistent with previous observations, did not correlate significantly with calcium (0.13) and only weakly with velocity (0.2).

Interestingly, we noted a general trend where changes in velocity tended to precede changes in calcium and neuromodulator activity, as indicated by positive cross-correlation peaks at 0.17s for serotonin, 0.34s for noradrenaline. This suggests a temporal lead of velocity changes over M2 activity. However, the cross-correlation plots for both serotonin and noradrenaline exhibited a characteristic shape with a negative lag dip followed by a positive peak, which can occur when signals generally rise and fall together over but there is some variation in which signal leads.

During air-puff delivery while in motion, we found that neither serotonin, noradrenaline, nor dopamine activities correlated well with calcium activity or with velocity, marking a distinct contrast from the responses observed during sucrose delivery.

Conditioned stimuli: auditory cues

During the OGNG task, auditory feedback cues preceded the US by two to four seconds for one second, but there were no significant fluorescence changes in the post-stimuli period for any of the cues, Figure 3.22. Auditory state cues were also presented, and again no significant fluorescence changes in the post-stimuli period (results not shown).

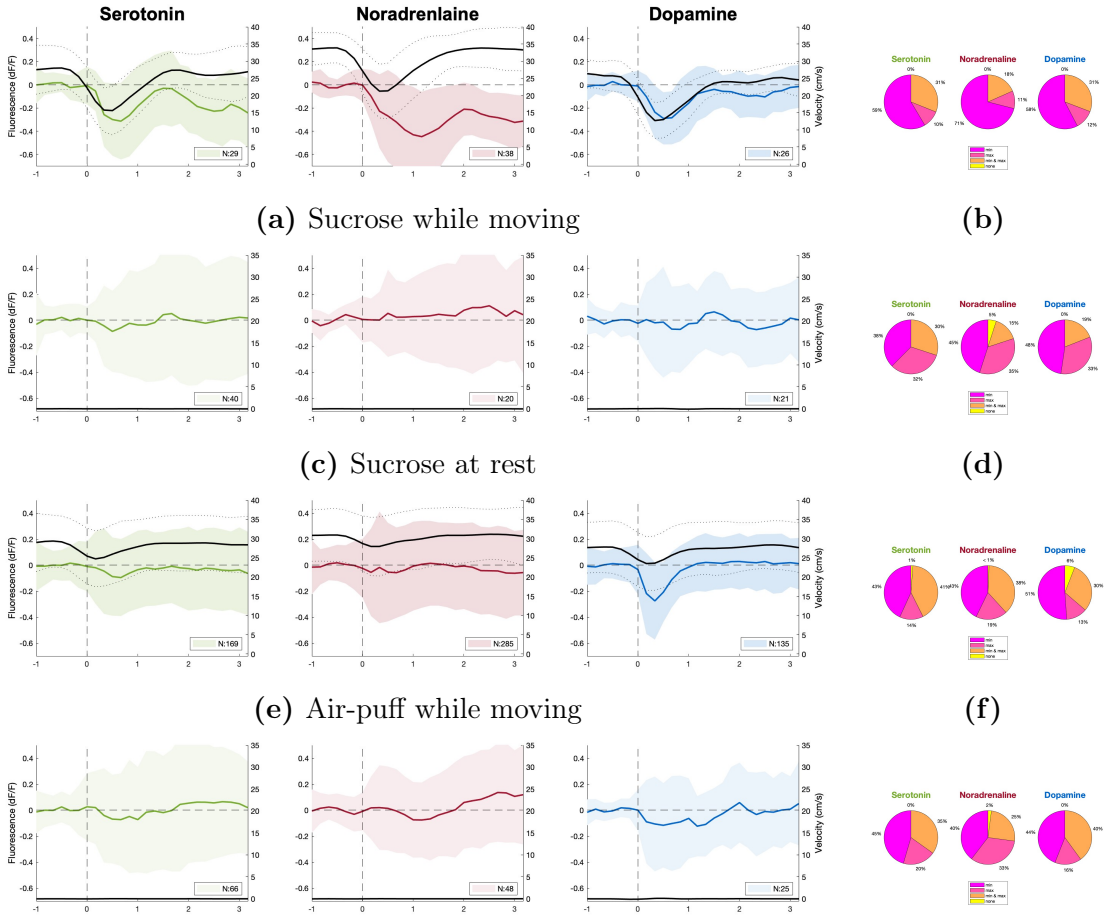

(a) Sucrose while moving **(b)**

(c) Sucrose at rest **(d)**

(e) Air-puff while moving **(f)**

(g) Air-puff at rest **(h)**

Figure 3.18: Subsection of data from Fig 3.12 in which mice are at rest or moving for the whole time course. Comparing 3.18a to 3.18c, and 3.18e to 3.18g, it appears that without locomotion serotonergic responses are muted and noradrenergic responses are abolished, while dopaminergic responses remain intact. Cued and uncued data were combined to increase the sample size. The mean and standard deviation fluorescence are indicated in colour, and the mean and standard deviation velocity of the treadmill is in black. The pre-stimulus period from -1 to 0 was used to calculate dF/F.

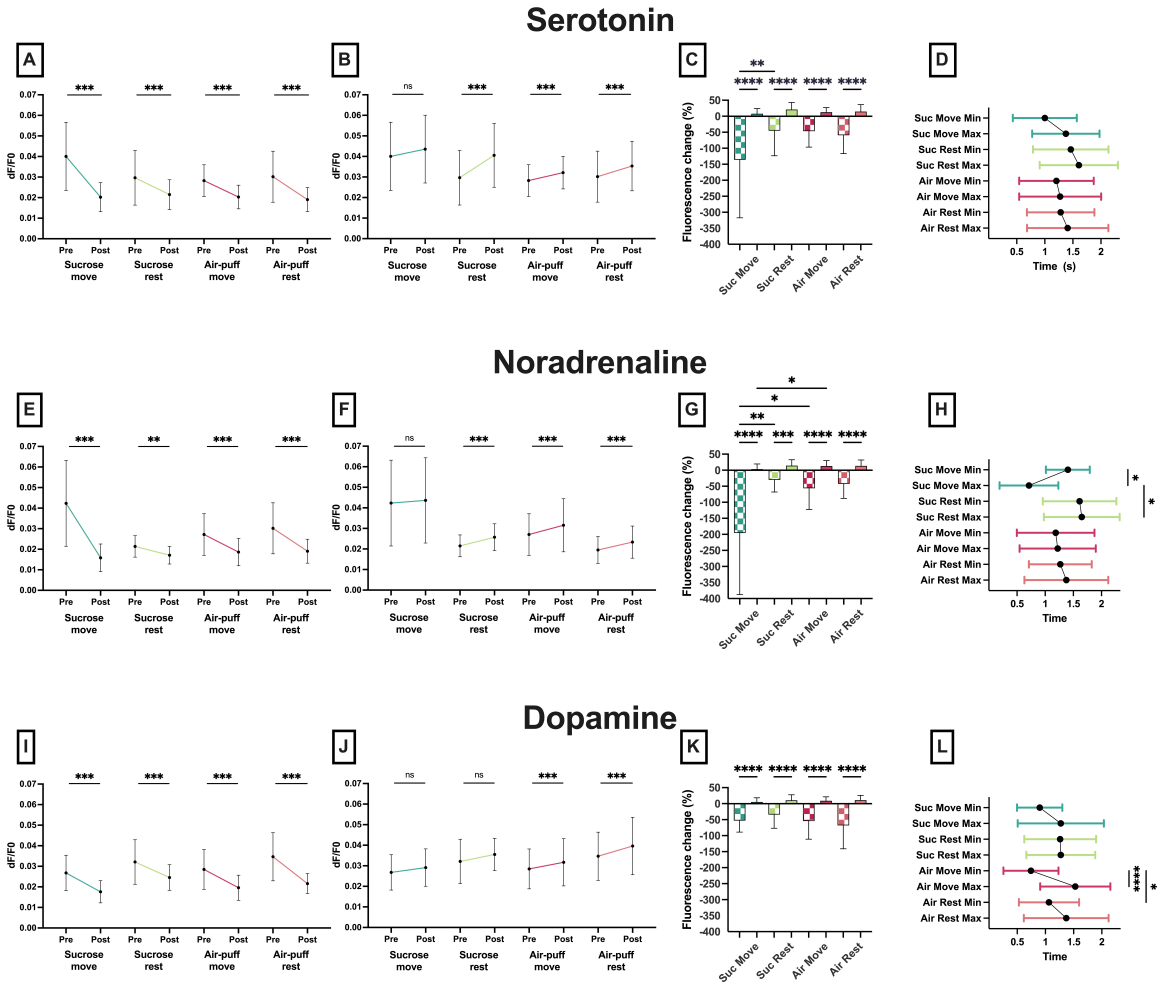

Figure 3.19: Analysis of the subsection of data presented in Fig 3.13 in which mice are at rest or moving throughout the peristimulus period. **A, E, I:** Peak post-stimulus minima with respect to pre-stimulus. Statistically significant differences seen for all ($p = 0.0024$) except noradrenaline sucrose at rest ($p = 0.0114$). **B, F, J:** Peak post-stimulus maxima with respect to pre-stimulus. Significant maxima seen for most ($p < 0.0096$ for all), but notably not for sucrose move for any neuromodulator, nor for sucrose rest for dopamine. **C, G, K:** Comparison between states; Minima chequered, maxima block colours. Serotonin and noradrenaline responses to move sucrose were significantly different from the other states. Whereas, dopamine responses to each of these US were of similar magnitude, since there was no statistical difference between states;**D, H, L:** Differences were found between the delays to peak fluorescence for noradrenaline sucrose move and dopamine air-puff move. Left, centre left used Wilcoxon paired tests p-values were corrected using the Holm-Šídák method; Centre right, right used Kruskal-Wallis tests.

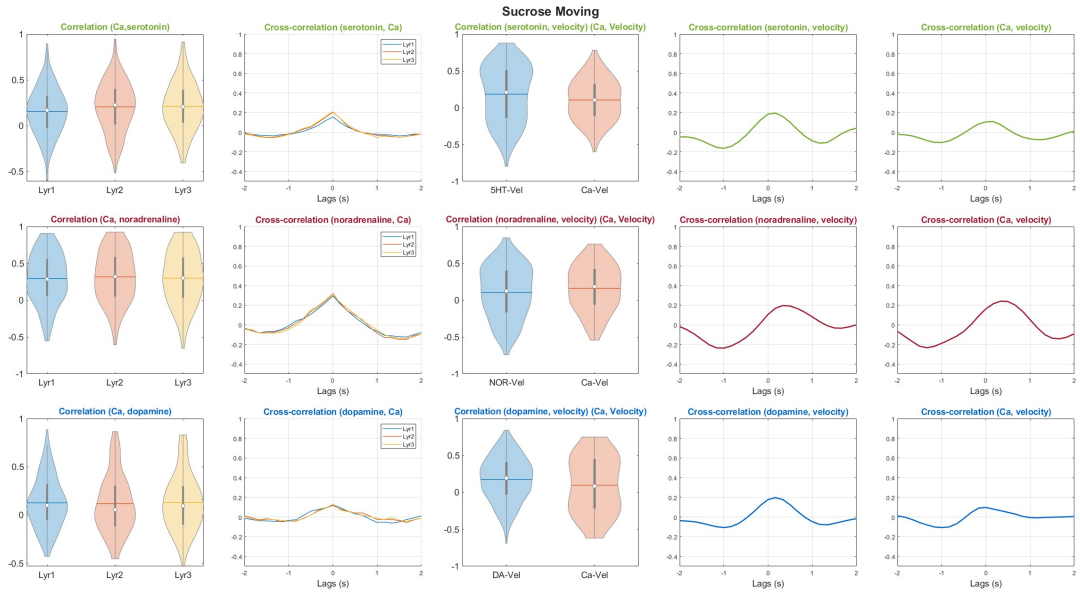

Figure 3.20: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion sucrose while moving. **Serotonin**: median correlation serotonin-calcium (0.21), cross-correlation peak serotonin-calcium (0.21 at 0s), median correlation serotonin-velocity (0.21), median correlation calcium-velocity (0.1), cross-correlation peak serotonin-velocity (0.2 at 0.17s), cross-correlation peak calcium-velocity (0.1 at 0.17s). **Noradrenaline**: median correlation noradrenaline-calcium (0.3), cross-correlation peak noradrenaline-calcium (0.3 at 0s), median correlation noradrenaline-velocity (0.12), median correlation calcium-velocity (0.18), cross-correlation peak noradrenaline-velocity (0.2 at 0.34s), cross-correlation peak calcium-velocity (0.24 at 0.34s). **Dopamine**: median correlation dopamine-calcium (0.09), cross-correlation peak dopamine-calcium (0.13 at 0s), median correlation dopamine-velocity (0.19), median correlation calcium-velocity (0.09), cross-correlation peak dopamine-velocity (0.2 at 0s), cross-correlation peak calcium-velocity (0.1 at 0s).

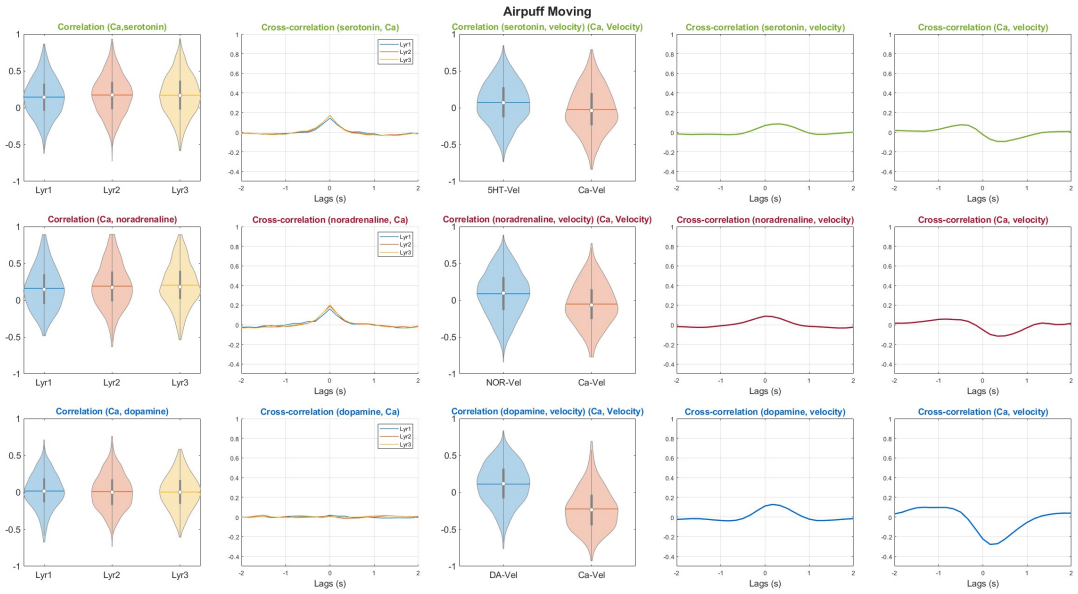

Figure 3.21: Correlation and cross-correlation between neuromodulator, calcium and velocity during locomotion air-puff while moving. **Serotonin**: median correlation serotonin-calcium (0.17), cross-correlation peak serotonin-calcium (0.17 at 0s), median correlation serotonin-velocity (0.07), median correlation calcium-velocity (-0.04), cross-correlation peak serotonin-velocity (0.08 at 0.34s), cross-correlation peak calcium-velocity (-0.09 at 0.34s). **Noradrenaline**: median correlation noradrenaline-calcium (0.17), cross-correlation peak noradrenaline-calcium (0.2 at 0s), median correlation noradrenaline-velocity (0.1), median correlation calcium-velocity (-0.07), cross-correlation peak noradrenaline-velocity (0.09 at 0s), cross-correlation peak calcium-velocity (-0.11 at 0s). **Dopamine**: median correlation dopamine-calcium (0.02), cross-correlation peak dopamine-calcium (0.02 at 0s), median correlation dopamine-velocity (0.12), median correlation calcium-velocity (-0.24), cross-correlation peak dopamine-velocity (0.13 at 0.17s), cross-correlation peak calcium-velocity (-0.28 at 0.17s).

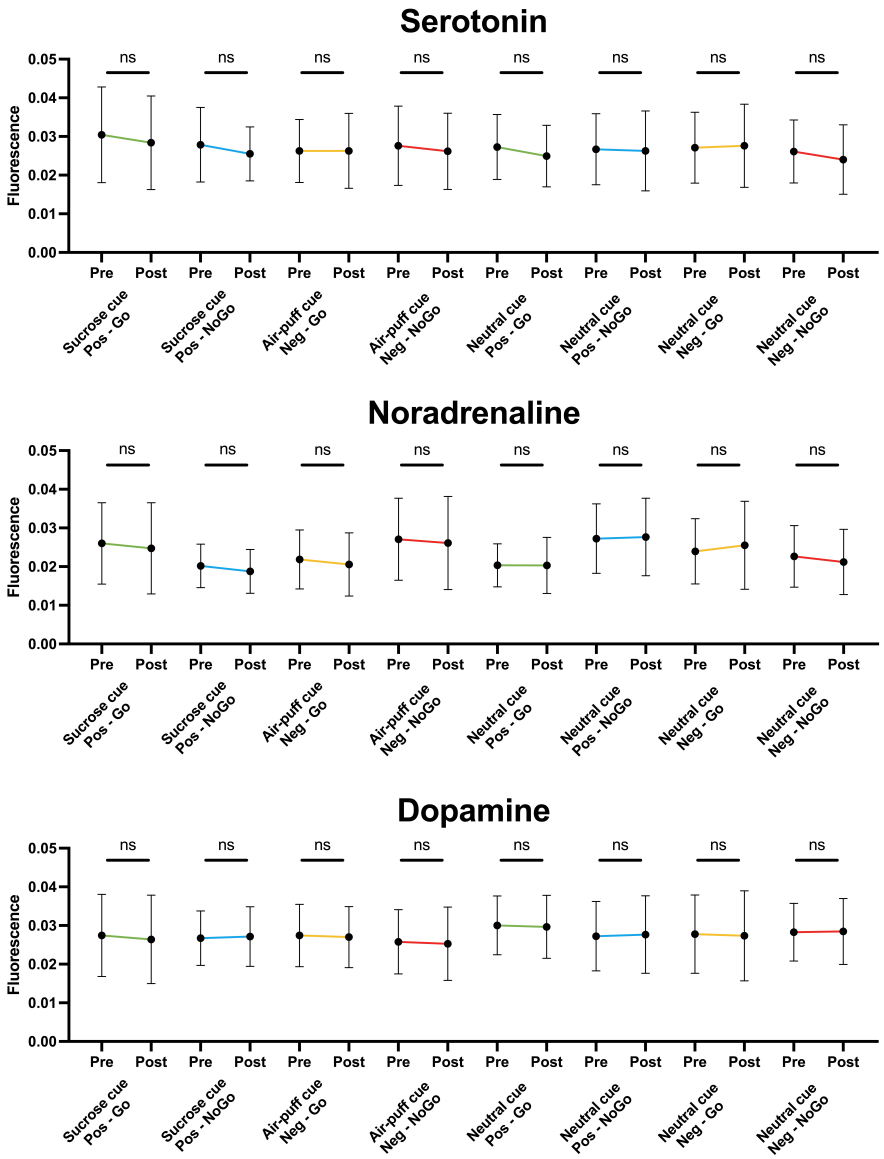

Figure 3.22: Analysis of peristimulus fluorescence in response to auditory cues during the OGNG task revealed no significant differences in median pre- and post-stimulus fluorescence for the neuromodulators. Pre-stimulus dF/F was based on the mean fluorescence from -1 to 0 seconds, and post-stimulus values were calculated as the mean 1 second around the peak fluorescence which occurred within 2 seconds of stimulus onset. Colour coding indicates the same behavioural state.

3.2.4 Spatial release patterns

To examine neuromodulator spatial release characteristics, we analysed the correlations between the time courses of pairs of hotspots (for definition see Section 2.9.2). Figure 3.23 illustrates the mean of pairwise correlations as a function of the distance between hotspots, comparing locomotion states (Fig. 3.23a), cued and uncued US delivery (Fig. 3.23b), and US delivery during locomotion or at rest (Fig. 3.23c).

We observed a consistent trend across neuromodulators during locomotion, with significant correlations at both onset and offset, but minimal during sustained locomotion or rest. Specifically, serotonergic activity showed a peak correlation of 0.57 at 20 µm during locomotion transitions, diminishing below the 0.3 correlation threshold at 170 and remaining around 0.28 at greater distances. Noradrenergic activity was highly correlated (0.65 at onset and 0.62 at offset) and linearly decreased with distance, falling below 0.3 at 220 µm. Dopaminergic activity showed the weakest correlation, peaking at 0.47 at onset and 0.42 at offset, and quickly diminishing beyond 60 µm. Notably, dopaminergic activity showed no significant correlation during sustained locomotion or rest.

Regarding responses to cued and uncued stimuli, serotonergic activity was similarly correlated (between 0.5 and 0.43 at 20 µm), decreasing sharply by 100 µm. Noradrenergic responses varied more distinctly between stimuli, with peak correlations for sucrose (0.61 cued, 0.58 uncued) and air-puffs (0.5 for both). Noradrenergic responses to sucrose were all broadly correlated as far as 340 /*mu*m, whereas responses to air-puffs were more localised, dropping off at 100 /*mu*m. Dopaminergic activity showed correlations for cued sucrose (0.42) and uncued air-puff (0.39), dropping off around 80 µm.

During sucrose delivery while moving, serotonin exhibited a U-shaped correlation pattern, peaking at 0.56 at 20 µm and 0.49 at 380 µm. This pattern was absent in other conditions, which peaked at 20 µm (0.42 sucrose rest, 0.38 air-puff rest, 0.31 air-puff move). Noradrenergic responses to sucrose while moving were notably high (0.62) and remained correlated across distances, while responses to air-puffs peaked at 0.45 and decayed rapidly. At rest, noradrenergic responses to stimuli only reached a 0.3 correlation at 20 µm. Dopaminergic responses under these conditions were not significantly correlated.

We utilised a two-way ANOVA to assess the impact of task state and distance on spatial correlation data in Figure 3.23a and Kruskall-Wallis tests to make group comparisons; p-values corrected using Holm-Šídák method. The large dataset yielded highly significant results for variance from both state and distance ($p = 0.0009$ for all). Examining the proportion of variation explained by distance or state revealed intriguing patterns.

Analysing activity in different locomotion states, we observed that for serotonin, 7% of the variation was attributable to distance, while 11% was linked to the state. In the case of noradrenaline, the state accounted for 11% of the variation, while distance explained 8%. Dopamine displayed the most pronounced difference, with 14% explained by state and only 3% by distance. Kruskal-Wallis testing confirmed the observed trends, indicating strong similarity between stable locomotion states (Rest and Move) and between state transitions (Onset and Offset), which were not statistically significant. In

contrast, all comparisons between stable states and transition states were statistically significant ($p = 0.0018$ for all).

Comparing responses to cued and uncued US, the variation in serotonin activity was mainly due to distance (7%), with state contributing only 0.5%. Air-puff Cued consistently exhibited lower correlation than Sucrose Cued ($p = 0.003$) and Sucrose Uncued ($p = 0.0088$). Noradrenergic activity showed higher variability with distance (6%) than with state (3%). Significant differences were observed between multiple groups: Sucrose Cued and Air-puff Cued ($p = 0.0109$), Sucrose Cued and Air-puff Cued ($p = 0.003$), Sucrose Uncued and Air-puff Cued ($p = 0.0039$), and Sucrose Uncued and Air-puff Uncued ($p = 0.0018$). Notably, there was no difference between Cued and Uncued Sucrose, or between Cued and Uncued Air-puff. For dopamine, a similar contribution of 3% was observed from distance and 2% from state. Figure 3.23b reveals a surprising similarity between Cued Sucrose and Uncued Air-puff, and between Uncued Sucrose and Cued Air-puff. However, the other group comparisons yielded significant differences: Cued Sucrose and Uncued Sucrose (both $p = 0.0018$), Cued Sucrose and Cued Air-puff ($p = 0.0018$), Uncued Sucrose and Uncued Air-puff ($p = 0.0223$), and Cued Air-puff and Uncued Air-puff ($p = 0.0084$).

Regarding US responses during rest or locomotion, serotonin fluorescence exhibited 9% variability attributed to the state and 7% to distance. This highlights the substantial contrast between sucrose during movement and the other states, with Sucrose Move being statistically significant from other groups (all $p = 0.0018$). In the case of noradrenaline, both the state and distance accounted for 7% of the variation. There were statistically significant differences between Sucrose Move and Sucrose Rest ($p = 0.0018$), Air-puff Move ($p = 0.0131$), and Air-puff Rest ($p = 0.0018$), as well as between Air-puff Move and Air-puff Rest ($p = 0.0131$). For dopamine, the distance contributed to 3% of the variation, while the state represented only 0.2%, marking a significant difference from the locomotion groups.

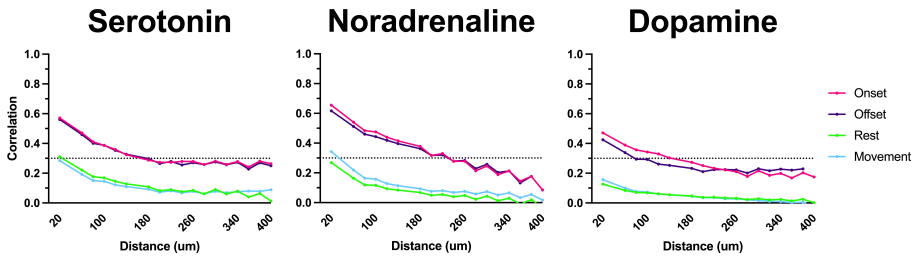

(a) Locomotion

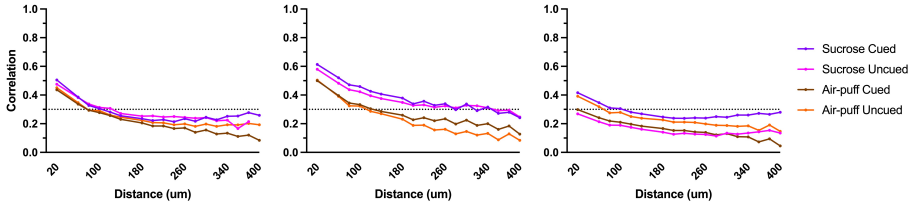

(b) Cued and uncued US responses

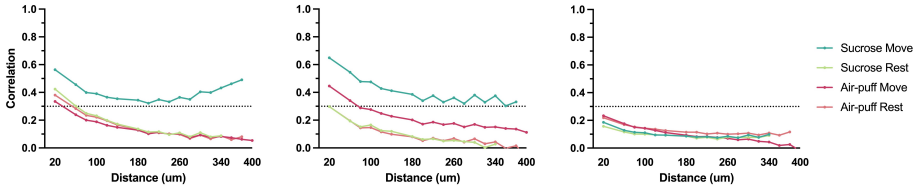

(c) US responses in motion and at rest

Figure 3.23: Depiction of correlated neuromodulator signals across M2 as a function of distance. Time courses of hot spots of activity were correlated with time courses from all other hot spots, see Section 2.9.2. Dashed line depicts correlation coefficient r= 0.3, set as threshold above which is considered correlated.

3.3 Pharmacology

Finally, we explore the responses of the novel GENIs to pharmacological manipulation. Specifically, we employed antagonists that target the GPCR-scaffold of each GENI to reduce GENI fluorescence activity. It's important to note that systemic drug administration affects both GENIs and neuromodulator-releasing cells, introducing a significant confounding factor that can lead to unexpected outcomes. Additionally, we utilized a noradrenaline reuptake inhibitor to enhance noradrenaline concentrations, a method that yields more straightforward and interpretable results. Furthermore, our findings indicate that neuromodulator systems can influence one another, with interventions in one system able to impact others.

We recorded data for 10 minutes at time points: 5 minutes before, and 15, 30 and 45 minutes following injection. Intradermal injection of the drug was administered to the mouse into the subcutaneous tissue located at the nape of the neck. This injection site was selected due to its accessibility while mice were head-fixed to the imaging setup.

For all analyses, we employed a mixed-effects model to evaluate the impact of "time"

after injection on mean fluorescence, as well as the magnitude of variation attributed to different "mice." Our results indicated that "time" was a significant source of variation in all pharmacological interventions, with the exception of "SCH23390 - dLight". Additionally, in all analyses involving multiple mice, differences between individual mice were also identified as a significant source of variability. To further assess differences between time points within the framework of our mixed-effects model, we conducted Dunnett's multiple comparisons test. This test facilitated comparisons against the baseline while controlling for the family-wise error rate, making it particularly suitable for our analysis.

The fluorescence of GRAB5HT, a 5-HT2C receptor-based chimera, was significantly altered upon the administration of SB242084, a 5-HT_{2C} antagonist. Following injection, 4 out of 5 mice showed decreased fluorescence 15 minutes post-injection ($p < 0.0001$). One mouse exhibited a substantial increase in fluorescence at 30 and 45 minutes. Both time ($p = 0.0007$) and individual mice ($p < 0.0001$) had a significant effect on fluorescence variance. The Dunnett's multiple comparison test indicated a significant difference from baseline at time points 15 ($p < 0.0001$) and 30 minutes ($p = 0.0019$).

The fluorescence of $GRAB_{NE}$, which utilises the ligand-binding domain of the endogenous noradrenergic $\alpha 2$ receptor, exhibited significant alterations upon the administration of Yohimbine, an $\alpha 2$ antagonist. Following injection, 3 out of 4 mice showed a decrease in fluorescence. Both time and individual mice significantly contributed to the variance of fluorescence ($p < 0.0001$ for both), as indicated by the mixed-effects model for both factors. All three time points were significantly different from baseline, with Dunnett's multiple comparison test showing $p < 0.0001$ for each. Additionally, the injection of the noradrenaline reuptake inhibitor Reboxetine induced a notable change in $GRAB_{NE}$ fluorescence. Two out of four mice exhibited increased fluorescence responses at 15 and 30 minutes post-injection. Both time ($p = 0.0004$) and individual mice ($p < 0.0001$) had a significant effect. Dunnett's multiple comparison confirmed a significant difference from baseline at 15 minutes ($p = 0.01$).

The response of dLight, which is based on the dopamine D1 receptor, to D1 antagonist SCH23390, varied strongly depending on the mouse ($p < 0.0001$). On average, the response did not vary significantly with time ($p = 0.0613$). However, as observed in the figure, injection induced a robust increase in two mice at different time points and a consistent decrease in fluorescence in one mouse across all time points.

In addition to the targeted pharmacological manipulations for each GENI presented in Fig 3.24, we subjected each GENI to the effects of all other drugs for exploratory purposes, illustrating the intricate nature of neuromodulator system manipulations. Our observations revealed that Yohimbine produced a decrease in fluorescence in all three GENIs. Reboxetine reduced the fluorescence responses of $GRAB_{5HT}$ and dLight, while increasing fluorescence in $GRAB_{NE}$. SB242084 caused a short-term decrease in fluorescence in $GRAB_{NE}$, a large increase in fluorescence in dLight, and an initial dip in $GRAB_{5HT}$ fluorescence followed, in some mice, by a rebound. SCH23390 induced a modest increase in fluorescence in GRAB5HT and GRABNE while causing significant bidirectional changes in dLight fluorescence.

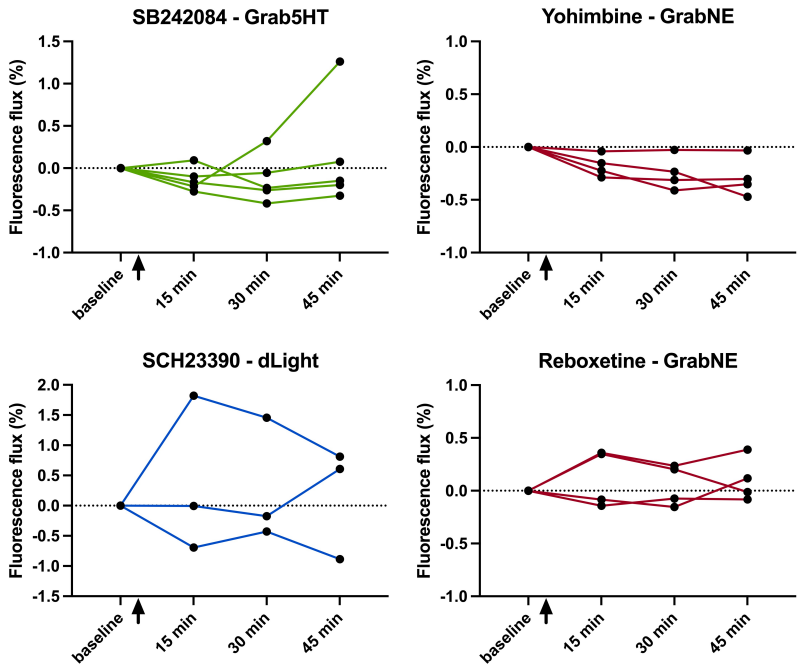

Figure 3.24: GENI-specific pharmacological manipulations. Recordings taken before and at 15-minute intervals following intradermal injection. Each line represents the mean fluorescence change from baseline for each mouse. Yohimbine ($\alpha 2$ antagonist) significantly altered GRABNE fluorescence at all time points compared to baseline ($p < 0.0001$). Reboxetine (NRI) induced a significant deviation from baseline at 15 minutes ($p = 0.01$). SB242084 (5-HT2C antagonist) caused a significant change at 15 minutes ($p < 0.0001$). SCH23390 induced substantial fluorescence changes in either direction depending on the mouse, resulting in an overall non-significant change from baseline.

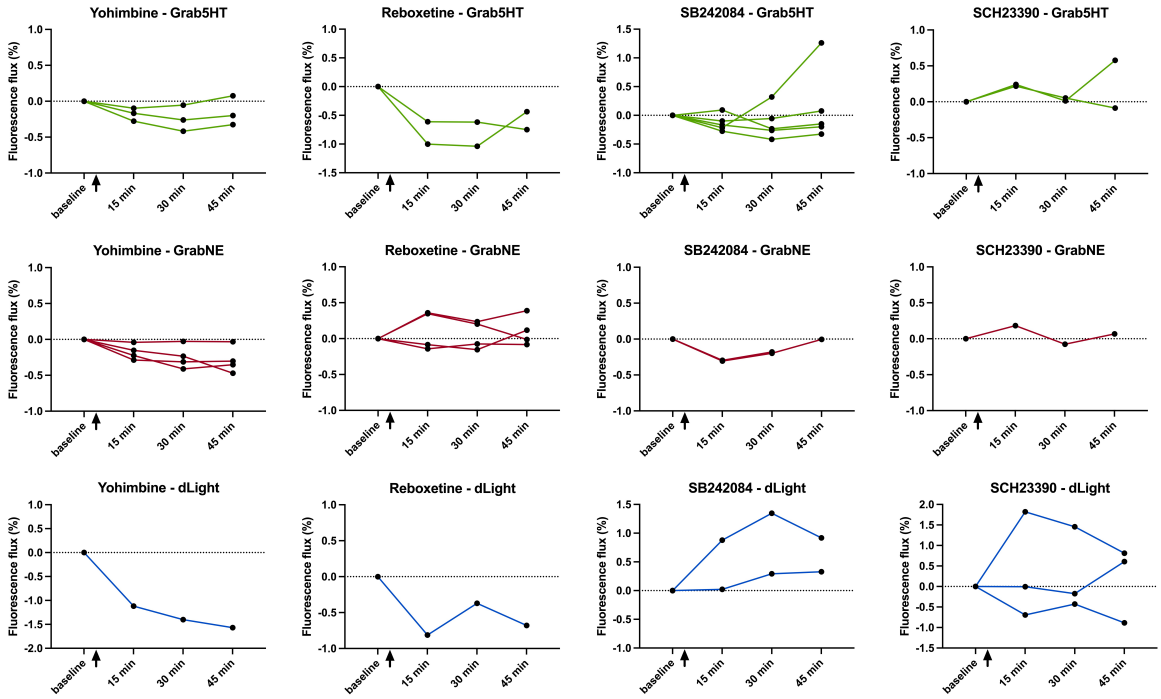

Figure 3.25: All pharmacological manipulations. Recordings taken before and at 15-minute intervals following intradermal injection. Recordings made of fluorescence from each GENI in response to each of the four pharmacological agents used in Fig 3.24. Each line represents the mean fluorescence change from baseline for each mouse.

Chapter 4

Discussion

The aim of this investigation was to characterise the release of the neuromodulators serotonin, noradrenaline, and dopamine in the secondary motor cortex. Characterisation took the form of analysing the temporal and spatial release dynamics during different behavioural states. The motivation behind the project was to use the novel GENIs to directly measure neuromodulator release in the cortex with the spatiotemporal precision that two-photon microscopy affords. In addition, we hoped to develop a novel OGNG task for use in mice, one which truly tested Go and No-Go behaviours in the more naturally relevant context of full body locomotion, rather than standard lick-based Go/NoGo tasks.

4.1 Neuromodulator release strongly correlated with locomotion transitions

One of our clearest findings was that all three neuromodulators showed robust activity changes in line with locomotion transitions. Serotonin, noradrenaline, and dopamine showed large increases in release as the mice transitioned from rest to locomotion (Fig 3.8). These were associated with simultaneous increases in calcium activity. Intracellular calcium fluctuations are widely used as a proxy for neuronal activation and signaling. Pertinently, we observed that changes in neuromodulator and calcium activity did not precede alterations in velocity (Fig. 3.10), contrary to what might be expected based on recordings in M2 [70, 79]. However, evidence suggests that during spontaneous locomotion, there is a balanced interaction between motor planning and sensory feedback activity in M2, resulting in a mean peak cross-correlation occurring within 0.1 s of 0 lags [35], which aligns with our findings. Additionally, our recordings were in layer 2/3, which is known to be involved with associating sensory information [91]. Recent research has suggested that rodent M2 primarily functions as a sensory-motor integration area for actions guided by sensory or motor cues rather than spontaneous locomotion onset [10].

Conversely, we saw decreases in fluorescence during transitions from locomotion to rest. These decreases were less rapid than for onset, which may reflect the various speeds and durations of running from which the mice were transitioning to rest. The higher positive lag values of the cross-correlation here suggest that the decrease in

locomotion occurred faster than the decrease in neuromodulator release and calcium activity. This may reflect the sensory feedback information associated with coming to rest. However, it should also be noted that the offset times for the fluorescent indicators τ_{off} is slower than the onset τ_{on} times for all GENIs, which may have contributed to the lingering signals (see Section 1.6).

In summary, noradrenergic activity aligns with locomotion, consistently tracking both its transitions and maintaining elevated levels during sustained movement. The broader standard deviation as compared to the other neuromodulators may reflect the small increases and decreases in locomotion during the sustained movement state. This relationship with locomotion aligns with existing literature that underscores the role of noradrenaline in facilitating mobile 'Go' actions and effortful tasks [5, 120]. Serotonergic activity varied with locomotion transitions, displaying higher average activity during sustained movement than at rest, although with a broad standard deviation that included negative values during locomotion. The serotonergic system's role in locomotion is complex. Activation of the dorsal raphe has been found to decrease locomotion in open field tests, whereas dorsal raphe GABAergic neurons can promote locomotion under escape conditions [26, 104]. Our findings of increased serotonin at locomotion onset and varied activity during movement are novel, diverging from previous results by Correia et al. [26], Seo et al. [104]. The observed patterns might be influenced by the stress of head-fixation, as discussed by [131]. Future experiments should investigate the role of serotonin in M2 in locomotion under non-stressful open field and stressful tail-suspension conditions used by Seo et al. [104], to provide clearer insights than our somewhat-stressful head-fixed setup allowed. Dopaminergic activity in the motor cortex during locomotion, consistent with findings from similar studies [86], is depicted in Figure 1.8.

4.2 Distinct and nuanced neuromodulator responses to appetitive and aversive stimuli

Among our observations, a particularly surprising finding stood out: a marked drop in fluorescence in response to US. As visualised in Figure 3.12, the administration of appetitive sucrose solution or aversive air-puffs led to reductions in mean fluorescence. Statistical analyses confirmed significant poststimulus minima for both US, regardless of whether they were cued or uncued (Fig. 3.13). Although the magnitude of responses to cued and uncued US remained largely comparable, their temporal patterns and response shapes exhibited notable differences. Furthermore, a pronounced standard deviation, extending from negative to positive, was consistent across all conditions. We plotted the standard deviation rather than the commonly used standard error of the mean to highlight the diversity of responses. Closer scrutiny of individual traces revealed a prevalence of biphasic responses, characterised by both peaks and troughs. To quantify these observations, we assessed the mean peak and trough for each time series. Significant peaks and troughs were evident in every condition. We then probed whether individual time series inherently exhibited both peaks and troughs or if this pattern emerged from the integration of bimodal data, as represented in the

pie chart figures. Given the pronounced responses to locomotion, we also investigated whether neuromodulator responses to US were influenced by locomotion-dependent factors (Figs. 3.18, 3.19). The cross-correlation between neuromodulator activity and velocity revealed that only during sucrose delivery, when mice consistently slowed to consume the reward, were the neuromodulators correlated with US delivery (Fig 3.20, 3.21).

Serotonin

In our examination of the mean serotonergic responses to sucrose (Fig 3.12), we found an initial decline followed by a subsequent increase in activity. Specifically, 39% of these responses showed solely a dip, 36% exhibited a biphasic nature, and 23% presented exclusively phasic peaks. Contrasting this, the serotonergic responses to uncued sucrose were predominantly characterised by dips, accounting for 55%; peaks were discerned in considerably smaller proportions (12% maxima-only and 28% mixed). This resulted in a marked and extended decline in fluorescence throughout the post-sucrose time course.

Interestingly, serotonergic reactions to both cued and uncued air-puff were strikingly similar, predominantly manifesting as a phasic dip. Yet, despite the lack of statistical difference between cued and uncued responses, the variability evident in the pie chart representations hint that statistical analysis of average activity might conceal a broader range of responses. Although it cannot be determined from our recordings, the distinct response types to cued and uncued air-puffs - 40-41% trough-exclusive, 20% peak-exclusive, and 35-36% mixed - may be produced by different serotonergic neurones, as individual neurones have been shown to be consistently excited or inhibited following US stimuli [24].

We then focused on data samples where the mice were consistently in locomotion or at rest (3.18). During movement, the troughs of serotonin activity deepened considerably, with an absence of significant peaks. Whereas, sucrose administration at rest led to a significantly less profound decline than during movement ($p = 0.0094$). Serotonergic responses to sucrose during locomotion were both stronger and more uniform than during rest. Specifically, when in motion, 59% of the results showed only decreases, 31% showed a mix of peaks and troughs, and a mere 10% exhibited peaks alone (Fig 3.19). This contrasts with the rest state, where there was a approximately even mix of time courses with minima and maxima. Thus, serotonergic responses to sucrose during locomotion appeared not only more potent but also more consistent than during rest. Intriguingly, it was during rest that the highest proportion of peaks in time courses was observed (32% peak-only, 30% peak and trough). Consistent increases in fluorescence in response to appetitive stimuli are commonly reported when recording from dorsal raphe nuclei [24, 68].

Examining mean velocity alongside time courses (Fig 3.18a), we noticed that mice consistently slowed down to consume the liquid, despite our data selection criteria requiring constant movement during the peristimulus period. This decrease in velocity seems to be a key factor influencing the fluorescence response profile, as indicated by the differences in responses observed between rest and movement periods. The cross-correlation peak at 0.17s confers that velocity changes tended to precede serotonin

fluctuations (Fig 3.20).

Comparatively, the distinction between air-puff administration during movement and rest was less pronounced, with no significant differences in response amplitude. The distribution of time courses showcasing peaks and troughs remained relatively stable (43-45% troughs only, 35-41% mixed, 14-20% peaks only), positioning these values between the figures for sucrose during movement and rest. An examination of the velocity plots (Fig 3.18e) revealed only a less consistent decrease in velocity in response to air-puffs, which may somewhat account for the less profound effects than seen with sucrose. However, the correlation between serotonin and velocity during sucrose but not during air-puffs, suggests serotonin is conveying a multifaceted signal.

In summary, in M2 serotonin signals both locomotion changes and the delivery of US. The distinct responses to cued versus uncued sucrose are noteworthy, in particular in contrast to the similar reactions to both cued and uncued air-puffs. This mirrors the findings of Cohen et al. [24], who observed greater responses to unexpected versus expected sucrose, but no such variation with air-puffs. They also noted that individual neurones consistently showed either increases or decreases in activity in response to the US. However, our results differ in that they reported a higher proportion of neurones exhibiting peaks rather than troughs in activity for both sucrose and air-puff scenarios; this may reflect a unique activity profile in M2. Furthermore, they found robust serotonergic and dopaminergic responses to US predicting cues, which we did not; this may simply reflect less strong associations made between the cues and US in our paradigm or a lack of cue related information being transferred to M2.

We found a strong distinction in serotonin responses between when mice were locomoting or at rest. During locomotion, serotonergic responses, especially in response to sucrose, were more pronounced and consistent than at rest. This activity likely represents both the deceleration of mice to consume sucrose and the act of consumption. Responses to sucrose at rest were more varied and included a higher proportion of activity peaks. The observed correlations between serotonin and velocity suggest that serotonin tracks both locomotion and sucrose delivery in M2. Phasic serotonin activity has been linked to both velocity decreases [26] and increases [104] in a state dependent manner. Our study provides a distinct perspective of neuromodulation in M2 and highlights the importance of complimenting nuclei recordings with target cortical area recordings.

Noradrenaline

When examining noradrenaline, both cued and uncued revealed a marked decrease in fluorescence. However, the responses to uncued sucrose were noticeably prolonged (Fig 3.12c), leading to a significantly delayed peak decrease compared to cued sucrose ($p = 0.0003$). Distinct differences were also observed between uncued sucrose and uncued air-puff (minima $p =< 0.0001$, maxima $p = 0.0142$), and between uncued and cued air-puffs (minima $p = 0.0009$). Notably, there were no significant disparities between the cued US. This points to unique noradrenergic responses to uncued US, consistent with its known role in signaling novelty [127].

Analysing response composition, uncued sucrose responses were predominantly minima-only (60%), contrasting with 51% for cued sucrose. The percentage of su-

crose uncued responses only with maximum was also smaller (9% vs. 19%). Responses to uncued air-puffs displayed greater variability (34% minima-only, 31% maxima-only, 32% mix), with greater proportion of time courses showing increases in activity than cued air-puffs (20%) and uncued sucrose (9%). This indicates that noradrenaline might exhibit different responses to uncued stimuli based on their valence, with more often decreases in activity for appetitive stimuli and a mix of responses to aversive stimuli.

When focussing on data during consistent locomotion or rest, a distinct picture emerged (Fig 3.18). During locomotion, we saw the same strong dip for sucrose that we found previously, with the majority (71%) being minima-only responses, and no significant maxima observed. However, the magnitude of responses to sucrose at rest was significantly smaller ($p = 0.0033$). Noradrenaline was correlated with calcium activity and velocity during sucrose delivery while moving. The peak in velocity cross-correlation occurred at 0.34s, meaning that noradrenaline faithfully tracked changes in velocity, but did not precede it. Mean noradrenaline responses to air-puffs while moving were smaller than for sucrose ($p = 0.0131$), and were not correlated. This may be attributable to the more diverse neuromodulator responses to air-puffs, in terms of minima and maxima. This also may reflect more diverse behavioural responses to air-puffs such as flinching and escape responses, as compared to more stereotyped sucrose consumption.

At rest, responses to sucrose and air-puff were statistically indistinguishable, and the response composition of minima and maxima for both stimuli was similar, with a higher degree of maxima than their respective moving responses. The average time course for sucrose at rest did not appear to deviate from baseline, mirroring findings by Xiang et al. [127]; although we show that averaging may be concealing varied underlying peaks and troughs. Intriguingly, Feenstra et al. [37] observed that aversive stimuli led to increased noradrenaline release in the PFC, however, the slow sampling rate of their micro-dialysis technique may be capturing elevated tonic activity due to stress [4] rather than the mix of phasic responses we detected.

Overall, it appears that noradrenergic activity changes in relation to appetitive and aversive stimuli are significantly influenced by two factors: expectation and locomotion. Our findings particularly highlight differences in responses to cued versus uncued stimuli, most notably with uncued sucrose. This is consistent with extensive literature linking noradrenaline to saliency [5, 98, 129]. In terms of locomotion, the types of noradrenergic responses varied between resting and moving states. During locomotion, sucrose delivery predominantly triggered minima-only responses, while air-puffs led to a mix of both minima and maxima. The invariance of responses to stimuli of opposing valence at rest, but not during locomotion, suggests that noradrenergic reactions to US are more influenced by changes in locomotion than by the stimuli themselves. This observation aligns with the work of Bouret and Sara [14] which showed that LC activation aligns more with the subsequent behavioural action than the stimulus onset. Furthermore, the noradrenaline consistently showed a higher correlation with velocity than the other neuromodulators.

Dopamine

Examining the time courses of dopamine activity, there was notable consistency in response magnitude and makeup between both cued and uncued US. Both sucrose and air-puffs exhibit a similar dip in activity (Fig 3.12). For sucrose, the majority of time courses displayed minima only (58% cued and 55% uncued), with 27% presenting mixed responses in both cases. For air-puffs, however, the minima-only responses were less frequent (44% cued and 34% uncued), with an increase in mixed responses (34% cued and 41% uncued).

A consistent bi-phasic pattern emerged when examining the delay-to-peaks. This pattern, characterised by an initial dip followed by a rebound, was evident for uncued sucrose ($p = 0.007$), cued air-puff ($p < 0.0001$), and uncued sucrose ($p = 0.013$). This bi-phasic response bears resemblance to the response pattern observed with serotonin. Such bi-phasic patterns might play a role in modulating oscillatory activity, which both dopamine and serotonin have been implicated in [18, 59].

When evaluating US responses during locomotion, dopamine again consistently showed an activity dip shortly after sucrose and air-puff delivery, followed by a rebound as confirmed by the delay-to-peak plots. Responses during locomotion were more temporally consistent than at rest, leading to a clearer mean response pattern (Fig 3.18). Yet, when observing the percentage change plots, the magnitude of dips and peaks between rest and locomotion was virtually indistinguishable. At rest, both sucrose and air-puff had a higher percentage of time courses displaying maxima, possibly obscuring patterns in the mean time courses. The resting state responses seem to feature a temporally varied positive component, which was subdued in the mean but becomes evident when dissecting the components. The correlation analysis showed that dopamine was weakly correlated with velocity during sucrose but not air-puff delivery.

In conclusion, dopamine activity also seems to convey a multifaceted signal, but with a greater focus on the US. The signal includes a component that aligns with locomotion changes and a bi-phasic component that responds to both appetitive and aversive stimuli. Interestingly, despite previous findings that dopamine neurones respond differently to cued versus uncued scenarios [24], we observed no such distinction in our study. Our findings bear significant similarities to those of Patriarchi et al. [86], reproduced in Fig 1.8. Their study underscored the variability of dopamine activity within the motor cortex, identifying regions of interest (ROIs) that were either responsive to locomotion or to rewards. On closely examining their time courses, a pattern emerges that is similar to ours: post-reward delivery, there is an initial decrease in activity, followed by a subsequent rebound. However, their plots primarily highlight the robust action-related response, which somewhat overshadows the dip in activity following sucrose delivery. In contrast, our peristimulus plots are exclusively focused on the responses to the US, offering a clearer view of this complex activity pattern.

4.2.1 Responses to cues

Despite our expectations, we did not observe any neuromodulator responses to cues predicting the US, as shown in Figure 3.22. This finding contrasts with previous findings of responses to US predicting cues from neuromodulator nuclei [24, 120]. One

possible explanation is that the mice did not learn the association between the cues and the outcomes. However, there is evidence indicating that the mice anticipated the stimuli in a state-dependent manner. The changes in velocity depicted in 3.18 suggest that mice often adjusted their speed just before the delivery of stimuli. This observation is further supported by the state-dependent velocity variations shown in Figure 3.4.

Another possibility is that the cue signals, while recognised by the nuclei, were not robustly relayed to the M2 cortex. This could be because these cues were not deemed essential for goal-oriented motor modulation. The cues analysed here were feedback cues which predicted the outcome. It may be that only relevant cues for instrumental tasks are transmitted to M2, but not those associated with Pavlovian conditioning. Considering the suboptimal performance in the instrumental aspect of our task, drawing definitive conclusions remains challenging. Further experiments are essential to definitively determine whether both Pavlovian and instrumental cues are equally transmitted to M2.

Conclusion

We observed that the activity profiles of serotonin and noradrenaline exhibited similarities in response to US. Both neuromodulators demonstrated pronounced decreases in activity upon cued sucrose delivery. Serotonin appeared to confer a multiplexed signal which responded most to changes in locomotion and sucrose delivery, but not to aversive air-puffs. For noradrenaline, the responses appear more closely tied to changes in locomotion rather than the US itself, as indicated by the distinct response profiles at rest which had a larger proportion of peaks in activity. Additionally, both serotonin and noradrenaline exhibited greater variability in response to sucrose at rest and to air-puffs under all conditions. In contrast, dopamine consistently responded to both US types in both cued and uncued conditions. As with the other neuromodulators, dopaminergic responses also presented a higher frequency of maxima when responding to sucrose at rest. Collectively, these findings suggest that all three neuromodulatory systems convey a complex blend of information regarding both locomotion and US, with distinct nuances evident in the response profile of each system.

4.3 Spatiotemporal release dynamics

To delineate the spatial characteristics of neuromodulator release, we examined the pairwise temporal correlations of activity by distance. This method proved especially effective when focussing exclusively on hotspots of activity (as described in 2.9.2). By doing so, we eliminated high noise prevalent in areas with low expression, which otherwise generated a myriad of misleading correlations and anti-correlations. We discovered that, during locomotion onset and offset, activity was broadly correlated across the tissue: spanning 170 μm for serotonin, 220 μm for noradrenaline, and 60 μm for dopamine. However, this expanse of correlated activity diminished considerably during sustained locomotion and rest, with only the nearest distances revealing correlations surpassing our 0.3 significance threshold. Statistical analyses of these groups

revealed no variances between locomotion transitions or between sustained locomotion and rest. Yet, there were marked differences between transitions and sustained states across all neuromodulators. This indicates that all three neuromodulators participate in a widespread, coordinated facilitation of locomotion transition. Nevertheless, we observed distinct differences in the behaviour of each neuromodulator system between states.

Serotonin

Turning our attention to serotonin, we identified a similar pattern of correlation by distance during both onset and offset. Intriguingly, the correlated activity consistently plateaued around 0.3 across all distances. This suggests the presence of a loosely synchronised signal spanning the entire tissue during locomotion transitions — a phenomenon absent during sustained locomotion or rest.

Furthermore, we observed statistically significant disparities between cued and uncued activity. The most prominent distinction was between cued air-puff and both sucrose conditions. This disparity seems driven by the elevated correlated activity over longer distances, a trend absent in cued air-puffs.

The strongest effect was observed in serotonergic responses to sucrose during movement. The elevated activity correlation appears to capture the pronounced response magnitudes for this state depicted in Figure 3.19. Remarkably, activity correlations at both 20 and 380 μm surged to approximately 50%, while intermediary distances saw correlations drop to 30%. Such data might be indicative of distinct modular activity — possibly originating from differences between the dorsal and medial raphe, or even within the raphe itself [96]. The observed discrepancy between cued and uncued scenarios might be a manifestation of the perceived controllability of the state, an aspect that the serotonergic system has been documented to discern [1, 104].

Noradrenaline

During both onset and offset, noradrenergic activity exhibited its strongest correlation locally, but which extended widely across the tissue. These correlations appear to diminish linearly with distance, tapering near zero for larger distances. During sustained locomotion, the correlation of noradrenergic activity was restricted to immediate localities. The swift decline with increasing distance seen in all states hints at spatial modularity within noradrenergic signals, a notion supported both functional [15, 118] and anatomical [20] studies.

A notable difference was evident between the effects of sucrose and air-puff, but not between cued and uncued, states. The correlation-by-distance measurements here therefore align more with the changes in response magnitude rather than response timing that were depicted in Figure 3.13. Further, when comparing responses to US during movement versus rest, noradrenergic activity appeared more coordinated during locomotion — more so for sucrose than air-puff. In comparison to Uematsu et al. [118], our findings do not lend evidence to their 'global broadcast signal' for aversive stimuli hypothesis, as our findings did not indicate correlated noradrenergic activity for air-puffs in a resting state.

Dopamine

Dopaminergic activity displayed the lowest correlation among the studied neuromodulators. During both locomotion onset and offset, only activity within a range of 70μm exhibited correlation. Remarkably, throughout sustained movement and rest, we observed no correlated dopaminergic activity across any measured distance. Similarly, no correlated activity emerged when comparing responses to US during both locomotion and rest, despite significant departures from baseline being evident in the analysis in Figure 3.18. Such findings resonate with the observed heterogeneity of dopamine activity within the cortex [71], wherein even proximate ROIs manifest diverse response patterns [86].

Surprisingly, our analysis of dopaminergic responses unveiled an intriguing pattern: cued sucrose and uncued air-puff exhibited greater similarity, as did uncued sucrose and cued air-puff. This marked distinction wasn't mirrored in our other metrics, hinting that a more in-depth exploration might yield insightful revelations.

Conclusion

Our investigation of the spatiotemporal dynamics of neuromodulator release highlighted distinct patterns in serotonin, noradrenaline, and dopamine. Specifically, while all three neuromodulators showcased widespread activity during locomotion transitions, their correlated activities varied, with serotonin displaying a loosely synchronised signal across the entire tissue and noradrenaline presenting spatial modularity in its signals. Dopamine, in contrast, showed the least correlated activity, emphasising its complex and diverse response patterns within the cortex. Collectively, these findings underscore the intricate nature of neuromodulatory systems in the brain, hinting at the need for more in-depth studies to fully understand their specific roles and interactions.

4.4 Orthogonalised Go/NoGo task

We did not find any performance improvement in any state (Fig 3.1), although individual mice appeared to show small improvements in specific states, see Table A.1. However, we saw clear changes in behaviour between the early and late sessions (Fig 3.4). In comparing velocities during the first and final 50 sessions, we can see that a significant increase in velocity in the the positive reinforcement of Go state ($p< 0.0001$), which was not seen in the positive reinforcement of NoGo. However, this change in behaviour was not reflected in an improvement in task performance, which suggests that mice did form Pavlovian approach behaviour, but that the specifics of the task requirements were too difficult.

Intriguingly, we found a significant increase in velocity Neg-NoGo state ($p< 0.0001$), despite this change in locomotion further increasing the probability of aversive stimuli. This suggests that the air-puff may have elicited a Pavlovian escape response. Mice also increased their velocity during Neg-Go states, however, this trend was not significant. The reason may be the much lower experience of this state, as most mice tended to run most of the time, they often experienced the air-puff negative reinforcement in the NoGo state, but not in the Go state.

Overall, our findings align with other studies using the OGNG task, which demonstrate that performance varies due to the effects of innate Pavlovian tendencies. We anticipated that mice would struggle with certain task states where Pavlovian and instrumental learning are in conflict. Humans also experience difficulty in some states, performing at near-chance levels in the Neg-Go state [50]. However, our data suggest that mice exhibit Pavlovian responses that are distinct from those of humans. While both species improved in the Pos-Go state, benefiting from Pavlovian approach behaviours that support Go actions, their responses to negative reinforcement was differed. Humans decreased locomotion in the face of aversive loss, impeding the Neg-Go state. Whereas mice increased their locomotion in the Neg-NoGo state, suggesting an escape response. This difference likely reflects the species-specific nature of these evolutionary ancient Pavlovian reactions Domjan [29].

Considering the differences in responses between cued and uncued conditions based on the state, it seems that the neuromodulator systems were able to differentiate between these task states. Despite this, there were no observable neuromodulator responses to either the feedback cues 3.22 or the state cues (data not shown). This suggests neuromodulator signals relating to sensory cues are not be transmitted to the M2 motor region. Given the low overall task performance, from our study alone we cannot definitively conclude that M2 does not receive neuromodulator responses to cues. However, this finding is consistent with other research employing GENI in motor regions, which also did not observe cue responses for dopamine [86] or noradrenaline [15].

Comparing directly to Jones et al. [62], the other example of an OGNG task performed in rodents, both our studies found performance particularly low. They used initial single-state training sessions whereby the mice learn one state until they reach criterion, then the other state, learning two states in total. We opted against this method to avoid the complexity of varying training session counts among mice. In retrospect, their approach may have been preferable as many mice eventually learned two states, whereas our mice failed to learn four states. In future, splitting the behavioural requirements into two groups would simplify the task for the mice and may improve task performance. Mice would simultaneously learn either Pavlovian congruent (Pos-Go and Neg-NoGo) or Pavlovian conflict (Neg-Go and Pos-NoGo) states. This approach would balance the simplicity of the method of Jones et al. [62] with our comprehensive design, without imposing performance thresholds, which can pose challenges for imaging studies with optimal periods for data collection.

Another difference was that in their study rodents were reinforced to stay in one location or move to another location, however were not penalised for moving freely within the specified zone. Thus NoGo did not likely evoke motor inhibition. Their task may be akin to a dynamic place-preference/place-avoidance task rather than a Go/NoGo task. In comparison, the our task required full body movement in the Go state. This distinction allowed us to tease apart the difference between neuromodulator responses to locomotion the responses the US. One limitation for us was that mice tended to run while on the setup, which meant that some mice simply did not experience the reinforcement associated with NoGo states.

4.5 Pharmacology

Our pharmacological findings yielded mixed success. For each GENI we administered antagonists to match the GPCR it was based on. With the administration of Yohimbine (an $\alpha 2$ antagonist), we observed a decrease in GRABNE fluorescence in 3 out of 4 animals. Conversely, after administering Reboxetine (NRI), which was anticipated to enhance noradrenaline availability and thus raise GRABNE fluorescence, we noted an increase in fluorescence in 2 out of 4 mice. The variance in responses might be attributed to disparities in intradermal injection targeting; different tissue types can significantly influence drug onset times.

Upon administering SB242084 (a 5-HT2C antagonist), most mice exhibited a marked decline in fluorescence. Yet, in one instance, there was a pronounced rebound in activity. Deciphering this outcome is challenging, but it echoes the complex role of $5HT_{2C}$ receptors as highlighted by [16]. An even more perplexing observation was made with SCH23390 (a D1 antagonist), which elicited pronounced fluorescence changes in either direction, contingent upon the individual mouse. Based on manual imaging video assessment, in the case of increased activity, there was an almost complete whiteout of dLight activity across the imaging region. This surge was especially surprising considering the assertion by Patriarchi et al. [86] that "The response to DA was abolished in the presence of the DRD1 antagonists SCH-23390".

Each data point in our graphs is derived from the average activity of 10 video segments, each 10 seconds long, captured during a 10-minute video session. All segments underwent rigorous and identical motion correction, significantly reducing the likelihood that our results are artefacts. However, to further validate our findings, additional experiments involving these drugs are necessary. One limitation of our study was the timing of the pharmacological interventions, which occurred after the optimal imaging window. This can lead to complications from overexpressing GENIs or GECIs, resulting in increased cellular dysfunction and mortality. In retrospect, it would have been more advantageous to perform pharmacological controls at the beginning of the study or in a separate cohort of mice.

4.6 Methodological limitations

In our study, we refrained from making direct statistical comparisons between neuromodulators due to differences in GENI dynamics and expression. Notably, the onset and offset dynamics of each GENI (τ_{on}, τ_{off}) and their dynamic ranges varies significantly (see 1.6), rendering the validity of direct comparisons questionable. However, by comparing the same animals across different states, we gained insights into state-dependent changes.

We also noted that distinct GENI expression patterns emerged, with apparent interactions between GENIs and GECIs expression. Specifically, while GRABNE expression was robust, tropism competition appears to have lead to distinct areas of expression for each with minimal overlap (Fig. 3.5). Despite dLight exhibiting low expression, its dynamic range was high. Among them, GRAB5HT demonstrated the most balanced GENI and GECI expression, promising valuable insights in follow-up

analysis.

The GENIs employed in this study are less established than calcium indicators and therefore warrant cautious interpretation of results. In particular, while the GENI are highly ligand-specific due to their inclusion of endogenous receptors, there remains the possibility of other confounding physiological phenomenon. Null-mutant controls, which are ligand-insensitive, would have been beneficial to rule out artefacts, such as haemodynamic changes or pH fluctuations. Recent findings suggest haemodynamic changes could significantly influence the GENI signal [83]. A static control fluorophore, such as a null-mutant or RFP, would have allowed for estimation of haemodynamic influences. Additionally, most GFP-based biosensors, including those used here, are known to be pH-sensitive [53], a factor not controlled for in this experiment nor addressed in the original development of these GENIs [38, 86, 110, 123]. Future work should include null-mutants to ascertain the impact of haemodynamic and pH changes on GENI signals.

An important limitation affecting the generalisability of our findings is the use of head fixation, which is discernibly stressful for mice. This is problematic for investigating neuromodulator release, as neuromodulators are deeply connected with stress responses [5, 104]. Stress can result in increased baseline activity and reduce performance-related phasic responses for noradrenaline, dopamine, and serotonin [3, 131]. Despite our rigorous efforts to acclimate mice to head fixation, it is worth noting that this may also have augmented neuromodulator responses and also contributed to slow task learning [80].

Addressing motion artefacts in our imaging was challenging. We meticulously selected only stable video sections and further eliminated components correlating with the motion correction algorithm. Initially, our SVD correction targeted each short video segment. But given the correlation between locomotion and motion correction, this approach was sometimes overly conservative. First applying the SVD method to the entire 20,000-frame unbinned video, followed by binning for analysis, proved more effective. This method distinguished between tissue motion and locomotion-induced fluorescence changes. However, as with all two-photon imaging, there is still a possibility of lingering motion artefacts.

It should be noted that the lack of performance improvement across all states might be due to inadequate training. Our regimen entailed two 10-minute sessions daily, amounting to about 300 trials per state. In comparison, Goltstein et al. [45] found that mice will perform up to 200 trials per day, while Patriarchi et al. [86] had mice perform 300-700 trials over 60-90 minutes per day. Repeating this study with longer sessions for each animal per day may result in improved performance which would better help answer our experimental aims.

We should also note the lower number of video sections available for certain analysis. This was the result of a combination of factors. Firstly, mouse behavioural tendencies, whereby certain mice almost never rested, resulting in very few sessions to examine US at rest. Secondly, after initial analysis produced some unexpectedly large results, we checked our recordings and found the persistence of some motion artefacts. As we were using a novel technique in an area that had very limited comparative studies in the field, we aired on the side of caution and excluded any video sections with signs of motion. This meant that a large number of the trials we recorded were excluded.

Subsequent investigations in M2 may benefit limiting the scope of their behavioural paradigm to only delivering US to mice at rest. Using a technique such as one-photon microscopy, which has some well-established and robust motion correction strategies, may also be fruitful.

4.7 Concluding thoughts on the open questions

What is the spatial extent of volume transmission?

To determine the spatial extent of neuromodulator release, we analysed activity correlations in relation to distance. We looked at two length scales, 5-10-15 μm in (e.g. Figure 3.10) and 25-400 $/mu$m scale in (e.g. Figure 3.23a). As noted in the methodological limitations, the relative novelty of the GENI make drawing strong conclusions difficult. However, our results provide some preliminary insights. We observed that the spatial extent varied with locomotion state and US delivery at the $>$ 25 μm scale. However, we did not observe differences at the 5 μm scale.

Serotonin exhibited an unexpected spatial release profile. During locomotion transitions, peak correlations were observed at shorter distances, with a consistent general correlation of around 0.3 observed throughout the tissue. Intriguingly, under the combination of sucrose delivery and locomotion, serotonin showed a more correlated activity across the tissue than during sucrose delivery at rest or air-puff delivery during movement.

Noradrenaline demonstrated the most widespread activity correlations, which generally decreased linearly with distance. Its release was notably widespread during locomotion onset, spanning up to 220 μm, and during sucrose delivery while moving, extending to at least 360 μm. In contrast, during sucrose delivery while at rest, correlated signals were confined to a mere 20 μm.

Dopamine presented the most heterogeneous signals among the neuromodulators. Correlated activity was predominantly observed at very short distances ($<$100 μm). Surprisingly, during US delivery, despite significant dips in mean hotspot fluorescence, spatial correlation was absent. This implies a heterogeneous release pattern into small pockets, marked by large amplitude fluctuations.

Future experiments could benefit from employing optogenetic stimulation at neuromodulator nuclei alongside two-photon recordings in the cortex. Applying diverse stimulation protocols, including inhibition, weak and tetanic stimulation, would determine spatial spread during minimum and maximum release. Additionally, utilising various lenses could prove beneficial. A lower magnification lens would assess if different cortical areas receive distinct, modular neuromodulator signals. Conversely, a higher magnification lens would help determine the presence of small-scale gradients in neuromodulator release.

Does *in vivo* neuromodulator release cause an increase or decrease in neuronal activity?

To evaluate the relationship between neuronal and neuromodulator activities, we analysed the correlation between calcium activity in soma and the surrounding neuromod-

ulator activity. During the onset of locomotion, all neuromodulator systems demonstrated a positive correlation with calcium activity. However, at the cessation of movement, although the average correlations remained positive across all systems, dopamine displayed a significant portion of activity that was anti-correlated. In response to US, serotonin and noradrenaline continued to show a positive correlation with calcium activity, but dopamine did not, particularly during air-puff delivery in motion, where the correlation between dopamine and calcium activity was notably absent.

Nonetheless, the axiom that correlation does not imply causation is pertinent here. Cross-correlation analysis revealed that changes in calcium activity aligned with neuromodulator changes, peaking simultaneously (0 lag), indicating that, on average, neither signal consistently preceded the other. The correlation between calcium and neuromodulator activities was moderate during the onset and offset of locomotion but fell below the 0.2 threshold in most other conditions, suggesting a non-causal correlation and the influence of common external variables, at least at this time scale.

Although these findings may seem to counter the hypothesis that noradrenaline acts as a network reset signal for initiating task-relevant actions [13], it is important to recognise that those results applied to task-specific actions. In our case, due to low task performance, it is likely that the observed actions were self-directed. Future research focusing on task-driven action initiation might uncover a preceding noradrenergic signal relative to both the action and calcium activity.

Overall, serotonin and noradrenaline activity tended to positively correlate with calcium activity. This correlation appears to vary depending on the state of locomotion, generally between 0.2 and 0.3. Whereas, dopamine activity did not consistently correlate with calcium activity, suggesting varied interactions between different neuromodulators and neuronal activity.

It is important to note that our findings should be interpreted with caution. The novel GENI used in this study were not robustly characterised here. Specifically, there were difficulties encountered with the pharmacological experiments, and there was a lack of controls for haemodynamic and pH confounds. This limits our ability to make definitive statements about the specific influences of each neuromodulator on neuronal calcium activity. Future experiments should use GENI in combination with a static indicator, such as GFP, as a robust control.

How do neurones response to the mixture in vivo neuromodulator signals?

We observed that all three neuromodulator systems responded similarly to locomotion transitions and while some differences emerged in their responses to US these results were also more similar than we had expected. This raises questions about the necessity and function of these overlapping neuromodulator systems. Is their redundancy strategic, designed to compensate for potential dysfunctions? Or does it reflect a nuanced, context-dependent orchestration of multiple neuromodulatory actions, activating simultaneously or selectively depending on the situation?

While there is some evidence of compensatory mechanisms, as noted in the introduction, a key characteristic of neuromodulator systems is the diversity of effects these systems can produce. Each neuromodulator system influences neuronal excitability in unique ways, determined by receptor locations, downstream effects, and the types

of neurones bearing these receptors. Moreover, instances of cooperative interactions between these systems have been documented. For example, Arnsten [4] described how noradrenaline and dopamine can concurrently enhance ensembles relating to correct responses and suppress competing ensembles, respectively. Our findings provide some conditions under which neuromodulators exhibit overlapping responses, such as locomotion transitions, and distinct responses, such as to air-puffs.

From our results it could be possible to analyse neuronal calcium activity in the context of specific neuromodulator milieu. However, the complexity of this becomes apparent when examining just three systems, not to mention the probable external influences also play a role. Additionally, there is an additional challenge of the weak causal link between *in vivo* neuromodulator release and changes in neuronal calcium activity, as suggested by our cross-correlation results. Consequently, a more extensive dataset is required to delineate how each neuromodulator milieu influences neuronal excitability. Employing more direct methods such as optogenetic stimulation or inhibition might provide clearer insights into how neurones react to mixed neuromodulator signals.

Our results underscore the importance of further experiments using tasks that require both movement and stationary states. Although our study highlighted differences in responses to US during locomotion and rest, the low task performance limited our ability to make definitive conclusions about task-related movements. Future studies should aim to delineate these neuromodulator responses more precisely in well-learned, tightly timed tasks, building on the preliminary findings presented here.

What is the overall role of each neuromodulator?

Our findings shed some light on the murky function of neuromodulators specifically in the M2 region. The cross-correlation analysis between neuromodulator release, calcium activity, and velocity indicates that changes in velocity often preceded alterations in both neuromodulator and calcium levels. Given that M2 is implicated in motor planning it might be expected activity should precede action. However, our analysis focused on spontaneous behaviour, and our recordings were from layer 2/3—which is known to receive sensory information from other cortical regions. In these cases it is probable that calcium signals relate to sensory feedback, which aligns with the known functionalities of murine M2. However, it is intriguing to note that neuromodulator activity also follows changes in velocity, and may be conveying sensory feedback information. Under the assumption that neuromodulator activity is consistent between layers, this sensory feedback signal would also be conveyed to the layer 5 output layers, which is not a commonly held hypothesis.

Delving into the specific roles of each neuromodulator, our data suggest that serotonin is more closely associated with responses to positive stimuli and locomotion, while noradrenaline exhibits a similar pattern but seems particularly linked to locomotion. Dopamine in the cortex, on the other hand, shows a high degree of spatial heterogeneity. This distinction highlights the complexity and specialisation of cortical neuromodulatory systems. To further understand these roles, follow-up experiments could involve cortical slice studies where a mixture of serotonin, noradrenaline, and dopamine are applied in varying ratios. Such experiments would provide a solid foun-

dation for answering questions about the specific functions of these neuromodulators.

Conclusion

Our investigation has provided in-depth insights into the spatiotemporal dynamics of serotonin, noradrenaline, and dopamine in the secondary motor cortex, revealing how these neuromodulators respond to various stimuli and behavioural states. Employing high-resolution two-photon microscopy and innovative GENI, we were able to elucidate the nuanced responses of these neuromodulators to different stimuli and in different contexts.

Our first significant observation was the robust correlation between neuromodulator activity and transitions in locomotion. All three neuromodulators showed robust increases in activity during locomotion onset and decreases during locomotion cessation. This was particularly evident in noradrenaline, which showed a strong coupling with locomotive states. Cross-correlation analysis suggests that neuromodulator activity generally follows changes in velocity, indicating an unexpected role in sensory feedback.

Another key finding was the unexpected decrease in neuromodulator activity in response to both appetitive (sucrose) and aversive (air-puff) stimuli, diverging from our initial expectations of opposing responses. However, this decline appears to be predominantly related to decreases in locomotion in response to US rather than to the US itself, as suggested by the clear distinctions between US delivery during locomotion and during rest. We also found nuanced differences between neuromodulator responses. Serotonin and noradrenaline displayed more potent and uniform responses during locomotion compared to rest, in particular in response to sucrose. Whereas, dopamine responses were more consistent across states and yet more heterogeneous in spatial release.

Furthermore, response patterns for all three systems were surprisingly varied, often manifesting as biphasic responses of minima and maxima which varied depending on the behavioural state and US. This indicates multifaceted neuromodulatory responses that influence locomotion behaviour and behavioural responses to US. This study was limited to the motor region. To gain further insight into spatiotemporal release characteristics in the cortex, future research should extend these experiments to other brain regions.

We developed a novel OGNG task for mice with full-body Go and NoGo action requirements. This approach allowed us to dissect how M2 responds to multifaceted requirements involving locomotion actions and US-related actions. Our finding that the Pavlovian approach and avoidance behaviours can impede performance align with the work of Guitart-Masip et al. [50] on the original OGNG task, however, with mice the influence of avoidance behaviours was particularly strong. Further studies with

more trials per day are necessary to optimise performance in this task.

Our analyses reveal that locomotion acts as a significant contextual factor in neuromodulator responses. Neuromodulators respond both to the US and the motor behaviours associated with them. This idea aligns with the ideas proposed by Guitart-Masip et al. [50] when they designed the original OGNG task. However, this interaction complicates the investigation of the role of neuromodulators in Pavlovian and instrumental conditioning. In may be useful for future studies to consider the influence the ability of subjects to move on their neuromodulatory responses.

Our findings contribute new insights into the complex interplay between neuromodulators and external stimuli, and suggest that the state of locomotion is an important contextual factor that modulates these responses. This work lays the groundwork for future research aimed at disentangling the multifaceted mechanisms that underlie neuromodulatory function in reaction to environmental cues and states. These findings have implications for theories of motor control, decision-making, and neuropsychiatric disorders involving the serotonin, noradrenaline, and dopamine systems.

Bibliography

[1] J. Amat, M. V. Baratta, E. Paul, S. T. Bland, L. R. Watkins, and S. F. Maier. Medial prefrontal cortex determines how stressor controllability affects behavior and dorsal raphe nucleus. *Nature neuroscience*, 8(3):365–371, Mar. 2005.

[2] C. Aoki and V. M. Pickel. Ultrastructural relations between beta-adrenergic receptors and catecholaminergic neurons. *Brain research bulletin*, 29(3-4):257–263, Sept. 1992.

[3] A. F. T. Arnsten. Stress signalling pathways that impair prefrontal cortex structure and function. *Nature reviews. Neuroscience*, 10(6):410–422, June 2009.

[4] A. F. T. Arnsten. Catecholamine influences on dorsolateral prefrontal cortical networks. *Biological psychiatry*, 69(12):e89–99, June 2011.

[5] G. Aston-Jones and J. D. Cohen. An intergrative theory of locus coeruleus - norepinepherine function: adaptive gain and optimal performance. *dx.doi.org*, 28 (1):403–450, July 2005.

[6] M. C. Avery and J. L. Krichmar. Neuromodulatory systems and their interactions: a review of models, theories, and experiments. *Frontiers in Neural Circuits*, 11:108, 2017.

[7] D. Avesar and A. T. Gulledge. Selective serotonergic excitation of callosal projection neurons. *Frontiers in Neural Circuits*, 6:12, 2012.

[8] O. M. Balcer, M. A. Seager, S. D. Gleason, X. Li, K. Rasmussen, J. K. Maxwell, G. Nomikos, A. Degroot, and J. M. Witkin. Evaluation of 5-HT7 receptor antagonism for the treatment of anxiety, depression, and schizophrenia through the use of receptor-deficient mice. *Behavioural brain research*, 360:270–278, Mar. 2019.

[9] S. J. Bang, P. Jensen, S. M. Dymecki, and K. G. Commons. Projections and interconnections of genetically defined serotonin neurons in mice. *The European journal of neuroscience*, 35(1):85–96, Jan. 2012.

[10] F. Barthas and A. C. Kwan. Secondary motor cortex: where 'sensory' meets 'motor' in the rodent frontal cortex. *Trends in neurosciences*, 40(3):181–193, Mar. 2017.

[11] M. Bogdanov, J. E. Timmermann, J. Gläscher, F. C. Hummel, and L. Schwabe. Causal role of the inferolateral prefrontal cortex in balancing goal-directed and habitual control of behavior. *Scientific reports*, 8(1):9382, June 2018.

[12] J. Bonaventura, C. Quiroz, N.-S. Cai, M. Rubinstein, G. Tanda, and S. Ferré. Key role of the dopamine D4 receptor in the modulation of corticostriatal glutamatergic neurotransmission. *Science advances*, 3(1):e1601631, Jan. 2017.

[13] S. Bouret and S. J. Sara. Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning. *The European journal of neuroscience*, 20(3):791–802, Aug. 2004.

[14] S. Bouret and S. J. Sara. Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends in neurosciences*, 28(11):574–582, Nov. 2005.

[15] V. Breton-Provencher, G. T. Drummond, J. Feng, Y. Li, and M. Sur. Spatiotemporal dynamics of noradrenaline during learned behaviour. *Nature*, 606(7915): 732–738, June 2022.

[16] R. L. Carhart-Harris and D. J. Nutt. Serotonin and brain function: a tale of two receptors. *Journal of psychopharmacology (Oxford, England)*, 31(9):1091–1120, Sept. 2017.

[17] G. V. Carr, F. Maltese, D. R. Sibley, D. R. Weinberger, and F. Papaleo. The dopamine D5 receptor is involved in working memory. *Frontiers in pharmacology*, 8:666, 2017.

[18] P. Celada, M. V. Puig, and F. Artigas. Serotonin modulation of cortical neurons and networks. *Frontiers in Integrative Neuroscience*, 7:25, 2013.

[19] J.-C. Cerpa, A. R. Marchand, and E. Coutureau. Distinct regional patterns in noradrenergic innervation of the rat prefrontal cortex. *Journal of chemical neuroanatomy*, 96:102–109, Mar. 2019.

[20] D. J. Chandler and B. D. P. D. Waterhouse. Evidence for broad versus segregated projections from cholinergic and noradrenergic nuclei to functionally and anatomically discrete subregions of prefrontal cortex. *Frontiers in behavioral neuroscience*, 6, May 2012.

[21] Chandler, Dan J, Jensen, Patricia, McCall, Jordan G, Pickering, Anthony E, Schwarz, Lindsay A, and Totah, Nelson K. Redefining noradrenergic neuromodulation of Behavior: impacts of a modular locus coeruleus architecture. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 39 (42):8239–8249, Oct. 2019.

[22] R. L. Clarkson, A. T. Liptak, S. M. Gee, V. S. Sohal, and K. J. Bender. D3 receptors regulate excitability in a unique class of prefrontal pyramidal cells. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 37 (24):5846–5860, June 2017.

[23] J. Y. Cohen and C. D. Grossman. Dorsal raphe serotonergic neurons regulate behavior on multiple timescales. *Handbook of the Behavioral Neurobiology of Serotonin*, 31:521–529, Jan. 2020.

[24] J. Y. Cohen, M. W. Amoroso, and N. Uchida. Serotonergic neurons signal reward and punishment on multiple timescales. *eLife*, 4:463, Feb. 2015.

[25] C. Colangelo, P. Shichkova, D. Keller, H. Markram, and S. Ramaswamy. Cellular, synaptic and network effects of acetylcholine in the neocortex. *Frontiers in Neural Circuits*, 13:24, 2019.

[26] P. A. Correia, E. Lottem, D. Banerjee, A. S. Machado, M. R. Carey, and Z. F. Mainen. Transient inhibition and long-term facilitation of locomotion by phasic optogenetic activation of serotonin neurons. *eLife*, 6, 2017.

[27] D. Datta, S.-T. Yang, V. C. Galvin, J. Solder, F. Luo, Y. M. Morozov, J. Arellano, A. Duque, P. Rakic, A. F. T. Arnsten, and M. Wang. Noradrenergic α1-adrenoceptor actions in the primate dorsolateral prefrontal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 39(14): 2722–2734, Apr. 2019.

[28] G. Di Giovanni, A. Chagraoui, R. Bharatiya, and P. De Deurwaerdère. Serotonergic control of excitability: from neuron to networks. *Handbook of the Behavioral Neurobiology of Serotonin*, 31:197–215, Jan. 2020.

[29] M. Domjan. The essentials of conditioning and learning. Pacific Grove, CA: Brooks, 1996.

[30] Dong, Hui, Wang, Juan, Yang, Yan-Fei, Shen, Yan, Qu, Wei-Min, and Huang, Zhi-Li. Dorsal striatum dopamine levels fluctuate across the sleep-wake cycle and respond to salient Stimuli in Mice. *Frontiers in neuroscience*, 13:242, 2019.

[31] K. Doya. Metalearning and neuromodulation. *Neural networks : the official journal of the International Neural Network Society*, 15(4-6):495–506, June 2002.

[32] K. Doya. Modulators of decision making. *Nature neuroscience*, 11(4):410–416, Apr. 2008.

[33] V. A. Doze, R. S. Papay, B. L. Goldenstein, M. K. Gupta, K. M. Collette, B. W. Nelson, M. J. Lyons, B. A. Davis, E. J. Luger, S. G. Wood, J. R. Haselton, P. C. Simpson, and D. M. Perez. Long-term α1A-adrenergic receptor stimulation improves synaptic plasticity, cognitive function, mood, and longevity. *Molecular Pharmacology*, 80(4):747–758, Oct. 2011.

[34] B. Engelhard, J. Finkelstein, J. Cox, W. Fleming, H. J. Jang, S. Ornelas, S. A. Koay, S. Y. Thiberge, N. D. Daw, D. W. Tank, and I. B. Witten. Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. *Nature*, 570(7762):509–513, June 2019.

[35] D. Eriksson, M. Heiland, A. Schneider, and I. Diester. Distinct dynamics of neuronal activity during concurrent motor planning and execution. *Nature communications*, 12(1):5390, Sept. 2021.

[36] L. L. Fan, Q. J. Zhang, J. Liu, J. Feng, Z. H. Gui, U. Ali, L. Zhang, C. Hou, T. Wang, Y. P. Hui, Y. N. Sun, and Z. H. Wu. In vivo effect of 5-HT7 receptor agonist on pyramidal neurons in medial frontal cortex of normal and 6-hydroxydopamine-lesioned rats: an electrophysiological study. *Neuroscience*, 190:328–338, Sept. 2011.

[37] M. Feenstra, M. Vogel, M. Botterblom, R. Joosten, and J. de Bruin. Dopamine and noradrenaline efflux in the rat prefrontal cortex after classical aversive conditioning to an auditory cue. *European Journal of Neuroscience*, 2001.

[38] J. Feng, C. Zhang, J. E. Lischinsky, M. Jing, J. Zhou, H. Wang, Y. Zhang, A. Dong, Z. Wu, H. Wu, W. Chen, P. Zhang, J. Zou, S. A. Hires, J. J. Zhu, G. Cui, D. Lin, J. Du, and Y. Li. A genetically encoded fluorescent sensor for rapid and specific in vivo detection of norepinephrine. *Neuron*, 102(4):745–761.e8, May 2019.

[39] S. B. Floresco. Prefrontal dopamine and behavioral flexibility: shifting from an "inverted-U" toward a family of functions. *Frontiers in neuroscience*, 7, Apr. 2013.

[40] A. Friedman, D. Homma, L. G. Gibb, K.-i. Amemori, S. J. Rubin, A. S. Hood, M. H. Riad, and A. M. Graybiel. A corticostriatal path Ttrgeting striosomes controls decision-making under conflict. *Cell*, 161(6):1320–1333, June 2015.

[41] K. Friston, P. Schwartenbeck, T. FitzGerald, M. Moutoussis, T. Behrens, and R. J. Dolan. The anatomy of choice: dopamine and decision-making. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 369(1655):20130481–20130481, Nov. 2014.

[42] A. Funamizu, B. Kuhn, and K. Doya. Neural substrate of dynamic Bayesian inference in the cerebral cortex. *Nature neuroscience*, 19(12):1682–1689, Dec. 2016.

[43] A. Furr, M. D. Lapiz-Bluhm, and D. A. Morilak. 5-HT2A receptors in the orbitofrontal cortex facilitate reversal learning and contribute to the beneficial cognitive effects of chronic citalopram treatment in rats. *The international journal of neuropsychopharmacology*, 15(9):1295–1305, Oct. 2012.

[44] K. Fuxe, L. F. Agnati, M. Marcoli, and D. O. Borroto-Escuela. Volume transmission in central dopamine and noradrenaline neurons and its astroglial targets. *Neurochemical research*, 40(12):2600–2614, Dec. 2015.

[45] P. M. Goltstein, S. Reinert, A. Glas, T. Bonhoeffer, and M. Hübener. Food and water restriction lead to differential learning behaviors in a head-fixed two-choice visual discrimination task for mice. *PLOS ONE*, 13(9):e0204066, 2018.

[46] N. M. Goodfellow, C. D. C. Bailey, and E. K. Lambe. The native serotonin 5-HT5A receptor: electrophysiological characterization in rodent cortex and 5-HT1A-mediated compensatory plasticity in the knock-out mouse. *The Journal*

of neuroscience : the official journal of the Society for Neuroscience, 32(17): 5804–5809, Apr. 2012.

[47] P. Greengard. The neurobiology of slow synaptic transmission. *Science*, 294 (5544):1024–1030, Nov. 2001.

[48] K. Grzelka, P. Kurowski, M. Gawlak, and P. Szulczyk. Noradrenaline modulates the membrane potential and holding current of medial prefrontal cortex pyramidal neurons via β1-adrenergic Receptors and HCN Channels. *Frontiers in cellular neuroscience*, 11:341, 2017.

[49] M. Guitart-Masip, Q. J. M. Huys, L. Fuentemilla, P. Dayan, E. Duzel, and R. J. Dolan. Go and no-go learning in reward and punishment: Interactions between affect and effect. *NeuroImage*, 62(1):154–166, Aug. 2012.

[50] M. Guitart-Masip, E. Duzel, R. Dolan, and P. Dayan. Action versus valence in decision making. *Trends in Cognitive Sciences*, 18(4):194–202, Apr. 2014.

[51] M. Guitart-Masip, M. Economides, Q. J. M. Huys, M. J. Frank, R. Chowdhury, E. Duzel, P. Dayan, and R. J. Dolan. Differential, but not opponent, effects of L -DOPA and citalopram on action learning with reward and punishment. *Psychopharmacology*, 231(5):955–966, Mar. 2014.

[52] N. Hansen and D. Manahan-Vaughan. Locus Coeruleus Stimulation Facilitates Long-Term Depression in the Dentate Gyrus That Requires Activation of β-Adrenergic Receptors. *Cerebral cortex (New York, N.Y. : 1991)*, 25(7):1889–1896, July 2015.

[53] S. Hario, G. N. T. Le, H. Sugimoto, K. Takahashi-Yamashiro, S. Nishinami, H. Toda, S. Li, J. S. Marvin, S. Kuroda, M. Drobizhev, T. Terai, Y. Nasu, and R. E. Campbell. High-performance genetically encoded green fluorescent biosensors for intracellular l-Lactate. *ACS central science*, 10(2):402–416, Feb. 2024.

[54] K. D. Harris. Top-Down Control of Cortical State. *Neuron*, 79(3):408–410, Aug. 2013.

[55] D. M. Herz, H. R. Siebner, O. J. Hulme, E. Florin, M. S. Christensen, and L. Timmermann. Levodopa reinstates connectivity from prefrontal to premotor cortex during externally paced movement in Parkinson's disease. *NeuroImage*, 90:15–23, Apr. 2014.

[56] H. Homayoun and B. Moghaddam. Differential representation of Pavlovian-instrumental transfer by prefrontal cortex subregions and striatum. *The European journal of neuroscience*, 29(7):1461–1476, Apr. 2009.

[57] J. A. Hosp, H. E. Nolan, and A. R. Luft. Topography and collateralization of dopaminergic projections to primary motor cortex in rats. *Experimental brain research*, 233(5):1365–1375, May 2015.

[58] K. Iigaya, M. S. Fonseca, M. Murakami, Z. F. Mainen, and P. Dayan. An effect of serotonergic stimulation on learning rates for rewards apparent after long intertrial intervals. *Nature communications*, 9(1):2477–10, June 2018.

[59] L. Iskhakova, P. Rappel, M. Deffains, G. Fonar, O. Marmor, R. Paz, Z. Israel, R. Eitan, and H. Bergman. Modulation of dopamine tone induces frequency shifts in cortico-basal ganglia beta oscillations. *Nature communications*, 12(1): 7026, Dec. 2021.

[60] M. Jing, P. Zhang, G. Wang, J. Feng, L. Mesik, J. Zeng, H. Jiang, S. Wang, J. C. Looby, N. A. Guagliardo, L. W. Langma, J. Lu, Y. Zuo, D. A. Talmage, L. W. Role, P. Q. Barrett, L. I. Zhang, M. Luo, Y. Song, J. J. Zhu, and Y. Li. A genetically encoded fluorescent acetylcholine indicator for in vitro and in vivo studies. *Nature biotechnology*, 36(8):726–737, Sept. 2018.

[61] M. Jing, Y. Zhang, H. Wang, and Y. Li. G-protein-coupled receptor-based sensors for imaging neurochemicals with high sensitivity and specificity. *Journal of neurochemistry*, 151(3):279–288, Nov. 2019.

[62] S. Jones, E. S. Paul, P. Dayan, E. S. J. Robinson, and M. Mendl. Pavlovian influences on learning differ between rats and mice in a counter-balanced Go/NoGo judgement bias task. *Behavioural brain research*, 331:214–224, July 2017.

[63] H. R. Kim, A. N. Malik, J. G. Mikhael, P. Bech, I. Tsutsui-Kimura, F. Sun, Y. Zhang, Y. Li, M. Watabe-Uchida, S. J. Gershman, and N. Uchida. A unified framework for dopamine signals across timescales. *Cell*, 183(6):1600–1616.e25, Dec. 2020.

[64] A. E. Kisilevsky and G. W. Zamponi. D2 dopamine receptors interact directly with N-type calcium channels and regulate channel surface expression levels. *Channels*, 2(4):269–277, July 2008.

[65] S.-H. Lee and Y. Dan. Neuromodulation of brain states. *Neuron*, 76(1):209–222, Oct. 2012.

[66] S. C. Leiser, Y. Li, A. L. Pehrson, E. Dale, G. Smagin, and C. Sanchez. Serotonergic regulation of prefrontal cortical circuitries involved in cognitive processing: a review of individual 5-HT receptor mechanisms and concerted effects of 5-HT receptors exemplified by the multimodal antidepressant vortioxetine. *ACS chemical neuroscience*, 6(7):970–986, Mar. 2015.

[67] J. J. Letzkus, S. B. E. Wolff, and A. Lüthi. Disinhibition, a Circuit Mechanism for Associative Learning and Memory. *Neuron*, 88(2):264–276, Oct. 2015.

[68] Y. Li, W. Zhong, D. Wang, Q. Feng, Z. Liu, J. Zhou, C. Jia, F. Hu, J. Zeng, Q. Guo, L. Fu, and M. Luo. Serotonin neurons in the dorsal raphe nucleus encode reward signals. *Nature communications*, 7:10503, Jan. 2016.

[69] C. Liu and P. S. Kaeser. Mechanisms and regulation of dopamine release. *Current opinion in neurobiology*, 57:46–53, Aug. 2019.

[70] J. Liu, D. Liu, X. Pu, K. Zou, T. Xie, Y. Li, and H. Yao. The secondary motor cortex-striatum circuit contributes to suppressing inappropriate responses in perceptual decision behavior. *Neuroscience bulletin*, 39(10):1544–1560, Oct. 2023.

[71] S. Lohani, A. K. Martig, K. Deisseroth, I. B. Witten, and B. Moghaddam. Dopamine modulation of prefrontal cortex activity is manifold and operates at multiple temporal and spatial scales. *Cell Reports*, 27(1):99–114.e6, Apr. 2019.

[72] M. L. Lőrincz and A. R. Adamantidis. Monoaminergic control of brain states and sensory processing: Existing knowledge and recent insights obtained with optogenetics. *Progress in neurobiology*, 151:237–253, Apr. 2017.

[73] S. Ma, H. Zhong, X. Liu, and L. Wang. Spatial distribution of neurons expressing single, double, and triple molecular characteristics of glutamatergic, dopaminergic, or GABAergic neurons in the mouse ventral tegmental area. *Journal of molecular neuroscience : MN*, pages 1–18, May 2023.

[74] S. Matias, E. Lottem, G. P. Dugué, and Z. F. Mainen. Activity patterns of serotonin neurons underlying cognitive flexibility. *eLife*, 6:365, Mar. 2017.

[75] A. M. Michaiel, P. R. L. Parker, and C. M. Niell. A hallucinogenic serotonin-2A receptor agonist reduces visual response gain and alters temporal dynamics in mouse V1. *Cell Reports*, 26(13):3475–3483.e4, Mar. 2019.

[76] K. Miyazaki, K. W. Miyazaki, A. Yamanaka, T. Tokuda, K. F. Tanaka, and K. Doya. Reward probability and timing uncertainty alter the effect of dorsal raphe serotonin neurons on patience. *Nature communications*, 9(1):2048, June 2018.

[77] K. W. Miyazaki, K. Miyazaki, K. F. Tanaka, A. Yamanaka, A. Takahashi, S. Tabuchi, and K. Doya. Optogenetic activation of dorsal raphe serotonin neurons enhances patience for future rewards. *Current biology : CB*, 24(17): 2033–2040, Sept. 2014.

[78] A. Mueller, S. B. Shepard, and T. Moore. Differential expression of dopamine D5 receptors across neuronal subtypes in macaque frontal eye field. *Frontiers in neural circuits*, 12:12, 2018.

[79] M. Murakami, M. I. Vicente, G. M. Costa, and Z. F. Mainen. Neural antecedents of self-initiated actions in secondary motor cortex. *Nature neuroscience*, 17(11): 1574–1582, Nov. 2014.

[80] A. Nasr, S. E. Dominiak, K. Sehara, M. A. Nashaat, R. N. S. Sachdev, and M. E. Larkum. Efficient training approaches for optimizing behavioral performance and reducing head fixation time. *PloS one*, 17(11):e0276531, 2022.

[81] K. M. Nautiyal and R. Hen. Serotonin receptors in depression: from A to B. *F1000Research*, 6, 2017.

[82] S. Nomura, M. Bouhadana, C. Morel, P. Faure, B. Cauli, B. Lambolez, and R. Hepp. Noradrenalin and dopamine receptors both control cAMP-PKA signaling throughout the cerebral cortex. *Frontiers in cellular neuroscience*, 8:247, 2014.

[83] G. Ocana-Santero, A. M. Packer, T. Sharp, and S. J. B. Butt. In vivo two-Photon microscopy reveals sensory-evoked serotonin (5-HT) release in adult mammalian neocortex. *ACS chemical neuroscience*, 15(3):456–461, Feb. 2024.

[84] M. Parent and L. Descarries. Ultrastructure of the serotonin innervation in mammalian central nervous system. *Handbook of the Behavioral Neurobiology of Serotonin*, 31:49–90, Jan. 2020.

[85] T. Parr and K. J. Friston. Uncertainty, epistemics and active inference. *Journal of the Royal Society, Interface*, 14(136):20170376, Nov. 2017.

[86] T. Patriarchi, J. R. Cho, K. Merten, M. W. Howe, A. Marley, W.-H. Xiong, R. W. Folk, G. J. Broussard, R. Liang, M. J. Jang, H. Zhong, D. Dombeck, M. von Zastrow, A. Nimmerjahn, V. Gradinaru, J. T. Williams, and L. Tian. Ultrafast neuronal imaging of dopamine dynamics with designed genetically encoded sensors. *Science*, 360(6396):eaat4422, June 2018.

[87] A. L. Pehrson, T. Jeyarajah, and C. Sanchez. Regional distribution of serotonergic receptors: a systems neuroscience perspective on the downstream effects of the multimodal-acting antidepressant vortioxetine on excitatory and inhibitory neurotransmission. *CNS spectrums*, 21(2):162–183, Apr. 2016.

[88] J. L. Pérez-López, R. Contreras-López, J. O. Ramírez-Jarquín, and F. Tecuapetla. Direct glutamatergic signaling from midbrain dopaminergic neurons onto pyramidal prefrontal cortex Neurons. *Frontiers in Neural Circuits*, 12: 70, 2018.

[89] B. U. Phillips and T. W. Robbins. The role of central serotonin in impulsivity, compulsivity, and decision-making: comparative studies in experimental animals and humans. *Handbook of the Behavioral Neurobiology of Serotonin*, 31:531–548, Jan. 2020.

[90] A. T. Popescu, M. R. Zhou, and M.-m. Poo. Phasic dopamine release in the medial prefrontal cortex enhances stimulus discrimination. *Proceedings of the National Academy of Sciences*, 113(22):E3169–E3176, May 2016.

[91] M. Quiquempoix, S. L. Fayad, K. Boutourlinsky, N. Leresche, R. C. Lambert, and T. Bessaih. Layer 2/3 pyramidal neurons control the gain of cortical output. *Cell reports*, 24(11):2799–2807.e4, Sept. 2018.

[92] G. Radnikow and D. Feldmeyer. Layer- and cell type-specific modulation of excitatory neuronal activity in the neocortex. *Frontiers in Neuroanatomy*, 12:1, 2018.

[93] Y. Ranjbar-Slamloo and Z. Fazlali. Dopamine and noradrenaline in the brain; overlapping or dissociate functions? *Frontiers in molecular neuroscience*, 12:334, 2019.

[94] P. Redgrave, N. Vautrelle, P. G. Overton, and J. Reynolds. Phasic dopamine signaling in action selection and reinforcement learning. *Handbook of Basal Ganglia Structure and Function, Second Edition*, 24:707–723, Jan. 2016.

[95] J. Ren, D. Friedmann, J. Xiong, C. D. Liu, B. R. Ferguson, T. Weerakkody, K. E. DeLoach, C. Ran, A. Pun, Y. Sun, B. Weissbourd, R. L. Neve, J. Huguenard, M. A. Horowitz, and L. Luo. Anatomically defined and functionally distinct dorsal raphe serotonin sub-systems. *Cell*, 175(2):472–487.e20, Oct. 2018.

[96] J. Ren, A. Isakova, D. Friedmann, J. Zeng, S. M. Grutzner, A. Pun, G. Q. Zhao, S. S. Kolluru, R. Wang, R. Lin, P. Li, A. Li, J. L. Raymond, Q. Luo, M. Luo, S. R. Quake, and L. Luo. Single-cell transcriptomes and whole-brain projections of serotonin neurons in the mouse dorsal and median raphe nuclei. *eLife*, 8, Oct. 2019.

[97] A. Rivera, A. Penafiel, M. Megias, L. F. Agnati, J. F. López-Téllez, B. Gago, A. Gutiérrez, A. de la Calle, and K. Fuxe. Cellular localization and distribution of dopamine D4 receptors in the rat cerebral cortex and their relationship with the cortical dopaminergic and noradrenergic nerve terminal networks. *Neuroscience*, 155(3):997–1010, Aug. 2008.

[98] A. C. Sales, K. J. Friston, M. W. Jones, A. E. Pickering, and R. J. Moran. Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: An Active Inference model. *PLoS computational biology*, 15(1):e1006267, Jan. 2019.

[99] N. Santana and F. Artigas. Laminar and cellular distribution of monoamine receptors in rat medial prefrontal cortex. *Frontiers in Neuroanatomy*, 11:160, Sept. 2017.

[100] S. J. Sara. The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience*, 10(3):211–223, Mar. 2009.

[101] D. Sargin, H.-S. Jeoung, N. M. Goodfellow, and E. K. Lambe. Serotonin regulation of the prefrontal cortex: cognitive relevance and the impact of developmental perturbation. *ACS chemical neuroscience*, 10(7):3078–3093, July 2019.

[102] M. Sarter and K. B. Phillips. The neuroscience of cognitive-motivational styles: Sign- and goal-trackers as animal models. *Behavioral neuroscience*, 132(1):1–12, Feb. 2018.

[103] A. Saunders, A. J. Granger, and B. L. Sabatini. Corelease of acetylcholine and GABA from cholinergic forebrain neurons. *eLife*, 4:3859, Feb. 2015.

[104] C. Seo, A. Guru, M. Jin, B. Ito, B. J. Sleezer, Y.-Y. Ho, E. Wang, C. Boada, N. A. Krupa, D. S. Kullakanda, C. X. Shen, and M. R. Warden. Intense threat switches

dorsal raphe serotonin neurons to a paradoxical operational mode. *Science*, 363 (6426):538–542, Feb. 2019.

[105] M. Soiza-Reilly and P. Gaspar. From B1 to B9: a guide through hindbrain serotonin neurons with additional views from multidimensional characterization. *Handbook of the Behavioral Neurobiology of Serotonin*, 31:23–40, Jan. 2020.

[106] C. K. Starkweather, S. J. Gershman, and N. Uchida. The medial prefrontal cortex shapes dopamine reward prediction errors under state uncertainty. *Neuron*, 98 (3):616–629.e6, May 2018.

[107] E. K. Stephens, A. L. Baker, and A. T. Gulledge. Mechanisms underlying serotonergic excitation of callosal projection neurons in the mouse medial prefrontal cortex. *Frontiers in Neural Circuits*, 12:2, 2018.

[108] M. Subramaniam and J. Roeper. Subtypes of midbrain dopamine neurons. *Handbook of Basal Ganglia Structure and Function, Second Edition*, 24:317–334, Jan. 2016.

[109] Sul, Jung Hoon, Jo, Suhyun, Lee, Daeyeol, and Jung, Min Whan. Role of rodent secondary motor cortex in value-based action selection. *Nature neuroscience*, 14 (9):1202–1208, Sept. 2011.

[110] F. Sun, J. Zeng, M. Jing, J. Zhou, J. Feng, S. F. Owen, Y. Luo, F. Li, H. Wang, T. Yamaguchi, Z. Yong, Y. Gao, W. Peng, L. Wang, S. Zhang, J. Du, D. Lin, M. Xu, A. C. Kreitzer, G. Cui, and Y. Li. A gnetically encoded fluorescent sensor enables rapid and specific detection of dopamine in flies, fish, and mice. *Cell*, 174 (2):481–496.e19, July 2018.

[111] O. K. Swanson, R. Semaan, and A. Maffei. Reduced dopamine signaling impacts pyramidal neuron excitability in mouse motor cortex. *bioRxiv*, 6: 2020.02.12.946301, Feb. 2020.

[112] S. Tanaka, K. Shishida, N. Schweighofer, Y. Okamoto, S. Yamawaki, and K. Doya. Serotonin affects association of aversive outcomes to past actions. *Journal of Neuroscience*, 2009.

[113] L. M. Teles-Grilo Ruivo, K. L. Baker, M. W. Conway, P. J. Kinsley, G. Gilmour, K. G. Phillips, J. T. R. Isaac, J. P. Lowry, and J. R. Mellor. Coordinated acetylcholine release in prefrontal cortex and hippocampus is associated with arousal and reward on distinct timescales. *Cell Reports*, 18(4):905–917, Jan. 2017.

[114] A. Thiele and M. A. Bellgrove. Neuromodulation of attention. *Neuron*, 97(4): 769–785, Feb. 2018.

[115] M. K. Tian, E. F. Schmidt, and E. K. Lambe. Serotonergic suppression of mouse prefrontal circuits implicated in task attention. *eNeuro*, 3(5):ENEURO.0269–16.2016, Sept. 2016.

[116] W. Timberlake and G. Wahl. Stimulus and response contingencies in the misbehavior of rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 1982.

[117] R. Tremblay, S. Lee, and B. Rudy. GABAergic interneurons in the neocortex: from cellular properties to circuits. *Neuron*, 91(2):260–292, July 2016.

[118] A. Uematsu, B. Z. Tan, E. A. Ycu, J. S. Cuevas, J. Koivumaa, F. Junyent, E. J. Kremer, I. B. Witten, K. Deisseroth, and J. P. Johansen. Modular organization of the brainstem noradrenaline system coordinates opposing learning states. *Nature neuroscience*, 20(11):1602–1611, Nov. 2017.

[119] C. M. Vander Weele, C. A. Siciliano, G. A. Matthews, P. Namburi, E. M. Izadmehr, I. C. Espinel, E. H. Nieh, E. H. S. Schut, N. Padilla-Coreano, A. Burgos-Robles, C.-J. Chang, E. Y. Kimchi, A. Beyeler, R. Wichmann, C. P. Wildes, and K. M. Tye. Dopamine enhances signal-to-noise ratio in cortical-brainstem encoding of aversive stimuli. *Nature*, 563(7731):397–401, Nov. 2018.

[120] C. Varazzani, A. San-Galli, S. Gilardeau, and S. Bouret. Noradrenaline and dopamine neurons in the reward/effort trade-off: a direct electrophysiological comparison in behaving monkeys. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 35(20):7866–7877, May 2015.

[121] E. M. Vazey, D. E. Moorman, and G. Aston-Jones. Phasic locus coeruleus activity regulates cortical encoding of salience information. *Proceedings of the National Academy of Sciences of the United States of America*, 115(40):E9439–E9448, Oct. 2018.

[122] J. Verharen, J. W. de Jong, T. J. M. Roelofs, C. F. M. Huffels, R. van Zessen, M. C. M. Luijendijk, R. Hamelink, I. Willuhn, H. E. M. den Ouden, G. van der Plasse, R. A. H. Adan, and L. J. M. J. Vanderschuren. A neuronal mechanism underlying decision-making deficits during hyperdopaminergic states. *Nature communications*, 9(1):731–15, Feb. 2018.

[123] J. Wan, W. Peng, X. Li, T. Qian, K. Song, J. Zeng, F. Deng, S. Hao, J. Feng, P. Zhang, Y. Zhang, J. Zou, S. Pan, M. Shin, B. J. Venton, J. J. Zhu, M. Jing, M. Xu, and Y. Li. A genetically encoded sensor for measuring serotonin dynamics. *Nature neuroscience*, 24(5):746–752, May 2021.

[124] X. Wang, P. Zhong, and Z. Yan. Dopamine D4 receptors modulate GABAergic signaling in pyramidal neurons of prefrontal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(21):9185–9193, Nov. 2002.

[125] M. Watabe-Uchida, N. Eshel, and N. Uchida. Neural circuitry of reward prediction error. *doi.org*, 40(1):373–394, Aug. 2017.

[126] C. M. V. Weele, C. A. Siciliano, and K. M. Tye. Dopamine tunes prefrontal outputs to orchestrate aversive processing. *Brain research*, 1713:16–31, June 2019.

[127] L. Xiang, A. Harel, H. Gao, A. E. Pickering, S. J. Sara, and S. I. Wiener. Behavioral correlates of activity of optogenetically identified locus coeruleus noradrenergic neurons in rats performing T-maze tasks. *Scientific reports*, 9(1):1361–13, Feb. 2019.

[128] C. R. Yang and J. K. Seamans. Dopamine D1 receptor actions in layers V-VI rat prefrontal cortex neurons in vitro: modulation of dendritic-somatic signal integration. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 16(5):1922–1935, Mar. 1996.

[129] A. J. Yu and P. Dayan. Uncertainty, neuromodulation, and attention. *Neuron*, 46(4):681–692, May 2005.

[130] P. Zhong and Z. Yan. Distinct physiological effects of dopamine D4 receptors on prefrontal cortical pyramidal neurons and fast-spiking interneurons. *Cerebral cortex (New York, N.Y. : 1991)*, 26(1):180–191, Jan. 2016.

[131] W. Zhong, Y. Li, Q. Feng, and M. Luo. Learning and stress shape the reward response patterns of serotonin neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 37(37):8863–8875, Sept. 2017.

[132] B. Zingg, H. Hintiryan, L. Gou, M. Y. Song, M. Bay, M. S. Bienkowski, N. N. Foster, S. Yamashita, I. Bowman, A. W. Toga, and H.-W. Dong. Neural networks of the mouse neocortex. *Cell*, 156(5):1096–1111, Feb. 2014.

Appendices

Mouse	Pos Go		Pos NoGo		Neg Go		Neg No Go	
	Odds Ratio β1 Slope up / down	Likelihood ratio p value	Odds Ratio β1 Slope up / down	Likelihood ratio p value	Odds Ratio β1 Slope up / down	Likelihood ratio p value	Odds Ratio β1 Slope up / down	Likelihood ratio p value
R3	1.004	0.1204	1.019	0.9968	1.009	1.0000	1.034	1.0000
B3	0.922	0.1204	0.922	0.4304	1.068	0.3736	0.9051	**0.0046 ****
G5	1.083	**0.0028 ****	0.999	1.0000	1.063	**0.0028 ****	0.9226	**0.0066 ****
R1	0.933	**0.0028 ****	1.021	1.0000	0.9942	**0.0028 ****	1.142	**0.0066 ****
B2	0.956	0.2498	0.988	1.0000	1	0.8695	1.001	**0.0066 ****
G3	0.996	**0.0076 ****	1.083	0.9635	0.9466	0.9999	0.9957	**0.0028 ****
L1	1.490	0.9999	0.994	0.9635	0.9028	0.9986	0.6805	0.6298
Mean change per session (%)	5.486		0.371		-0.234		-4.559	
Median change per session (%)	-0.4		-0.1		0		-0.43	

Table A.1: Summary table of data from Fig 3.1. Odds ratios exceeding 1 indicate performance enhancement over time are highlighted in green. Odds ratios below 1 indicate worsening performance are highlighted in red. Likelihood test p-values, which evaluate whether the rate of performance change significantly deviates from zero, are also coloured to depending on the direction of performance change. In the final row the mean change in performance is converted to percentage for clarity. To control for multiple comparisons and reduce the risk of Type I errors, the Holm-Šídák method was used to correct the p-values obtained from the likelihood ratio tests.

A.1 Preliminary experiments

Injection pilots (2020)

In 2020 I performed a variety of virus injection pilot experiments to confirm whether I could successfully transfect the novel GENI, and determine the optimal procedure.

1. **Injection site pilot**: 2020/05. 1 mouse, 3 locations with $GRAB_{NE}$ & jRGECO1a. Confirmed positioning of the cranial window and accessibility of injection sites.

2. **$GRAB_{NE}$ pilot:** 2020/07. 4 mice. 1 site per hemisphere.
   Right hemisphere: Different ratios of $GRAB_{NE}$ and jRGECO1a. The virus containing $GRAB_{NE}$ appears to repulse the jRGECO1a virus quite strongly, resulting in lower expression of jRGECO1a. A ratio of 2:1 jRGECO1a to $GRAB_{NE}$ was able to somewhat overcome this, although the balance of expression was still less consistent than with the other GENI.
   Left hemisphere: Functional and static indicators: either $GRAB_{NE}$ & GFP, or jRGECO1a & TurboRFP. Confirmed the functionality of $GRAB_{NE}$ and jRGECO1a, i.e. that they both produce dynamic signals. At this time we were only in the possession of $GRAB_{NE}$.

3. **$GRAB_{NE}$ & dLight pilot:** 2020/10. 6 mice. 1 site per hemisphere.
   Each mouse was injected with a mixture of one GENI and jRGECO1a producing viruses in each hemisphere, at varying ratios to determine optimal expression.
   First use of the newly built behavioural setup. Mice were habituated to being head-fixed and exposed to uncued appetitive (sucrose solution) and aversive (air-puff) stimuli.
   Pharmacology pilot:

   - SCH23390 D1 antagonist to inactivate dLight (1mg/kg).
   - Yohimbine $\alpha 2$ antagonist to inactivate $GRAB_{NE}$ (0.1mg/kg).

Behaviour pilot experiment

Prior to the results presented in this thesis, I performed a earlier version of the experiment on 7 mice. However, upon analysing the data the mice did not appear to learn the task; Mice performance did not consistently improve through over the 24 days, with 30 minutes training per day. We therefore decided to repeat the experiment with a simplified schedule and with some improvements to the behavioural setup as described below and depicted in Figure A.1.

Task feedback tone

The feedback cue strategy may have been difficult to learn. In the 2021 version there were only 2 feedback tones, correct and incorrect. This may seem simpler, yet these tones indicated 3 possible outcomes. For example, correct performance in the 'positive reinforcement of go' state and 'negative reinforcement of go' state would both receive a 'correct performance' tone, however, these would be followed by sucrose delivery and air-puff omission, respectively. In the 2022 experiment I added one more tone allowing for direct correspondence between the feedback and the outcome, i.e. one tone each for sucrose, air-puff and neutral. Thereby, correct performance in the 'positive reinforcement of go' state would receive an 'Appetitive sucrose' tone and 'negative reinforcement of go' state would receive a 'Neutral' tone. This strategy does not require the mice to fully understand the paradigm. Moreover, if the mice do not learn the instrumental requirements, the tone-outcome pairing can be analysed as a Pavlovian paradigm, as was the case here.

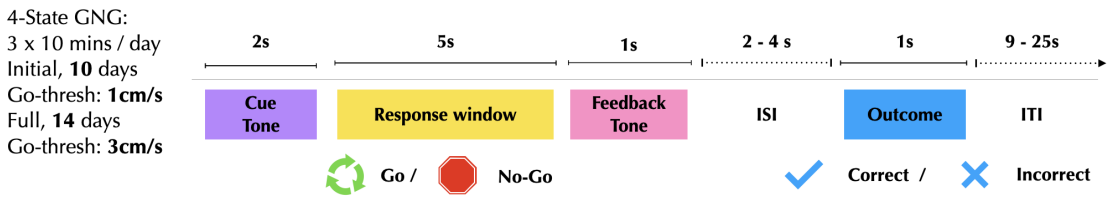

Figure A.1: 2021 schedule: 2 feedback cues, 2 Go-action thresholds

Task action requirements

In the 2021 version the task was divided into easier and harder parts with Go-action thresholds of 1cm/s and then 3cm/s, respectively. The lower threshold was selected to increase the chances of sampling the correct Go-action in the initial stages of learning. However, it became clear during the experiment that when mice decided to move they generally started running faster than 5cm/s, and so the easy and hard separation was inconsequential. In the 2022 version the two parts were revised to Pavlovian and Instrumental stages. In the Pavlovian paradigm mice were only exposed to the 3 feedback tones and corresponding stimuli. In the Instrumental paradigm the mice were also exposed to the 4 state cues and action window as well.

Lickometer

In 2021 I used an optical lickometer to track licking behaviour online. However, the optical lickometer was difficult to position and often recorded false positives, such as nose twitches and grooming, and false negatives, when the lickometer was positioned too far or at an angle. In the 2022 experiment I chose to record licking behaviour post-hoc using facially-directed video footage and pose-tracking software SLEAP.