

Modelling the mechanisms underlying variable spatiotemporal cortical response dynamics

Rishikesan Maran

Faculty of Science, The University of Sydney

A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy

2026

Abstract

Elucidating how the brain’s diverse repertoire of neural dynamics emerges from its fixed anatomical structure, known as the ‘structure–function relationship’, remains a challenge in macroscale neuroscience. While progress has been made in predicting resting-state (spontaneous) functional connectivity from structural connectivity, this static paradigm often fails to capture the mechanisms that shape dynamics over faster (sub-second) timescales, or how the spatiotemporal properties of brain dynamics can reconfigure over time under a fixed structural connectome. Accordingly, this thesis investigates the mechanisms that underlie spatiotemporal variability in stimulus-evoked cortical dynamics over fast (sub-second) timescales. Its main contribution is a timely investigation into the open question of why the geometry of brain anatomy can successfully capture key statistical properties of spontaneous brain dynamics, while ignoring the highly specific arrangements of the various long-distance inter-regional fibres that support global brain communication. Specifically, this thesis demonstrates, through newly constructed models, that long-distance connectivity is essential for capturing fast-timescale interactions between specific remote populations, but these interactions are obscured in slower order-of-seconds fluctuations of spontaneous activity. Furthermore, these fast dynamics are shown to be contingent on the spatial proximity of the driving stimulus, and critically gated by the simultaneous burst-like firing of subcortical arousal nuclei, which is often ignored in the modelling literature. Collectively, this thesis challenges the notion of a fixed structure–function relationship, showing that brain dynamics over fast timescales, despite evolving on a static anatomical backbone, can exhibit a diverse range of spatiotemporal patterns, mediated by variations in the spatial profile of driving stimuli, and fluctuations in internal arousal.

Statement of originality

This is to certify that the content of this thesis is my own work. This thesis has not been submitted for any other degree or purpose. I certify that the intellectual content of this thesis is the product of my own work, and that all assistance received in preparing this thesis and all sources have been acknowledged.

Rishikesan Maran

Candidate

May 23, 2026

Included papers and attribution

Listed below are details of the papers on which Chapters 2, 3, and 4 are based.

Chapter 2	Analyzing the brain’s dynamic response to targeted stimulation using generative modeling. Maran, R, Müller, EJ, Fulcher, BD Published in <i>Network Neuroscience</i> 9: 237–258, 2025. I contributed to conceptualising the review, created visual schematics, and wrote the drafts of the manuscript.
Chapter 3	Modelling the influences of non-local connectomic projections on geometrically constrained cortical dynamics Maran, R, Müller, EJ, Fulcher, BD In Submission. I constructed the model, developed corresponding numerical algorithms for simulation, conducted the investigations, analysed the data, visualised the results, and wrote the drafts of the manuscript.
Chapter 4	Dynamic routing on a static structure: noradrenergic gating of cortical wave propagation Maran, R, Fulcher, BD, Müller, EJ In Preparation. I constructed the model, developed corresponding numerical algorithms for simulation, conducted the investigations, analysed the data, visualised the results, and wrote the drafts of the manuscript.

Rishikesan Maran

Candidate

May 23, 2026

Associate Professor Ben D Fulcher

Primary Supervisor

May 23, 2026

Artificial Intelligence Statement

During the preparation of the thesis, the author used ChatGPT, Claude and Google Gemini for the purposes of grammar and text enhancement. The author confirms that where text was modified by generative AI, the content was reviewed for possible errors, inaccuracies, and bias. The author takes full responsibility for the submitted thesis and ensures the work is their own and has used generative AI within the parameters of use.

Australian Government Support

This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

Acknowledgements

I'd like to thank my supervisors Ben Fulcher and Eli Müller for taking me on as a student, and for their continual support and guidance throughout the past few years; the Australian Government Research Training Program (RTP) for financial support; and all fellow colleagues who have helped me along this journey. Despite this immense support and guidance, this thesis's errors and shortcomings are my own.

Contents

1	Introduction	8
1.1	The structure–function relationship	8
1.2	Key themes of this thesis	9
1.3	Organisation of this thesis	13
2	Analysing the brain’s dynamic response to targeted stimulation using generative modelling	15
2.1	Introduction	16
2.2	A dynamical systems framing of targeted brain stimulation	18
2.3	Targeted stimulation experiments	20
2.4	Generative models of brain activity	24
2.5	Conclusion	31
3	Modelling the influences of non-local connectomic projections on geometrically constrained cortical dynamics	34
3.1	Introduction	35
3.2	Materials and methods	37
3.3	Results	41
3.4	Discussion	55
4	Dynamic routing on a static structure: noradrenergic gating of cortical wave propagation	60
4.1	Introduction	61
4.2	Materials and methods	64
4.3	Results	68
4.4	Discussion	78
5	Concluding summary and future directions	82
5.1	Summary of findings	82
5.2	Implications for future work	83
	References	86
A	Supporting Information to Chapter 3	108
B	Supporting Information to Chapter 4	119

1. Introduction

1.1 The structure–function relationship

The brain is a complex organ that is responsible for receiving, processing and transmitting signals to other organs. In humans, it consists of billions of neurons (nerve cells), which communicate via electrical and chemical processes. While the physiological mechanisms underpinning neuron-to-neuron dynamical interactions are relatively well understood [1], it is much less clear how the brain’s relatively fixed structure of many interacting neurons, connecting axonal fibres, and neurotransmitter receptors, can together give rise to the various dynamical patterns that underlie its functions [2]. Elucidating the correspondence between the brain’s fixed anatomical structure and its spatially distributed macroscale dynamics, commonly referred to as the ‘structure–function relationship’, remains a fundamental goal in neuroscience, and is a core overarching motivation behind the work presented in this thesis.

To date, progress in understanding the structure–function relationship in mammalian brains has benefited from technological advances in structural and functional brain imaging. Structural imaging techniques such as tractography (with diffusion weighted imaging) and viral tract tracing allow us to map the brain’s structural connectivity, which we define as the arrangement of fibres between pairs of neurons, or as the density of fibres between pairs of brain regions [3]. In parallel, whole-brain functional imaging techniques such as functional magnetic resonance imaging (fMRI) or magnetoencephalography (MEG) can measure brain activity across regions and over time, which can be analysed to examine dynamic coordination between regions [4]. Access to structural and functional imaging datasets have driven substantial efforts to investigate whether the statistical properties of brain dynamics can be reproduced from simulations of generative or dynamical models incorporating structural connectivity data. Notably, the widely investigated statistic of resting-state functional connectivity can successfully be predicted by structural connectivity-constrained models, achieving accuracies or $r \approx 0.4$ in human brains [5–9], $r \approx 0.5$ in macaque [10, 11], and $r \approx 0.5$ in mouse [12, 13], where r is the Pearson correlation between empirical and simulated functional connectivity values. These findings demonstrate that the spatiotemporal patterns of spontaneous macroscale brain dynamics are to some extent constrained by the static arrangement of inter-regional structural fibres across subjects and species.

While structural and functional connectivity have been important measurements in providing evidence for the structure–function relationship, they still remain incomplete

representations of the brain’s structure and dynamics. Maps of structural connectivity have traditionally recorded the densities of fibres between brain regions in the cortex (outer layer of the brain) (Fig. 1.1A(i)), while neglecting the arrangements of fibres which connect subcortical structures such as the thalamus and brainstem. This tendency to reduce brain structure to just the corticocortical fibre system may also be a contributing factor to why the correlation between structural and functional connectivity appears to cap at only $r \approx 0.5$ [14]. Likewise, functional connectivity summarises the spatiotemporal patterns of unperturbed spontaneous fluctuations (Fig. 1.1A(ii)), and is hence limited in its ability to reveal the dynamics exhibited when the brain is perturbed by a stimulus. Motivated by the above limitations, this thesis explores how models can be used to illuminate the structure–function relationship beyond structural and functional connectivity, by incorporating a wider range of the brain’s structural attributes, and making predictions of a wider range of dynamical phenomena. Below, a case is made for the three particular directions that are explored in this thesis.

1.2 Key themes of this thesis

1.2.1 From spontaneous dynamics to stimulus-evoked dynamics

In the macroscale brain modelling literature, dynamics has typically been studied in a spontaneous setting, where fluctuations of neural activity are observed without engagement with an explicit stimulus imposed by the experiment (Fig. 1.1A(ii)) [15]. To simulate spontaneous fluctuations, models frequently incorporate a stochastic process as a driving input, to induce noisy perturbations in the model’s dynamics about its steady state (dynamics in the absence of input) [16, 17]. This stochastic input reflects the stimuli processed by the brain that are uncontrolled by the experimenter, including external sensory fluctuations and internal signals from other organs [1]. The spatiotemporal properties of the simulated fluctuations are then summarised into a statistic, such as the prevalently studied functional connectivity, or dynamic functional connectivity [18].

A major appeal of studying the brain’s structure–function relationship in a spontaneous setting is practical, in that activity can be measured without simultaneous experimental manipulation. Despite this practical benefit, it is known that the spatiotemporal patterns of brain dynamics can be altered under motor and cognitive tasks [19–21], different pharmacological interventions [22–24], or under sustained electrical neuromodulation [25, 26]. Because the spontaneous paradigm does not explicitly characterise driving stimuli, it provides limited insight into how the spatiotemporal patterns of brain dynamics are influenced by such external factors. This limitation leaves an opportunity for future studies to test structure-constrained models on such non-spontaneous forms of dynamics that the




Figure 1.1: **The relationship between the brain's anatomical structure and its spatially distributed macroscale dynamics, also known as the 'structure–function relationship', can be examined more deeply beyond the prevalently studied relationship between structural and functional connectivity.** **A:** Structural connectivity encodes the density of interregional axonal fibres, mostly in the cortex (**i.**), and functional connectivity is the pairwise correlation structure of the brain's spontaneous fluctuations, which is activity observed in the absence of an explicit stimulus (**ii.**). **B:** Instead of spontaneous dynamics, structure-constrained models can aim to predict the spatiotemporal properties of stimulus-evoked brain dynamics (**i.**), which is the brain's dynamic response to an impulse stimulus. If brain dynamics is conceptualised as a trajectory in a state space of relevant neural variables (**ii.**), then stimulus-evoked dynamics (red trajectory) will operate in states far from the attractor on which spontaneous dynamics (blue trajectory) will operate. **C:** Many topological properties of structural connectivity can be captured with a spatially homogeneous, isotropic geometric connection principle, known as the exponential distance rule. Founded on this principle, simpler geometric models have shown promise in successfully capturing the spatiotemporal properties of brain dynamics just from the geometry of brain anatomy (**i.**), without incorporating individual fibre densities. These successes motivate the need to investigate the dynamical contribution of specific long-distance fibres between remote populations (**ii.**), whose arrangements are unaccounted for by geometry. **D:** Brain dynamics is mediated not only by structural fibres between regions, but also modulated by arousal systems in the brainstem and forebrain (**i.**). These systems project widely to numerous regions (**ii.**), and alter the gain of efferent target regions. Modulating gain allows spatiotemporal dynamics to reconfigure over time under a fixed structural connectome.

brain can exhibit.

One particular class of non-spontaneous dynamics, which forms the focus of this thesis, is that of stimulus-evoked activity, which is the brain’s response to a transient, spatially localised stimulus (Fig. 1.1B(i)). This impulse stimulus can be sensory, such as a visual flash, or it can be artificially generated through targeted stimulation techniques such as optogenetics [27] or transcranial magnetic stimulation [28]. One benefit of studying the structure–function relationship via the stimulus-evoked paradigm, over the spontaneous paradigm, is that it presents an opportunity to reveal novel structural mechanisms that are difficult to detect in spontaneous settings. As visualised in Fig. 1.1B(ii), if brain dynamics is conceptualised as a trajectory in a state space of relevant neural variables, then the trajectories of spontaneous dynamics will be concentrated near an attractor of the state space, whereas the trajectories of stimulus-evoked dynamics will operate far from the attractor, revealing non-linear behaviours over faster timescales. Testing models against stimulus-evoked dynamics may drive the discovery of the subtle mechanisms underpinning these rapid non-linearities, that are less apparent when observing spontaneous dynamics. Furthermore, since the experimenter controls where, when, and how the stimulus is applied, the spatiotemporal profile of the model’s driving input can be designed to correctly reflect the experimental setup. With control over this spatiotemporal profile, the stimulus-evoked paradigm provides a framework for models to make predictions of a wider range of spatiotemporal patterns of brain dynamics than spontaneous settings can offer.

1.2.2 The brain as a geometric physical object

For a mammalian brain with millions or billions of neurons, structural imaging techniques are not yet capable of recording every axonally connected pair of neurons at whole-brain coverage. For this reason, it is standard practice for structural connectivity maps to discretise brain tissue into a disjoint set of spatially contiguous regions, and represent connectivity as a mathematical graph—called the connectome—in which the edges encode inter-regional fibre densities (Fig. 1.1A(i)) [29–32]. Consequently, the vast majority of models constrained by structural connectivity spatially discretise the brain in the same manner, and thus depict distributed brain activity as a discrete network of dynamic nodes—each representing a distinct region—that are coupled in accordance to the weights of the connectome’s edges [17].

One limitation of the aforementioned topological approach to building connectomic models, is that it does not account for how the connectome’s edge topology is shaped by geometry or spatial embedding, an intrinsically simpler structural attribute of brain anatomy. It is well established that corticocortical structural connectivity tends to follow a spatially homogeneous, isotropic geometric connection principle—also known as the ‘exponential distance rule’ [33–36]. Founded on this physical principle, a simpler but lesser appreciated class of models, known as neural field models, have shown promise in capturing

many dynamical properties of macroscale cortical dynamics from just the geometry of brain anatomy, while ignoring any connectome-specific information (Fig. 1.1C(i)) [37–39]. A recent comparative study has in fact shown neural field models to make predictions of functional connectivity and task activation maps of comparable accuracies to more detailed connectomic models [37]. This finding motivates the need for future models to address the dynamical contribution of the more topologically complex patterns of inter-regional connectivity that exists beyond geometry.

A particular aspect of connectivity unaccounted for by geometry, are the various long-distance fibres which connect specific remote populations (Fig. 1.1C(ii)) [40]. Since the arrangements of long-distance fibres are heterogeneous and anisotropic in space [36, 41–44], they are not captured by the geometric connectivity assumed by neural field models. Understanding the dynamical role of long-distance fibres is important because they are widely believed to be a core structural mechanism supporting the global integration of information in the brain [2]. To study their specific dynamical contribution, this thesis will develop a model which allows long-distance fibres to be added on top of geometric connectivity, by enhancing the connectivity strengths and conduction speeds of fibres at specific locations. While geometric connectivity facilitates travelling wave dynamics, the long-distance fibres will facilitate rapid propagation between selected remote populations, thereby inducing perturbations in the wave dynamics. With this construction, the model provides a controlled way to isolate the dynamical role of specifically arranged long-distance fibres separately from the influence of geometry, allowing us to determine when and how such fibres meaningfully reshape macroscale brain dynamics.

1.2.3 Neuromodulatory control of flexible dynamics

A core motivation behind scientific interest in the structure–function relationship is to understand the neural correlates of the brain’s functions, such as memory, perception, language, cognition, and consciousness. One remarkable feature of the brain’s complex array of behaviours, which distinguishes it from that of many non-biological communication systems, is that they are predominantly learned through experience with the external environment, rather than hardwired from birth [45]. Furthermore, the brain is able to integrate new information with existing knowledge, and flexibly switch its learnt functional configurations between different tasks [46]. Mounting empirical evidence suggests that this unique adaptive flexibility of brain function is supported by the various neuromodulatory systems of the subcortical brainstem and forebrain (Fig. 1.1D(i)) [47, 48]. These systems send axonal projections diffusely across the brain, including the cortex (Fig. 1.1D(ii)), and alter the excitability or gain of neurons in efferent target populations through the release of neuromodulators including dopamine, serotonin, noradrenaline, and acetylcholine [49]. Through this mechanism, the brain can rapidly reconfigure its spatiotemporal dynamics on demand, without physically modifying its static structural backbone.

To date, a substantial amount of research has gone into understanding the circuit mechanisms that govern the firing behaviours of subcortical arousal nuclei (in the brainstem and forebrain) [50]. However, modelling the modulatory effects of arousal systems on the spatiotemporal patterns of macroscale neural dynamics have only more recently been attempted. Several modelling studies have attempted to investigate this relationship by coupling existing neural dynamics with gain dynamics [51–53], but such formulations assume a constant afferent input from arousal nuclei over time, thus providing limited insight into how arousal systems can actively reconfigure activity across the brain in response to changing environmental demands or task requirements. Accordingly, future models can provide a more biophysically realistic and behaviourally relevant account of modulatory influences of arousal systems on spatiotemporal neural dynamics, by including the temporal dynamics of arousal nuclei in response to external stimuli.

One behaviourally relevant temporal pattern of arousal nuclei firing, is the tendency of certain arousal nuclei to respond strongly to unexpected changes in the external environment [50]. In particular, the locus coeruleus of the brainstem, which constitutes the primary source of the noradrenergic arousal system (the system that releases noradrenaline), tends to display phasic bursts of activity in response to task-relevant stimuli [54]. To investigate the effect of phasic noradrenergic modulation on neural dynamics, this thesis will develop a novel computational model that explicitly tracks the dynamics of the locus coeruleus, and couples its dynamics to the gain dynamics of efferent target populations. With this mechanism, the model provides a more explicit structural link between locus coeruleus firing and macroscale neural dynamics, and provides a more realistic framework to investigate how the locus coeruleus controls the brain’s functional configuration in accordance to external stimuli.

1.3 Organisation of this thesis

The work presented in this thesis consists of three chapters, each of which builds upon previous ones to address the key themes mentioned above:

- Chapter 2 presents an argument, based in dynamical systems theory, on how new structural mechanisms underlying brain dynamics over fast sub-second timescales can be uncovered by modelling stimulus-evoked dynamics. This chapter also serves as the literature review of this thesis, detailing the various stimulation and measurement techniques that can be utilised in empirical settings to generate datasets of stimulus-evoked dynamics, and providing a novel organisation of different models of brain activity found in the literature and the specific types of mechanisms they encode.
- Chapter 3 develops a new model of macroscale cortical dynamics, that can simulta-

neously capture travelling wave dynamics shaped by cortical geometry, and rapid communication dynamics facilitated by specifically arranged long-distance connectivity. This model is used to investigate how adding long-distance fibres to geometric connectivity perturbs the travelling wave dynamics, and the contexts in which these perturbations are more or less meaningful. Importantly, the findings in this chapter support the arguments presented in Chapter 2, as simulations of stimulus-evoked dynamics (instead of spontaneous dynamics) were used to uncover the fast-timescale influences of long-distance fibres on cortical dynamics. To keep the presentation of this chapter focused on the key findings, the model equations are described and summarised, with details of their derivation (and associated numerical algorithm for solving the equations) in Appendix A.

- Chapter 4 develops a novel model of corticothalamic dynamics, focused on isolating the neuromodulatory influence of the locus coeruleus on spatiotemporal brain dynamics, from the influences of existing structural fibres between regions. This physiological model extends upon the phenomenological model developed in Chapter 3, introducing more biophysically realistic elements such as corticothalamic interactions, synaptodendritic dynamics, inhibitory mechanisms, more realistic long-distance connectivity strengths, and the aforementioned locus coeruleus. With this model, it is shown how the timing and topography of the locus coeruleus' engagement can actively steer the spatial propagation of stimulus-evoked wave dynamics. It also builds on the findings of Chapter 3, showing how the fast-timescale dynamics mediated by long-distance fibres are crucially contingent on sufficient neuromodulatory influence from the locus coeruleus. Similar to Chapter 3, the model equations and numerical methods are summarised in the chapter and detailed in Appendix B.

Each chapter in this thesis corresponds to a self-contained journal article, resulting in some repetition of introductory material. To articulate the links between the chapters in addressing the key goals of this thesis, each chapter begins with a 'chapter summary', which contextualises the work within the broader aims of the thesis, and establishes its connection to preceding chapters.

2. Analysing the brain’s dynamic response to targeted stimulation using generative modelling

CHAPTER SUMMARY

The aim of this chapter is to provide a theoretical motivation for studying the structure–function relationship through the lens of stimulus-evoked dynamics; i.e., the dynamic response to an impulse focal stimulus. It utilises concepts from dynamical systems theory to argue that stimulus-evoked dynamics may be shaped by distinct mechanisms that are otherwise obscured during spontaneous activity. The chapter then reviews the various emerging targeted stimulation techniques capable of generating datasets of stimulus-evoked dynamics, and the opportunities that integrating targeted stimulation experiments with different types of models can provide in expanding our understanding of the various mechanisms underlying brain dynamics. This framework establishes the conceptual basis for the modelling investigations conducted in the subsequent chapters.

Abstract

Generative models of brain activity have been instrumental in testing hypothesised mechanisms underlying brain dynamics against experimental datasets. Beyond capturing the key mechanisms underlying spontaneous brain dynamics, these models hold an exciting potential for understanding the mechanisms underlying the dynamics evoked by targeted brain-stimulation techniques. This paper delves into this emerging application, using concepts from dynamical systems theory to argue that the stimulus-evoked dynamics in such experiments may be shaped by new types of mechanisms distinct from those that dominate spontaneous dynamics. We review and discuss: (i) the targeted experimental techniques across spatial scales that can both perturb the brain to novel states and resolve its relaxation trajectory back to spontaneous dynamics; and (ii) how we can understand these dynamics in terms of mechanisms using physiological, phenomenological, and data-driven models. A tight integration of targeted stimulation experiments with generative quantitative modelling provides an important opportunity to uncover novel mechanisms of

brain dynamics that are difficult to detect in spontaneous settings.

2.1 Introduction

Statistical analyses of experimental neuroimaging data have long demonstrated that the brain supports a rich repertoire of spatiotemporal dynamics [55–58]. To better understand the intrinsic mechanisms which give rise to the brain’s dynamical properties, generative models of brain activity—also known as computational models, dynamic models, or mathematical models—can be used to test hypothesised mechanisms against properties of brain dynamics that they aim to explain or predict. While the term ‘generative model’ can have a range of interpretations across fields, here we use it as shorthand for ‘generative model of brain dynamics’, to describe a model that can generate (through simulation) a set of time-series of brain activity that can be compared to data measured experimentally [59]. Existing generative models allow us to simulate brain activity based on a wide variety of encoded mechanisms across spatial scales [16, 17, 59–62], allowing the hypothesised mechanisms to be directly evaluated against the dynamics of activity accessed in experimental neural recordings.

A crucial advantage that generative models provide is the ability to rigorously test candidate mechanisms against datasets from different experimental settings. Early generative models were used to test mechanisms underlying dynamics of brain activity recorded during cognitive or motor tasks, or in response to sensory stimuli [63–72]. Rich dynamical properties were later found even in *spontaneous* brain activity, i.e., activity not attributable to any externally applied task or stimulus [15]. Many of these properties have hypothesised mechanistic explanations that have been tested through generative models [5, 73–79]. As datasets from different experimental settings reveal novel spatiotemporal properties of brain dynamics, generative models present new opportunities to refine and hypothesise new underlying mechanistic accounts that can accurately capture them.

A new opportunity of growing interest to generative modelling, is to test our understanding of the dynamics of brain activity evoked by targeted stimulation. Unlike traditionally used sensory stimulation paradigms, targeted stimulation techniques directly target neurons or neural populations using concentrated external inputs of energy. This approach bypasses the constraints imposed by the brain’s natural sensory pathways, offering significantly greater spatial precision in target selection and flexibility in stimulation strength. Commonly used targeted stimulation techniques include optogenetics [27], electrode stimulation [80, 81], transcranial magnetic stimulation [28], and chemogenetics [82], all of which are applicable *in vivo*, and allow for simultaneous measurement of evoked responses with compatible measurement techniques [83–87]. This review focuses on the new opportunities that datasets of activity measured from targeted stimulation experiments

present in validating and refining our generative models of brain activity, and thereby better understanding the important mechanisms which underpin brain dynamics.

Why should one be interested in the mechanisms underlying the brain’s response to targeted stimulation? There are both theoretical and practical motivations. Theoretically, despite the high-dimensionality of brain activity recordings, brain dynamics of spontaneous activity is known to be constrained to a lower-dimensional manifold [19, 20, 88–91]. This implies the existence of a vast space of unexplored states that the brain does not typically access, unless perturbed with a sufficiently strong stimulus. Targeted stimulation techniques enable us to access states that are not only inaccessible during spontaneous activity, but also difficult to access during sensory-evoked activity, hence initiating new forms of artificially evoked activity that transcend the brain’s natural activation patterns. Consequently, generative models present an important opportunity to test and calibrate mechanisms on datasets obtained through a unique experimental paradigm, potentially driving the discovery of new mechanisms of brain dynamics distinct from those that sufficiently capture spontaneous dynamics. Practically, there is growing interest in developing clinical neuromodulation technologies using targeted stimulation to address a wide range of neurodegenerative and neurological disorders [92]. While trial-and-error methods for determining effective stimulation parameters have shown promising results for specific patients [93, 94], a wider deployment of neuromodulation necessitates a systematic approach that can more efficiently explore the vast stimulation parameter space and tailor the technology to each patient’s needs [95, 96]. An enhanced mechanistic understanding of the brain’s response to the targeted stimulation techniques behind these technologies can therefore accelerate the development of more systematic personalised protocols, by guiding the interpretation of individual responses to different stimulation parameter combinations.

In this review, we focus on the theoretically motivated opportunities in understanding the brain’s response to targeted stimulation (for a review on the practically motivated opportunities mentioned above, see [62, 92, 95, 97] for examples). With reference to a theoretical abstracted picture of brain dynamics, we argue for the existence of hidden nonlinear mechanisms of brain dynamics which are challenging to detect from spontaneous dynamics. We highlight three available tools that we argue are essential to uncover these mechanisms: (1) targeted stimulation techniques to make precise perturbations that can transition brain activity away from spontaneous patterns; (2) measurement techniques to simultaneously record the brain’s response to the perturbation as it reverts to spontaneous patterns; and (3) generative models to test mechanisms that we hypothesise to capture the dynamics of recorded activity evoked by targeted stimulation. Anchored around this argument, we explore in detail the opportunities that integrating these tools—that is, integrating different targeted stimulation and measurement techniques with different physiological, phenomenological, and data-driven models—provide in expanding

our understanding of the various mechanisms underlying brain dynamics.

2.2 A dynamical systems framing of targeted brain stimulation

To concretely establish our theoretical motivation for understanding the brain’s response to targeted stimulation, we present a geometric depiction of brain dynamics which abstracts the complexities of brain activity measurements across different spatial scales and species, shown in Fig. 2.1. With reference to dynamical systems theory, we use this picture to argue that complex brain dynamics are likely underpinned by mechanisms that are not clearly apparent in the dynamics of commonly analysed spontaneous activity, and that targeted stimulation techniques and measurement techniques are necessary tools to generate brain dynamics that manifest these hidden mechanisms. We also argue that the dynamics evoked by targeted stimulation are likely to be nonlinear, and that to rigorously test and validate the hidden mechanisms underlying these dynamics necessitates the use of generative modelling and simulation.




Figure 2.1: **An abstracted picture of brain dynamics that conceptualises the dynamics measured from distinct experimental settings, while simplifying the complexities of brain activity measurements across other experimental conditions.** **A:** Brain activity dynamics can be represented as the evolution of a set of key state variables, such as: **i.** the membrane potentials of individual neurons in a network; or **ii.** the activation of distributed modes of whole-brain activity. **B:** In this abstracted picture, spontaneous brain dynamics can be viewed as stochastic fluctuations about an attractor (a manifold that attracts nearby trajectories). Common attractor classes include: **i.** fixed points [98, 99], and **ii.** limit cycles [19, 20]. **C:** In contrast to spontaneous dynamics, stimulus-evoked brain dynamics consists of: **i.** a perturbation to a point that is typically far from the attractor, followed by **ii.** a dynamic relaxation back to the attractor. Stimulus-evoked dynamics can exhibit interesting nonlinearities (**iii.**), where the response (red) qualitatively differs with the strength of the perturbation (blue).

In this abstracted picture, the brain’s activity is characterised as the evolution of a brain state over time. Brain states can be encapsulated by a set of time-varying

state variables, $\{x_1(t), x_2(t), x_3(t), \dots, x_N(t)\}$, and brain dynamics thus correspond to trajectories, $\mathbf{x}(t) = [x_1(t), x_2(t), \dots, x_N(t)]$, through the state space. A specific spatial scale of interest can be addressed through appropriate selection of the state variables, from the activities of individual neurons in a population [100] (depicted in Fig. 2.1A(i)), to the spatially distributed functional gradients or modes which describe patterns of population-scale activity over the whole brain [37, 101] (depicted in Fig. 2.1A(ii)).

By abstracting brain dynamics as trajectories through a state space, we can conceptualise the dynamics of brain activity measured from qualitatively distinct experimental settings. Consider the commonly examined setting of spontaneous activity, which is measured when the brain is at rest and absent from any explicit external stimulus. Measurements indicate that spontaneous brain states are concentrated at a subset of the state space with dimension much less than that of the state space, and do not deviate significantly from this subset [19, 20, 88–90]. Consistent with these properties, spontaneous dynamics in our abstracted picture can be conceptualised as trajectories concentrated near *attractors*, which are low-dimensional subsets of the state space that attract neighbouring trajectories [102]. Additional sensory inputs and background biological processes can also cause trajectories to temporarily deviate from and return to the attractor, effectively yielding stochastic fluctuations about the attractor. Attractors, which define spontaneous dynamics across different experimental conditions, can be further categorised into classes of similar geometric features; such as stable fixed points [98, 99] (depicted in Fig. 2.1B(i)), or flows along a more complex low-dimensional manifold such as a limit cycle [19, 20, 88] (depicted in Fig. 2.1B(ii)). Through classes of attractors, we can concisely capture the intrinsic dynamical properties of spontaneous brain activity in this abstracted picture, whilst aggregating the impact of other experimental conditions.

In addition to the low-dimensional attractor where trajectories of spontaneous dynamics are concentrated, are a myriad of possible trajectories of brain dynamics through states far from the attractor. Unfortunately, as brain states move further from the attractor, they become less likely to be observed in spontaneous settings. However, our abstracted picture can conceptualise an alternative *stimulus-evoked* setting that employs external intervention to reach these distant states, as shown in Fig. 2.1C. Initially positioned near the attractor (depicted here as a fixed point), the brain is first perturbed with a sufficiently strong stimulus to a distant state (Fig. 2.1C(i)), and subsequently relaxes back to the attractor following stimulus termination (Fig. 2.1C(ii)). According to nonlinear dynamical systems theory, a central characteristic of stimulus-evoked dynamics is the nonlinear nature of the relaxation trajectory [102], where its shape is expected to be highly sensitive to the magnitude of the perturbation, as illustrated in Fig. 2.1C(iii). However, while such nonlinearities can produce a diverse range of trajectories far from the attractor, their effects tend to vanish as the trajectory approaches the attractor governing spontaneous dynamics, implying that spontaneous dynamics can be sufficiently approximated by simpler linear

models. Generative modelling studies have indeed shown that linear models can capture properties of macroscale spontaneous brain dynamics as effectively as nonlinear models, such as the established resting state functional connectivity benchmark [8, 79, 103]. These findings, supported by predictions of nonlinear dynamical systems theory, suggest that the nonlinear mechanisms shaping stimulus-evoked brain dynamics in this dynamical systems picture are difficult to be detected in spontaneous settings. Thus, to uncover and understand these nuanced hidden mechanisms of stimulus-evoked dynamics in practice (beyond this abstracted picture), we require an experimental paradigm which allows us to perturb the brain away from its spontaneous attractor (Fig. 2.1C(i)) and resolve the nonlinear relaxation trajectory (Fig. 2.1C(ii)), alongside a model which allows us to question and validate the nonlinear mechanisms which underpin these trajectories.

To this end, we bring attention to a practically feasible strategy that integrates targeted stimulation techniques and measurement techniques with generative models. Targeted stimulation techniques can perturb the brain from the attractor to distant states, with spatial coverage and precision that surpass traditional sensory stimuli [27, 28, 82, 104]. High precision selection of states is crucial, as the nonlinear dynamics of stimulus-evoked activity would be highly sensitive to perturbation strengths. Compatible measurement techniques play an indispensable role, by simultaneously tracking the brain’s relaxation trajectory back to the attractor at sufficient spatial and temporal resolution. Finally, the simulation approach of generative models is essential for linking hypothesised mechanisms to the dynamics observed in measurements; as a mechanistic understanding of nonlinear properties is challenging to pursue from static statistical measures alone [105].

Targeted stimulation techniques, measurement techniques, and generative models are therefore three important tools which, when integrated, can uncover and provide understanding of new mechanisms of brain dynamics that are concealed in spontaneous settings. By conducting experiments with stimulation and measurement techniques, we can generate datasets of stimulus-evoked dynamics with potentially new forms of nonlinearities. By constraining generative models with datasets from these experiments, we can rigorously test and validate hypotheses about the novel mechanisms which drive their nonlinear dynamics. To explore the opportunities that emerge from this integration, we first delve into the different experiments that can be executed.

2.3 Targeted stimulation experiments

Through an abstracted picture of brain dynamics, illustrated in Fig. 2.1, we have argued why targeted stimulation experiments play an important role in revealing novel mechanisms of brain dynamics in measured brain activity. This picture serves another purpose: it describes the fundamental dynamical properties of activity in any targeted stimulation experiment,

while abstracting away details of the specific techniques used. This interpretation bridges a wide body of stimulation and measurement techniques developed from disjoint fields, each with their own mechanics of action, spatial scale, and species compatibility. Here, we review a diverse range of key stimulation and measurement techniques that can be used in a targeted stimulation experiment, contextualising their features within our unified picture.

In a targeted stimulation experiment, a stimulation technique perturbs the brain to a state distant from the attractor representing spontaneous dynamics (Fig. 2.1(i)), and a measurement technique captures the brain's subsequent relaxation trajectory back to the attractor (Fig. 2.1(ii)). Key targeted stimulation and measurement techniques are illustrated in Fig. 2.2. Based on the spatial scale in which brain states are defined, it is crucial to select an appropriate targeted stimulation technique to elicit a precise perturbation in this state space, and an appropriate measurement technique to accurately capture the relaxation.




Figure 2.2: **Stimulus-evoked brain activity can be generated through a diverse range of targeted stimulation techniques and measured by a range of measurement techniques at different spatial scales.** **A:** Targeted stimulation techniques can perturb the brain to novel brain states at different degrees of spatial precision. From neuron-scale techniques, such as optogenetics; to mesoscale techniques, such as ICMS and DBS; to macroscale techniques, such as chemogenetics and TMS. **B:** Different measurement techniques can measure the brain's response to targeted stimulation at different spatial resolutions. From neuron-resolution modalities, such as Neuropixels electrodes, patch clamps, and calcium imaging; to mesoscale modalities, such as iEEG systems; to macroscale modalities, such as EEG and fMRI.

2.3.1 Targeted stimulation techniques

Different targeted stimulation techniques can be selected to elicit perturbations at a spatial precision of choice. Microscale techniques can precisely stimulate individual neurons [27], while meso-to-macroscale techniques can stimulate neural populations ranging from millimetres to centimetres in size [28, 80, 81]. Figure 2.2A illustrates key targeted stimulation techniques, ordered by degrees of spatial precision. We review these techniques below.

At the microscale, techniques can stimulate a single target neuron or a set of neurons. The most prevalent microscale stimulation technique is optogenetics, in which target neurons are genetically modified with injected photoreceptor proteins called opsins, and later exposed to light of a specified wavelength to generate action potentials [27]. Two-photon microscopy allows the light beam to be made sufficiently narrow to be directed exclusively to the target neurons [106–108].

At the mesoscale, neural populations can be similarly stimulated using optogenetics, or via electrical methods such as intracortical microstimulation (ICMS) and deep brain stimulation (DBS). A compelling benefit of using optogenetics at this scale is its ability to preferentially target specific cell-types, cortical layers, and projection targets through viral-vector gene-delivery techniques, although such targeting is not always perfectly selective in practice [109]. The alternative methods, ICMS and DBS, are electrical techniques which stimulate local populations in the vicinity of an implanted electrode using electric fields [80, 81]. While these methods lack the forms of specificity that can be achieved with optogenetics, they can be applied in human brains as it does not require genetic modification of the target neural tissue. Local populations near the cortical surface can be targeted by ICMS through a chronically implanted multielectrode array [81], whereas populations of varying depth down to subcortical structures can also be targeted by DBS through electrodes implanted deep in the brain [80].

Finally, macroscale neural populations can be stimulated using electrical methods such as transcranial magnetic stimulation (TMS), or chemogenetics which is a genetic modification technique similar to optogenetics. TMS uses an external electromagnetic coil to generate a magnetic field, which then induces an electric field over the target population [28, 110]. TMS is a non-invasive technique, as the generated magnetic field passes across the scalp and skull, however the difficulty in downsizing the coil has mainly restricted its use to primate brains [111]. In smaller rodent brains, macroscale populations can be stimulated using chemogenetics, in which the target population is injected with an engineered protein called Designer Receptors Exclusively Activated by Designer Drugs (DREADDs), and consequently activated by an orally administered designer drug [82]. Similar to optogenetics, chemogenetics benefits from the forms of specificity provided by viral-vector techniques, but it is able to stimulate much larger neural populations, as the

designer drug diffuses throughout the brain when orally administered [112]. Chemogenetic stimulations also have the unique ability to modulate brain activity for sustained periods ranging from hours to weeks [113].

2.3.2 Measurement techniques

Following the perturbation elicited by the selected targeted stimulation technique, we can simultaneously measure the brain’s relaxation trajectory back to the attractor of spontaneous activity. Collections of such measurements form a dataset of stimulus-evoked activity, which we have argued to reveal undiscovered mechanisms of brain dynamics upon further analysis. To capture the relaxation, the selected measurement technique must measure the exhibited brain activity at a matching spatial resolution. Microscale measurement techniques can resolve the response from individual neurons; while mesoscale and macroscale measurement techniques capture aggregated activity of neural populations. Figure 2.2B illustrates key measurement techniques, ordered by degrees of spatial resolution. We review these techniques below.

At a microscale resolution, single-unit recording methods capture the spiking activity of individual neurons. The classic single-unit recording method is whole-cell patch-clamp electrophysiology, which reads the intracellular potential from an electrode of a micropipette tightly suctioned to the target neuron membrane [114]. The spatial coverage of patch clamps is however limited to a few neurons. An alternate emerging modality is calcium imaging, which measures changes in levels in calcium concentration as an indirect proxy for spiking activity [115]. Modern microscopy techniques can measure calcium levels of hundreds of neurons, recently measuring activity across all neurons in *C. elegans* in response to neuron scale optogenetics [116]. With emerging microelectrode arrays, such as the recently developed Neuropixels probes [117], it is now possible to directly measure neuronal spiking with a spatial coverage comparable to that of calcium imaging.

Activity of a mesoscale population is conventionally measured by the local field potential (LFP) of its extracellular medium, which aggregates the spiking of its constituent neurons [118]. LFPs can be recorded by intracranial electroencephalographic (iEEG) electrodes penetrating the cortex. Commonly used iEEG electrode designs include subdural strip or grid electrodes, which can monitor LFPs from multiple positions from the cortical surface, or depth electrodes, which can measure at multiple depths at a single surface position [119].

Activity of a macroscale neural population could be measured by simultaneously tracking mesoscale populations with multiple iEEG systems, however this option can only be implemented in rare surgical scenarios (for example, the F-TRACT project [120]), as many electrodes become excessively invasive. Instead, there exist non-invasive functional neuroimaging techniques which make proxy measurements at this spatial scale. Electroencephalography (EEG) can measure electrical potentials from electrodes attached

around the scalp at high temporal resolution, making it suitable for measuring transient responses to stimulation such as TMS [28]. Complementing EEG is functional magnetic resonance imaging (fMRI), which measures the blood oxygen level dependent (BOLD) signals tracking underlying neural activity. Despite its limited temporal resolution, fMRI with its high spatial coverage can be used to measure responses to a macroscale modulatory stimulation method such as chemogenetics [22, 23, 84].

In summary, there are a range of sophisticated targeted stimulation techniques that can perturb precise neurons or broad neural populations; and measurement techniques which can simultaneously resolve the subsequent relaxation of neuronal spiking activity or aggregated neural population activity. Thanks to these technological advancements, datasets of stimulus-evoked activity can be generated through targeted stimulation experiments at any spatial scale of our interest: micro-, meso-, or macroscale. As argued through an abstracted picture of brain dynamics (Fig. 2.1), stimulus-evoked brain activity may manifest highly nonlinear dynamics, driven by unexplored mechanisms. However, integrating these experiments with generative models allows us to rigorously test hypotheses about these nonlinear mechanisms at any spatial scale. In the next section, we detail how different generative models can be utilised to harness this significant opportunity.

2.4 Generative models of brain activity

Generative models provide a mechanistic account of the spatiotemporal dynamics of brain activity observed under any experimental setting. Mechanisms are encoded as a set of dynamical rules, such as difference equations or differential equations, which govern the evolution of brain activity over time. By simulating these governing rules, we can generate a spatially distributed time series of brain activity, which can be compared against activity from experimental data. Mechanisms of generative models can also be designed to generate nonlinear time series, allowing us to capture the nonlinear dynamical properties of stimulus-evoked activity. Aside from the ability to generate neuronal dynamics, the hypothesised mechanisms of generative models can also be assessed and validated by integrating the model with a target experimental dataset. This process may involve model selection and model calibration, where the model’s parameters and possible structural characteristics are optimised to explain the dataset as accurately as possible [59]. One can additionally assess the model’s predictions of the target dynamics through an uncertainty quantification phase, in which statistical methods quantify the reliability of model’s parameter estimates and predictions [121, 122], and sensitivity analyses evaluate how robust the output dynamics are to small changes in the inferred parameters [123, 124]. Through the generation of nonlinear dynamics, and the integration with targeted stimulation datasets, generative models are therefore indispensable tools in the endeavour to infer mechanisms of brain

dynamics revealed in the activity evoked by targeted stimulation.




Figure 2.3: **Generative models come in a variety of flavours, and can be positioned along a spectrum between three extremes, depending on the type of mechanistic account of brain dynamics they provide.** **A:** Physiological models—such as biological neural networks, neural mass models, and neural field models—model physiological processes of brain activity with physically measurable biological quantities [16, 17]. **B:** Phenomenological models—such as linear network models, coupled oscillator models, and conductance-based models—model abstracted mechanisms that can capture essential dynamical properties of brain activity with minimal physiological detail [60, 61]. **C:** Data-driven models—such as models constructed through system identification or artificial neural networks—learn statistical associations of brain dynamics from sample data to make accurate predictions of brain dynamics [59, 62].

Generative models encode various forms of mechanisms to provide different forms of accounts of brain dynamics. Figure 2.3 illustrates a spectrum of key generative models, organised by the type of mechanisms they encode. By choosing and employing an appropriate model, we can test and validate the novel mechanisms that we hypothesise to drive the dynamics evoked by targeted stimulation. *Physiological models*, for example, use physically measurable quantities to quantitatively encapsulate complex physiological processes in the brain [16, 17]. In contrast, *phenomenological models* use mathematical abstractions to capture the core spatiotemporal properties of brain dynamics hypothesised by the mechanism [60, 61]. Finally, we distinguish *data-driven models*, which learn statistical associations from sample data to make accurate predictions of brain dynamics [59,

62]. Despite their varying mechanisms, all generative models in Fig. 2.3 can simulate both spontaneous and stimulus-evoked activity as abstracted by the dynamical systems picture in Fig. 2.1. The relevant output variables construct a state space where brain activity follows a trajectory of brain states (Fig. 2.1A). Introducing stochastic noise as an input variable concentrates the trajectory around an attractor, reflecting the model’s depiction of spontaneous dynamics at the corresponding spatial scale (Fig. 2.1B). Contrastingly, setting the input variable to a deterministic spatiotemporal profile aligned with the experimental stimulus would perturb the brain state away from the attractor before drawing it back via the relaxation trajectory (Fig. 2.1C).

In the subsections below, we review how physiological, phenomenological, and data-driven models are constructed, and how the specific types of mechanisms they encode can contribute to our understanding of the dynamics evoked by targeted stimulation. For each model category, we provide examples of modelling studies which have tested mechanisms underlying brain dynamics through an integration with datasets of stimulus-evoked activity, with a primary focus on emerging datasets of targeted stimulation. We also highlight how findings of some studies support our core argument presented through the abstracted picture in Fig. 2.1: that stimulus-evoked dynamics could be driven by novel mechanisms of brain dynamics which play a more peripheral role in shaping spontaneous dynamics.

2.4.1 Physiological models

In physiological models, variables and parameters represent physical quantities such as cell-body potentials or synaptic coupling strengths, and interact in accordance to experimentally supported physiological processes [16, 17, 125]. Physically measurable variables and parameters assist in interpreting the functional significance of individual variables, and the effects of changes in static parameters on the resultant dynamics. Crucially, the physiological formulation allows predictions of the model to be tested not just against the model’s ability to capture the predicted brain dynamics, but also against physical measurements from physiological experiments. Physiological models can study the dynamics of responses to a wide range of targeted stimulation techniques, with input variables quantifying the technique’s specific mechanics of action. For example, a physiological model could be formulated to directly compare responses to optogenetics, which directly depolarises neuron membranes, with responses to electrode stimulation, which feeds a current into the extracellular medium. We can construct physiological models to test mechanisms of physiological processes at different spatial scales, ranging from the microscale to whole brain scale.

To capture microscale brain dynamics, we can use *biological neural networks* which simulate the activity of interconnected neurons [125]. Separate single-neuron models, such as the Hodgkin–Huxley model [126], are coupled to capture neuron–neuron interactions reflecting underlying synaptic processes. By assigning separate models to each neuron,

biological neural networks can simulate dynamics with realistic heterogeneity among neurons. Specific parameters can be assigned to heterogeneous classes of neurons, such as excitatory and inhibitory neurons [127–131], or classes of intricate morphological features [132, 133]. Ongoing projects to create expansive classification databases such as the Blue Brain Project [134] and the Allen Cell Types Database [135], alongside technological advances in computing power, suggest that future biological neural networks could simulate brain dynamics at neuronal resolution whilst spanning large spatial scales.

Zooming out from the microscale, mean-field population models instead track the mean activity of all neurons in a population, without tracking each constituent neuron [16, 17]. Crucially, this mean-field approach to model construction requires assumptions on the connectivities between populations in space. For example, there are *neural field models* [64, 65], which assume isotropic connectivity between populations to allow for efficient computation of population-scale activity across continuous space [16, 74, 136, 137]. Another class of mean-field models across populations are *neural mass models* or brain-network models [17, 138], which coarse-grain neural populations into discrete spatial regions, and then couple these regions to one another in a network, using measured structural connectome data [7, 139, 140]. Notwithstanding the limitations in accuracy and physical realism due to mean-field assumptions [141], the significant reduction in the number of variables allows mean-field population models to be computed much more efficiently than the equivalent hypothetical biological neural network of the same spatial scale.

By designing stimulation inputs with unique spatiotemporal profiles, physiological models integrated with targeted stimulation datasets can help investigate the mechanisms behind varying responses to different stimulation forms. For instance, a biological neural network of the rat somatosensory cortex revealed how morphological features like axon length shape neuron activation under ICMS and TMS under various stimulation parameters [132]. Another biological neural network of the mouse and monkey somatosensory cortex investigated recurrent mechanisms underpinning the difference in responses between optogenetic and visual sensory stimulation [85]. Additionally, physiological models can account for nonlinearities observed in experimental responses to targeted stimulation. For example, a biological neural network of the macaque motor cortex demonstrated potential nonlinear mechanisms underpinning distortive effects of ICMS on the dynamics of ongoing task-related activity [83]. At the macroscale, a neural field model with nonlinear synaptic mechanisms studied the calcium-dependent mechanisms underlying the nonlinear effects of TMS-induced plasticity [142]. A similar neural field model demonstrated significant differences in local synaptic strength parameters when fitted to spontaneous EEG versus evoked potentials [72]. This discrepancy was reconciled with a dynamic synaptic strength governed by a nonlinear gain modulation mechanism [143]. Finally, a neural mass model was used to demonstrate how recurrent mechanisms can give rise to spatial gradients in

excitability, which lead to spatial variations in the strengths of intracranially evoked EEG responses [144].

2.4.2 Phenomenological models

While physiological models offer realistic biophysically constrained explanations of how brain activity stems from independently measurable physiological processes, it is also useful to avoid excessive physiological details that complicate the model’s dimensionality and human interpretability [145]. An alternative approach, taken by phenomenological models, is to replace a full physiological description with simpler canonical dynamical systems, which preserve the dynamics hypothesised to be exhibited by the underlying complex physiological processes [60, 61]. This abstraction is supported by the observation that many physiologically realistic models of stable brain dynamics exhibit canonical dynamical structures such as fixed points and limit cycles [12, 59, 77]. While the variables and parameters of phenomenological models are not physiologically interpretable (and thus not experimentally measurable), they simplify the model’s complexity and the systems-level interpretation of the resulting simulated dynamics. Phenomenological models are thus valuable for investigating the core spatially distributed properties of the dynamics of networks of neurons or brain regions in response to targeted stimulation, without comprehensive detail of the local physiological mechanisms that underpin each neuron or brain region’s dynamics.

Phenomenological models vary in the sophistication of dynamics exhibited by their canonical structures. A simple example is a *linear network model*, which uses a linear system of equations to model the dynamics of a network of nodes representing neurons or neural populations [61, 97]. Each node’s dynamics are driven by a linear combination of activities from coupled nodes, and any external perturbation such as targeted brain stimulation. Linear models are a means to capture correlations between the activity of coupled neurons or neural populations with the simplest form of node-to-node coupling interactions [146, 147]. They may also allow for closed-form expressions for network responses to inputs, mathematically linking the system dynamics directly to its underlying coupling structure [148, 149].

Generalising the linear setting, there are nonlinear network models governed by nonlinear systems of equations, which can capture more sophisticated forms of local dynamics. A simple example is linear threshold models [150], which extend linear models by incorporating an additional threshold to the linear combination input. More sophisticated *coupled oscillator models* aim to capture oscillatory properties of brain dynamics as observed in recordings of brain activity across spatial scales [76, 151]. Commonly used oscillators include the Kuramoto oscillator which oscillates with constant amplitude [6, 152], and the Stuart–Landau (Hopf) oscillator which can also change in amplitude [76, 153, 154]. A more physiologically inspired nonlinear model is a *conductance-based model*, which models

population-level activity abstracted as a simpler biological neural network [5, 155, 156]. Here, the aggregate activity of each population is tracked by the dynamics of a single neuron model, which assumes a sufficiently strong coherence between constituent neurons of each population [17]. Despite this strong physiological assumptions, conductance-based models are capable of exhibiting population-scale spatiotemporal dynamics with interesting properties that are commonly observed in biological neural networks.

In comparison to physiological models, the abstractions incorporated in the local dynamics of phenomenological models limit their ability to distinguish between the dynamics evoked by different targeted stimulation techniques. However, phenomenological models can be used to more directly investigate how the brain’s underlying structure shapes the spatially distributed properties of its stimulus-evoked dynamics under an arbitrary stimulus input, as they allow the experimenter to parameterise or minimise complex local dynamics and thereby isolate the contribution of large-scale structural connectivity to the resulting activity patterns. For example, a coupled oscillator model integrated with mouse optogenetic stimulation data demonstrated how heterogeneous long-range structural connections allow stimulus-evoked activity to reveal a rich repertoire of novel dynamical responsive functional networks, extending the established repertoire of resting state networks observed in spontaneous dynamics [157, 158]. Phenomenological models can also investigate the extent to which deviations in spatiotemporal dynamics between stimulus-evoked and spontaneous activity are influenced by the stimulus position. For example, a coupled oscillator model demonstrated how TMS-induced changes in functional connectivity varied between stimulating highly connected hub regions in hub regions and stimulating weakly connected peripheral regions [152]. Another coupled oscillator model studied how the influence of stimulus position varied between different global brain states such as wakefulness and deep sleep [153].

2.4.3 Data-Driven models

The governing rules of phenomenological and physiological models summarised above reflect qualitatively interpretable mechanisms underlying brain dynamics, such as the internal physiological processes of neurons in biological neural networks, or the oscillatory behaviours of population-scale neural recordings in a whole-brain oscillator model. In contrast, data-driven models encode statistical mechanisms which can describe the quantitative structure of the governing rules without a qualitative interpretation [59, 62]. Through the use of statistical techniques, data-driven models are capable of exploring a wide range of possible governing rules which reflect encoded statistical mechanisms, enabling highly accurate predictions of brain dynamics that are unconstrained by predefined physiological or phenomenological mechanisms. This flexibility is particularly advantageous when studying responses to targeted stimulation, where underlying qualitative mechanisms may be difficult to hypothesise in advance. At the same time, choosing rule structures requires

careful judgement: models need enough coefficients to predict complex dynamics, but not too many under-constrained coefficients which undermine the model’s generalisation ability.

The range of different possible statistical properties of brain activity that a data-driven model can capture, varies with the range of governing rules that it can explore. For example, there are data-driven models which use the process of *system identification* to infer a set of mathematical equations that govern the dynamics of brain activity from the dataset [62]. While system identification algorithms are limited to exploring and discovering equation-based governing rules, these equations enable the interpretation of the learned statistical properties of individual time series, and pairwise dependencies between time series. Commonly used system identification procedures include linear time series models which learn simple linear equation structures [159]; and procedures which can estimate nonlinear equation structures [160–162]. Comparing different system identification procedures therefore provides an opportunity to detect nonlinearities in stimulus-evoked dynamics, by comparing the capacities of linear and nonlinear equations in predicting responses [163–165].

There are also more flexible data-driven models such as *artificial neural networks* (not to be confused with biological neural networks), which can be used to solve a more general problem of learning a predictive mapping between an input and output variable from a dataset of observations [62]. The majority of artificial neural network architectures, including the simplest feed-forward architecture, learn nonlinear mappings between variables as a sequence of layers of units, including the input layer, output layer, and hidden layers [166]. Each unit receives inputs from units of the previous layer, then feeds the input into an activation rule before sending it to units of the next layer. While less interpretable than the systems of equations discovered by system identification methods, artificial neural networks are more flexible in replicating a wider range of governing rule structures, including those with high orders of nonlinearities that would otherwise be computationally expensive to be captured as a set of equations [167, 168]. This flexibility potentially enables artificial neural networks to create future personalised models that make accurate predictions of individual responses to targeted stimulation, adapting important individual-specific structural and functional characteristics which are difficult to encode as physiological or abstracted mechanisms [169].

A systematic comparison of data-driven models can be used to test whether hypothesised statistical mechanisms are necessary in capturing responses to targeted stimulation. For example, a comparative study of responses to DBS found significant differences in fitted parameters of a linear autoregressive with exogenous input (ARX) model between spontaneous and stimulus-evoked dynamics [163]. This finding suggests that a single linear model may be insufficient to capture both regimes, and instead motivates the use of a switched linear model as a simple nonlinear model that switches parameter sets in

the presence of targeted stimulation. Similar predictive powers were also found between switched linear models and less interpretable feedforward neural networks, demonstrating that switched linear models sufficiently explained responses to DBS without significantly trading off predictive capabilities [163]. Another avenue of systematic comparison involves an architecturally distinct family of recurrent neural networks, which use recurrent connections to accumulate memory of an indefinite number of past inputs and observed states [166]. While recurrent neural networks have mostly focused on modelling dynamics during task stimuli [168, 170], they may also investigate how dynamic working-memory mechanisms can impact variations in time-dependent responses to sequences of targeted or sensory stimulation. For example, a comparison of recurrent neural networks with other memoryless neural network architectures demonstrated the necessity of internal memory mechanisms in predicting responses to visual stimuli sequences of varying images [171].

In summary, in this section we have explored the variety of different generative models, as well as their different aims, advantages, and disadvantages for modelling dynamics evoked by targeted stimulation. Choosing an appropriate model is crucial for investigating different types of questions about this form of dynamics: testing mechanisms related to the physiology of neurons would benefit from a biological neural network, or exploring the role of the connectome in the spatially distributed response to TMS could benefit from a phenomenological linear network model, or a project to accurately predict the firing rates of single neurons in response to optogenetic stimulation could benefit from system identification models or artificial neural network models. Integrated with datasets of activity from targeted stimulation experiments, generative models complete our toolkit for uncovering the mechanisms underlying the complex, high-dimensional dynamics measured from this unique experimental paradigm.

2.5 Conclusion

Emerging targeted stimulation techniques such as optogenetics and electrode stimulation, combined with measurement techniques such as calcium imaging and EEG, now allow us to measure the intricate spatiotemporal responses of neurons and neural populations to a range of perturbations with unprecedented precision [27, 28, 82, 104]. Using an abstracted picture of brain dynamics, and drawing on results from dynamical systems theory, we presented an argument for why the dynamics from novel states accessed in these targeted stimulated experiments may be driven by new types of mechanisms, distinct to those underpinning the simpler low-dimensional dynamics of spontaneous activity. The mechanisms which govern dynamics evoked by targeted stimulation are also functionally significant, as they reflect how the brain recruits diverse neural populations while processing complex streams of sensory inputs. Fortunately, through an integration

of targeted stimulation and measurement techniques with different generative models, we can uncover, analyse, and better understand these mechanisms at a chosen spatial scale. Targeted stimulation techniques allow us to precisely perturb the brain to a range of states distant from the attractor manifold on which spontaneous dynamics unfold. Measurement techniques can simultaneously capture the brain’s trajectory as it relaxes back to the attractor, generating datasets of activity that could reveal the novel (and likely highly nonlinear) mechanisms that govern the relaxation’s dynamics. Finally, by integrating datasets of activity evoked by targeted stimulation with generative models, it is possible to rigorously test and validate mechanisms that we hypothesise to drive the dynamics. Different generative modelling approaches are suited to testing different types of hypothesised mechanisms, from different physiological processes that can be incorporated in physiological models [16, 17], to the abstracted dynamical mechanisms treated in phenomenological models [60, 61], and the statistical mechanisms empirically learned using data-driven models [159, 166]. Our exploration touches on some of the many opportunities to uncover mechanistic insights into complex brain dynamics, ultimately guiding the practical implementation of emerging neuromodulation technologies in both research and clinical settings.

While our dynamical systems framework provides a foundation for interpreting the unique characteristics of brain dynamics evoked by targeted stimulation compared to spontaneous brain activity, there are alternative ways to conceptualise these dynamics. One such approach involves the separation of timescales [172]. Some mechanisms may act too quickly to significantly influence the trajectories of fluctuations near the spontaneous activity attractor but become crucial in capturing the far-from-attractor dynamics evoked by targeted stimulation. Brain activity may therefore possibly be modelled as a multi-timescale system, with mechanisms operating over various timescales in response to targeted stimulation [139, 143]. Targeted stimulation can also trigger new forms of non-equilibrium dynamics if multiple attractors exist in the state space, where specific perturbations can cause the relaxation trajectory to migrate to different attractors. Each attractor may represent unique forms of spontaneous brain activity, such as the states of deep sleep versus wakefulness [173], or sequential discrete thoughts during cognition [174]. Generative models of targeted stimulation can therefore inform the development of neuromodulation treatments to restore healthy brain function, by designing optimal perturbations that can force brain states to transition from ‘diseased’ to ‘healthy’ attractor dynamics [175–177].

As a novel experimental paradigm, datasets generated from targeted stimulation experiments allow us to not only test newly hypothesised mechanisms of brain dynamics, but also to refine existing mechanisms which have been validated mostly on datasets of spontaneous activity. Integrating existing models with new datasets can potentially contribute towards resolving standing issues of degeneracy in computational neuroscience, where generative models founded on differing mechanisms can lead to similar dynamics

of spontaneous activity [178]. It is possible that the relaxation dynamics measured in targeted stimulation experiments, given their potential nonlinearities, can be used to arbitrate between different generative quantitative models of brain dynamics, beyond the relatively simple statistical properties of spontaneous activity that both forms of models are known to adequately capture, such as their pairwise linear correlation structure (as ‘functional connectivity’) [8, 79].

Looking at and beyond the targeted stimulation setting, we suggest that more synergistic interactions between experimentalists and theoreticians can accelerate progress in our mechanistic understanding of brain dynamics. As experimentalists forge more sophisticated paradigms that each generate unique forms of brain dynamics, theoreticians can build more robust generative models that can explain or predict the dynamics of brain activity from a wider range of experimental settings. Conversely, refined generative models motivate the development of further experimental designs tailored to test and validate new theoretical accounts and predictions. This reciprocal relationship between experimental precision and theoretical rigour can uncover novel hypotheses and validate theoretical predictions, fostering a deeper and more integrative comprehension of the intricate mechanisms underlying brain function.

3. **Modelling the influences of non-local connectomic projections on geometrically constrained cortical dynamics**

CHAPTER SUMMARY

The previous chapter argued, using dynamical systems theory, that modelling the brain’s response to an impulse stimulus can reveal mechanisms of brain dynamics operating over fast (sub-second) timescales that are obscured in spontaneous settings. Building on this opportunity, this chapter aims to investigate the dynamical role of specifically arranged long-distance connections—referred to as fast-conducting non-local projections (FNPs)—and the contexts in which they provide a distinct contribution beyond dynamics obtained from geometrically constrained connectivity. The motivation behind studying this particular structural mechanism is to address an open question in macroscale neuroscience, of why simple geometric theories can successfully capture spontaneous dynamics over long (order-of-seconds) timescales, despite neglecting the spatial heterogeneity of FNP connectivity. These findings demonstrate that, while geometry is sufficient to constrain slow spontaneous fluctuations, FNP connectivity is essential for capturing dynamic interactions between remote populations that occur over fast timescales. Note that the chapter summarises the model equations and numerical details, with full information (including model derivation) provided in Appendix A.

Abstract

The function and dynamics of the cortex are fundamentally shaped by the specific wiring configurations of its constituent axonal fibres, also known as the connectome. However, many dynamical properties of macroscale cortical activity are well captured by instead describing the activity as propagating waves across the cortical surface, constrained only by the surface’s two-dimensional geometry. It thus remains an open question why the local geometry of the cortex can successfully capture macroscale cortical dynamics, despite

neglecting the specificity of Fast-conducting, Non-local Projections (FNPs) which are known to mediate the rapid and non-local propagation of activity between remote neural populations. Here we address this question by developing a novel mathematical model of macroscale cortical activity in which cortical populations interact both by a continuous sheet and by an additional set of FNPs wired independently of the sheet’s geometry. By simulating the model across a range of connectome topologies, external inputs, and timescales, we demonstrate that the addition of FNPs strongly shape the model dynamics of rapid, stimulus-evoked responses on fine millisecond timescales ($\lesssim 30$ ms), but contribute relatively little to slower, spontaneous fluctuations over longer timescales (> 30 ms), which increasingly resemble geometrically constrained dynamics without FNPs. Our results suggest that the discrepant views regarding the relative contributions of local (geometric) and non-local (connectomic) cortico-cortical interactions are context-dependent: While FNPs specified by the connectome are needed to capture rapid communication between specific distant populations (as per the rapid processing of sensory inputs), they play a relatively minor role in shaping slower spontaneous fluctuations (as per resting-state functional magnetic resonance imaging).

3.1 Introduction

Cortical function emerges from complex interactions between spatially distributed neural populations, mediated by the structural axonal fibres that connect them [179]. While the distribution of all fibres in the cortex is skewed towards shorter lengths [33], some fibres are genetically encoded to connect specific remote populations [40, 180], sometimes forming highly myelinated bundles [181]. Despite their substantial energy requirements for development and maintenance [182], neurological experiments from the past century have demonstrated the critical role of these specifically positioned fibres—which we hereafter call fast-conducting non-local projections (FNPs)—in supporting cognition and behaviour, by revealing the functional deficits that occur when they are damaged or removed [183–190]. Stimulus–response experiments also demonstrate that FNPs are specialised to propagate macroscale cortical activity at speeds much faster than travelling waves on the continuous cortical surface, thus enabling rapid, non-local interactions between specific populations that bypass the constraints of the cortex’s local geometry [191–195]. Based on empirical observations including those above, FNPs are widely considered to be a core structural mechanism supporting the global integration of information in the brain [2, 196]. Accordingly, several high-profile experimental efforts have aimed to develop comprehensive connectomes across species, which detail the specific populations that FNPs (as well as other fibres) connect [29–31], and hence significantly accelerate our understanding of the structural underpinnings of cortical function and dysfunction [197, 198].

Despite extensive evidence demonstrating the key role of FNPs in supporting cortical function and shaping cortical dynamics, recent analyses have indicated that many dynamical properties of macroscale activity, on millimetre lengthscales and above, can be well captured by geometrically constrained models of neural activity that neglect the specificity of connectome-informed FNP connectivity [37–39]. This body of work suggests that macroscale cortical activity is well approximated as the superposition of travelling waves on a physically continuous cortical sheet, such that the waves are constrained only by the local two-dimensional geometry of the cortical surface rather than the specific wiring positions of the cortex’s constituent axonal fibres. This geometric description of cortical activity is derived by a broad class of generative models called neural field models [64, 65, 136, 199] which assume a mean approximation rule of cortical structural connectivity that is spatially homogeneous (exhibiting the same properties at all points) and spatially isotropic (exhibiting the same properties in all directions). It remains an unresolved and fundamental problem in macroscale neuroscience how neural field models, which commonly assume homogeneous and isotropic FNP connectivity, in contrast to the reliable heterogeneity observed in the connectome [36, 41, 42, 200], can nevertheless generate predictions of cortical dynamics that closely match many experimental observations.

Here we seek to mechanistically reconcile two seemingly conflicting empirical findings: (i) direct evidence of the role of FNPs in facilitating the non-local dynamics of stimulus-evoked whole-brain responses measured with voltage-sensitive dye (VSD) imaging [191–193], electroencephalography (EEG) [201, 202], and wide-field calcium imaging [194]; and (ii) the predictive successes of geometrically constrained neural field models, which ignore the specificity of FNP connectivity informed by connectomic data, in capturing the spatiotemporal dynamics of resting-state functional magnetic resonance imaging (fMRI) [37–39], the spatial patterns of fMRI task activation maps [37, 38], and resting-state magnetoencephalography (MEG) [38]. To this end, we develop a new mathematical model of macroscale cortical activity in which interactions between cortical populations are mediated both by: (i) a continuous sheet of fibres that are homogeneously and isotropically distributed in space; and also by (ii) a set of specific FNPs whose wiring positions are independent of the geometry of the sheet. Our phenomenological model is a simple extension of a neural field model, in which the additionally wired FNPs perturb the dynamics of the travelling waves along the cortical surface by behaving as ‘topological shortcuts’ that facilitate rapid propagation between specific distant populations (at speeds faster than the travelling waves). By simulating the model over a variety of FNP connectome topologies, we find that the influence of FNPs on the model dynamics is strongly dependent on two key factors: (i) the spatial precision of the external driving input (e.g., stimulus-evoked responses versus spontaneous activity with no explicit external stimulus); and (ii) the timescales over which spatiotemporal activity is resolved (e.g., the subsecond timescales of VSD versus the second timescales of fMRI). Our results thus provide a plausible explanation

for the empirical findings above, demonstrating that the connectome (which informs the specificity of FNP connectivity) and cortical geometry (which spatially aggregates this specificity) are best suited to describe and capture macroscale cortical dynamics from contrasting experimental settings and modalities.

3.2 Materials and methods

In order to investigate the additional contribution of fast non-local interactions to geometrically constrained cortical dynamics, we developed a mathematical model of macroscale cortical population-level activity (on millimetre lengthscales and above). Our model isolates the role of specific arrangements of FNPs in perturbing cortical dynamics, separately from the influence of the sheet geometry. The model’s structure is a two-dimensional continuous sheet representing the cortical surface, augmented with a set of FNPs that connect pairs of points on the sheet. While there are existing models that combine geometric and connectomic mechanisms of interaction [157, 158, 203–205], the model introduced here supports finite-speed travelling waves on a two-dimensional cortical sheet with multiple FNPs, thereby overcoming key limitations that render prior models unsuited to our investigations (discussed further in Discussion).

3.2.1 Model description

Figure 3.1A illustrates the architecture of our model—containing both geometrically constrained and fast non-local connectomic mechanisms of cortico-cortical interaction—which we describe below. The cortical sheet of a cerebral hemisphere is represented as a continuous two-dimensional surface, denoted Ω . Cortical populations are identified by their spatial location $\mathbf{r}$ on Ω , and their neural activity over time is represented by the field variable $\phi(\mathbf{r}, t)$. The spatiotemporal evolution of $\phi(\mathbf{r}, t)$ over time $t \in [0, \infty)$ is governed by the non-local partial differential equation:

$$\mathcal{D}_t \phi - r^2 \nabla^2 \phi - C \circ \phi = f, \quad \text{for all } (\mathbf{r}, t) \in \Omega \times [0, \infty); \quad (3.1)$$

$$\phi(\mathbf{r}, 0) = \dot{\phi}(\mathbf{r}, 0) = 0. \quad (3.2)$$

In Eq. (3.1), $\mathcal{D}_t$ is a second-order linear differential operator in time, and the term $r^2 \nabla^2 \phi(\mathbf{r}, t)$ captures local spatial relations on Ω for different points $\mathbf{r}$. Both terms together capture local propagation as travelling waves along Ω , which arises from a geometric (spatially homogeneous and isotropic) connectivity rule (illustrated in Fig. 3.1B) [64, 65, 136]. The term $(C \circ \phi)(\mathbf{r}, t)$ captures non-local propagation between specific points on Ω facilitated by FNPs (forming ‘shortcuts’ through Ω , as illustrated in Fig. 3.1C). Each FNP is a unidirectional projection from an assigned source point (denoted ‘ $\mathbf{a}$ ’ in Fig. 3.1A) to a

target point (denoted ‘**b**’ in Fig. 3.1A), through which activity propagates more rapidly than by a travelling wave along Ω . The term $f(\mathbf{r}, t)$ is the external input that drives the model’s activity. In the absence of input, $f \equiv 0$, the initial conditions in Eq. (3.2) define zero activity at all points $\mathbf{r}$ and for all time t . Further details on the above terms, and a derivation of Eq. (3.1) from physiological principles, are in Appendix A, Sec. A.1.




Figure 3.1: **Schematic of our mathematical model of macroscale cortical activity with local geometric (blue) and non-local connectomic (green) mechanisms of cortico-cortical interaction.** The model equations are defined in Eq. (3.1). **A.** The two-dimensional continuous surface Ω represents the cortical sheet of a cerebral hemisphere, on which neural populations are represented as points. The mean population-level activity at each point, on millimetre lengthscales and above, is described by variable ϕ . This model includes both local non-local mechanisms of interaction between points on Ω . Local geometric propagation of activity takes the form of travelling waves along Ω , while non-local connectomic propagation takes the form of rapid, non-local interactions between spatially distant pairs of points on Ω (an example of which is illustrated as **a** and **b**). **B.** Geometric propagation (travelling waves) is derived from a spatially homogeneous and isotropic connectivity rule [64, 65, 136]. **C.** Connectomic propagation between specific pairs of points is facilitated by fast-conducting, non-local projections (FNPs), which are axonal fibres that propagate activity faster than travelling waves along Ω , and is considered to underpin macroscale cortical function by facilitating rapid inter-communication between spatially distributed neural systems. FNPs introduce spatial inhomogeneities and anisotropies to the otherwise geometric connectivity.

A notable property of the non-local term $C \circ \phi$ is that, for any given input f , the total activity in the system integrated over space and time is preserved after the addition of any FNP to the structural connectivity. This property prevents non-oscillatory instabilities in the dynamics, which can occur when activity is regenerated from FNPs at a faster rate than activity globally dissipates [203, 204], and therefore allows multiple FNPs to be added with arbitrary connectivity strengths. In alternative models, total activity is not preserved (e.g., the neural field model by Jirsa and colleagues [203–205]) — in such models, the system can achieve non-oscillatory instabilities with sufficiently large FNP

connectivity strengths [203–205]; thus, the connectivity strengths of each FNP must be progressively reduced to maintain linear stability.

The model defined by Eq. (3.1) will be used to simulate cortical dynamics for different external inputs $f(\mathbf{r}, t)$, so that the perturbatory influence of FNPs on the dynamics can be evaluated separately for each input setting. For example, to simulate stimulus-evoked dynamics in response to an impulse stimulus, f is set as an impulse-like input that is Gaussian in both space and time [72, 206]:

$$f^{\text{stim}}(\mathbf{r}, t) \propto \exp\left(\frac{-\|\mathbf{r} - \mathbf{r}_{os}\|^2}{2\sigma_x^2}\right) \exp\left(\frac{-(t - t_{os})^2}{2\sigma_t^2}\right);$$

$$\int_0^\infty dt \int_\Omega d^2\mathbf{r} f^{\text{stim}}(\mathbf{r}, t) = 1; \quad (3.3)$$

where $\mathbf{r}_{os} \in \Omega, t_{os} > 0$ are the position and onset time of the stimulus, respectively, and σ_x, σ_t are the spatial and temporal width of the stimulus, respectively. In contrast, to simulate noise-driven dynamics—that is, temporally unstructured inputs—we set f as a continuous space–time stochastic process:

$$f^{\text{noise}}(\mathbf{r}, t) = g(\mathbf{r})\dot{W}(\mathbf{r}, t);$$

$$\mathbb{E}(\dot{W}(\mathbf{r}, t)\dot{W}(\mathbf{r}', t')) = \delta(\mathbf{r} - \mathbf{r}')\delta(t - t'); \quad (3.4)$$

where $\dot{W}$ on $\Omega \times [0, \infty)$ is a space-time white noise process [39, 207], and $g(\mathbf{r})$ on Ω is a smooth deterministic function which controls the variance of the process at different points on Ω .

3.2.2 Model implementation

To numerically solve the model defined by Eq. (3.1), we developed a finite-difference scheme used for non-local partial differential equations, in which derivatives in time and space are approximated by centred finite differences and definite integrals are approximated using a quadrature rule [208]. For our investigations, Ω is taken to be a square with periodic boundary conditions, making its geometry equivalent to a flat torus. A torus has a closed surface, like a more physically realistic sphere [209], but permits more efficient numerical simulations on a square grid. The side length of the square was set to $L = 0.4$ m, so that the area of Ω is approximately that of a human cortical hemisphere [136]. We then uniformly partitioned a grid of points from Ω with grid spacing Δx , and time points from $[0, \infty)$ with timestep Δt . With this time and space partitioning, we used a finite-difference scheme to numerically solve Eq. (3.1). Further details and the derivation of this scheme, including the chosen grid spacing and timestep for our investigations, can be found in Appendix A, Sec. A.2.

The parameters of Eq. (3.1) (detailed in Eqs (A.18), (A.19), (A.20) in Appendix A,

Sec. A.1) and their values used in our investigations are given in Table 3.1. For clarity, we divide the parameters into three groups: (i) parameters that govern geometric propagation through travelling waves: r , γ , ν_0 ; (ii) parameters that govern interaction via FNPs: $\{c_m\}_m$, $\{\tau_m\}_m$; and (iii) parameters that govern the spatiotemporal properties of the external input f : σ_x, σ_t . The value of parameter r , the characteristic lengthscale of the geometric connectivity rule, and the value of γ , the travelling wave speed divided by r , were both taken from the two-dimensional neural field model developed by Robinson et al. for the human cortex [136, 210]. The value of ν_0 , which is the effective rate at which activity at a given position is regenerated, was taken from the excitatory-only neural field model developed by Jirsa and Haken [65].

Table 3.1: **Nominal model parameter values for numerical experiments presented here.**

Parameter	Description	Value
Topology		
L	Length of Ω (as a square)	0.4 m
Local Geometric Propagation		
r	Lengthscale of geometric connectivity	0.086 m
γ	Wave speed divided by r	116 s^{-1}
ν_0	Cortical gain	0.756
Non-local Connectomic Propagation		
c_m	Connectivity strength of m th FNP	$7.396 \times 10^{-3} \text{ m}^2$ for all m
τ_m	Conduction time of m th FNP	0 s for all m
Input		
σ_x	Spatial width of stimulus input	$4 \times 10^{-3} \text{ m}$
σ_t	Temporal width of stimulus input	$6 \times 10^{-4} \text{ s}$

We set c_m , which is the connectivity strength of the m th FNP, to equate the strength of local and non-local propagation from the source of the FNP. Accordingly, we set $c_m = r^2$ for each m , so that when the model is numerically treated, the connectivity strength of the FNP ($c_m/(\Delta x)^2$ in Eq. (A.24) in Appendix A, Sec. A.2) is equal to the strength of coupling between geometrically nearest neighbours on the grid from the local geometry ($r^2/(\Delta x)^2$ in Eq. (A.23) in Appendix A, Sec. A.2). We set the value of τ_m , which represents the conduction time of the m th FNP, so that it is always smaller than the conduction time of a travelling wave between the FNP source and target, but is also small enough to avoid oscillatory instabilities independently of the chosen connectome configuration [203, 204]. However, the conduction speeds of FNPs vary widely from factors other than their length [211], and the critical values of the τ_m values separating stability from instability depends on the spatial configurations of the FNPs [203]. For this reason, we fixed $\tau_m = 0$ for each m . The values of σ_x, σ_t (in Eq. (3.3)), the spatial and temporal widths of the impulse stimulus, respectively, were set to sufficiently small values such that the input mimics a

delta impulse function without causing spurious oscillations in the numerical scheme.

3.3 Results

Using the phenomenological model defined in Eq. (3.1), we aim to shed light on the variable role of FNPs across different experimental settings, with specific individual FNPs playing a clear role in facilitating corticocortical interactions that underpin distributed information processing in some settings [187, 189, 191, 192], but being apparently negligible relative to idealised geometric connectivity rules in others [37–39].

3.3.1 Model demonstration

We first introduce the behaviour of the model, through a simple demonstration of how the incorporation of FNPs to the model’s structural connectivity perturbs its stimulus-evoked dynamics (i.e., evoked response to an impulse stimulus, cf. Eq. (3.3)). Two configurations of structural connectivity were considered. The geometric model, depicted in Fig. 3.2A, is governed by the spatially homogeneous and isotropic structural connectivity rule of neural field models, thus completely constrained by the geometry of Ω and absent of any FNPs [136]. The hybrid model, depicted in Fig. 3.2B, additionally incorporates 50 FNPs whose source and target positions were each sampled uniformly at random from Ω . We then computed the response of both models to an impulse stimulus at position $\mathbf{r}_{os} = (L/2, L/2)$, shown with a red dot in Figs 3.2A and B. The stimulus-evoked responses of the geometric and hybrid models are plotted in Figs 3.2C and D, respectively, as heatmaps of their spatial activity profiles $\phi(\mathbf{r}, t)$, at each of six time points following stimulus onset. Consistent with the connectivity rule of neural field models, the stimulus-evoked response of the geometric model is a travelling wave emanating from $\mathbf{r}_{os}$. By contrast, the response of the hybrid model (containing FNPs) also comprises the activation of specific non-local points due to non-local activity propagation via FNPs. These non-local activations lead to additional travelling wave fronts under the geometric connectivity, yielding dynamics comprising a superposition of multiple travelling waves. In summary, FNPs yield more complex stimulus-evoked dynamics by non-locally propagating activity between spatially distant populations at speeds faster than the travelling waves constrained by cortical geometry.

3.3.2 Influence of a single FNP on cortical dynamics resolved across varying timescales

We aim to understand how a single FNP’s perturbation on the model’s stimulus-evoked dynamics evolves over different timescales. We then use this distinction of perturbation by timescale to shed light on how the influence of FNPs on measured macroscale cortical dy-




Figure 3.2: **FNPs shape stimulus-evoked dynamics by enabling rapid, non-local propagation of activity between neural populations.** We used the model defined in Eq. (3.1) to examine the influence of FNPs on stimulus-evoked dynamics, by comparing the model’s stimulus-evoked dynamics under two structural connectivity configurations, in response to the same impulse stimulus. **A.** The geometric model is governed by a spatially homogeneous and isotropic connectivity rule, and is purely constrained by the geometry of Ω (shown as the square). **B.** The hybrid model’s structural connectivity additionally contains 50 FNPs, depicted as arrows from their source to target. The impulse stimulus applied to both models was positioned at the point $\mathbf{r}_{os}$, annotated with a red dot. **C.** The geometric model’s stimulus-evoked response is illustrated as heatmaps of activity profile $\phi(\mathbf{r}, t)$ at time points, $t = 0, 2, 4, 6, 8, 10$ ms, measured relative to stimulus onset. **D.** The hybrid model’s stimulus-evoked response is illustrated in a similar manner as in **C.** Specific FNPs of the hybrid connectivity (depicted in **B**) facilitate rapid activity propagation from their source to target at different times and are annotated manually in each heatmap. Note that we have imposed an upper limit of $\phi = 650$ on the colour scale for all heatmaps in **C** and **D**, to facilitate clear visibility across all time points.

namics could differ between activity resolved over fast sub-second timescales by modalities such as voltage-sensitive dye (VSD) imaging, EEG or calcium imaging [191–194, 201, 202]; and activity resolved over slower, second timescales by modalities such as fMRI [37–39].

We compared the stimulus-evoked response of the geometric model depicted in Fig. 3.3A(i), and a hybrid model containing a single additional FNP from $\mathbf{p} \rightarrow \mathbf{q}$, depicted in Fig. 3.3A(ii). The length of this FNP ($L\sqrt{2}/4 \approx 14$ cm) is comparable to the longest intra-hemispheric FNPs measured in human cortex [212]. The impulse stimulus applied to both model was set to position $\mathbf{p}$, shown as a red dot in Figs 3.3A(i) and (ii), which was selected to ensure that the FNP $\mathbf{p} \rightarrow \mathbf{q}$ would propagate activity from $\mathbf{p}$ to $\mathbf{q}$ immediately at stimulus onset. The stimulus-evoked responses of both the geometric and hybrid model to this impulse stimulus are shown in Figs 3.3A(iii) and (iv). As expected, the FNP perturbs the stimulus-evoked dynamics by non-locally activating $\mathbf{q}$, which results in an additional travelling wave emanating from $\mathbf{q}$. However, from the point that the two travelling waves begin to superimpose, from $t \approx 10$ ms onwards, the geometric and hybrid model’s response becomes progressively similar. At the later time points shown in Fig. 3.3A, $t = 30, 50$ ms, we see that the two models exhibit very similar spatial patterns in their responses, with the hybrid model’s spatial activity pattern displaying a very small perturbation, near $\mathbf{p}$ and $\mathbf{q}$, due to the perpetual propagation of activity between these two points via the FNP. To quantify the timescale dependency of the FNP’s perturbation on the dynamics, we measured the cosine dissimilarity metric (or cosine distance) between

the spatial activity patterns of the responses of the two models over time t , denoted as $C(t)$ (see Eq. (A.30) in Appendix A, Sec. A.3 for mathematical formulation). A plot of $C(t)$ in Fig. 3.3B shows that the FNP maximally perturbs the dynamics at time $t \approx 8$ ms, which closely corresponds to the time at which the two waves from $\mathbf{p}$ and $\mathbf{q}$ meet. We thus find that a single FNP has a strongly timescale-dependent perturbation on the model's dynamics, being substantial over short, millisecond timescales ($\lesssim 30$ ms) but minimal on longer timescales (> 30 ms).

To understand why the stimulus-evoked dynamics are most strongly influenced by a stimulus-proximate FNP over short, millisecond timescales ($\lesssim 30$ ms), we observed the evoked response of the geometric and hybrid models at two specific points $\mathbf{r}_1$ (in Fig. 3.3C(i)) and $\mathbf{r}_2$ (in Fig. 3.3C(ii)), which are positioned symmetrically on the interval from $\mathbf{p}$ to $\mathbf{q}$. At both $\mathbf{r}_1$ and $\mathbf{r}_2$, the FNP's perturbation on the response commences at time $t \approx 4$ ms (illustrated by the solid red line), through a reduction in activity at $\mathbf{r}_1$ and an increase in activity at $\mathbf{r}_2$. The increase in activity at $\mathbf{r}_2$ arises from the additional travelling wave emanating from $\mathbf{q}$, which arrives at $\mathbf{r}_2$ before the wave emanating from $\mathbf{p}$. Contrastingly, the decrease in activity at $\mathbf{r}_1$ arises from the total activity preservation property of the model, where the non-local propagation of activity from $\mathbf{p}$ to $\mathbf{q}$ reduces the amplitude of the travelling wave emanating from $\mathbf{p}$. However, the perturbation induced by the FNP is mostly 'cancelled' once the waves emanating from $\mathbf{p}$ and $\mathbf{q}$ begin to interfere. For $\mathbf{r}_1$ and $\mathbf{r}_2$, this cancellation occurs at $t \approx 11$ ms (illustrated by the dotted red line in Figs 3.3C(i) and (ii)), where the negative perturbation at $\mathbf{r}_1$ is compensated for by the superposition of activity arriving from the secondary wave emanating from $\mathbf{q}$, and the positive perturbation at $\mathbf{r}_2$ is cancelled by a weaker incident primary travelling wave from $\mathbf{p}$ from $t \approx 11$ ms. This experiment indicates that, at any given point $\mathbf{r}$ on Ω , the FNP can only substantially perturb the stimulus-evoked response over a specific time interval: from the time at which the first wavefront (from either $\mathbf{p}$ or $\mathbf{q}$) arrives at the point, to the time at which the second wavefront (from either $\mathbf{q}$ or $\mathbf{p}$, respectively) arrives. By generalising this result across all points on Ω , we therefore expect that the perturbation $C(t)$ in Fig. 3.3B would be maximal at a time soon after $t = (L\sqrt{2}/8)/(r\gamma) \approx 7$ ms, when the two waves in the hybrid dynamics superimpose; but would become minimal from time $(L\sqrt{2}/2)/(r\gamma) \approx 28$ ms, when the travelling wave from $\mathbf{p}$ reaches every point on Ω . These simple calculations capture the simulation results of Fig. 3.3B well and explain why an FNP strongly affects spatial activity patterns only over a relatively short time interval on short (millisecond) timescales.

The findings above suggest that the perturbation of the FNP $\mathbf{p} \rightarrow \mathbf{q}$ on the model's stimulus-evoked dynamics would reduce in magnitude as the resolved dynamics becomes increasingly restricted to longer timescales, such as the order-of-seconds timescales resolved by fMRI [213]. We explicitly tested this prediction by examining the perturbation of the FNP on the model's BOLD response to the impulse stimulus, i.e., the BOLD signal




Figure 3.3: The addition of a single fast-conducting non-local projection (FNP) predominantly shapes stimulus-evoked dynamics on short timescales $\lesssim 30 \text{ ms}$. **A.** We compared the stimulus-evoked dynamics of: **i.** the geometric model, and **ii.** a hybrid model whose structural connectivity additionally contains a single FNP from $\mathbf{p} = (3L/8, 3L/8)$ to $\mathbf{q} = (5L/8, 5L/8)$. An impulse stimulus was applied to both models at position $\mathbf{p}$, shown by the red dot. **iii.** and **iv.** show heatmaps of the stimulus-response dynamics $\phi(\mathbf{r}, t)$ of the geometric and hybrid models for $t = 0, 5, 10, 20, 30, 50 \text{ ms}$. Note that the first snapshot, labelled $t \approx 0$, was taken at a time slightly earlier than stimulus onset (-1.33 ms), to indicate that $\mathbf{q}$ is activated after $\mathbf{p}$ in the hybrid dynamics. The colour scale for all heatmaps is shown on the right of the figure. **B.** The cosine dissimilarity (or cosine distance) between the spatial activity patterns of the geometric and hybrid model’s stimulus-evoked dynamics over time, denoted $C(t)$ (Eq. (A.30) in Appendix A, Sec. A.3), quantifies the FNP’s perturbation on the model’s stimulus-evoked dynamics over time. **C.** We examined the difference between the geometric and hybrid model’s stimulus-evoked response at two locations: **i.** $\mathbf{r}_1 = (7L/16, 7L/16)$, and **ii.** $\mathbf{r}_2 = (9L/16, 9L/16)$. Insets in the top right corner of each plot show the positions of $\mathbf{r}_1$ and $\mathbf{r}_2$. The solid and dashed red lines in both plots represent the time at which the perturbation of the FNP on the dynamics commences ($t \approx 4 \text{ ms}$) and ends ($t \approx 11 \text{ ms}$), respectively.

observation of the stimulus-evoked response. At each $\mathbf{r}$, the BOLD response was approximated as $z(\mathbf{r})$, the stimulus-evoked response's zero temporal frequency component [213], or equivalently, the time-integrated response as depicted schematically in Fig. 3.4A (see Eq. (A.28) in Appendix A, Sec. A.2 for details on how $z(\mathbf{r})$ is computed numerically). This linear approximation closely aligns with, and hence can be generalised to BOLD responses computed with more complex nonlinear hemodynamic forward models [214].




Figure 3.4: **When neural dynamics are resolved only over long timescales (such as those accessible using fMRI), a single FNP has a much smaller impact on measured neural dynamics.** **A.** The modelled BOLD response (BOLD signal observation of the stimulus-evoked response), denoted z , is computed as the zero-frequency temporal Fourier component of the model's stimulus-evoked response. **i.** To demonstrate the contrast between the stimulus-evoked response and the BOLD response, we apply an impulse stimulus (positioned at the red dot) to the geometric model. **ii.** The stimulus-evoked response of the geometric model, $\phi(\mathbf{r}, t)$, is then integrated over time to approximate the slow BOLD response, shown by the heatmap of $z(\mathbf{r})$. **B.** We modelled the evoked BOLD response of: **i.** the geometric model, and **ii.** the hybrid model containing the FNP $\mathbf{p} \rightarrow \mathbf{q}$, when the stimulus is positioned at $\mathbf{p}$ (indicated by a red dot). **iii.** and **iv.** plot the BOLD response $z(\mathbf{r})$ as a heatmap for the geometric and hybrid models, respectively. **C.** The cosine dissimilarity between the BOLD responses of the geometric and hybrid models, denoted C_z (Eq. (A.31) in Appendix A, Sec. A.3), quantifies the perturbation of the FNP on the evoked BOLD response. C_z is illustrated as the horizontal grey line, $C(t)$ from Fig. 3.3B is illustrated as the black curve, and $C_{\max} = \max_t \{C(t)\}$ is illustrated as the dotted black line.

Figures 3.4B(i) and (ii) illustrate again the geometric and hybrid model respectively from Figs 3.3A(i) and (ii), and the impulse stimulus positioned at $\mathbf{p}$. Figures 3.4B(iii) and (iv) are heatmaps of the evoked BOLD response of both models to this stimulus, i.e., the time-integrated spatial activity pattern of the stimulus-evoked responses in Figs 3.3A(iii) and (iv). The FNP's perturbation on the modelled BOLD response is concentrated near $\mathbf{p}$ and $\mathbf{q}$, decreasing the response near $\mathbf{p}$ and increasing it near $\mathbf{q}$, which is consistent

with the same FNP’s perturbation on the model’s stimulus-evoked dynamics on longer timescales (> 30 ms) in our experiments above (Figs 3.3A(iii) and (iv)). To quantify this perturbation, we measured C_z , the cosine dissimilarity between the BOLD responses of the geometric and hybrid model (see Eq. (A.31) in Appendix A, Sec. A.3 for mathematical detail). A comparison between C_z and $C(t)$ in Fig. 3.4C shows that $C_z = 0.016$ is larger than the limiting value of $C(t)$ as $t \rightarrow \infty$, meaning that the FNP’s strong perturbations over short timescales ($\lesssim 30$ ms) are partially captured by the BOLD response. However, C_z is less than a fifth of $C_{\max} = \max_t\{C(t)\} = 0.085$, the maximum attained value of $C(t)$ at time $t \approx 8$ ms, indicating the short-timescale perturbation captured by the BOLD response is still substantially lower than what would be measured if the dynamics were properly time-resolved. Hence, restricting our observations of the activity dynamics to the longer, order-of-seconds timescales captured by fMRI indeed limits our ability to access the shorter timescales over which FNPs most prominently shape the stimulus-evoked dynamics. These findings provide a potential explanation for the discrepant empirical findings on the influence of FNPs on cortical dynamics across differing timescales: An FNP facilitates rapid non-local activity propagation on the sub-second timescales resolved by VSD imaging, EEG and calcium imaging [191–194, 201, 202], but has a weaker influence on the longer timescales captured by fMRI, in which the idealised geometric connectivity rule of neural field models becomes a more valid approximation [37, 38].

3.3.3 Influence of a single FNP on cortical dynamics driven by inputs of varying spatial precision

We next investigated how the perturbation of a single FNP on the model dynamics depends on the spatial precision of the driving input. We particularly explored perturbations on noise-driven dynamics (dynamics driven by noisy or temporally unstructured input), so that we could directly measure the FNP-induced perturbations to spontaneous activity, which is a commonly studied form of noise-driven dynamics where the input is spatially uniform [15, 215].

The noise-driven linear spatiotemporal dynamics of cortical activity can be decomposed as a superposition of linear evoked responses to spatially uncorrelated impulse stimuli across the cortical surface [141, 213, 216]. Hence, the perturbation of an FNP on noise-driven dynamics can be interpreted as its perturbation on stimulus-evoked dynamics, aggregated across all stimulus positions on the cortical surface, and weighted in accordance to the variance of the noisy input at each position. The previous experiments in Fig. 3.3 characterised the perturbation to geometric stimulus-evoked dynamics caused by a single FNP when the stimulus is incident at the FNP source location. However, a sweep of all possible stimulus positions on Ω shows that the extent of this perturbation (quantified by both $C_{\max} = \max_t\{C(t)\}$ and C_z) decreases with its distance from the source of the FNP (Appendix A, Fig. A.1), indicating that the FNP’s non-local propagation is highly

specific to activity that is proximate to its source location. Under our aforementioned interpretation, this result in the stimulus-evoked setting therefore suggests that a single FNP's perturbation to noise-driven dynamics would depend not only on the configuration of the FNPs, but also on the spatial precision of the driving noisy input. In particular, we predict that the FNP's perturbation on noise-driven dynamics increases as the input becomes more spatially localised at the FNP source, and decreases as the input becomes more spatially uniform in variance per spontaneous dynamics.

To test and quantify this prediction, we investigated how the FNP $\mathbf{p} \rightarrow \mathbf{q}$ perturbs the dynamics of a range of noise-driven activity—from activity driven by noisy input spatially localised at $\mathbf{p}$, to spontaneous activity (activity driven by spatially uniform noisy input)—and compared the magnitude of the perturbation of the FNP across this variation of input precision. As illustrated in Fig. 3.5A, we formulated the noisy input f^{stim} in Eq. (3.4) by setting $g(\mathbf{r})$, the input variance, as a Gaussian with fixed center $\mathbf{p}$ (Fig. 3.5A(i)). By varying the spatial width σ_n , we could assign different spatial profiles to the variance of the noisy inputs, ranging from spatially localised input at $\mathbf{p}$ ($\sigma_n \rightarrow 0$), to spatially uniform input (uniformly distributed in the limit $\sigma_n \rightarrow \infty$). For any choice of σ_n , we then summarised the statistics of the resultant noise-driven activity via the spatial correlation function from reference point $\mathbf{p}$, denoted $\gamma_{\mathbf{p}}(\mathbf{r})$ (Fig. 3.5A(ii)), which measures the Pearson correlation between the activities at $\mathbf{p}$ and $\mathbf{r}$ for each $\mathbf{r} \in \Omega$ (see Appendix A, Sec. A.3 for details on computation). The reference point $\mathbf{p}$ was natural to focus on, as it allows us to understand how the FNP is perturbing the dynamics by propagating activity from $\mathbf{p}$ (to $\mathbf{q}$).

Figure 3.5B plots heatmaps of $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$, the spatial correlation function of the hybrid model, for $\sigma_n = 3L/80, L/8, L/4$, and $\sigma_n \rightarrow \infty$ (corresponding to space-time white noise). We first note for all σ_n the expected decay of spatial correlation as a function of distance from $\mathbf{p}$, mediated by local geometric propagations of activity. The FNP also expectedly perturbs the geometry-mediated noise-driven dynamics by mediating long-range correlations between $\mathbf{p}$ and $\mathbf{q}$, the FNP target, due to non-local propagation between the two points. To quantify this perturbation, we compared $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$ with $\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{r})$, the spatial correlation function of the geometric model with reference point $\mathbf{p}$, for the same selections of σ_n (shown in Appendix A, Fig. A.2). Comparisons of the elements of $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$ and $\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{r})$ for each σ_n are plotted as scatters in Figs 3.5C(i)–(iv) for each σ_n value. These plots show that the FNP's perturbation on noise-driven dynamics is observable in all cases, with the largest deviation in the linear fit between $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$ and $\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{r})$ occurring at $\mathbf{q}$, denoted by the green dot. However, the Pearson correlation between the elements of $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$ and $\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{r})$, denoted r , increases with σ_n . The noise-driven dynamics of the geometric and hybrid models therefore becomes increasingly similar in correlation structure as the noisy input becomes more spatially uniform, characteristic of spontaneous activity. This finding is in line with our expectation and suggests that differences in the spatial precisions of




Figure 3.5: **The influence of an FNP on noise-driven dynamics is strongest for spatially precise input drives centred on the FNP source.** **A.** We investigated the influence of a single FNP $\mathbf{p} \rightarrow \mathbf{q}$ on noise-driven dynamics for noisy inputs for varying levels of spatial precision. **i.** We formulated the noisy input in Eq. (3.4) by setting $g(\mathbf{r}) = \exp(-\|\mathbf{r} - \mathbf{p}\|^2/2\sigma_n^2)$, a Gaussian function with centre $\mathbf{p}$ and spatial width σ_n (indicated by the radius of the red circle). **ii.** We summarised the spatial correlation structure of the noise-driven model dynamics via the spatial correlation function to the reference point $\mathbf{p}$ (with other points $\mathbf{r}$), which we denote as $\gamma_{\mathbf{p}}(\mathbf{r})$. **B:** A heatmap of $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$, the spatial correlation function with reference point $\mathbf{p}$ of the hybrid model containing the $\mathbf{p} \rightarrow \mathbf{q}$ FNP, for four different spatial widths of noisy input: **i.** $\sigma_n = 3L/80$; **ii.** $\sigma_n = L/8$; **iii.** $\sigma_n = L/4$; and **iv.** spatially uniform input ($\sigma_n \rightarrow \infty$). The σ_n values are indicated by the radius of an annotated red circle for i., ii., and iii. For each heatmap, the maximum value of the colour scale is the maximum attained value of $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$. **C:** Scatter plots of the elements of the functions $\gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{r})$ and $\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{r})$ (the spatial correlation function with reference point $\mathbf{p}$ of the geometric model) for the four σ_n values. In all scatter plots, the green dot corresponds to the point $(\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{q}), \gamma_{\mathbf{p}}^{\text{hyb}}(\mathbf{q}))$, the grey line is the linear fit between the elements, and r is the Pearson correlation.

input may also contribute to conflicting empirical findings about the influence of FNPs on cortical dynamics: while an FNP exerts its strongest influence when the input is spatially localised at its source [191–194], it has a weaker influence on spontaneous dynamics driven by spatially uniform inputs, which can be captured by a simpler geometric connectivity rule [37–39].

3.3.4 Influence of a complex network of FNPs on cortical dynamics

The results above demonstrate a single added FNP perturbs the model dynamics most strongly when the dynamics are observed on relatively short timescales ($\lesssim 30$ ms) and when the driving input is spatially localised near the source of the FNP. We now extend to a more realistic setting in which a connectome of multiple FNPs are distributed across the cortical sheet, facilitating rapid and non-local interactions across a complex distributed network. This setting allows us to explore specific network properties of FNPs that violate

the spatially homogeneous and isotropic connectivity rule commonly assumed by neural field models [36, 41, 42, 200].

The perturbation of a single FNP on stimulus-evoked dynamics being strongest over timescales $\lesssim 30$ ms (in Fig. 3.3) could be explained by a reduction in perturbation over longer timescales (> 30 ms) following the superposition of travelling waves from the source and target of the FNP. However, when multiple FNPs are present, the superposition of waves from the source and target of one FNP can be followed by a subsequent perturbation by an FNP at another location, as demonstrated in Fig. 3.2 (using an example connectome of 50 random FNPs). Motivated by this possibility, here we aimed to quantitatively test whether the timescale-dependent perturbation of a single FNP, as demonstrated above, can be extended to a connectome of multiple FNPs.

For a fixed number of FNPs N , we generated an ensemble of 200 connectomes, each with N randomly positioned FNPs with source and target positions sampled uniformly on Ω . An example of a generated connectome for each of $N = 10, 20, 50, 100$ FNPs is illustrated in Fig. 3.6A. As above, we measured the perturbation of each connectome on the stimulus-evoked dynamics using $C(t)$ (cf. Eq. (A.30), Appendix A, Sec. A.3), in a stimulus-response setting using an impulse stimulus positioned at the source location of the first randomly sampled FNP. This stimulus position was selected to ensure that the perturbation of all sampled connectomes begins from the same time (the time of stimulus onset), so that the $C(t)$ curves align and can be compared straightforwardly.

The perturbation curves $C(t)$ across ensembles of connectomes for each of $N = 10, 20, 50, 100$ are shown in Fig. 3.6B. For all t , we find that $C(t)$ tends to increase with the number of FNPs N . This finding is consistent with our expectation, since increasing N increases the expected number of rapid non-local activations during the response. $C(t)$ for all N also tends to increase to a maximum value at a time between 10 ms and 20 ms. Note that this time to maximum is larger than that attained with a single FNP (≈ 8 ms, see Fig. 3.3B), reflecting the different times from stimulus onset at which each FNP perturbs the dynamics. However, after reaching the maximum, $C(t)$ for all N tends to reduce to a smaller limiting value over longer timescales (> 30 ms). This subsequent behaviour of $C(t)$ for each N therefore demonstrates the timescale-dependent perturbation of a connectome of multiple randomly positioned FNPs on cortical dynamics, with the most salient perturbation occurring over shorter timescales ($\lesssim 30$ ms).

In our investigations of noise-driven dynamics (in Fig. 3.5), a single FNP was shown to contribute a relatively small perturbation to geometrically mediated spontaneous dynamics (driven by a noisy input of spatially uniform variance). This was explained by the inverse relationship between the perturbation ($C_{\max}$ or C_z) and the distance between the stimulus position and the source of the FNP (Appendix A, Fig. A.1). However, connectomic data indicates that FNPs tend to organise into non-random topological configurations, leading to spatially inhomogeneous and anisotropic connectivity patterns that violate the average




Figure 3.6: **A connectome of multiple randomly positioned FNPs perturbs geometric dynamics more strongly with the number of FNPs N over timescales $\lesssim 30$ ms.** **A.** For each of four different numbers of FNPs: $N = 10, 20, 50, 100$, we generated an ensemble of 200 connectomes, with source and target locations sampled uniformly at random from Ω . **B.** Realisations of $C(t)$ for each connectome in each of the four ensembles, with each ensemble’s average shown as a thicker curve (computed as the ensemble average of $C(t)$ for a fixed t).

connectivity rule typically used in neural field models [36]. A well-described characteristic of cortical structural connectivity is the non-uniform distribution of FNPs, which tends to concentrate connectivity on a subset of highly connected cortical regions called ‘hubs’ [40–42, 200]. Hub–hub connectivity makes the existence of chains of inter-linked FNPs (i.e., where the target of one FNP is proximate to the source of another FNP) more likely and thus facilitating distributed, multi-FNP interactions. Such effects would contribute differently in hub areas (where multi-FNP chains are more likely than in random networks) than in non-hub areas (where multi-FNP chains are less likely than in random networks). Accordingly, we wanted to investigate how the concentration of connectivity on specific spatial locations may lead to higher-order network effects (beyond that of a single FNP) that result in different behaviour to that observed for single FNPs in isolation.

Relative to random (unconstrained) connectivity, we explored the influence on spontaneous dynamics (quantified as $C_{\max}$) of three key connectomic constraints: the exponential distance rule (indexed by parameter λ_e), hub specificity (indexed by parameter λ_h), and rich-club specificity (indexed by parameter λ_r). For a given λ_e , λ_h , or λ_r , connectomes were generated with a rejection sampling algorithm (see Appendix A, Sec. A.3 for details on each algorithm). We explain each constraint in turn.

The exponential distance rule (EDR) captures the empirical finding that the probability of a FNP existing between two given cortical populations decays exponentially as a function

of the length of the projection [33, 34, 217]. To generate a connectome with a given EDR decay rate, we used a decay parameter λ_e to control the rate at which the probability of an FNP existing between two given points on Ω decays exponentially with its (Euclidean) separation distance (accounting for periodic boundary conditions) (see Eq. (A.32) in Appendix A, Sec. A.3 for formulation of this probability). Examples of three different connectomes with 10 FNPs and varying EDR decay rates, generated from $\lambda_e = 0, 0.5, 1$, are shown in Fig. 3.7A(i), demonstrating the increasing penalisation of long-range connections with increasing λ_e .

Hub specificity captures the empirical observation that specific hubs of the cortex tend to be particularly strongly connected and are thought to play a major role in mediating distributed information transfer across the cortex [41, 42, 218]. To generate a connectome with a given level of hub specificity, we used a scaling parameter λ_h to assign the probability that each FNP connects a population inside a hub to (or from) a population outside the same hub; ranging from the probability attained by chance from uniform sampling ($\lambda_h = 0$) to certainty ($\lambda_h = 1$) (see Eq. (A.33) in Appendix A, Sec. A.3 for formulation of this probability). For simplicity, we modelled the case of four hub areas (which also accords with some early work on hub identification in the human cortex [219, 220]) which were chosen as square areas on the domain Ω . As depicted in Figs 3.7A(ii), hub areas were allocated with centres $(L/2, L/2) + (\pm L/4, \pm L/4)$ to maximise spacing between hubs, and side length $\sqrt{34}$ since the cortex of each cerebral hemisphere was parcellated into 34 regions [219, 220]. Examples of three different connectomes with 10 FNPs for $\lambda_h = 0, 0.5, 1$ are illustrated in Fig. 3.7A(ii). At $\lambda_h = 0$ we see spatially uniform random FNP connectivity unaffected by the assignment of hub areas, whereas at $\lambda_h = 1$ we see that all FNPs connect with a hub.

Rich-club specificity captures the empirical phenomenon of dense inter-hub connectivity, resulting in a rich-club core of the connectome [40, 200, 221]. To generate a connectome with a given level of rich-club specificity, we used a parameter λ_r to assign the probability that each FNP connects two distinct hubs; ranging from the probability attained by chance from spatially uniform random connectivity ($\lambda_r = 0$) to certainty ($\lambda_r = 1$) (see Eq. (A.34) in Appendix A, Sec. A.3 for formulation of this probability). Example realisations of connectomes containing 10 FNPs for $\lambda_r = 0, 0.5, 1$ are plotted in Fig. 3.7A(iii).

We investigated how incorporating each of these three connectomic constraints (EDR, λ_e ; hub specificity, λ_h ; or rich-club specificity, λ_r) affects the perturbation to the model's spontaneous dynamics. We generated an ensemble of 200 connectomes each containing $N = 10, 20, 50, 100$ FNPs, now repeating for different values of $\lambda_e \in [0, 1]$ (for the EDR constraint), $\lambda_h \in [0, 1]$ (for the hub specificity constraint), or $\lambda_r \in [0, 1]$ (for the rich-club specificity constraint). We then quantified the average perturbation of each ensemble on the modelled spontaneous dynamics. Specifically, we measured the average perturbation of each connectome in the ensemble on stimulus-evoked dynamics, for a




Figure 3.7: The influence of a connectome of FNPs on spontaneous dynamics is either unaffected or reduced by the incorporation of distance-dependent or positional specificity-based connectomic constraints. **A.** We examine the contribution of three connectivity constraints using a parameter that quantify the extent of the constraint: $\lambda_e, \lambda_h, \lambda_r$. **i.** An exponential distance rule (EDR) constraint is captured by the parameter λ_e , which controls the rate at which connection probability decays with separation distance. **ii.** A hub specificity constraint is captured by the parameter λ_h , which controls the probability that a given FNP is incident with a hub area (coloured orange). **iii.** A rich-club specificity constraint is captured by the parameter λ_r , which controls the probability that a given FNP connects a pair of hubs. **B.** The distribution of $C_{\max}$ over an ensemble of 200 connectomes containing N FNPs each and 200 uniformly sampled stimulus positions, for increasing values of: **i.** λ_e ; **ii.** λ_h ; and **iii.** λ_r . Each distribution is represented with error bars of one standard deviation from the mean. **C.** We compared the distribution of $C_{\max}$ for different stimulus positions, across three sampled connectomes: **i.** A random (unconstrained) connectome; **ii.** A maximally hub constrained connectome ($\lambda_h = 1$); and **iii.** A maximally rich-club constrained connectome ($\lambda_r = 1$); each containing $N = 100$ FNPs. **iv.**, **v.**, and **vi.** show heatmaps of $C_{\max}$ by stimulus position of the sampled random, hub-constrained, and rich-club-constrained connectomes shown in **i.**, **ii.**, and **iii.** respectively. The stimulus positions were a 40×40 grid of equally spaced points on Ω , $\{(iL/40, jL/40) : i, j = 1, \dots, 40\}$. Source positions of all FNPs are indicated by black dots. **D.** Scatter plots of the sampled values of $C_{\max}$ and the distance between the sampled stimulus and centre of the nearest hub, for an ensemble of 200 maximally hub constrained ($\lambda_h = 1$) connectomes and 200 maximally rich-club constrained ($\lambda_r = 1$) connectomes. The mean of both distributions is annotated using a dotted horizontal line, and the average value of $C_{\max}$ over an ensemble of 200 random connectomes (with no hubs). Exponential fits to the data points from the maximally hub and rich-club constrained ensemble are annotated by the solid curves.

uniformly sampled stimulus position on Ω . This quantification approach, which draws on the relationship between noise-driven activity and stimulus-evoked dynamics demonstrated in the previous subsection, is computationally faster than the conventional approach of computing correlation structures of spontaneous dynamics (see Appendix A, Fig. A.3 for similar results of the experiments using correlation structures for a smaller ensemble sample size). Accordingly, for each sampled connectome we sampled an impulse stimulus uniformly at random from Ω , and measured $C_{\max} = \max_t \{C(t)\}$ (cf. Eq. (A.30) in Appendix A, Sec. A.3), the maximum perturbation that the connectome induces on the evoked response to this stimulus over time if the dynamics was resolved to sufficiently fine timescales.

Distributions of $C_{\max}$ as a function of λ_e , λ_h , and λ_r are plotted in Figs 3.7B(i) and (ii). For a fixed λ_e , λ_h , or λ_r , we see an increase in the average $C_{\max}$ with N (the number of FNPs), matching the increase in perturbation with N across timescales shown above for the random connectome (Fig. 3.6B). The decreasing in $C_{\max}$ for increasing λ_e —i.e., as progressively shorter FNPs are favoured—matches the expectation that fast-conducting perturbations induced by FNPs increase when the FNPs propagate activity over longer distances. Indeed, in the limiting case of $\lambda_e \rightarrow \infty$, the average length of each sampled FNP approaches zero, and in turn $C_{\max}$ approaches zero.

Interestingly, the average $C_{\max}$ did not increase with λ_h , nor with λ_r , remaining relatively constant for small $N = 10, 20$, and decreasing for larger $N = 50, 100$ (Figs 3.7B(ii) and (iii)). To better understand this result, we investigated how λ_h and λ_r affected the dependence of $C_{\max}$ on the position of the stimulus. To this end, we sampled and studied three connectomes: a random (unconstrained) connectome (Fig. 3.7C(i)); a maximally hub constrained connectome, $\lambda_h = 1$, in which every FNP connects a hub (Fig. 3.7C(ii)); and a maximally rich-club constrained connectome, $\lambda_r = 1$, in which every FNP connects a pair of hubs (Fig. 3.7C(iii)). We constructed all three connectomes with $N = 100$ FNPs, since $C_{\max}$ was most sensitive to λ_h and λ_r in Figs 3.7B(ii) and (iii) for this value of N . Figures 3.7C(iv)–(vi) show heatmaps of $C_{\max}$ by stimulus position for the random, hub-constrained ($\lambda_h = 1$), and rich-club constrained ($\lambda_r = 1$) connectomes, respectively. Compared to random connectivity, the hub constrained and rich-club constrained connectomes exhibit a more spatially heterogeneous distribution of $C_{\max}$, with higher $C_{\max}$ values for stimuli targeted on or near hub areas, and lower $C_{\max}$ values as the stimulus becomes more distant from the hubs. This spatial non-uniformity can be explained by the fact that a single FNP perturbs stimulus-evoked dynamics most strongly when the stimulus is positioned at the FNP source location (Appendix A, Fig. A.1), with the non-random connectomes displaying a higher density of FNP sources (indicated by the black dots) at hub areas. Figure 3.7D demonstrates more concretely that the spatial nonuniformity of $C_{\max}$ attained from hub specificity and rich-club specificity persists across an ensemble of 200 maximally hub constrained connectomes ($\lambda_h = 1$) and 200 maximally rich-club constrained ($\lambda_r = 1$) connectomes. Consistent with the concentration of high $C_{\max}$ near hub locations, the plot

explicitly demonstrates the decreasing relationship between $C_{\max}$ and the distance between the stimulus and the nearest hub for the hub constrained and rich-club constrained models (shown by solid curves). However, when compared to a random ensemble of connectomes (dashed horizontal lines in Fig. 3.7D) the non-random connectomes tend to have larger values of $C_{\max}$ for stimulus positions proximate to a hub (stimulus–hub distances < 0.03 m), and smaller values of $C_{\max}$ for stimulus positions distant from a hub (stimulus–hub distances > 0.03 m). Aggregating across all stimulus locations results in the hub and rich-club constrained ensembles having a lower average $C_{\max}$ than the random ensemble, indicating that the decrease in $C_{\max}$ for stimuli distant from hub locations outweighs the increase in $C_{\max}$ for stimuli proximate to hubs.

The results from Fig. 3.7C and D in the stimulus-evoked setting suggest that the hub and rich-club constraints result in a corresponding spatial specificity of the connectomic perturbation to spontaneous dynamics, which becomes concentrated on hubs. However these spatially localised perturbations are largely ‘washed out’ over larger lengthscales characteristic of the entire cortical surface, which is likely attributable to the fact that the hubs, where the perturbations increase, form a small proportion of the cortical surface ($4/34 \approx 12\%$ in our experiments). Taken together, while the hub and rich-club constraints can amplify the perturbation of FNPs on evoked responses to (specifically) hub stimulation, it has a comparatively weaker effect on the perturbation to spontaneous dynamics. This suggests that the impact of hub and rich-club connectivity on model dynamics depends on the spatial precision of the input driving the dynamics.

In summary, a connectome of multiple FNPs, that facilitates rapid and distributed communication across the cortex, perturbs the geometric dynamics more strongly with an increased number of FNPs but—as with prior results using a single FNP—these perturbations remain most prominent on timescales $\lesssim 30$ ms. Furthermore, relative to randomly embedded networks, simple models of spatially and topologically constrained connectomes (including hub and rich-club specificity) tend to concentrate the perturbative impact of FNPs on spontaneous dynamics at specific hub regions, but did not tend to elevate the aggregate perturbation over the entire cortical surface (and actually tended to reduce the spatially aggregated perturbation when the connectome contained sufficiently large numbers of FNPs, $N = 50, 100$). Our findings provide a plausible and testable account for why the spatially homogeneous and isotropic average connectivity rule of neural field models can approximate resting-state fMRI dynamics (spontaneous dynamics over timescales $\gtrsim 1$ s), despite the presence of a spatially embedded and topologically complex network of FNPs.

3.4 Discussion

This work aims to address two lines of evidence regarding the functional role of fast-conducting non-local projections (FNPs) on cortical dynamics. One range of experiments demonstrate the critical role of specific, genetically inherited FNPs in shaping the intricate spatiotemporal dynamics of cortical activity, such as the clear non-local activations seen in voltage-sensitive dye (VSD) imaging of stimulus-evoked cortical responses in mouse [191–194, 201, 202], or long-range correlations between cortical populations mediated by structural connectivity [5, 192, 222]. These empirical results, across scales and species, demonstrate the important role of FNP networks in facilitating the complex, spatially distributed cortico-cortical communication thought to underpin the brain’s remarkable cognitive functionality, despite their costs to develop, maintain and operate [40, 217, 221, 223]. On the other hand, a second collection of results demonstrates the success of neural field models which, via a spatially homogeneous and isotropic connectivity rule assumption *neglect* specific FNPs, but can nevertheless predict many aspects of cortical dynamics accurately, particularly the spontaneous fluctuations observed on the order-of-seconds timescales of fMRI [37–39]. This work sought to address the question: how can models of cortical dynamics that ‘smear out’ the anatomical specificity of FNPs to a geometric average nevertheless generate such accurate predictions of cortical dynamics?

To this end, we introduced a new mathematical model of macroscale cortical activity (on millimetre lengthscales and above) which incorporates both continuous geometric propagation (as travelling waves along the cortical surface and discrete connectomic propagation (as rapid, non-local interactions between specific remote populations mediated by FNPs). Our experiments demonstrate that the perturbations that FNPs induce on the model’s geometrically-constrained cortical dynamics is both timescale dependent—being strongest on short timescales $\lesssim 30$ ms—and dependent on the spatial precision of the input drive—being strongest for spatially localised inputs that are proximate to FNPs. These settings in which FNPs strongly shape cortical dynamics—fast dynamical responses to spatially targeted inputs—suggest rapid non-local interactions as a core mechanism for facilitating the fast and spatially distributed information processing of specific sensory input drives (e.g., specific cortical inputs from ‘core’ thalamic cells [224]). Conversely, in settings such as resting-state fMRI, which measures spontaneous dynamics over slower order-of-seconds timescales, the cortical dynamics increasingly resemble that of geometric propagation (i.e., as if there were no FNPs). Our findings thus provide a plausible mechanistic explanation for why different experiments (on different spatial and temporal scales, and examining spontaneous versus stimulus-response dynamics) can yield different conclusions regarding the role of FNPs, which our experiments indicate is highly dependent on the timescale resolved by measurement and the spatial precision of the input

drive.

Our results broadly match qualitative intuitions from the statistical mechanics of physical systems, which provides a way to understand how simple macroscopic behaviour can emerge from complex microscopic interactions. This includes the phenomenon of ‘universality’, in which physical systems with different microscopic interactions nevertheless display common macroscale behaviour (which can therefore exhibit surprising robustness to variations in microscopic physics) [225]. A commonly studied example is the relatively simple Navier–Stokes equations, which emerge from extremely complex interactions between fluid molecules at the microscale, and which well-describe the macroscopic behaviour of fluids regardless of the specific type of molecules (and their interactions). In qualitative accordance with this loose expectation, here we find that as cortical inputs become less precise in space (as in spontaneous dynamics) and are less accurately resolved in time (as in modalities like fMRI), the specific microscopic wiring configurations of individual structural fibres of the cortex matter less than the spatially averaged rules that govern them. Future work analysing multiscale patterns in neural systems [226], including relationships between scales, could investigate these scale-dependent relationships more directly and in more detail.

Our work suggests a potential reconciliation of conflicting perspectives on how best to build dynamical models of macroscale cortical activity constrained by underlying structural connectivity [16, 17, 227]. On one hand, connectome-based neural mass models treat the cortex as a discrete network of lumped nodes, representing parcellated regions, connected by edges that accumulate all structural fibres connecting the corresponding regions [57, 76, 139, 140, 146, 228]. This representation allows for explicit incorporation of specific structural connections from connectomic data, but artificially enforces a discretisation (via a parcellation) with hard boundaries on a physically continuous object. By contrast, neural field models treat the cortex as a two-dimensional continuum, and typically embed a geometric (spatially homogeneous and isotropic) structural connectivity onto this continuum [136, 199, 229]. This representation retains the continuity of cortical tissue at the millimetre lengthscale without enforcing spatial discretisation into parcels [141], but its idealised distance-dependent connectivity assumption cannot fully capture the specificity of structural connectivity, particularly that of FNP connectivity. The findings of this work suggest that the continuous geometric approximation of structural connectivity (akin to neural field models) becomes a stronger approximation for capturing the spatiotemporal patterns of spontaneous cortical dynamics on longer timescales (like that of resting-state fMRI), while the additional non-local specificity of the connectome (which can be incorporated in neural mass models) becomes more crucial for capturing the cortex’s more rapid processing of spatially localised stimuli. Note that this does not imply that the specificity of the connectome is unimportant for capturing resting-state fMRI dynamics—indeed, the further constraining of geometric connectivity with specific connections from connectomic

data can improve predictions of resting-state fMRI dynamics [196, 230, 231]. Rather, our findings highlight stimulus-evoked response dynamics as an experimentally challenging but important regime for developing and refining future dynamical models of macroscale brain activity — a regime in which the specificity of structural connectivity may play a more significant role than in resting-state fMRI [232].

An important finding of this work is that a single FNP most strongly perturbed the model’s geometric dynamics over short millisecond timescales (< 10 ms), and weakened in influence on longer timescales. This finding may have direct implications for how FNPs perturb the spatial patterns of the geometric eigenmodes of cortical activity—the natural dynamical modes of activity derived from the cortex’s geometric or spatial structure [233]. If FNPs primarily influence the model’s faster dynamics, we predict that the addition of FNPs (from connectomic data) to the cortex’s geometric structure would induce perturbations to the geometric eigenmodes that increase with eigenmode order. This prediction accords with the results of another mathematical investigation, which modelled the perturbation of a single FNP in a way that—like our model—preserved the spacetime-integrated propagated activity [234]. However, since the geometric eigenmodes of higher order contribute less power to spontaneous activity [235], we suspect that the improvements that the connectome-perturbed eigenmodes provide in reconstructing measured activity would be dominated by physiological noise and measurement artifacts, both of which are major sources of error carried by resting-state fMRI measurements [236]. If true, it would be a challenging feat to identify an alternative basis of connectome-perturbed eigenmodes that significantly improves the reconstruction accuracy of resting-state fMRI from simpler geometric eigenmodes, which is found to be the case in ongoing work [230, 237–239]. Nevertheless, it remains important for future work to investigate how the addition of FNPs (or heterogeneities from other structural brain properties) perturb the spatial patterns of geometric eigenmodes, as a means to mechanistically interpret how these perturbations impact the accuracy of eigenmode-based reconstruction of measured activity.

Our results also highlight how a fixed architecture of structural connectivity in the cortex is capable of supporting a vast array of spontaneous dynamical regimes, depending on how the system is driven. This finding underscores an underappreciated specificity of the contribution of a given FNP in shaping cortical dynamics on the spatial properties of the cortical input drive (e.g., from the thalamus). For the commonly studied case of spontaneous dynamics, typically modelled as spatially uniform cortical input [207, 228], the specific (anisotropic and inhomogenous) configurations of FNPs make surprisingly little effect on the resulting spontaneous dynamics when the variance of the noisy input drive is spatially uniform. However, our modelling highlights how the same, fixed underlying structural connectivity can support a much wider repertoire of dynamical regimes, accessible by controlling the spatial precision of the input drive. In this way, some FNPs can be relatively ‘silent’ in shaping to the dynamics to some (distal) focal inputs, despite

corresponding to major fibre tracts with ‘high weight’ in a connectomic representation. This finding also accords with mouse experiments showing that long-range correlations between FNP-connected cortical populations were amplified after stimulation of the FNP source, but quietened after stimulation of a population distant from the FNP source [192]. Our modelling thus suggests a more fluid picture of ‘the structure-function relationship’ in which controllable inputs to the cortex can be used to facilitate a broad array of different types of distributed information processing that may be functionally suited to a given situation (e.g., balancing integration versus segregation on different timescales). One possible way in which the brain may take advantage of this flexibility is by adjusting the spatial profiles of core (spatially precise) and matrix (spatially diffuse) thalamocortical projections [224, 240, 241] to bias cortical computation towards regimes that are functionally suited to a given environmental situation.

The results of this work were obtained from simulations of a novel mathematical model of macroscale activity which incorporates both geometric and discrete connectomic structural mechanisms within a single modelling framework (provided in open code that accompanies this work). The motivation to develop a new model is to improve upon existing models in literature, which face limitations in simultaneously capturing travelling wave dynamics along a continuous sheet and rapid communication dynamics through FNPs. For example, one commonly used model is the neural mass model by Spiegler and colleagues, which simulates whole-brain stimulus-evoked responses [157, 158]. Our model improves upon this model by generating finite-speed travelling waves (rather than infinite-speed diffusion), and enabling simulation of dynamics in multiple contexts other than stimulus-evoked responses. Another commonly used model is the neural field model by Jirsa et al. [203–205] which simulates cortical activity on a one-dimensional continuum (line) with a single embedded FNP between two points. While this model bears similarities to ours in generating travelling waves on a continuum, our model extends the one-dimensional geometry to a two-dimensional surface, to more accurately representing the geometry of the cortical sheet and enabling the model’s structural connectivity to have multiple FNPs—consistent with the real cortex—without facing dynamical instabilities. There is much scope for future work to further develop and refine models of brain activity that can simultaneously capture multiple mechanisms of cortico-cortical interactions, in order to better understand the types of interaction that shape cortical dynamics most strongly.

A notable limitation of this work is that our modelling results, and their interpretation, rest on specific assumptions made in constructing the model. One particular notable property of the model is its preservation of the total system activity over time (e.g., in ‘transporting’ activity from one location to another, FNPs preserve total activity). This property stands in contrast to other models, in which FNPs can propagate activity to multiple target locations and increase total activity, and nonlinearities in the local dynamics prevent dynamical instabilities arising from this increase [196]. The validity of

the activity-preserving assumption that underlies the construction of our model (and the particular mathematical treatment of this assumption, cf. [234]) could be interrogated further in future work. Our model also assumes a system gain (model parameter ν_0) that is constant over time, however the perturbations that FNPs induce on long-timescale linear cortical dynamics can be shown to be amplified when the cortical gain is increased [234]. Accordingly, future work studying FNP perturbations on cortical dynamics could investigate how these perturbations are shaped by inputs from ascending arousal systems, which are known to modulate cortical gain [48, 242].

Other simplifying assumptions include setting the connectivity strength of each FNP, c_m , to a fixed value of r^2 . Future work could better calibrate this value (e.g., using structural connectivity measurements). We also made the major simplifying assumption of modelling FNPs in the limit of instantaneous conduction, $\tau_m = 0$, which neglects time delays due to physical conduction. Although an idealisation, this choice allowed us to more straightforwardly model and qualitatively understand the qualitative effects of rapid non-local projections that arrive at some distant target location much faster than a travelling wave. Human diffusion MRI data reveals that the ratio between a structural fibre’s geodesic length (along the cortical surface) and three-dimensional length through the white matter is approximately 7 for the fibres with the largest geodesic lengths, but approximately 2.5 for the average connection [44]. Since the axonal conduction velocity in neural field models, $v = r\gamma$, is parameterised to incorporate this ratio, a reasonable choice for the time delay of a given FNP $\mathbf{a}_m \rightarrow \mathbf{b}_m$ is $\tau_m = d_m/v \times (2.5/7)$, where $d_m = \|\mathbf{a}_m - \mathbf{b}_m\|$ is the geodesic length of the FNP. Finally, the cortical surface Ω was set to have a toroidal geometry, and nonrandom FNPs on Ω were sampled algorithmically under idealised topological constraints, rather than from connectomic data. Future work could thus investigate the effects of instead modelling the real geometry Ω of a cerebral hemisphere [209] and/or aligning the wiring configurations of all FNPs with connectomic data as pursued in previous work [196, 230].

In summary, we have developed a new model of cortical activity that allowed us to investigate the role of rapid, non-local projections in perturbing the otherwise geometrically constrained dynamics. We find that the role of FNPs in shaping the dynamics depends strongly on the timescales resolved and the spatial precision of the driving input, suggesting that FNPs play a crucial role in shaping rapid stimulus-evoked response dynamics, but can be approximately neglected in favour of simpler geometric rules in slower resting-state fMRI settings. These results provide a potential explanation for why some prior results demonstrate a strong role for FNPs in mediating non-local activations, while neglecting them in favour of an isotropic geometric rule becomes a stronger assumption in others. Our paper is accompanied by open code to reproduce all results presented here [243], and we hope it will be developed further in future work to better understand different mechanisms of cortical interaction.

4. Dynamic routing on a static structure: noradrenergic gating of cortical wave propagation

CHAPTER SUMMARY

The previous chapter demonstrated that spatiotemporal variability in stimulus-evoked dynamics can arise from the position of the stimulus—specifically, depending on whether the stimulus is sufficiently proximate to engage long-distance fibres. Building on this finding, this chapter aims to address how variability in responses can arise even from a stimulus of the same position, by integrating fluctuations in the brain’s internal arousal state. It develops a novel dynamic gain model of corticothalamic dynamics to investigate how the strength, timing and topography of neuromodulatory signals from the locus coeruleus, a subcortical structure that promotes arousal, can steer the spatial propagation of stimulus-evoked waves on the cortex. Furthermore, this chapter builds upon the findings in the previous chapter, demonstrating that the influence of long-distance connectivity on stimulus-evoked dynamics is critically gated by arousal levels at stimulus onset. Note that details of the model analysed in this chapter, including its derivation and numerical treatment, are in Appendix B.

Abstract

Stimulus-evoked cortical responses range from tightly localised activations to spatially extended travelling waves, yet the mechanisms through which these distinct regimes are flexibly recruited remain unclear. Based on neuroanatomical and neuroimaging evidence, we hypothesise that the noradrenergic (NA) arousal system, originating in the locus coeruleus (LC), dynamically modulates cortical gain to shape spatiotemporal propagation under a fixed anatomical backbone. To test this hypothesis, we extend an empirically grounded corticothalamic neural field model with an LC population that delivers modulatory (gain-changing) rather than driving inputs to the cortex, and we examine evoked responses under phasic versus tonic NA modulation, under a distance-dependent geometric connectivity plus additive fast-conducting non-local projections (FNPs) over long distances. Phasic LC activation produces spatiotemporally sustained travelling waves of cortical activity,

whereas tonic activation yields temporally prolonged but spatially localised responses; without modulation, waves are weak and dissipative. Timing and topography of LC engagement gate wave-onset latency and impose anisotropy, restricting propagation to modulated territories and halting it across unmodulated gaps. In the presence of FNPs, modulation enables rapid non-local activation of remote targets and sharply lowers the gain required for long-distance routing between remote populations. These results identify NA phasic gain control as a parsimonious mechanism for trial-to-trial variability in evoked dynamics and for context-dependent privileging of specific communication pathways over a largely static connectome.

4.1 Introduction

The cerebral cortex exhibits a wide variability in its responses to incoming sensory or artificial stimuli, ranging from spatially localised activations to spatially diffuse travelling wave activations [191, 244–248]. A key question is how this diversity of stimulus responses can arise from a relatively static structural connectivity of inter-regional axonal fibres. A possible explanation to this phenomenon is a dynamic modulatory mechanism—complementing static structural connectivity—that can control the gain of excitability of constituent neural populations [49, 249]. Here, we explore whether the noradrenergic arousal system, a known chemical neuromodulatory system originating in the brainstem, can serve as a particular gain modulation mechanism that underpins dynamic variability in stimulus responses.

The propagation of stimulus-evoked activity through the axonal fibres of macroscale brain tissue has been observed in a range of experimental stimulation paradigms, including whisker deflections [191, 250], visual stimuli [251–253], transcranial magnetic stimulation [201, 254], and optogenetics [192, 253]. In the cerebral cortex, fibres tend to be distributed more densely over shorter distances [33], causing propagation to be spatiotemporally organised as travelling waves of successively depolarising populations [255–257]. There are also fibres that are distributed more sparsely over longer distances [41, 42] which can induce perturbations from a wave propagation dynamic by propagating activity to remote populations, from which additional waves are initiated [192, 194, 201, 250]. Together, these corticocortical fibres of varying lengths constitute the structural (or anatomical) connectivity of the cortex, whose reorganisation over the organism’s lifespan is negligible on the sub-day timescales of a single stimulation experiment [60, 258].

Although the cortex maintains an effectively static structural connectivity, stimulation experiments indicate that its neural response to repeated presentations of the same stimulus can exhibit substantial variability in spatiotemporal patterns between trials [191, 244–248]. For example, whisker-deflection-evoked responses can exhibit sustained spreading

processes (travelling waves and rapid long-range activations) in some trials, but can also exhibit weaker, spatially localised activations in others [191, 250]. Spreading variability has also been observed in the consciousness literature, where the sustained propagation of activity across the brain is termed ‘ignition’ and found concomitant with states of consciousness [246, 247, 259]. Furthermore, the spreading processes (when they occur) vary between trials in propagation direction [252, 257, 260–262] and latency of activation from stimulus onset [251]. Given that the structural connectivity of the cortex is relatively stable between trials, these studies therefore suggest that the brain’s evoked responses to stimuli are constrained not only by structural connectivity, but by an additional time-varying mechanism that mediates the extent to which the response spreads across the cortical surface.

Here, we propose the noradrenergic (NA) ascending arousal system as a possible mechanism that can give rise to substantial spatiotemporal diversity in evoked cortical responses to an identical input stimulus. This diversity includes localised versus spreading stimulus-evoked responses, and variability in wave propagation direction and latency. The NA system is a key brain structure for promoting states of arousal and exploratory behaviour [263], and comprises a range of nuclei in the locus coeruleus (LC), located in the pons of the brainstem, which project widely throughout the brain, including the thalamus and cerebral cortex (Fig. 4.1A). At their target sites, LC neurons can increase the *gain* or excitability of their target neurons without necessarily eliciting an action potential [48, 49]. By modulating the gain of neurons in the cortex, the LC may be able to control the variability between strong, spatially sustained responses with high gain, and weak, localised responses with low unmodulated gain. In addition, while the LC collectively targets the cortex in a diffuse manner, the specific neurons of the LC that are activated, and the times of activation, vary significantly between trials [264, 265]. Based on these empirical observations, the NA system may also support spatiotemporal selectivity of activity propagation, which results in cross-trial variability in spreading times and directions.

Understanding the contribution of the NA system to the spatiotemporal properties of stimulus-evoked dynamics is limited by the difficulty of stimulating and recording activity from the locus coeruleus (LC, the primary structure of the NA system) as well as its diffuse projection targets across the cerebral cortex. Computational modelling offers a powerful solution to this barrier, by allowing the experimenter to simulate an approximation of neural activity both with and without the influence of the NA system. In addition, these models can be used to elucidate how contributions of the NA system can shape, or are shaped by other structural mechanisms within the brain, which may lead to interesting hypotheses that can be tested in future empirical experiments.

In this study, we build on an empirically validated corticothalamic model, extending it to include the NA system which is capable of modulating the gain of the existing system




Figure 4.1: **Schematic of our novel dynamic gain corticothalamic model, which we use to investigate the role of the noradrenergic (NA) ascending arousal system on macroscale corticothalamic dynamics.** **A.** The **i.** locus coeruleus (LC) projects to **ii.** the thalamus, and **iii.** the cerebral cortex in a diffuse manner. **B.** The model tracks corticothalamic activity as a network of corticothalamic nodes. Each node tracks the mean activities of the excitatory (e) and inhibitory neurons (i) in distinct cortical regions, and the mean activities of the reticular (r) and relay (s) nuclei in their topographically aligned thalamic regions. **C.** The nodes are arranged into a circular ring, and are inter-coupled to each other, including **i.** stronger connections over shorter distances per the exponential distance rule (EDR), as well as **ii.** fast-conducting non-local projections (FNPs), which are long-distance exceptions with faster conduction than other connections, and stronger connectivity than predicted by the EDR. The additional locus coeruleus (LC) population (**iii.**) then projects to all corticothalamic nodes, with such projections comprising the NA system. **D.** Each population e, i, r, s in each node has an activation function, which relates its mean membrane potential V to its mean firing rate Q . Driver inputs to a population increase its mean membrane potential, whereas modulator inputs increase its gain (the slope of the activation function about its steady state).

in the face of an incoming stimulus. Through a range of stimulus-response experiments, we show that variability in both the NA system’s spatial targets and timing of modulation yield corresponding variability in cortical responses to the same stimulus, where particular directions of activity propagation from the stimulus site are prioritised over others. Taken together, our simulations demonstrate how the NA system as a biological mechanism can support the brain’s diverse repertoire of spatiotemporal dynamics under a fixed connectivity structure. In addition, we demonstrate how phasic gain modulation can both promote stimulus propagation along specific fibres whose endpoints lie within the modulated regions and, conversely, silence propagation along fibres that do not. In this way, the brain can privilege some physical connections over others in a context-dependent manner, thus positioning dynamic arousal not only as a mechanism for variability seen in stimulus-response experiments, but also as a putative mechanism for how the brain may vary information propagation streams vital for flexible, context-dependent computation [46].

4.2 Materials and methods

4.2.1 Model description

To investigate the role of the NA system on cortical activity, we extend a previous biophysical model of the corticothalamic system [74], to include input from the noradrenergic locus coeruleus (LC) (Fig. 4.1A-B). We provide a description of this novel model construction below, and detail the complementing model equations in Appendix B, Sec. B.1.

In the existing corticothalamic model, the cortex and thalamus are parcellated across space into a network of nodes, each of which has an architecture illustrated in Fig. 4.1B. Each node tracks the mean activities of excitatory (e) and inhibitory (i) populations in a distinct cortical region, and the mean activities of reticular (r) and relay (s) nuclei in the topographically aligned thalamic region. These populations can then project via axonal fibres to other populations in the same node, illustrated as arrows in Fig. 4.1B. The excitatory projections of each node, coloured blue in Fig. 4.1B, can additionally project to other nodes. These excitatory projections have strengths which reflect the structural connectivity between the associated cortical regions, and time delays which reflect the finite velocity at which activity propagates axonally between the two regions. Such projections tend to be stronger over shorter ranges (Fig. 4.1C(i)), a phenomenon quantified by the exponential distance rule (EDR) [33, 34]. However, certain specifically positioned projections over longer distances exhibit strengths orders of magnitude larger than predicted by the EDR due to fibre bundling, and faster conduction velocities due to myelination (Fig. 4.1C(ii)) [36, 41, 42]. We call such projections fast-conducting non-local projections (FNPs). Note that the EDR also captures projections over long distances, but

with much weaker connectivity strengths and slower conduction speeds.

Here we extend the existing corticothalamic model to include an additional single locus coeruleus (LC) population (Fig. 4.1C(iii)), which projects to all existing corticothalamic nodes. These projections, which collectively represent the noradrenergic (NA) system, provide modulator inputs to the neural populations in other nodes (dotted lines in Fig. 4.1B), which are distinct from the driver inputs provided by the existing projections in the model (solid lines in Fig. 4.1B). As shown in Fig. 4.1D, a driver input to a population increases its mean membrane potential, which in turn increases its mean firing rate in accordance with its activation function. By contrast, a modulator input to a population increases its *gain*, which we quantify as the slope of its activation function about steady state [49, 266]. Modulator inputs hence have an indirect effect on a population’s mean firing rate, effectively amplifying the impact of any simultaneous driver input.

4.2.2 Model Implementation

The parameters of the model and their values used in our investigations are given in Table 4.1. For clarity, we group the parameters based on the aspects of the model they govern: (i) the model topology: N, R ; (ii) axonal propagation: λ, γ, v, t_0 ; (iii) the dendritic response: $\alpha, \beta, \{\nu_{ab}\}$; (iv) firing rates: $Q_{\max}, \theta, \sigma$, and; (v) the modulation dynamics of population gain by the NA system: $\eta, \{g_a\}, \tau_{el}$. Existing model parameters were previously fit to power spectra of human electrophysiological recordings in an ‘eyes-open’ state [267]. Below, we describe all newly incorporated parameters (highlighted bold in Table 4.1), and justify the choice of their assigned values.

To simplify model complexity in our investigations, the corticothalamic nodes are arranged with uniform spacing into a ring—a one-dimensional simplification of a single closed cortical hemisphere (Fig. 4.1C), and assigned a default geometric structural connectivity via an exponential distance rule (EDR), where the connectivity strength between two nodes decays exponentially with their Euclidean distance in two-dimensional coordinate space. We fit this distance decay to empirical connectivity estimates described below. The choice of a ring with geometric connectivity is to conveniently visualise the model’s travelling wave propagation of activity over time, and to isolate the roles of specifically arranged FNPs separately from the influence of geometry. Qualitatively similar results can be attained if the nodes are arranged in a more realistic spherical surface, in which the travelling waves are now two-dimensional (Appendix B, Fig. B.1). We also measured internode distances in two-dimensional space, instead of along the ring, since the EDR in previous human brain modelling studies considered distances between regions in physical three-dimensional space, rather than along the cortical surface (i.e., geodesically) [196].

Using the Schaefer-400 cortical parcellation scheme, a common spatial basis for large-scale recordings and modelling [268], we assigned the ring a total of $N = 200$ nodes, each representing a parcellated region of a single cortical hemisphere. The ring was set to radius

Table 4.1: Nominal model parameter values for numerical experiments presented here. Newly incorporated parameters in the model are highlighted bold.

Parameter	Description	Value
Topology		
N	Number of Nodes	200
R	Radius of Ring	0.08 m
Axonal Propagation		
λ	EDR decay rate	101.79 m^{-1}
γ	Propagation damping rate	116 s^{-1}
v_{ij}	Default maximum axonal conduction velocity from node i to j	2 ms^{-1}
$t_0/2$	Corticothalamic / Thalamocortical Delay	0.0425 s
Dendritic Response		
α	Postsynaptic potential decay rate	83 s^{-1}
β	Postsynaptic potential rise rate	769 s^{-1}
ν_{ab}	Mean time-integrated potential of population a per unit input from b	
ν_{ee}, ν_{ie}		$7.85 \times 10^{-3} \text{ Vs}$
ν_{ei}, ν_{ii}		$-9.88 \times 10^{-3} \text{ Vs}$
ν_{es}, ν_{is}		$0.90 \times 10^{-3} \text{ Vs}$
ν_{re}		$0.21 \times 10^{-3} \text{ Vs}$
ν_{rs}		$0.06 \times 10^{-3} \text{ Vs}$
ν_{se}		$2.68 \times 10^{-3} \text{ Vs}$
ν_{sr}		$-1.31 \times 10^{-3} \text{ Vs}$
ν_{sx}		$6.60 \times 10^{-3} \text{ Vs}$
ϕ_x	Steady state of input stimulus	1 s^{-1}
Firing Response		
$Q_{\max}$	Maximum firing rate	340 s^{-1}
θ	Mean firing threshold	$12.9 \times 10^{-3} \text{ V}$
$\bar{\sigma}$	Firing threshold spread	$3.8 \times 10^{-3} \text{ V}$
Gain Modulation Dynamics		
η	Adrenoceptor activation timescale	20 s^{-1}
g_a	Mean time-integrated modulation of population a per unit input from locus coeruleus (LC)	
g_e		0 s under no gain modulation
g_i, g_r, g_s		0 s
τ_{el}	LC to Excitatory Population Conduction time	0.0425 s
ρ_l	Gain of LC population	$2.09 \times 10^3 \text{ V}^{-1} \text{ s}^{-1}$
ν_{lx}	Mean time-integrated potential of LC per unit input from stimulus	$6.60 \times 10^{-3} \text{ Vs}$ under phasic LC firing, 0 Vs under tonic LC firing

$R = 0.08$ m, half the Euclidean length of the longest structural connections seen in human empirical connectivity estimates within the left hemisphere (Appendix B, Fig. B.2A) which couple posterior to frontal regions (Appendix B, Fig. B.2B). We fit an exponential distance rule to these structural estimates within the left hemisphere using a nonlinear least-squares method to obtain an exponential decay rate of $\lambda = 101.79 \text{ m}^{-1}$ to 5 significant figures (Appendix B, Fig. B.2A), consistent with previous literature [37]. By default, the connection strength between all pairs of nodes (j, k) was set to $c_{jk} = \exp(-\lambda d_{jk})$, where d_{jk} is their Euclidean distance on the ring, and the maximum conduction velocity was set to $v_{jk} = 2 \text{ ms}^{-1}$, following reported velocities of local travelling waves in the human cortex [255]. However, in the case nodes (j, k) are connected by an FNP, c_{jk} and v_{jk} are increased (by several orders of magnitude, depending on the experiment).

The gain modulation parameters govern how the NA system, originating in the LC, shapes the gain of each population over time. At the microscale, LC neurons send noradrenaline to the adrenoceptors of efferent target neurons [242], which in turn changes the firing thresholds of the efferent neurons [269]. Aggregating this microscale process over a population, we assumed that the change in each target population's gain is proportional to the change in concentration of activated adrenoceptors, implying that the rate at which the gain changes is equal to the rate of adrenoceptor activation [49]. We set this rate to $\eta = 20 \text{ s}^{-1}$, following reports that the timescale of adrenoceptor activation is approximately 0.05 s [270]. g_a is the strength at which population a is modulated. Although the LC as an aggregate has widespread projections across the brain [271], our aim was to investigate the effect of the NA system on the spatial propagation of activity over the cerebral cortex, which is in turn facilitated by the excitatory propagators of each node. Thus, we allowed only $g = g_e$ to change, and fixed $g_i = g_r = g_s = 0$. Finally, since the LC is a subcortical structure, we assumed that the conduction time from the LC to excitatory populations is equal to the thalamocortical delay ($\tau_{el} = t_0/2$), and that the gain of the LC is equal to the steady-state gain of the relay population ($\rho_l = \rho_s$).

To simulate the model's stimulus-evoked response, we set the stimulus propagator, denoted by the red arrow in Fig. 4.1B, as an impulse (delta function). This input was fed to the relay population of one node of the ring to mimic a sensory-like stimulus that reaches a single cortical region via the thalamus, such as the input of a visual stimulus to primary visual cortex [259], or a whisker stimulation input to the corresponding barrel cortex [191]. Although the LC receives afferents from many brain regions, including the medial prefrontal cortex and hypothalamus [271], empirical observations have shown phasic activations of the LC that are time-locked to task specific stimuli during behaviour [264, 272]. Here, we approximate this relationship between the LC's phasic firing behaviour and stimuli, by coupling the stimulus to the LC, denoted by the orange arrow in Fig. 4.1B. This coupling has the same synaptic strength as the relay population of the stimulus node ($\nu_{lx} = \nu_{sx}$), aligned with the model's random connectivity assumption [267]. By contrast,

to model tonic modulation, the LC is decoupled from the stimulus ($\nu_{lx} = 0$), and set to a constant firing rate.

All model equations were solved numerically using a finite difference scheme. Second-order ordinary differential equations were solved using centred finite difference methods, while first-order equations were solved using the ‘Backward Euler’ method [273]. Details about these numerical schemes are in Appendix B, Sec. B.3.

4.3 Results

Using our novel corticothalamic model, we investigate how modulating cortical gain through the NA system—tuned by strength, cortical targets, and timing—can affect the spatiotemporal properties of the observed stimulus-evoked responses. In addition, we explore how this dynamic gain modulation impacts the contributions of long-distance connectivity (FNPs) to cortical dynamics in space and time.

4.3.1 Effect of phasic and tonic modulation on stimulus-evoked dynamics

The NA system in our model exhibits two distinct modulation mechanisms: (i) phasic modulation, where the LC transiently increases its firing rate in response to a task-relevant stimuli, and; (ii) tonic modulation, where the LC maintains an increased constant firing rate uncoupled from the stimulus [54, 274]. Here, we investigate how the model’s stimulus-evoked cortical response, denoted by variable $\phi^j(t)$ (Appendix B, Sec. B.1), is shaped by either modulation mechanism. We fixed a default geometric structural connectivity (a fully connected network with edge weights set by the EDR) and, in the presence of phasic modulation, set the stimulus to project to the LC and (the thalamic relay population of the) target node simultaneously.

We first examined the stimulus-evoked response when there was no modulation mechanism, indicated by the absence of the LC (Fig. 4.2A(i)). Figure 4.2A(ii) illustrates the evoked response ($\phi^j(t)$ for each node $j = 0, \dots, 199$ over time t) from the time of stimulus onset, when the impulse stimulus was applied to node 100. We also plot the modulation factor over time, which is the gain of the excitatory population divided by its pre-stimulus(unmodulated) value. The evoked response takes the form of a highly dissipative travelling wave emanating from the stimulated node. In addition, this travelling wave rebounds approximately every $t_0 = 85$ ms, which arises from activity propagating to the thalamus and back to the cortex via the corticothalamic loop resonance. Figure 4.2A(iii) illustrates the peak activation (maximum activity) of each node, which notably decays rapidly with distance from the stimulus node due to the dissipative nature of the travelling wave.

Next, the effect of a phasic modulation mechanism was investigated by coupling




Figure 4.2: **Phasic modulation shapes stimulus-evoked responses by enabling spatiotemporally sustained cortical travelling waves.** We examined the effect of phasic and tonic modulation on the model's stimulus-evoked dynamics, by comparing the response to the same stimulus under three different NA system configurations: **A.** The NA system was absent, so no modulation took place; **B.** The LC fires phasically in response to the stimulus, and modulates the cortex with strength $g = 0.01$ s; **C.** The LC fires tonically at a constant firing rate $Q_l = 10$ s $^{-1}$, uncoupled from the stimulus, and modulates the cortex with strength $g = 0.01$ s. The nodes were numbered from 0 to 199 clockwise on the ring, and the stimulus was applied to node 100. For all three configurations: **i.** shows the wiring configuration between populations in each node; **ii.** shows heatmaps of the stimulus-evoked response ($\phi^j(t)$ (s $^{-1}$) for each node $j = 0, \dots, 199$ over time t) and the modulation factor, both from the time of stimulus onset; and **iii.** shows the peak activation of each node j (the maximum attained value of $\phi^j(t)$ over t). In **ii.**, the top-left inset depicts all nodes arranged in a ring (blue), the stimulus node (red) and an indication of modulation applied to the nodes (gray). In **ii.** and **iii.**, $\phi_{\max}$ is the maximum attained value of $\phi^j(t)$ across j, t .

the stimulus to the LC (population l), which in turn modulated the excitatory cortical population with strength $g = 0.01\text{ s}$ (Fig. 4.2B(i)). The evoked response under this configuration (Fig. 4.2B(ii)) is more sustained over time, and propagates over a larger proportion of the ring. Figure 4.2B(iii)) shows that although peak activation is still maximum at the stimulus node, this decays more slowly with distance from this node.

Last, we investigated how tonic modulation impacts evoked responses. As illustrated in Fig. 4.2C(i), the mean firing rate of the LC was set to 10 s^{-1} , and was closest to the largest possible firing rate at which the steady state of the model remained dynamically stable (results for higher firing rates can be found in Appendix B, Fig. B.3). We found that tonic modulation caused the evoked response to sustain over time, but was not a sufficient mechanism in allowing activity to propagate through space (Fig. 4.2C(ii)). Thus, the evoked response remained spatially localised (Fig. 4.2C(iii)).

In summary, the model’s travelling wave response is relatively weak and localised in the absence of NA modulation, strong and spatiotemporally sustained under phasic modulation, and temporally sustained but spatially localised under tonic modulation. These results support phasic modulation as a candidate neurobiological mechanism for facilitating spatially sustained travelling wave responses on short sub-second timescales, while maintaining system stability over longer timescales.

4.3.2 Effect of LC time of firing and cortical targets on evoked dynamics

In the previous section, we explored the effects of phasic LC activations with a fixed timing relationship to the stimulus onset and, importantly, a uniform spatial modulation of gain, i.e., all cortical nodes were targeted. However, individual LC neurons show variable phasic activation relative to stimulus onset in empirical recordings across repeated trials [264, 265, 275, 276], as well as substantial heterogeneity in their projection profiles to the cerebral cortex, targeting unique cortical regions [277]. Together, these key features align with empirically observed variability of the directions of propagation of travelling wave responses [252, 260, 261], and variability in the timing of the responses relative to stimulus onset [251]. We next investigated whether these empirical spatiotemporal variations can be captured by adjusting the LC’s time of phasing firing and spatial distribution of its cortical projections.

First, we examined how the timing of LC firing relative to stimulus onset affected the temporal properties of the stimulus-evoked response. We introduced a delay between the time of LC phasic firing and stimulus onset (parameter τ_{lx}). As shown in Figs 4.3A(i) and A(ii), delays of 25 ms and 50 ms resulted in corresponding delays in the time at which the wave begins to travel from the stimulus node, despite a fixed stimulus onset time. However, the peak activation of each node also tends to reduce with increasing delays, and with sufficiently large delays of 400 ms and above, the modulator inputs eventually misses the stimulus (Appendix B, Fig. B.4), consistent with empirically observed timescales of

LC influence on cortical activity [274]. This result suggests that the empirically observed variability in the timing of LC phasic firing can contribute to variable temporal shifts in propagating travelling waves in evoked responses to the same stimulus.




Figure 4.3: Temporal and spatial specificity of LC modulation of cortex shapes the spatiotemporal properties of evoked cortical travelling waves. **A.** We examine how the time at which the LC fires relative to stimulus onset affected the timing of evoked travelling waves. Heatmaps of the stimulus-evoked responses and the modulation factor over time are illustrated when the time of LC firing (relative to stimulus onset) was delayed by: **i.** 25 ms; and **ii.** 50 ms. **B.** We examine how restricting the cortical targets of the LC to a subset of cortical nodes would affect the propagation of a cortical travelling wave. The modulated region of nodes is highlighted gray on the vertical axis. For all four simulations, the LC modulated the cortex with strength $g = 0.01$. The stimulus location is shown with a red dot (node 100) on the vertical axis of all plots, as is the maximum response of any node, $\phi_{\max}$ is annotated.

We next examined how varying the set of specific nodes that the LC projects to can give rise to spatial bias in the direction of propagation of the evoked response. In Fig. 4.3B(i), we restricted the set of nodes that were targeted by LC modulation to the set $\{0, 1, \dots, 120\}$. We find that, compared to the response in Fig. 4.2B(ii) (where all nodes

were modulated), the evoked travelling wave is restricted to the modulated region, and stops propagating at the boundary of this region. Hence, despite isotropic corticocortical structural connectivity, the spatially heterogeneous distribution of LC projections to the cortex introduces anisotropy to the propagation of travelling waves. Specifically, activity propagates preferentially through NA-modulated cortical areas, capturing a key observation found in experimental animal studies [260, 261]. Furthermore, as depicted in Fig. 4.3B(ii), when the modulated region was $\{0, 1, \dots, 60\} \cup \{80, \dots, 120\}$, so that a spatial gap is introduced between nodes 60 and 80, the evoked travelling wave fails to propagate in a sustained manner across this unmodulated gap, and instead dissipates at the boundaries of the modulated patches. Hence, the spatial contiguity of LC modulation is an additional determinant of the extent to which wave propagation is sustained in space. Conversely, disjointedness between modulated regions, which may result from the regional specificity of LC projection distributions to the cortex [277, 278], can give rise to compression of evoked travelling waves, which has been empirically observed at boundaries between cortical regions [252].

4.3.3 Effect of LC modulation on long-distance cortical propagation

In addition to travelling waves, spreading processes during stimulus-evoked responses can also occur along fast-conducting non-local projections (FNPs) between remote cortical populations [191, 250]. These rapid propagations can be interpreted as a perturbation of the travelling wave dynamics, by introducing non-local pairwise interactions that cannot be accounted for by the local cortical geometry [196, 279]. Here, we test whether LC gain modulation can explain this non-local perturbation during stimulus-evoked responses, by introducing a single FNP to the existing geometric structural connectivity. We set the FNP to bidirectionally connect nodes 50 and 150, so that its length ($d_{50,150} = 2R = 0.16$ m) is comparable to empirical estimates of the longest intra-hemispheric white-matter fibres in human (highlighted green in Appendix B, Fig. B.2A) which couple posterior and frontal cortical regions (visualised in Appendix B, Fig. B.2B). Starting from the geometric connectivity strength after normalisation, the strength of this bidirectional FNP was increased by a factor of 10^5 . This is because the average normalised strength of empirical estimates of the posterior-frontal fibres was found to be 8.9×10^{-4} , which is approximately 10^5 times the normalised strength of the projection between the two nodes under geometric connectivity ($\exp(-\lambda \cdot d_{50,150}) / \sum_{j=0}^{199} \exp(-\lambda \cdot d_{50,j}) \approx 10^{-8}$). Due to the uncertainty in empirical measurements of conduction times of long-distance fibres in human brains [211], we set the minimum conduction time of the FNP to the idealised case of 0 s. While propagation delays could be investigated, this approximation is sufficient for our purposes of investigating the qualitative perturbatory effect of a non-local interaction (on a timescale much faster than that of geometric propagation) on the model dynamics.

An impulse stimulus was applied to node 50 (labelled red on the top-left ring insets in

Figs 4.4A and B), an endpoint node of the FNP, to maximise the perturbation that the FNP induces on the consequent dynamics. We then compared the response between two structural connectivity configurations: geometric connectivity without the FNP (top-left inset of Fig. 4.4A(i)); and an alternate hybrid connectivity containing the additional FNP (top-left inset of Fig. 4.4B(i)). Figures 4.4A(i) and B(i) indicate that, when all N nodes are modulated with strength $g = 0.01$, the evoked responses to the stimulus are similar under both connectivities: the added FNP minimally perturbs the dynamics in this uniform low-modulation setting. When the modulation strength is increased to $g = 0.02$ or $g = 0.03$ (Figs 4.4A(ii-iii) and B(ii-iii)), the FNP perturbs the dynamics more noticeably, initiating an additional travelling wave from node 150, the other endpoint of the FNP, at time $t \gtrsim 100$ ms. To quantify the FNP's perturbation on the evoked dynamics for each value of g , we quantified the cosine dissimilarity (or cosine distance) between the spatial activity patterns of the geometric and FNP perturbed responses over time t , denoted as $C(t)$ (see Eq. (B.28) in Appendix B, Sec. B.3 for formulation) [279]. A plot of $C(t)$ in Fig. 4.4C shows that the FNP's perturbation remains minimal for $g = 0.01$, but deviates from zero for $g = 0.02, 0.03$ at time $t \approx 100$ ms, which closely corresponds to the time at which node 150 activates. These results demonstrate that LC modulation enables not only sustained travelling wave activity propagation via geometric connectivity, but also rapid non-local propagations via spatially heterogeneously distributed FNP.

Figure 4.4C shows that the FNP-induced perturbation, $C(t)$, exhibits a sharp transition from $C(t) \approx 0$ at $g = 0.01$ (i.e., the FNP has a negligible on the dynamics), to non-trivial $C(t) > 0$ dynamics at $g = 0.02$ (i.e., the FNP substantially perturbs the model dynamics), suggesting a threshold-like dependence on modulation strength. To systematically characterise this property, we examined how the temporal dynamics of $C(t)$ varies across a more finely resolved spectrum of g values. Figure 4.5A shows that $C(t)$ sharply deviates from near-zero behaviour at a threshold of $g \approx 0.015$. Above this threshold, modulation amplifies activity at the FNP's distal endpoint to detectable levels before the travelling wave arrives (e.g., Figs 4.4B(ii) and (iii)); below this threshold, the FNP continues to propagate activity but at levels too small to distinguish from travelling wave dynamics under geometric connectivity alone (e.g., Fig. 4.4B(i)). This sharp, threshold-like dependence of the functional impact of an FNP on LC modulation strength suggests that the detectability of FNP-mediated non-local propagations depends critically on whether modulation supplied to the cortex exceeds some minimum level.

While sufficient modulation strengths are needed to amplify the functional relevance of FNP, Fig. 4.4C shows that $C(t)$ decays to zero more rapidly at strength $g = 0.03$ than at $g = 0.02$, indicating that the FNP's perturbation of dynamics is more persistent over longer timescales ($t > 100$ ms) at $g = 0.02$ than at $g = 0.03$, i.e., intermediate gain rather than extreme values. This observation can be explained by examining the spatial patterns of the model's evoked response under geometric and hybrid connectivity. For




Figure 4.4: **LC modulation enables rapid non-local activity propagation facilitated by fast-conducting non-local projections (FNPs).** We examined the influence of a single FNP on stimulus-evoked dynamics for different modulation strengths, by comparing the evoked dynamics under two structural connectivity configurations. **A.** Under default geometric connectivity (top-left inset of **i.**), the evoked response to an impulse stimulus at node 50 is simulated and shown over time for three modulation strengths: **i.** $g = 0.01$, **ii.** $g = 0.02$, and **iii.** $g = 0.03$. **B.** Under an alternative hybrid connectivity that additionally contains a bidirectional FNP between nodes 50 and 150 (top-left inset of **i.**), the evoked response to an impulse stimulus at node 50 is simulated and shown over time for the same three modulation strengths. The FNP has connectivity strength increased by a factor 10^5 , and minimum conduction time of 0s. **C.** For each g , the cosine dissimilarity (or cosine distance) between the spatial activity patterns of the geometric and perturbed dynamics over time, denoted $C(t)$ (Eq. (B.28) in Appendix B, Sec. B.3), quantifies the FNP's perturbation on the geometric dynamics over time.

$g = 0.02$, the travelling wave response under geometric connectivity settles to a spatially nonuniform pattern that decays with distance from the stimulus (Fig. 4.4A(ii)), while the response under hybrid connectivity reaches a spatially uniform activation across all nodes

(Fig. 4.4B(ii)). By contrast, for $g = 0.03$, the high modulation strength drives the response under geometric connectivity to fill the entire spatial extent (Fig. 4.4A(iii)), rendering its spatial pattern over longer timescales identical to that of the response under hybrid connectivity (Fig. 4.4B(iii)). This comparison suggests a non-monotonic relationship between modulation strength and the perturbation's persistence over long timescales.

To quantify this relationship across the full spectrum of modulation strengths shown in Fig. 4.5A, we quantified the cosine dissimilarity between the zero-frequency components of the geometric and perturbed responses, denoted by C_z (Eq. (B.29) in Appendix B, Sec. B.3) [279]. The zero-frequency component of each response is a linear approximation of the response when resolved over slow, order-of-seconds timescales by modalities such as functional magnetic resonance imaging (fMRI) [214]. Figure 4.5B shows that C_z remains close to zero for $g \lesssim 0.015$, increases sharply to a maximum value after crossing this threshold, before decaying back to zero as g increases further. This non-monotonic behaviour of C_z indicates that a FNP's perturbation is most persistent over long timescales within a narrow range of modulation strengths, beyond which the perturbation becomes increasingly restricted to shorter timescales $t < 100$ ms.




Figure 4.5: Moderate LC modulation maximises the dynamical effects of FNPs over order-of-seconds timescales. **A.** The perturbation of the FNP between nodes 50 and 150 on the evoked response to a stimulus at node 50, measured as $C(t)$, is plotted over time for a finely resolved set of modulation strengths g . **B.** For each g , the cosine dissimilarity between the zero-frequency components of the geometric and perturbed response patterns, denoted C_z (Eq. (B.29) in Appendix B, Sec. B.3), quantifies the FNP's perturbation on the evoked dynamics (relative to geometric connectivity) over slow, order-of-seconds timescales.

4.3.4 Effect of heterogeneous LC modulation in the presence of long-distance connections

The previous results demonstrate two mechanisms that can shape the spatial patterns of travelling wave-evoked responses under modulation: (i) assigning spatial specificity to the modulation introduces anisotropy to the travelling wave; and (ii) adding FNPs to the structural connectivity enables additional faster propagation between remote populations (beyond what would result from geometric connectivity). Here we investigate whether, by leveraging both mechanisms, activity could propagate over long distances by modulating only the endpoints of a single FNP rather than the entire network. Such a selective routing mechanism would allow activity to reach its intended remote target without spreading to intermediate regions, providing both metabolic efficiency and communication specificity that could support adaptive cognition.

We first tested whether modulating only two nodes of the network was sufficient to drive activity propagation between them, and then tested the impact of connecting them with an FNP in addition to the modulation. Figure 4.6A illustrates the evoked response under three conditions: (i) no modulation; (ii) both nodes modulated with strength $g = 0.05$, the largest strength considered in the previous experiment; and (iii) both nodes are similarly modulated and additionally connected by a FNP of the same enhanced properties as the previous experiment. Figure 4.6A(ii) shows that while modulation amplifies the activity at the stimulus node and proximate nodes (relative to Fig. 4.6A(i)), at this strength it is insufficient to evoke a sustained response at the target node. By contrast, Fig. 4.6A(iii) shows that combining the same modulation with an FNP successfully drives sustained activation at the target and its proximate nodes, demonstrating the critical role that the FNP plays in enabling propagation at this distance and modulation strength.

To understand why the FNP was necessary to enable propagation between the two modulated nodes, we first investigated why geometric connectivity alone was insufficient to facilitate a sustained activation of the target node in Figs 4.6A(i) and (ii). Figure 4.6B(i) shows how evoked activity of the target node under geometric connectivity alone varies as a function of: D , the distance between the stimulus and target node along the ring (in units of nodes, with $D = 1$ corresponding to being adjacent to the target node); and g , the modulation strength. Here, the node activity tends to decrease with increasing distance D and increase with g , which is expected since pairwise geometric connectivity reduces with distance, and modulation amplifies driving inputs. Interestingly, nodal activity also transitions sharply from weak and to strong sustained activity at a clear distance boundary on the $D - g$ parameter space. By plotting this boundary across a wider range of g values in Appendix B, Fig. B.5, we find that at the maximum distance $D = 100$, the threshold strength was $g \approx 0.75$. Thus, this explains why $g = 0$ and $g = 0.05$ in Fig. 4.6A(i) and (ii), both indicated by the points on this surface, resulted in no sustained activation, for




Figure 4.6: Combining spatially precise modulation with FNPs enables selective long-distance routing. **A.** The evoked response to a stimulus (node 50) was simulated under three conditions: **i.** No modulation or FNP; **ii.** Stimulus node and target node (node 150) modulated with strength $g = 0.05$; and **iii.** Stimulus and target nodes modulated at $g = 0.05$, while also being connected by an FNP (with same strength and conduction time used in previous experiments). **B. i.** The peak activation of the target node (maximum attained activity over time) in the absence of the FNP is plotted as a function of its distance from the stimulus node along the ring D and the strength of modulation applied to both nodes g . The parameter combinations considered in **A(i)** and **(ii)** are annotated. **ii.** The peak activation of the target node, when connected to the stimulus node by an FNP, is plotted as a function of the FNP’s enhanced strength and conduction velocity (relative to geometric connectivity), when the stimulus and target node as set as nodes 50 and 150 respectively, and the modulation strength applied to both nodes is fixed to $g = 0.05$. The parameter combination considered in **A(iii)** is denoted as an arrow (velocity $\rightarrow \infty$).

both values of g are below this threshold.

We next examined what FNP properties—specifically connection strength and conduction velocity—enable the long-distance propagation observed in Fig. 4.6A(iii) with a modulation strength of $g = 0.05$. Figure 4.6B(ii) shows the target node activation as a function of both the FNP connection strength (relative to geometric connectivity strength) and conduction speed (relative to baseline conduction velocity). Similar to Fig. 4.6B(i), this surface exhibits a boundary in parameter space separating regions of weak and sustained target activation, and the parameter combination used in Fig. 4.6A(iii)—indicated by the arrow on the surface—is in the sustained activation region. Hence, the FNP reduces the

required modulation strength from $g \approx 0.75$ (for geometric connectivity alone) to $g \leq 0.05$. This suggests that the formation of FNPs, despite their developmental and metabolic costs, may in fact be necessary to enable otherwise metabolically expensive long-distance selective routing with physiologically plausible levels of modulation.

4.4 Discussion

The cortex shows substantial spatiotemporal variability in its evoked responses to repeated presentations of the same stimulus, ranging from weak localised responses to strong spatially sustained responses of varying propagation directions. This work provides a biophysical account of a dynamic neuromodulatory mechanism that can give rise to this variability, whilst evolving on a static structural connectome. In doing so, it demonstrates that even with a spatially homogeneous and isotropic (geometric) structural connectivity rule (which captures much of overall corticocortical connectivity), anisotropic activity propagation can be achieved with spatially heterogeneous modulation. Furthermore, our results reconcile the unique contributions of geometric and fast long-distance connectivity to this propagation, showing how these contributions are modulated by time-varying dynamics from the arousal system. Taken together, this work provides a putative mechanism through which the brain may leverage unique aspects of its physical connectivity, allowing for the formation of privileged streams of activity propagation.

This work contributes to an emerging paradigm combining insights in corticocortical interactions with subcortical neurobiological mechanisms that can shape and constrain their ongoing spatiotemporal properties [46, 48, 241, 280, 281]. Historically, a corticocentric approach has revealed how macroscale interactions across the cerebral cortex relate to structural networks via white-matter fibre connectivity [9, 13, 57, 60]. In parallel, an evolving body of work has focused on subcortico-cortical interactions, i.e., via the thalamus and brainstem, as a primary driver of ongoing cortical dynamics [280, 282]. Our results combine these distinct perspectives, revealing how both elements of brain structure contribute to spatiotemporal activity propagation across the cerebral cortex. In particular, it shows how the significance of long-distance corticocortical connectivity is regulated by noradrenergic signalling from the brain-stem, allowing specific aspects of static structure to be dominant over others and this dominance to vary in time.

We showed that long-distance connections (FNPs) have a minimum modulation threshold at which activity can propagate through. This threshold nonlinearity bears striking similarity with ignition, the empirical phenomenon that prefrontal cortical areas tend to activate in an all-or-none fashion in response to sensory stimuli [283]. While previous mechanistic explanations attribute ignition to nonlinear interactions between the stimulus and spontaneous cortical activity prior to stimulus onset [246, 259, 283, 284], our model

positions the noradrenergic system as a neurobiological mechanism distinct from the cortex that may contribute to ignition phenomena. This suggests that both the noradrenergic system and spontaneous cortical activity may simultaneously gate ignition and poses several interesting questions about which mechanism predominates in different contexts, providing a key future direction for work linking theory to empirical observations of response variability.

The threshold nonlinearity of FNP propagation also demonstrates the importance of modulation in supporting the dynamical relevance of FNPs. Recent analyses have indicated that the dynamical properties of macroscale spontaneous (resting-state) cortical activity can be well captured by a simple distance-dependent or geometric connectivity rule that neglects the specificity of FNPs [37–39, 285]. The results of this work suggest that one possible reason contributing to the aforementioned successes of geometric models is that FNPs require sufficient modulation to magnify their dynamical effects to observable levels, and this modulation is seldom attained in spontaneous settings. However, we also found that the long-timescale perturbation induced by the FNP on cortical dynamics decreases with modulation strength past this threshold. This nonlinear relationship between long-timescale perturbation and modulation strength suggests that previous discrepant hypotheses on the timescales over which FNPs influence dynamics may be modulation strength-dependent; namely that for modulation strengths near the critical threshold, the role is maximised over all timescales [234, 240], and for modulation strengths past the threshold, the role becomes increasingly restricted to shorter millisecond timescales [279]. Thus, the level of neuromodulation may be an underappreciated determinant of whether and how long FNPs shape cortical dynamics.

Beyond spatially diffuse modulation, our model predicts that sparse modulation of specific cortical regions can selectively promote propagation of activity between those regions, enabling targeted routing of activity over both short and long distances [46]. In this way, particular information streams can be privileged over others in a context dependent manner, which may be vital for adaptive computation. While we manually specified which nodes to modulate in our simulations, anatomical evidence suggests this capacity can be achieved in the physical brain by activating the numerous identified LC neurons whose axon collaterals have been found to project to sparsely distributed sets of cortical targets [277, 278]. Interestingly, a recent study reveals that the cortical targets of these sparsely projecting LC neurons exhibit functional correlations [286], consistent with the aforementioned hypothesis that these sparsely projecting LC neurons play an important role in privileging specific pairwise interactions.

The investigations in this work focused particularly on the modulatory effects of LC projections on excitatory populations in the cortex, as the internode projections that govern spatial propagation of activity in our model are excitatory. However, the cortical targets of LC neurons, given their diffuseness, are not confined to a specific neuronal type

[271]. Furthermore, while we assumed that modulation results in a positive change in the gain of target populations ($g > 0$), empirical experiments suggest that certain classes of adrenoceptors can have a negative effect [269, 287]. For example, it has been hypothesised that the α_2 adrenoceptor can silence inhibitory neurons, leading to an excitation of total activity via disinhibition [288]. Future work can therefore consider assigning differing modulation strengths (positive or negative) and adrenoceptor activation timescales to both excitatory and inhibitory populations, and investigate their consequent effects on evoked cortical dynamics. Another direction is to investigate the effects of modulating thalamic populations, given the large inputs they receive from the LC [271]. Since the thalamic populations in our model project to cortical populations in the same node, it is likely that modulating them will alter the dominant temporal frequencies of the evoked response, but introducing more realistic thalamocortical heterogeneity into the model may also allow diverse spatial effects [289].

In this model, we quantified neuromodulation as the change in population gain, which we defined as the slope of the population’s activation function about steady state. However, neuromodulation can be defined more generally between pairs of populations, as the change in the synaptic strength parameters $\{\nu_{ab}\}$. The benefit of this generalisation is that it allows feedback pathways ($r \rightarrow s, e \rightarrow s, e \rightarrow r$) and feedforward pathways ($s \rightarrow e, s \rightarrow i, s \rightarrow r$) to be modulated independently [290]. With this ability, this model can be used to shed light on how the emphasis on internal prior predictions (via feedback modulation) versus external attention (via feedforward modulation) in the face of different stimuli is reflected quantitatively in the corresponding evoked response dynamics [143]. This balance may be of importance in understanding the relationship between modulation dynamics and more subtle features of sensory stimuli that are not easily quantifiable; for example, the novelty of a stimulus in comparison to its precedents in a sequence [291].

The central motivation of this work was to put forward the noradrenergic system as a possible neurobiological mechanism underpinning trial-to-trial variability in evoked cortical activity patterns. However, the brain consists of many such subcortical systems that similarly project to the cortex in a relatively diffuse fashion [46]. For one, the noradrenergic system is one of many ascending neuromodulatory arousal systems that project to the cortex, each involving a unique neurotransmitter [292]. While all arousal systems follow similar general principles of modulation, some can project directly onto others; for example, the locus coeruleus (origin of the noradrenergic system) projects directly to the basal forebrain (origin of the cholinergic system) [293]. Arousal systems can also interact indirectly at overlapping cortical targets, wherein the joint concentration of their respective released neurotransmitters at the cortex exhibit nonlinear coupled dynamics [294, 295]. It can therefore be argued that each arousal system’s influence on cortical dynamics cannot be understood completely in isolation from the others, calling for future modelling studies to investigate the effects of multiple interacting arousal systems on

dynamics. Aside from arousal, the cortex also receives driving (and partially modulatory) inputs from diffusely projecting matrix populations in the thalamus [289], which differ from the core thalamic cells of our model which project strictly to the topographically aligned cortical region. Matrix projections can broadcast information from the thalamus to multiple cortical regions simultaneously, thus have been hypothesised by both experimental and theoretical studies to support functional integration across large spatiotemporal scales necessary for conscious awareness [240, 241, 281, 296]. Future work could investigate the subtle differences in the dynamical effects of matrix thalamic inputs versus ascending modulator inputs, and the behavioural significance of these differences.

A key limitation of our model is that the input–output organisation of the LC was manually specified to create response variability, rather than determined in accordance to explicit dynamical rules. Specifications included turning on the driving input from the external stimulus to generate phasic firing following relevant stimuli. While the mechanistic factors that control stimulus-dependent phasic activation remain incomplete [297], there is still scope for future modelling work to investigate how this activation may be influenced by the diverse range of inputs that the LC receives, apart from sensory stimuli [271]. Of potential importance are the inputs from the prefrontal cortex, which have been hypothesised to control the switching of the LC between phasic and tonic firing behaviours [274, 298–300].

Other simplifying assumptions for the sake of mechanistic understanding include setting the cortex to have a one dimensional ring geometry and assigning distance-dependent structural connectivity. Future work could investigate more realistic geometries and incorporate pairwise connectivity strengths from tractography data, so that the effects of modulation on evoked responses can be discerned between different target corticothalamic regions. Furthermore, aside from the last experiment, the FNP in our investigations was modelled in the limit of instantaneous conduction, by setting the conduction time to zero. Despite being an idealisation, this choice allowed us to model FNPs distinctly from EDR-predicted long-distance connections, without having to consider the uncertainty in empirical measurements of FNP conduction velocities [211]. Finally, the modulation strength in our experiments was assumed to be uniform across all nodes in the modulated region. More realistic spatial distributions of modulation targets can be modelled by the incorporation of adrenoceptor density maps [301], giving rise to further target-dependent effects of modulation.

5. Concluding summary and future directions

5.1 Summary of findings

An overall focus of this thesis has been on investigating how the brain’s diverse range of functionally relevant, spatially distributed macroscale dynamics can arise from an approximately fixed structure of neurons, axonal fibres, and neurotransmitter receptors. To date, this structure–function relationship has been largely studied by aiming to make model-based predictions of functional connectivity from corticocortical structural connectivity. Motivated by the limitations of this static paradigm in capturing the complexity of brain dynamics, the works of this thesis have constructed new models of brain dynamics, with the aim of interrogating how specific attributes of brain structure interact with driving stimuli and subcortical arousal to drive spatiotemporal variability in stimulus-evoked cortical dynamics on a fixed structural backbone.

Chapter 3 addresses the ongoing debate regarding the degree to which geometric constraints (relative to full connectomic information) are sufficient to capture the key properties of macroscale cortical dynamics. By developing a model that integrates geometric (spatially homogeneous and isotropic) connectivity with specific long-distance connectivity unaccounted for by geometry, this work delineates the limits of the geometric theory. Long-distance fibres were shown to drive non-geometric spatial patterns in fast, sub-second stimulus-evoked dynamics, but only when the stimulus was proximate to a long-distance connection. These findings reveal that, despite structural connectivity being largely static, the spatial patterns of measured brain dynamics may or may not exhibit rapid interactions between remote populations, contingent upon the spatial profile of the external input.

Furthermore, Chapter 4 demonstrates that, even for a fixed stimulus location, the resulting dynamics are further governed by the strength, timing, and spatial topography of neuromodulatory signals from subcortical arousal nuclei. By targeting cortical populations in a spatially heterogeneous manner, the arousal system was shown to introduce anisotropy to stimulus-evoked travelling waves, even if the underlying structural connectivity is spatially isotropic. Crucially, with sufficient modulation strength, arousal systems were found to act as a gate that enables remote interactions mediated by long-distance connections. Consequently, the interaction between the fixed connectome and the fluctuating arousal state generates a diverse functional repertoire, explaining how a single structural backbone can produce highly variable patterns of spatiotemporal responses to a given input stimulus.

Collectively, these results suggest a more nuanced view of the structure–function relationship, where the dynamical role of the brain’s different structural attributes is dependent on the timescale of interest, the spatial profile of the driving stimulus, and the internal arousal state. Further investigations into the dynamical influences of such factors beyond structural connectivity can potentially drive the development of more accurate quantitative models that explain the brain’s diverse dynamical repertoire on a static anatomical backbone.

5.2 Implications for future work

Beyond the specific contributions outlined above, the findings of this thesis also have interesting implications on how future studies can use models to study the mechanisms of, or make predictions of macroscale brain dynamics. Several implications are summarised below.

5.2.1 The mechanistic and computational value of modelling responses to stimulation

A key divergence of this thesis from modelling studies relating structural to functional connectivity, is the shift from modelling stochastic spontaneous activity to modelling deterministic stimulus-evoked responses. As elaborated in Chapter 2, this focus was motivated by the insight that the contribution of certain model mechanisms to brain dynamics—particularly those operating at fast timescales—are often obscured in a spontaneous setting, but are more clearly observed when the brain is perturbed by an impulse stimulus. Indeed, it was through this paradigm that Chapter 3 demonstrated the fast-timescale contribution of long-distance fibres in response to a proximate impulse stimulus, in contrast to the slower timescale contribution of geometry that dominates when the stimulus is more diffuse in space per spontaneous dynamics. Similarly, Chapter 4 showed that short intermittent (phasic) firing of the locus coeruleus in response to transient stimuli was shown to shape the spatiotemporal patterns of cortical dynamics over fast timescales, whereas this contribution would be ‘washed-out’ over time if the investigations focused on spontaneous dynamics. Furthermore, the stimulus-evoked approach offers computational advantages for future research. Generating a functional connectivity matrix requires simulating brain activity over long periods (e.g., 864 seconds in Pang et al. [37]), whereas simulating each evoked response in this thesis required simulation of activity over no more than 1.5 seconds. This efficiency allows for rapid exploration of parameter spaces, providing a more scalable framework for interrogating the dynamical role of different mechanisms. These advantages suggest the opportunity for future modelling studies to embrace the utility of modelling responses to stimulation in service of studying the structure–function relationship.

5.2.2 The utility of simple phenomenological models for mechanistic understanding

Throughout this thesis, a conscious choice was made to develop a minimal, phenomenological model that clearly targets the key mechanisms of interest, before extending to a more complex, physiologically realistic model constrained by a larger set of mechanisms and parameters. The benefit of this approach is demonstrated by the robustness of the findings across chapters. Investigations with the simple phenomenological model in Chapter 3 identified that long-distance fibres drive fast dynamics over sub-second timescales, and the far more complex physiological model in Chapter 4, with more biophysically realistic elements, confirmed this same principle. If the investigations of Chapter 3 used the physiological model instead of the phenomenological model, it would have been far more difficult to disentangle the dynamical influence of long-distance connectivity from that of the arousal state. For instance, since long-distance fibres were shown to require sufficient arousal neuromodulation to impact dynamics, using the physiological model at baseline arousal might have led to the erroneous conclusion that long-distance connectivity plays no dynamical role. By first establishing the dynamical capacity of long-distance fibres in the phenomenological model of Chapter 3, it became possible to interpret the complex model's behaviour in Chapter 4 not as an absence of influence of this structural mechanisms, but as a state-dependent gating of that influence. Thus, this work demonstrates that complexity, while often viewed as a requirement for reliable investigation in computational neuroscience (e.g., the Blue Brain Project [134]), comes at not only at the cost of interpretability [145], but also the risk of misinterpretation. Instead, a potentially more constructive approach to dynamic brain modelling is to start with the simplest mathematical form necessary to generate the phenomenon—allowing for a clear attribution of mechanism—and subsequently utilise more physiologically realistic models to reveal the physiological conditions required to recruit that mechanism.

5.2.3 Using exploratory modelling to generate hypotheses for future experiments

Models in this thesis were used in an exploratory sense, where investigations made novel, falsifiable predictions of dynamics that can be tested in future empirical experiments. Exploratory modelling stands in contrast to the predictive application of models in structural-functional connectivity studies, where models are fitted to existing empirical data to assess whether they can reproduce the dynamics [5, 8, 13]. While it cannot serve as a substitute for predictive modelling, exploratory modelling can be a powerful framework for generating new scientific hypotheses that can be explicitly evaluated and potentially falsified through future empirical experiments [302]. Several findings in this thesis have produced new specific, testable predictions. For example, the results of Chapter 3 propose

the hypothesis that the influence of long-distance fibres on cortical dynamics will vary with the timescales resolved by the measurement modality; for example, electrophysiological recordings measuring millisecond timescale activity should reveal stronger influences of long-range connectivity to stimulus-evoked responses than haemodynamic recordings resolving slower order-of-seconds timescale activity. This prediction could be falsified if no consistent differences are observed across such modalities of different timescale resolutions. Similarly, the hypotheses of Chapter 4 predict that trial-to-trial variability in spatiotemporal patterns of evoked cortical responses will co-vary with fluctuations in arousal system activity, such that increased firing in arousal nuclei corresponds to measurable changes in the sustainment of propagation of cortical responses. This prediction would be falsified if variability in cortical responses is found to be independent of measured arousal-related activity. In both cases, these predictions establish clear criteria for empirical validation and provide a basis for direct comparison between model-derived dynamics and observed neural data.

5.2.4 Translating models of brain stimulation to clinical application

While modelling stimulus-evoked responses in this thesis had theoretical motivations—aimed at uncovering the mechanisms underpinning brain dynamics—it also has practical motivations surrounding its translational value in clinical settings. Models of responses to stimulation can guide the deployment of more effective clinical neuromodulation technologies for treating neurodegenerative and neurological disorders, as they can help elucidate the causal relationship between stimulation parameters and their effects on patient behaviour [92]. One useful conceptualisation of this relationship is a chain of events: the initial excitation of neural elements from the stimulation device, the subsequent evoked neural response to this excitation, and finally the mapping from neural response to behavioural outcome [303]. While there is comprehensive literature on physics-based modelling of the initial excitation (reviewed in [304–307]), there has been relatively less progress in predicting the downstream neural response, largely because neural outcomes tend to vary widely between trials and subjects even with identical stimulation parameters [308–311]. Recent literature has begun to address this hurdle, by showing how the variability in responses are constrained by corresponding variability in the spontaneous state prior to stimulus onset [259, 312]. The work of this thesis extends this emerging perspective, by demonstrating that the arousal state also acts as a critical constraint, acting as a gain control mechanism that drives further variability in signal propagation. In this light, to attain more accurate predictions of the neural response to stimulation, future clinical studies could integrate real-time monitoring of arousal states (e.g., pupil diameter [282]), and incorporate the arousal state at the moment of stimulation to predict the response.

Bibliography

- [1] E. R. Kandel, J. Koester, S. Mack, and S. Siegelbaum, editors. *Principles of Neural Science*. McGraw Hill, sixth edition, 2021. ISBN 978-1-259-64223-4.
- [2] H.-J. Park and K. Friston. Structural and Functional Brain Networks: From Connections to Cognition. *Science*, 342(6158):1238411, 2013. doi: 10.1126/science.1238411.
- [3] O. Sporns. *Networks of the Brain*. MIT Press, 2016. ISBN 978-0-262-52898-6.
- [4] H. Duffau, editor. *Brain Mapping*. Springer Vienna, 2011. doi: 10.1007/978-3-7091-0723-2.
- [5] C. J. Honey, O. Sporns, L. Cammoun, et al. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009. doi: 10.1073/pnas.0811168106.
- [6] J. Cabral, E. Hugues, O. Sporns, and G. Deco. Role of local network oscillations in resting-state functional connectivity. *NeuroImage*, 57(1):130–139, 2011. doi: 10.1016/j.neuroimage.2011.04.010.
- [7] G. Deco, M. L. Kringelbach, A. Arnatkeviciute, et al. Dynamical consequences of regional heterogeneity in the brain’s transcriptional landscape. *Science Advances*, 7(29):eabf4752, 2021. doi: 10.1126/sciadv.abf4752.
- [8] A. Messé, D. Rudrauf, A. Giron, and G. Marrelec. Predicting functional connectivity from structural connectivity via computational models using MRI: An extensive comparison study. *NeuroImage*, 111:65–75, 2015. doi: 10.1016/j.neuroimage.2015.02.001.
- [9] M. Demirtaş, J. B. Burt, M. Helmer, et al. Hierarchical Heterogeneity across Human Cortex Shapes Large-Scale Neural Dynamics. *Neuron*, 101(6):1181–1194.e13, 2019. doi: 10.1016/j.neuron.2019.01.017.
- [10] G. Deco, A. R. McIntosh, K. Shen, et al. Identification of Optimal Structural Connectivity Using Functional Connectivity and Neural Modeling. *Journal of Neuroscience*, 34(23):7910–7916, 2014. doi: 10.1523/JNEUROSCI.4423-13.2014.
- [11] P. Ritter, M. Schirner, A. R. McIntosh, and V. K. Jirsa. The Virtual Brain Integrates Computational Modeling and Multimodal Neuroimaging. *Brain Connectivity*, 3(2):121–145, 2013. doi: 10.1089/brain.2012.0120.
- [12] P. H. Siu, E. Müller, V. Zerbi, et al. Extracting Dynamical Understanding From Neural-Mass Models of Mouse Cortex. *Frontiers in Computational Neuroscience*, 16, 2022. doi: 10.3389/fncom.2022.847336.
- [13] F. Melozzi, E. Bergmann, J. A. Harris, et al. Individual structural features constrain the mouse functional connectome. *Proceedings of the National Academy of Sciences*, 116(52):26961–26969, 2019. doi: 10.1073/pnas.1906694116.

- [14] L. E. Suárez, R. D. Markello, R. F. Betzel, and B. Misic. Linking Structure and Function in Macroscale Brain Networks. *Trends in Cognitive Sciences*, 24(4):302–315, 2020. doi: 10.1016/j.tics.2020.01.008.
- [15] M. D. Fox and M. E. Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8(9):700–711, 2007. doi: 10.1038/nrn2201.
- [16] G. Deco, V. K. Jirsa, P. A. Robinson, et al. The Dynamic Brain: From Spiking Neurons to Neural Masses and Cortical Fields. *PLOS Computational Biology*, 4(8):e1000092, 2008. doi: 10.1371/journal.pcbi.1000092.
- [17] M. Breakspear. Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20(3):340–352, 2017. doi: 10.1038/nn.4497.
- [18] R. M. Hutchison, T. Womelsdorf, E. A. Allen, et al. Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, 80:360–378, 2013. doi: 10.1016/j.neuroimage.2013.05.079.
- [19] J. M. Shine, M. Breakspear, P. T. Bell, et al. Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nature Neuroscience*, 22(2):289–296, 2019. doi: 10.1038/s41593-018-0312-0.
- [20] J. M. Shine, L. J. Hearne, M. Breakspear, et al. The Low-Dimensional Neural Architecture of Cognitive Complexity Is Related to Activity in Medial Thalamic Nuclei. *Neuron*, 104(5):849–855.e3, 2019. doi: 10.1016/j.neuron.2019.09.002.
- [21] Q. K. Telesford, M.-E. Lynall, J. Vettel, et al. Detection of functional brain network reconfiguration during task-driven cognitive states. *NeuroImage*, 142:198–210, 2016. doi: 10.1016/j.neuroimage.2016.05.078.
- [22] V. Zerbi, A. Floriou-Servou, M. Markicevic, et al. Rapid Reconfiguration of the Functional Connectome after Chemogenetic Locus Coeruleus Activation. *Neuron*, 103(4):702–718.e5, 2019. doi: 10.1016/j.neuron.2019.05.034.
- [23] M. Markicevic, O. Sturman, J. Bohacek, et al. Neuromodulation of striatal D1 cells shapes BOLD fluctuations in anatomically connected thalamic and cortical regions. *eLife*, 12:e78620, 2023. doi: 10.7554/eLife.78620.
- [24] G. Rabuffo, H.-A. Lokossou, Z. Li, et al. Mapping global brain reconfigurations following local targeted manipulations. *Proceedings of the National Academy of Sciences*, 122(16):e2405706122, 2025. doi: 10.1073/pnas.2405706122.
- [25] R. Polanía, M. A. Nitsche, and W. Paulus. Modulating functional connectivity patterns and topological functional organization of the human brain with transcranial direct current stimulation. *Human Brain Mapping*, 32(8):1236–1249, 2011. doi: 10.1002/hbm.21104.
- [26] B. Sehm, A. Schäfer, J. Kipping, et al. Dynamic modulation of intrinsic functional connectivity by transcranial direct current stimulation. *Journal of Neurophysiology*, 108(12):3253–3263, 2012. doi: 10.1152/jn.00606.2012.
- [27] K. Deisseroth. Optogenetics: 10 years of microbial opsins in neuroscience. *Nature Neuroscience*, 18(9):1213–1225, 2015. doi: 10.1038/nn.4091.

- [28] N. C. Rogasch and P. B. Fitzgerald. Assessing cortical network properties using TMS–EEG. *Human Brain Mapping*, 34(7):1652–1669, 2012. doi: 10.1002/hbm.22016.
- [29] S. W. Oh, J. A. Harris, L. Ng, et al. A mesoscale connectome of the mouse brain. *Nature*, 508(7495):207–214, 2014. doi: 10.1038/nature13186.
- [30] R. Kötter. Online retrieval, processing, and visualization of primate connectivity data from the CoCoMac Database. *Neuroinformatics*, 2(2):127–144, 2004. doi: 10.1385/NI:2:2:127.
- [31] D. C. Van Essen, S. M. Smith, D. M. Barch, et al. The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80:62–79, 2013. doi: 10.1016/j.neuroimage.2013.05.041.
- [32] O. Sporns, G. Tononi, and R. Kötter. The Human Connectome: A Structural Description of the Human Brain. *PLOS Computational Biology*, 1(4):e42, 2005. doi: 10.1371/journal.pcbi.0010042.
- [33] S. Horvát, R. Gămănuț, M. Ercsey-Ravasz, et al. Spatial Embedding and Wiring Cost Constrain the Functional Layout of the Cortical Network of Rodents and Primates. *PLOS Biology*, 14(7):e1002512, 2016. doi: 10.1371/journal.pbio.1002512.
- [34] M. Ercsey-Ravasz, N. T. Markov, C. Lamy, et al. A Predictive Network Model of Cerebral Cortical Connectivity Based on a Distance Rule. *Neuron*, 80(1):184–197, 2013. doi: 10.1016/j.neuron.2013.07.036.
- [35] J. A. Henderson and P. A. Robinson. Geometric Effects on Complex Network Structure in the Cortex. *Physical Review Letters*, 107(1):018102, 2011. doi: 10.1103/PhysRevLett.107.018102.
- [36] J. A. Henderson and P. A. Robinson. Using Geometry to Uncover Relationships Between Isotropy, Homogeneity, and Modularity in Cortical Connectivity. *Brain Connectivity*, 3(4):423–437, 2013. doi: 10.1089/brain.2013.0151.
- [37] J. C. Pang, K. M. Aquino, M. Oldehinkel, et al. Geometric constraints on human brain function. *Nature*, 618(7965):566–574, 2023. doi: 10.1038/s41586-023-06098-1.
- [38] J. A. Henderson, K. M. Aquino, and P. Robinson. Empirical estimation of the eigenmodes of macroscale cortical dynamics: Reconciling neural field eigenmodes and resting-state networks. *Neuroimage: Reports*, 2(3):100103, 2022. doi: 10.1016/j.ynirp.2022.100103.
- [39] P. A. Robinson, J. A. Henderson, N. C. Gabay, et al. Determination of Dynamic Brain Connectivity via Spectral Analysis. *Frontiers in Human Neuroscience*, 15:655576, 2021. doi: 10.3389/fnhum.2021.655576.
- [40] A. Arnatkeviciute, B. D. Fulcher, S. Oldham, et al. Genetic influences on hub connectivity of the human connectome. *Nature Communications*, 12(1):4237, 2021. doi: 10.1038/s41467-021-24306-2.
- [41] N. T. Markov, M. Ercsey-Ravasz, C. Lamy, et al. The role of long-range connections on the specificity of the macaque interareal cortical network. *Proceedings of the National Academy of Sciences*, 110(13):5187–5192, 2013. doi: 10.1073/pnas.1218972110.
- [42] R. F. Betzel and D. S. Bassett. Specificity and robustness of long-distance connections in weighted, interareal connectomes. *Proceedings of the National Academy of Sciences*, 115(21):E4880–E4889, 2018. doi: 10.1073/pnas.1720186115.

- [43] M. P. van den Heuvel and O. Sporns. Rich-Club Organization of the Human Connectome. *Journal of Neuroscience*, 31(44):15775–15786, 2011. doi: 10.1523/JNEUROSCI.3539-11.2011.
- [44] J. A. Henderson and P. A. Robinson. Relations Between the Geometry of Cortical Gyrification and White-Matter Network Architecture. *Brain Connectivity*, 4(2):112–130, 2014. doi: 10.1089/brain.2013.0183.
- [45] J. Mei, A. Rodriguez-Garcia, D. Takeuchi, et al. Improving the adaptive and continuous learning capabilities of artificial neural networks: Lessons from multi-neuromodulatory dynamics. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2501.06762.
- [46] E. J. Müller, B. R. Munn, and J. M. Shine. The brain that controls itself. *Current Opinion in Behavioral Sciences*, 63:101499, 2025. doi: 10.1016/j.cobeha.2025.101499.
- [47] J. M. Shine. Neuromodulatory control of complex adaptive dynamics in the brain. *Interface Focus*, 13(3):20220079, 2023. doi: 10.1098/rsfs.2022.0079.
- [48] B. R. Munn, E. J. Müller, G. Wainstein, and J. M. Shine. The ascending arousal system shapes neural dynamics to mediate awareness of cognitive states. *Nature Communications*, 12(1):6016, 2021. doi: 10.1038/s41467-021-26268-x.
- [49] J. M. Shine, E. J. Müller, B. Munn, et al. Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics. *Nature Neuroscience*, 24(6):765–776, 2021. doi: 10.1038/s41593-021-00824-6.
- [50] M. C. Avery and J. L. Krichmar. Neuromodulatory Systems and Their Interactions: A Review of Models, Theories, and Experiments. *Frontiers in Neural Circuits*, 11, 2017. doi: 10.3389/fncir.2017.00108.
- [51] G. Deco, J. Cruzat, J. Cabral, et al. Whole-Brain Multimodal Neuroimaging Model Using Serotonin Receptor Maps Explains Non-linear Functional Effects of LSD. *Current Biology*, 28(19):3065–3074.e6, 2018. doi: 10.1016/j.cub.2018.07.083.
- [52] M. L. Kringelbach, J. Cruzat, J. Cabral, et al. Dynamic coupling of whole-brain neuronal and neurotransmitter systems. *Proceedings of the National Academy of Sciences*, 117(17):9566–9576, 2020. doi: 10.1073/pnas.1921475117.
- [53] D. Depannemaecker, C. Duprat, G. Casagrande, et al. A next generation neural mass model with neuromodulation. *Journal of Computational Neuroscience*, 54(1):23–43, 2026. doi: 10.1007/s10827-026-00920-1.
- [54] G. Aston-Jones and J. D. Cohen. An Integrative Theory of Locus Coeruleus-Norepinephrine Function: Adaptive Gain and Optimal Performance. *Annual Review of Neuroscience*, 28(1):403–450, 2005. doi: 10.1146/annurev.neuro.28.061604.135709.
- [55] F. Lopes da Silva, A. van Rotterdam, P. Barts, et al. Models of Neuronal Populations: The Basic Mechanisms of Rhythmicity. In *Progress in Brain Research*, volume 45, pages 281–308. Elsevier, 1976. doi: 10.1016/S0079-6123(08)60995-4.
- [56] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, 34(4):537–541, 1995. doi: 10.1002/mrm.1910340409.

- [57] C. J. Honey, R. Kötter, M. Breakspear, and O. Sporns. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences*, 104(24):10240–10245, 2007. doi: 10.1073/pnas.0701519104.
- [58] D. Vidaurre, S. M. Smith, and M. W. Woolrich. Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, 114(48):12827–12832, 2017. doi: 10.1073/pnas.1705120114.
- [59] M. Ramezani-Panahi, G. Abrevaya, J.-C. Gagnon-Audet, et al. Generative Models of Brain Dynamics. *Frontiers in Artificial Intelligence*, 5, 2022. doi: 10.3389/frai.2022.807406.
- [60] J. Cabral, M. L. Kringelbach, and G. Deco. Functional connectivity dynamically evolves on multiple time-scales over a static structural connectome: Models and mechanisms. *NeuroImage*, 160:84–96, 2017. doi: 10.1016/j.neuroimage.2017.03.045.
- [61] J. Z. Kim and D. S. Bassett. Linear Dynamics and Control of Brain Networks. In B. He, editor, *Neural Engineering*, pages 497–518. Springer International Publishing, 2020. doi: 10.1007/978-3-030-43395-6_17.
- [62] G. Acharya, S. F. Ruf, and E. Nozari. Brain modeling for control: A review. *Frontiers in Control Engineering*, 3, 2022. doi: 10.3389/fcteg.2022.1046764.
- [63] V. K. Jirsa, R. Friedrich, H. Haken, and J. A. S. Kelso. A theoretical model of phase transitions in the human brain. *Biological Cybernetics*, 71(1):27–35, 1994. doi: 10.1007/BF00198909.
- [64] V. K. Jirsa and H. Haken. Field Theory of Electromagnetic Brain Activity. *Physical Review Letters*, 77(5):960–963, 1996. doi: 10.1103/PhysRevLett.77.960.
- [65] V. K. Jirsa and H. Haken. A derivation of a macroscopic field theory of the brain from the quasi-microscopic neural dynamics. *Physica D: Nonlinear Phenomena*, 99(4):503–526, 1997. doi: 10.1016/S0167-2789(96)00166-2.
- [66] D. Zipser. Recurrent Network Model of the Neural Mechanism of Short-Term Active Memory. *Neural Computation*, 3(2):179–193, 1991. doi: 10.1162/neco.1991.3.2.179.
- [67] N. Brunel and X.-J. Wang. Effects of Neuromodulation in a Cortical Network Model of Object Working Memory Dominated by Recurrent Inhibition. *Journal of Computational Neuroscience*, 11(1):63–85, 2001. doi: 10.1023/A:1011204814320.
- [68] B. H. Jansen and V. G. Rit. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biological Cybernetics*, 73(4):357–366, 1995. doi: 10.1007/BF00199471.
- [69] C. J. Rennie, P. A. Robinson, and J. J. Wright. Unified neurophysical model of EEG spectra and evoked potentials. *Biological Cybernetics*, 86(6):457–471, 2002. doi: 10.1007/s00422-002-0310-9.
- [70] O. David, J. M. Kilner, and K. J. Friston. Mechanisms of evoked and induced responses in MEG/EEG. *NeuroImage*, 31(4):1580–1591, 2006. doi: 10.1016/j.neuroimage.2006.02.034.
- [71] R. J. Moran, S. J. Kiebel, K. E. Stephan, et al. A neural mass model of spectral responses in electrophysiology. *NeuroImage*, 37(3):706–720, 2007. doi: 10.1016/j.neuroimage.2007.05.032.

- [72] C. C. Kerr, C. J. Rennie, and P. A. Robinson. Physiology-based modeling of cortical auditory evoked potentials. *Biological Cybernetics*, 98(2):171–184, 2008. doi: 10.1007/s00422-007-0201-1.
- [73] P. A. Robinson, P. N. Loxley, S. C. O’Connor, and C. J. Rennie. Modal analysis of corticothalamic dynamics, electroencephalographic spectra, and evoked potentials. *Physical Review E*, 63(4):041909, 2001. doi: 10.1103/PhysRevE.63.041909.
- [74] P. A. Robinson, C. J. Rennie, and D. L. Rowe. Dynamics of large-scale brain activity in normal arousal states and epileptic seizures. *Physical Review E*, 65(4):041924, 2002. doi: 10.1103/PhysRevE.65.041924.
- [75] M. Breakspear, J. A. Roberts, J. R. Terry, et al. A Unifying Explanation of Primary Generalized Seizures Through Nonlinear Brain Modeling and Bifurcation Analysis. *Cerebral Cortex*, 16(9):1296–1313, 2006. doi: 10.1093/cercor/bhj072.
- [76] G. Deco, M. L. Kringelbach, V. K. Jirsa, and P. Ritter. The dynamics of resting fluctuations in the brain: Metastability and its dynamical cortical core. *Scientific Reports*, 7(1):3095, 2017. doi: 10.1038/s41598-017-03073-5.
- [77] F. Freyer, J. A. Roberts, R. Becker, et al. Biophysical Mechanisms of Multistability in Resting-State Cortical Rhythms. *Journal of Neuroscience*, 31(17):6353–6361, 2011. doi: 10.1523/JNEUROSCI.6693-10.2011.
- [78] F. Abdelnour, M. Dayan, O. Devinsky, et al. Functional brain connectivity is predictable from anatomic network’s Laplacian eigen-structure. *NeuroImage*, 172:728–739, 2018. doi: 10.1016/j.neuroimage.2018.02.016.
- [79] E. Nozari, M. A. Bertolero, J. Stiso, et al. Macroscopic resting-state brain dynamics are best described by linear models. *Nature Biomedical Engineering*, pages 1–17, 2023. doi: 10.1038/s41551-023-01117-y.
- [80] M. L. Kringelbach, N. Jenkinson, S. L. F. Owen, and T. Z. Aziz. Translational principles of deep brain stimulation. *Nature Reviews Neuroscience*, 8(8):623–635, 2007. doi: 10.1038/nrn2196.
- [81] S. N. Flesher, J. L. Collinger, S. T. Foldes, et al. Intracortical microstimulation of human somatosensory cortex. *Science Translational Medicine*, 8(361):361ra141–361ra141, 2016. doi: 10.1126/scitranslmed.aaf8083.
- [82] B. L. Roth. DREADDs for Neuroscientists. *Neuron*, 89(4):683–694, 2016. doi: 10.1016/j.neuron.2016.01.040.
- [83] D. J. O’Shea, L. Duncker, W. Goo, et al. Direct neural perturbations reveal a dynamical mechanism for robust computation. *bioRxiv preprint*, 2022. doi: 10.1101/2022.12.16.520768.
- [84] M. Markicevic, B. D. Fulcher, C. Lewis, et al. Cortical Excitation:Inhibition Imbalance Causes Abnormal Brain Network Dynamics as Observed in Neurodevelopmental Disorders. *Cerebral Cortex*, 30(9):4922–4937, 2020. doi: 10.1093/cercor/bhaa084.
- [85] A. Sanzeni, A. Palmigiano, T. H. Nguyen, et al. Mechanisms underlying reshuffling of visual responses by optogenetic stimulation in mice and monkeys. *Neuron*, 111(24):4102–4115.e9, 2023. doi: 10.1016/j.neuron.2023.09.018.

- [86] M. Bonnard, S. Chen, J. Gaychet, et al. Resting state brain dynamics and its transients: A combined TMS-EEG study. *Scientific Reports*, 6(1):31220, 2016. doi: 10.1038/srep31220.
- [87] J. Luboeinski and T. Tchumatchenko. Nonlinear response characteristics of neural networks and single neurons undergoing optogenetic excitation. *Network Neuroscience*, 4(3):852–870, 2020. doi: 10.1162/netn_a_00154.
- [88] M. M. Churchland, J. P. Cunningham, M. T. Kaufman, et al. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012. doi: 10.1038/nature11129.
- [89] R. Chaudhuri, B. Gerçek, B. Pandey, et al. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, 2019. doi: 10.1038/s41593-019-0460-x.
- [90] J. P. Cunningham and B. M. Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, 2014. doi: 10.1038/nn.3776.
- [91] M. D. Humphries. Strong and weak principles of neural dimension reduction. *Neurons, Behavior, Data analysis, and Theory*, 5(2), 2021. doi: 10.51628/001c.24619.
- [92] D. L. Kurtin, V. Giunchiglia, J. Vohryzek, et al. Moving from phenomenological to predictive modelling: Progress and pitfalls of modelling brain stimulation in-silico. *NeuroImage*, 272:120042, 2023. doi: 10.1016/j.neuroimage.2023.120042.
- [93] E. J. Cole, K. H. Stimpson, B. S. Bentzley, et al. Stanford Accelerated Intelligent Neuromodulation Therapy for Treatment-Resistant Depression. *American Journal of Psychiatry*, 177(8):716–726, 2020. doi: 10.1176/appi.ajp.2019.19070720.
- [94] R. Ramasubbu, S. Lang, and Z. H. T. Kiss. Dosing of Electrical Parameters in Deep Brain Stimulation (DBS) for Intractable Depression: A Review of Clinical Studies. *Frontiers in Psychiatry*, 9:302, 2018. doi: 10.3389/fpsy.2018.00302.
- [95] Y. Wang, F. Hutchings, and M. Kaiser. Chapter 9 - Computational modeling of neurostimulation in brain diseases. In S. Bestmann, editor, *Progress in Brain Research*, volume 222 of *Computational Neurostimulation*, pages 191–228. Elsevier, 2015. doi: 10.1016/bs.pbr.2015.06.012.
- [96] M. Capogrosso and S. F. Lempka. A computational outlook on neurostimulation. *Bioelectronic Medicine*, 6(1):10, 2020. doi: 10.1186/s42234-020-00047-3.
- [97] E. Tang and D. S. Bassett. *Colloquium* : Control of dynamics in brain networks. *Reviews of Modern Physics*, 90(3):031003, 2018. doi: 10.1103/RevModPhys.90.031003.
- [98] M. P. van den Heuvel and H. E. Hulshoff Pol. Exploring the brain network: A review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, 20(8):519–534, 2010. doi: 10.1016/j.euroneuro.2010.03.008.
- [99] V. Sip, M. Hashemi, T. Dickscheid, et al. Characterization of regional differences in resting-state fMRI with a data-driven network model of brain dynamics. *Science Advances*, 9(11):eabq7547, 2023. doi: 10.1126/sciadv.abq7547.
- [100] G. Buzsáki. Large-scale recording of neuronal ensembles. *Nature Neuroscience*, 7(5):446–451, 2004. doi: 10.1038/nn1233.

- [101] J. M. Huntenburg, P.-L. Bazin, and D. S. Margulies. Large-Scale Gradients in Human Cortical Organization. *Trends in Cognitive Sciences*, 22(1):21–31, 2018. doi: 10.1016/j.tics.2017.11.002.
- [102] S. H. Strogatz. *Nonlinear Dynamics and Chaos with Student Solutions Manual: With Applications to Physics, Biology, Chemistry, and Engineering*. CRC press, 2018. doi: 10.1201/9780429492563.
- [103] Y. Hosaka, T. Hieda, K. Hayashi, et al. Linear models replicate the energy landscape and dynamics of resting-state brain activity. *bioRxiv preprint*, 2024. doi: 10.1101/2024.05.21.595246.
- [104] C. J. Keller, C. J. Honey, L. Entz, et al. Corticocortical Evoked Potentials Reveal Projectors and Integrators in Human Brain Networks. *Journal of Neuroscience*, 34(27):9152–9163, 2014. doi: 10.1523/JNEUROSCI.4289-13.2014.
- [105] Y. J. John, K. S. Sawyer, K. Srinivasan, et al. It’s about time: Linking dynamical systems with human neuroimaging to understand the brain. *Network Neuroscience*, 6(4):960–979, 2022. doi: 10.1162/netn_a_00230.
- [106] J. P. Rickgauer and D. W. Tank. Two-photon excitation of channelrhodopsin-2 at saturation. *Proceedings of the National Academy of Sciences*, 106(35):15025–15030, 2009. doi: 10.1073/pnas.0907084106.
- [107] L. Tong, S. Han, Y. Xue, et al. Single cell in vivo optogenetic stimulation by two-photon excitation fluorescence transfer. *iScience*, 26(10):107857, 2023. doi: 10.1016/j.isci.2023.107857.
- [108] O. A. Shemesh, D. Tanese, V. Zampini, et al. Temporally precise single-cell-resolution optogenetics. *Nature Neuroscience*, 20(12):1796–1806, 2017. doi: 10.1038/s41593-017-0018-8.
- [109] O. Yizhar, L. E. Fenno, T. J. Davidson, et al. Optogenetics in Neural Systems. *Neuron*, 71(1):9–34, 2011. doi: 10.1016/j.neuron.2011.06.004.
- [110] S. Tremblay, N. C. Rogasch, I. Premoli, et al. Clinical utility and prospective of TMS–EEG. *Clinical Neurophysiology*, 130(5):802–844, 2019. doi: 10.1016/j.clinph.2019.01.001.
- [111] I. Alekseichuk, K. Mantell, S. Shirinpour, and A. Opitz. Comparative modeling of transcranial magnetic and electric stimulation in mouse, monkey, and human. *NeuroImage*, 194:136–148, 2019. doi: 10.1016/j.neuroimage.2019.03.044.
- [112] S. M. Sternson and B. L. Roth. Chemogenetic Tools to Interrogate Brain Functions. *Annual Review of Neuroscience*, 37(1):387–407, 2014. doi: 10.1146/annurev-neuro-071013-014048.
- [113] J. Muir, J. Lopez, and R. C. Bagot. Wiring the depressed brain: Optogenetic and chemogenetic circuit interrogation in animal models of depression. *Neuropsychopharmacology*, 44(6):1013–1026, 2019. doi: 10.1038/s41386-018-0291-6.
- [114] C. L. Hill and G. J. Stephens. An Introduction to Patch Clamp Recording. In M. Dallas and D. Bell, editors, *Patch Clamp Electrophysiology: Methods and Protocols*, Methods in Molecular Biology, pages 1–19. Springer US, 2021. doi: 10.1007/978-1-0716-0818-0_1.
- [115] C. Grienberger and A. Konnerth. Imaging Calcium in Neurons. *Neuron*, 73(5):862–885, 2012. doi: 10.1016/j.neuron.2012.02.011.
- [116] F. Randi, A. K. Sharma, S. Dvali, and A. M. Leifer. Neural signal propagation atlas of *Caenorhabditis elegans*. *Nature*, 623(7986):406–414, 2023. doi: 10.1038/s41586-023-06683-4.

- [117] N. A. Steinmetz, C. Aydin, A. Lebedeva, et al. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372(6539):eabf4588, 2021. doi: 10.1126/science.abf4588.
- [118] G. Buzsáki, C. A. Anastassiou, and C. Koch. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420, 2012. doi: 10.1038/nrn3241.
- [119] J. P. Lachaux, D. Rudrauf, and P. Kahane. Intracranial EEG and human brain mapping. *Journal of Physiology-Paris*, 97(4):613–628, 2003. doi: 10.1016/j.jphysparis.2004.01.018.
- [120] M. Jedynak, A. Boyer, B. Chanteloup-Forêt, et al. Variability of Single Pulse Electrical Stimulation Responses Recorded with Intracranial Electroencephalography in Epileptic Patients. *Brain Topography*, 36(1):119–127, 2023. doi: 10.1007/s10548-022-00928-7.
- [121] J. Jha, M. Hashemi, A. N. Vattikonda, et al. Fully Bayesian estimation of virtual brain parameters with self-tuning Hamiltonian Monte Carlo. *Machine Learning: Science and Technology*, 3(3):035016, 2022. doi: 10.1088/2632-2153/ac9037.
- [122] D. R. Penas, M. Hashemi, V. K. Jirsa, and J. R. Banga. Parameter estimation in a whole-brain network model of epilepsy: Comparison of parallel global optimization solvers. *PLOS Computational Biology*, 20(7):e1011642, 2024. doi: 10.1371/journal.pcbi.1011642.
- [123] M. Hashemi, A. Ziaemehr, M. M. Woodman, et al. Simulation-based inference on virtual brain models of disorders. *Machine Learning: Science and Technology*, 5(3):035019, 2024. doi: 10.1088/2632-2153/ad6230.
- [124] N. Baldy, M. Breyton, M. M. Woodman, et al. Inference on the Macroscopic Dynamics of Spiking Neurons. *Neural Computation*, pages 1–43, 2024. doi: 10.1162/neco_a_01701.
- [125] J. Ma and J. Tang. A review for dynamics in neuron and neuronal network. *Nonlinear Dynamics*, 89(3):1569–1578, 2017. doi: 10.1007/s11071-017-3565-3.
- [126] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952. doi: 10.1113/jphysiol.1952.sp004764.
- [127] R. Gast, S. A. Solla, and A. Kennedy. Neural heterogeneity controls computations in spiking neural networks. *Proceedings of the National Academy of Sciences*, 121(3):e2311885121, 2024. doi: 10.1073/pnas.2311885121.
- [128] R. A. Stefanescu and V. K. Jirsa. A Low Dimensional Description of Globally Coupled Heterogeneous Neural Networks of Excitatory and Inhibitory Neurons. *PLOS Computational Biology*, 4(11):e1000219, 2008. doi: 10.1371/journal.pcbi.1000219.
- [129] I. D. Landau, R. Egger, V. J. Dercksen, et al. The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks. *Neuron*, 92(5):1106–1121, 2016. doi: 10.1016/j.neuron.2016.10.027.
- [130] A. Arkhipov, N. W. Gouwens, Y. N. Billeh, et al. Visual physiology of the layer 4 cortical circuit in silico. *PLOS Computational Biology*, 14(11):e1006535, 2018. doi: 10.1371/journal.pcbi.1006535.
- [131] Y. N. Billeh, B. Cai, S. L. Gratiy, et al. Systematic Integration of Structural and Functional Data into Multi-scale Models of Mouse Primary Visual Cortex. *Neuron*, 106(3):388–403.e18, 2020. doi: 10.1016/j.neuron.2020.01.040.

- [132] A. S. Aberra, A. V. Peterchev, and W. M. Grill. Biophysically realistic neuron models for simulation of cortical stimulation. *Journal of Neural Engineering*, 15(6):066023, 2018. doi: 10.1088/1741-2552/aadbb1.
- [133] A. S. Aberra, B. Wang, W. M. Grill, and A. V. Peterchev. Simulation of transcranial magnetic stimulation in head model with morphologically-realistic cortical neurons. *Brain Stimulation*, 13(1):175–189, 2020. doi: 10.1016/j.brs.2019.10.002.
- [134] H. Markram. The Blue Brain Project. *Nature Reviews Neuroscience*, 7(2):153–160, 2006. doi: 10.1038/nrn1848.
- [135] N. W. Gouwens, J. Berg, D. Feng, et al. Systematic generation of biophysically detailed models for diverse cortical neuron types. *Nature Communications*, 9(1):710, 2018. doi: 10.1038/s41467-017-02718-3.
- [136] P. A. Robinson, C. J. Rennie, and J. J. Wright. Propagation and stability of waves of electrical activity in the cerebral cortex. *Physical Review E*, 56(1):826–840, 1997. doi: 10.1103/PhysRevE.56.826.
- [137] P. A. Robinson. Propagator theory of brain dynamics. *Physical Review E*, 72(1):011904, 2005. doi: 10.1103/PhysRevE.72.011904.
- [138] A. Pathak, D. Roy, and A. Banerjee. Whole-Brain Network Models: From Physics to Bedside. *Frontiers in Computational Neuroscience*, 16:866517, 2022. doi: 10.3389/fncom.2022.866517.
- [139] R. Chaudhuri, K. Knoblauch, M.-A. Gariel, et al. A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron*, 88(2):419–431, 2015. doi: 10.1016/j.neuron.2015.09.008.
- [140] G. Deco, V. Jirsa, A. R. McIntosh, et al. Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings of the National Academy of Sciences*, 106(25):10302–10307, 2009. doi: 10.1073/pnas.0901831106.
- [141] P. A. Robinson. Physical brain connectomics. *Physical Review E*, 99(1):012421, 2019. doi: 10.1103/PhysRevE.99.012421.
- [142] P. K. Fung and P. A. Robinson. Neural field theory of calcium dependent plasticity with applications to transcranial magnetic stimulation. *Journal of Theoretical Biology*, 324:72–83, 2013. doi: 10.1016/j.jtbi.2013.01.013.
- [143] T. Babaie-Janvier and P. A. Robinson. Neural Field Theory of Evoked Response Potentials With Attentional Gain Dynamics. *Frontiers in Human Neuroscience*, 14, 2020. doi: 10.3389/fnhum.2020.00293.
- [144] D. Momi, Z. Wang, S. Parmigiani, et al. Stimulation mapping and whole-brain modeling reveal gradients of excitability and recurrence in cortical networks. *Nature Communications*, 16(1):3222, 2025. doi: 10.1038/s41467-025-58187-6.
- [145] P. A. Robinson. Ten rules for effective modeling. *NeuroImage*, 263:119622, 2022. doi: 10.1016/j.neuroimage.2022.119622.
- [146] F. Abdelnour, H. U. Voss, and A. Raj. Network diffusion accurately models the relationship between structural and functional brain connectivity networks. *NeuroImage*, 90:335–347, 2014. doi: 10.1016/j.neuroimage.2013.12.039.

- [147] A. Raj, A. Kuceyeski, and M. Weiner. A Network Diffusion Model of Disease Progression in Dementia. *Neuron*, 73(6):1204–1215, 2012. doi: 10.1016/j.neuron.2011.12.040.
- [148] L. Parkes, J. Z. Kim, J. Stiso, et al. Asymmetric signaling across the hierarchy of cytoarchitecture within the human connectome. *Science Advances*, 8(50):eadd2185, 2022. doi: 10.1126/sciadv.add2185.
- [149] S. Gu, R. F. Betzel, M. G. Mattar, et al. Optimal trajectories of brain state transitions. *NeuroImage*, 148:305–317, 2017. doi: 10.1016/j.neuroimage.2017.01.003.
- [150] B. Mišić, R. F. Betzel, A. Nematzadeh, et al. Cooperative and Competitive Spreading Dynamics on the Human Connectome. *Neuron*, 86(6):1518–1529, 2015. doi: 10.1016/j.neuron.2015.05.035.
- [151] M. Breakspear, S. Heitmann, and A. Daffertshofer. Generative Models of Cortical Oscillations: Neurobiological Implications of the Kuramoto Model. *Frontiers in Human Neuroscience*, 4, 2010. doi: 10.3389/fnhum.2010.00190.
- [152] L. L. Gollo, J. A. Roberts, and L. Cocchi. Mapping how local perturbations influence systems-level brain dynamics. *NeuroImage*, 160:97–112, 2017. doi: 10.1016/j.neuroimage.2017.01.057.
- [153] G. Deco, J. Cabral, V. M. Saenger, et al. Perturbation of whole-brain dynamics in silico reveals mechanistic differences between brain states. *NeuroImage*, 169:46–56, 2018. doi: 10.1016/j.neuroimage.2017.12.009.
- [154] G. Deco, J. Cabral, M. W. Woolrich, et al. Single or multiple frequency generators in on-going brain activity: A mechanistic whole-brain model of empirical MEG data. *NeuroImage*, 152:538–550, 2017. doi: 10.1016/j.neuroimage.2017.03.023.
- [155] M. Breakspear, J. R. Terry, and K. J. Friston. Modulation of excitatory synaptic coupling facilitates synchronization and complex dynamics in a nonlinear model of neuronal dynamics. *Neurocomputing*, 52–54:151–158, 2003. doi: 10.1016/S0925-2312(02)00740-3.
- [156] R. Larter, B. Speelman, and R. M. Worth. A coupled ordinary differential equation lattice model for the simulation of epileptic seizures. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 9(3):795–804, 1999. doi: 10.1063/1.166453.
- [157] A. Spiegler, E. C. A. Hansen, C. Bernard, et al. Selective Activation of Resting-State Networks following Focal Stimulation in a Connectome-Based Network Model of the Human Brain. *eNeuro*, 3(5):ENEURO.0068–16.2016, 2016. doi: 10.1523/ENEURO.0068-16.2016.
- [158] A. Spiegler, J. K. Abadchi, M. Mohajerani, and V. K. Jirsa. In silico exploration of mouse brain dynamics by focal stimulation reflects the organization of functional networks and sensory processing. *Network Neuroscience*, 4(3):807–851, 2020. doi: 10.1162/netn_a_00152.
- [159] L. Ljung. System Identification. In A. Procházka, J. Uhlíř, P. W. J. Rayner, and N. G. Kingsbury, editors, *Signal Analysis and Prediction*, pages 163–173. Birkhäuser, 1998. doi: 10.1007/978-1-4612-1768-8_11.
- [160] S. A. Billings. *Nonlinear System Identification: NARMAX Methods in the Time, Frequency, and Spatio-Temporal Domains*. John Wiley & Sons, 2013. ISBN 9781118535561.

- [161] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.
- [162] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A Data-Driven Approximation of the Koopman Operator: Extending Dynamic Mode Decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015. doi: 10.1007/s00332-015-9258-5.
- [163] G. Acharya, K. A. Davis, and E. Nozari. Predictive modeling of evoked intracranial eeg response to medial temporal lobe stimulation in patients with epilepsy. *Communications Biology*, 7(1):1210, 2024. doi: 10.1038/s42003-024-06859-2.
- [164] Y. Yang, S. Qiao, O. G. Sani, et al. Modelling and prediction of the dynamic responses of large-scale brain networks during direct electrical stimulation. *Nature Biomedical Engineering*, 5(4):324–345, 2021. doi: 10.1038/s41551-020-00666-w.
- [165] J.-Y. Chang, A. Pigorini, M. Massimini, et al. Multivariate autoregressive models with exogenous inputs for intracerebral responses to direct electrical stimulation of the human brain. *Frontiers in Human Neuroscience*, 6, 2012. doi: 10.3389/fnhum.2012.00317.
- [166] L. Ljung, C. Andersson, K. Tiels, and T. B. Schön. Deep Learning and System Identification. *IFAC-PapersOnLine*, 53(2):1175–1181, 2020. doi: 10.1016/j.ifacol.2020.12.1329.
- [167] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. doi: 10.1016/0893-6080(89)90020-8.
- [168] S. Vyas, M. D. Golub, D. Sussillo, and K. V. Shenoy. Computation Through Neural Population Dynamics. *Annual Review of Neuroscience*, 43(1):249–275, 2020. doi: 10.1146/annurev-neuro-092619-094115.
- [169] J. Misra, S. G. Surampudi, M. Venkatesh, et al. Learning brain dynamics for decoding and predicting individual differences. *PLOS Computational Biology*, 17(9):e1008943, 2021. doi: 10.1371/journal.pcbi.1008943.
- [170] V. Mante, D. Sussillo, K. V. Shenoy, and W. T. Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, 2013. doi: 10.1038/nature12742.
- [171] U. Güçlü and M. A. J. van Gerven. Modeling the Dynamics of Human Brain Activity with Recurrent Neural Networks. *Frontiers in Computational Neuroscience*, 11, 2017. doi: 10.3389/fncom.2017.00007.
- [172] C. Kuehn. Introduction. In C. Kuehn, editor, *Multiple Time Scale Dynamics*, Applied Mathematical Sciences, pages 1–17. Springer International Publishing, 2015. doi: 10.1007/978-3-319-12316-5_1.
- [173] G. Deco, J. Cruzat, J. Cabral, et al. Awakening: Predicting external stimulation to force transitions between different brain states. *Proceedings of the National Academy of Sciences*, 116(36):18088–18097, 2019. doi: 10.1073/pnas.1905534116.
- [174] M. J. Spivey and R. Dale. Continuous Dynamics in Real-Time Cognition. *Current Directions in Psychological Science*, 15(5):207–211, 2006. doi: 10.1111/j.1467-8721.2006.00437.x.

- [175] S. F. Muldoon, F. Pasqualetti, S. Gu, et al. Stimulation-Based Control of Dynamic Brain Networks. *PLOS Computational Biology*, 12(9):e1005076, 2016. doi: 10.1371/journal.pcbi.1005076.
- [176] Y. S. Perl, C. Pallavicini, I. P. Ipiña, et al. Perturbations in dynamical models of whole-brain activity dissociate between the level and stability of consciousness. *PLOS Computational Biology*, 17(7):e1009139, 2021. doi: 10.1371/journal.pcbi.1009139.
- [177] M. L. Kringelbach and G. Deco. Brain States and Transitions: Insights from Computational Neuroscience. *Cell Reports*, 32(10):108128, 2020. doi: 10.1016/j.celrep.2020.108128.
- [178] C. Bernard. Brain’s Best Kept Secret: Degeneracy. *eNeuro*, 10(11), 2023. doi: 10.1523/ENEURO.0430-23.2023.
- [179] S. L. Bressler and V. Menon. Large-scale brain networks in cognition: Emerging methods and principles. *Trends in Cognitive Sciences*, 14(6):277–290, 2010. doi: 10.1016/j.tics.2010.04.004.
- [180] P. M. Thompson, T. D. Cannon, K. L. Narr, et al. Genetic influences on brain structure. *Nature Neuroscience*, 4(12):1253–1258, 2001. doi: 10.1038/nm758.
- [181] A. Schüz and V. Braitenberg. The Human Cortical White Matter: Quantitative Aspects of Cortico-Cortical Long-Range Connectivity. In *Cortical Areas*. CRC Press, 2002. ISBN 978-0-429-21910-8.
- [182] E. Bullmore and O. Sporns. The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336–349, 2012. doi: 10.1038/nrn3214.
- [183] R. W. Sperry. Cerebral Organization and Behavior. *Science*, 133(3466):1749–1757, 1961. doi: 10.1126/science.133.3466.1749.
- [184] N. Geschwind. Disconnexion syndromes in animals and man. I. *Brain: A Journal of Neurology*, 88(2):237–294, 1965. doi: 10.1093/brain/88.2.237.
- [185] C. Wernicke. *Der aphasische symptomengruppe*. 1874. ISBN 978-3-642-65950-8.
- [186] K. Goldstein. *The Organism: A Holistic Approach to Biology Derived from Pathological Data in Man*. The Organism: A Holistic Approach to Biology Derived from Pathological Data in Man. Zone Books, 1995. ISBN 978-0-942299-96-0.
- [187] M. Catani and D. H. ffytche. The rises and falls of disconnection syndromes. *Brain: A Journal of Neurology*, 128(Pt 10):2224–2239, 2005. doi: 10.1093/brain/awh622.
- [188] J. X. O’Reilly, P. L. Croxson, S. Jbabdi, et al. Causal effect of disconnection lesions on interhemispheric functional connectivity in rhesus monkeys. *Proceedings of the National Academy of Sciences*, 110(34):13982–13987, 2013. doi: 10.1073/pnas.1305062110.
- [189] C. M. Filley and R. D. Fields. White matter and cognition: Making the connection. *Journal of Neurophysiology*, 116(5):2093–2104, 2016. doi: 10.1152/jn.00221.2016.
- [190] D. S. Grayson, E. Bliss-Moreau, C. J. Machado, et al. The Rhesus Monkey Connectome Predicts Disrupted Functional Networks Resulting from Pharmacogenetic Inactivation of the Amygdala. *Neuron*, 91(2):453–466, 2016. doi: 10.1016/j.neuron.2016.06.005.
- [191] I. Ferezou, F. Haiss, L. J. Gentet, et al. Spatiotemporal Dynamics of Cortical Sensorimotor Integration in Behaving Mice. *Neuron*, 56(5):907–923, 2007. doi: 10.1016/j.neuron.2007.10.007.

- [192] M. H. Mohajerani, A. W. Chan, M. Mohsenvand, et al. Spontaneous cortical activity alternates between motifs defined by regional axonal projections. *Nature Neuroscience*, 16(10):1426–1435, 2013. doi: 10.1038/nm.3499.
- [193] D. H. Lim, M. H. Mohajerani, J. LeDue, et al. In vivo Large-Scale Cortical Mapping Using Channelrhodopsin-2 Stimulation in Transgenic Mice Reveals Asymmetric and Reciprocal Relationships between Cortical Areas. *Frontiers in Neural Circuits*, 6, 2012. doi: 10.3389/fncir.2012.00011.
- [194] Y. Liu, P. Bech, K. Tamura, et al. Cell class-specific long-range axonal projections of neurons in mouse whisker-related somatosensory cortices. *eLife*, 13:RP97602, 2024. doi: 10.7554/eLife.97602.
- [195] S. Hubatz, G. Hucher, D. E. Shulz, and I. Férézou. Spatiotemporal properties of whisker-evoked tactile responses in the mouse secondary somatosensory cortex. *Scientific Reports*, 10(1):763, 2020. doi: 10.1038/s41598-020-57684-6.
- [196] G. Deco, Y. Sanz Perl, P. Vuust, et al. Rare long-range cortical connections enhance human information processing. *Current Biology*, 31(20):4436–4448.e5, 2021. doi: 10.1016/j.cub.2021.07.064.
- [197] C. I. Bargmann and E. Marder. From the connectome to brain function. *Nature Methods*, 10(6):483–490, 2013. doi: 10.1038/nmeth.2451.
- [198] A. Fornito and E. T. Bullmore. Connectomics: A new paradigm for understanding brain disease. *European Neuropsychopharmacology*, 25(5):733–748, 2015. doi: 10.1016/j.euroneuro.2014.02.011.
- [199] P. L. Nunez and B. A. Cutillo, editors. *Neocortical Dynamics and Human EEG Rhythms*. Oxford University Press, 1995. ISBN 978-0-19-505728-7.
- [200] M. P. Van Den Heuvel and O. Sporns. Rich-Club Organization of the Human Connectome. *The Journal of Neuroscience*, 31(44):15775–15786, 2011. doi: 10.1523/JNEUROSCI.3539-11.2011.
- [201] M. Massimini, F. Ferrarelli, S. Sarasso, and G. Tononi. Cortical mechanisms of loss of consciousness: Insight from TMS/EEG studies. *Archives Italiennes de Biologie*, 150(2/3):44–55, 2012. doi: 10.4449/aib.v150i2.1361.
- [202] C. Seguin, M. Jedynak, O. David, et al. Communication dynamics in the human connectome shape the cortex-wide propagation of direct electrical stimulation. *Neuron*, 2023. doi: 10.1016/j.neuron.2023.01.027.
- [203] M. R. Qubbaj and V. K. Jirsa. Neural Field Dynamics with Heterogeneous Connection Topology. *Physical Review Letters*, 98(23):238102, 2007. doi: 10.1103/PhysRevLett.98.238102.
- [204] V. K. Jirsa. Neural Field Dynamics with Local and Global Connectivity and Time Delay. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 367(1891):1131–1143, 2009. doi: 10.1098/rsta.2008.0260.
- [205] V. K. Jirsa and J. A. S. Kelso. Spatiotemporal pattern formation in neural systems with heterogeneous connection topologies. *Physical Review E*, 62(6):8462–8465, 2000. doi: 10.1103/PhysRevE.62.8462.
- [206] K. N. Mukta, X. Gao, and P. A. Robinson. Neural field theory of evoked response potentials in a spherical brain geometry. *Physical Review E*, 99(6):062304, 2019. doi: 10.1103/PhysRevE.99.062304.

- [207] P. Sanz-Leon, P. A. Robinson, S. A. Knock, et al. NfTsim: Theory and Simulation of Multiscale Neural Field Dynamics. *PLOS Computational Biology*, 14(8):e1006387, 2018. doi: 10.1371/journal.pcbi.1006387.
- [208] N. I. Kavallaris and T. Suzuki. *Non-Local Partial Differential Equations for Engineering and Biology: Mathematical Modeling and Analysis*, volume 31 of *Mathematics for Industry*. Springer International Publishing, 2018. doi: 10.1007/978-3-319-67944-0.
- [209] N. C. Gabay and P. A. Robinson. Cortical geometry as a determinant of brain activity eigenmodes: Neural field analysis. *Physical Review E*, 96(3):032413, 2017. doi: 10.1103/PhysRevE.96.032413.
- [210] P. Robinson, C. Rennie, D. Rowe, and S. O’Connor. Estimation of multiscale neurophysiologic parameters by electroencephalographic means. *Human Brain Mapping*, 23(1):53–72, 2004. doi: 10.1002/hbm.20032.
- [211] H. A. Swadlow and S. G. Waxman. Axonal conduction delays. *Scholarpedia*, 7(6):1451, 2012. doi: 10.4249/scholarpedia.1451.
- [212] J. A. Roberts, A. Perry, A. R. Lord, et al. The contribution of geometry to the human connectome. *NeuroImage*, 124:379–393, 2016. doi: 10.1016/j.neuroimage.2015.09.009.
- [213] P. A. Robinson, S. Sarkar, G. M. Pandejee, and J. A. Henderson. Determination of effective brain connectivity from functional connectivity with application to resting state connectivities. *Physical Review E*, 90(1):012707, 2014. doi: 10.1103/PhysRevE.90.012707.
- [214] P. A. Robinson, P. M. Drysdale, H. Van der Merwe, et al. BOLD responses to stimuli: Dependence on frequency, stimulus form, amplitude, and repetition rate. *NeuroImage*, 31(2):585–599, 2006. doi: 10.1016/j.neuroimage.2005.12.026.
- [215] L. Q. Uddin. Bring the Noise: Reconceptualizing Spontaneous Neural Activity. *Trends in Cognitive Sciences*, 24(9):734–746, 2020. doi: 10.1016/j.tics.2020.06.003.
- [216] P. A. Robinson. Interrelating anatomical, effective, and functional brain connectivity using propagators and neural field theory. *Physical Review E*, 85(1):011912, 2012. doi: 10.1103/PhysRevE.85.011912.
- [217] B. D. Fulcher and A. Fornito. A transcriptional signature of hub connectivity in the mouse connectome. *Proceedings of the National Academy of Sciences*, 113(5):1435–1440, 2016. doi: 10.1073/pnas.1513302113.
- [218] R. F. Betzel, A. Avena-Koenigsberger, J. Goñi, et al. Generative models of the human connectome. *NeuroImage*, 124:1054–1064, 2016. doi: 10.1016/j.neuroimage.2015.09.041.
- [219] P. Hagmann, L. Cammoun, X. Gigandet, et al. Mapping the Structural Core of Human Cerebral Cortex. *PLOS Biology*, 6(7):e159, 2008. doi: 10.1371/journal.pbio.0060159.
- [220] Y. Iturria-Medina, R. C. Sotero, E. J. Canales-Rodríguez, et al. Studying the human brain anatomical network via diffusion-weighted MRI and Graph Theory. *NeuroImage*, 40(3):1064–1076, 2008. doi: 10.1016/j.neuroimage.2007.10.060.
- [221] M. P. van den Heuvel, R. S. Kahn, J. Goñi, and O. Sporns. High-cost, high-capacity backbone for global brain communication. *Proceedings of the National Academy of Sciences*, 109(28):11372–11377, 2012. doi: 10.1073/pnas.1203593109.

- [222] M. D. Greicius, K. Supekar, V. Menon, and R. F. Dougherty. Resting-State Functional Connectivity Reflects Structural Connectivity in the Default Mode Network. *Cerebral Cortex*, 19(1):72–78, 2009. doi: 10.1093/cercor/bhn059.
- [223] P. E. Vértes, T. Rittman, K. J. Whitaker, et al. Gene transcription profiles associated with inter-modular hubs and connection distance in human functional magnetic resonance imaging networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1705):20150362, 2016. doi: 10.1098/rstb.2015.0362.
- [224] J. M. Shine, L. D. Lewis, D. D. Garrett, and K. Hwang. The impact of the human thalamus on brain-wide information processing. *Nature Reviews Neuroscience*, 24(7):416–430, 2023. doi: 10.1038/s41583-023-00701-0.
- [225] J. P. Sethna. *Statistical Mechanics: Entropy, Order Parameters, and Complexity*. Oxford Master Series in Physics. Oxford University Press, second edition, 2021. ISBN 978-0-19-886525-4.
- [226] B. R. Munn, E. J. Müller, I. Favre-Bulle, et al. Multiscale organization of neuronal activity unifies scale-dependent theories of brain function. *Cell*, 187(25):7303–7313.e15, 2024. doi: 10.1016/j.cell.2024.10.004.
- [227] J. D. Griffiths, S. P. Bastiaens, and N. Kaboodvand. Whole-Brain Modelling: Past, Present, and Future. In M. Giugliano, M. Negrello, and D. Linaro, editors, *Computational Modelling of the Brain: Modelling Approaches to Cells, Circuits and Networks*, pages 313–355. Springer International Publishing, 2022. doi: 10.1007/978-3-030-89439-9_13.
- [228] A. Ghosh, Y. Rho, A. R. McIntosh, et al. Noise during Rest Enables the Exploration of the Brain’s Dynamic Repertoire. *PLOS Computational Biology*, 4(10):e1000196, 2008. doi: 10.1371/journal.pcbi.1000196.
- [229] V. Jirsa, K. Jantzen, A. Fuchs, and J. Kelso. Spatiotemporal forward solution of the EEG and MEG using network modeling. *IEEE Transactions on Medical Imaging*, 21(5):493–504, 2002. doi: 10.1109/TMI.2002.1009385.
- [230] J. Vohryzek, Y. Sanz-Perl, M. L. Kringelbach, and G. Deco. Human brain dynamics are shaped by rare long-range connections over and above cortical geometry. *Proceedings of the National Academy of Sciences*, 122(1):e2415102122, 2025. doi: 10.1073/pnas.2415102122.
- [231] G. Deco, Y. S. Perl, and M. L. Kringelbach. Non-local Schrödinger diffusion model reveals mechanisms of critical brain dynamics. *Cell Reports Physical Science*, page 102663, 2025. doi: 10.1016/j.xcrp.2025.102663.
- [232] R. Maran, E. J. Müller, and B. D. Fulcher. Analyzing the brain’s dynamic response to targeted stimulation using generative modeling. *Network Neuroscience*, 9(1):237–258, 2025. doi: 10.1162/netn_a_00433.
- [233] P. Robinson, X. Zhao, K. Aquino, et al. Eigenmodes of brain activity: Neural field theory predictions and comparison with experiment. *NeuroImage*, 142:79–98, 2016. doi: 10.1016/j.neuroimage.2016.04.050.
- [234] P. A. Robinson. Near-critical corticothalamic eigenmodes: Effects of nonuniform connectivity on modes, activity, and communication channels. *Physical Review E*, 111(1):014404, 2025. doi: 10.1103/PhysRevE.111.014404.

- [235] E. J. Müller, B. R. Munn, K. M. Aquino, et al. The music of the hemispheres: Cortical eigenmodes as a physical basis for large-scale brain activity and connectivity patterns. *Frontiers in Human Neuroscience*, 16:1062487, 2022. doi: 10.3389/fnhum.2022.1062487.
- [236] T. T. Liu. Noise contributions to the fMRI signal: An overview. *NeuroImage*, 143:141–151, 2016. doi: 10.1016/j.neuroimage.2016.09.008.
- [237] L. Sina Mansour, H. Behjat, D. V. De Ville, et al. Eigenmodes of the brain: revisiting connectomics and geometry. *bioRxiv preprint*, 2024. doi: 10.1101/2024.04.16.589843.
- [238] J. Faskowitz, D. Moyer, D. A. Handwerker, et al. Commentary on Pang et al. (2023) Nature. *bioRxiv preprint*, 2023. doi: 10.1101/2023.07.20.549785.
- [239] K. R. Patil, K. Jung, and S. B. Eickhoff. Commentary on Pang et al. (2023) Nature. *bioRxiv preprint*, 2023. doi: 10.1101/2023.10.06.561240.
- [240] E. J. Müller, B. R. Munn, and J. M. Shine. Diffuse neural coupling mediates complex network dynamics through the formation of quasi-critical brain states. *Nature Communications*, 11(1):6337, 2020. doi: 10.1038/s41467-020-19716-7.
- [241] E. J. Müller, B. R. Munn, M. J. Redinbaugh, et al. The non-specific matrix thalamus facilitates the cortical information processing modes relevant for conscious awareness. *Cell Reports*, 42(8), 2023. doi: 10.1016/j.celrep.2023.112844.
- [242] I. F. Orlando, J. M. Shine, T. W. Robbins, et al. Noradrenergic and cholinergic systems take centre stage in neuropsychiatric diseases of ageing. *Neuroscience & Biobehavioral Reviews*, 149:105167, 2023. doi: 10.1016/j.neubiorev.2023.105167.
- [243] geometricFNPmodel. <https://github.com/DynamicsAndNeuralSystems/geometricFNPmodel>.
- [244] M. N. Shadlen and W. T. Newsome. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *Journal of Neuroscience*, 18(10):3870–3896, 1998. doi: 10.1523/JNEUROSCI.18-10-03870.1998.
- [245] V. de Lafuente and R. Romo. Neuronal correlates of subjective sensory experience. *Nature Neuroscience*, 8(12):1698–1703, 2005. doi: 10.1038/nm1587.
- [246] B. van Vugt, B. Dagnino, D. Vartak, et al. The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science*, 360(6388):537–542, 2018. doi: 10.1126/science.aar7186.
- [247] A. D. Cul, S. Baillet, and S. Dehaene. Brain Dynamics Underlying the Nonlinear Threshold for Access to Consciousness. *PLOS Biology*, 5(10):e260, 2007. doi: 10.1371/journal.pbio.0050260.
- [248] C. Sergent, S. Baillet, and S. Dehaene. Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, 8(10):1391–1400, 2005. doi: 10.1038/nm1549.
- [249] J. M. Shine, M. J. Aburn, M. Breakspear, and R. A. Poldrack. The modulation of neural gain facilitates a transition between functional segregation and integration in the brain. *eLife*, 7:e31130, 2018. doi: 10.7554/eLife.31130.
- [250] R. Aronoff, F. Matyas, C. Mateo, et al. Long-range connectivity of mouse primary somatosensory barrel cortex. *European Journal of Neuroscience*, 31(12):2221–2233, 2010. doi: 10.1111/j.1460-9568.2010.07264.x.

- [251] L. Muller, A. Reynaud, F. Chavane, and A. Destexhe. The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nature Communications*, 5(1):3675, 2014. doi: 10.1038/ncomms4675.
- [252] W. Xu, X. Huang, K. Takagaki, and J.-y. Wu. Compression and Reflection of Visually Evoked Cortical Waves. *Neuron*, 55(1):119–129, 2007. doi: 10.1016/j.neuron.2007.06.016.
- [253] M. H. Mohajerani, D. A. McVea, M. Fingas, and T. H. Murphy. Mirrored Bilateral Slow-Wave Cortical Activity within Local Circuits Revealed by Fast Bihemispheric Voltage-Sensitive Dye Imaging in Anesthetized and Awake Mice. *The Journal of Neuroscience*, 30(10):3745–3751, 2010. doi: 10.1523/JNEUROSCI.6437-09.2010.
- [254] M. Massimini, F. Ferrarelli, M. J. Murphy, et al. Cortical reactivity and effective connectivity during REM sleep in humans. *Cognitive Neuroscience*, 1(3):176–183, 2010. doi: 10.1080/17588921003731578.
- [255] L. Muller, F. Chavane, J. Reynolds, and T. J. Sejnowski. Cortical travelling waves: Mechanisms and computational principles. *Nature Reviews Neuroscience*, 19(5):255–268, 2018. doi: 10.1038/nrn.2018.20.
- [256] J. Cruddas, J. C. Pang, and A. Fornito. Cortical traveling waves in time and space: Physics, physiology, and psychology. *Neuron*, 114(6):985–1005, 2026. doi: 10.1016/j.neuron.2025.12.019.
- [257] J.-Y. Wu, X. Huang, and C. Zhang. Propagating Waves of Activity in the Neocortex: What They Are, What They Do. *The Neuroscientist*, 14(5):487–502, 2008. doi: 10.1177/1073858408317066.
- [258] R. F. Betzel, L. Byrge, Y. He, et al. Changes in structural and functional connectivity among resting-state networks across the human lifespan. *NeuroImage*, 102:345–357, 2014. doi: 10.1016/j.neuroimage.2014.07.067.
- [259] U. Klatzmann, S. Froudast-Walsh, D. P. Bliss, et al. A dynamic bifurcation mechanism explains cortex-wide neural correlates of conscious access. *Cell Reports*, 44(3):115372, 2025. doi: 10.1016/j.celrep.2025.115372.
- [260] M. T. Lippert, K. Takagaki, W. Xu, et al. Methods for Voltage-Sensitive Dye Imaging of Rat Cortical Activity With High Signal-to-Noise Ratio. *Journal of Neurophysiology*, 98(1):502–512, 2007. doi: 10.1152/jn.01169.2006.
- [261] A. Grinvald, E. E. Lieke, R. D. Frostig, and R. Hildesheim. Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *Journal of Neuroscience*, 14(5):2545–2568, 1994. doi: 10.1523/JNEUROSCI.14-05-02545.1994.
- [262] T. M. Patten, C. J. Rennie, P. A. Robinson, and P. Gong. Human Cortical Traveling Waves: Dynamical Properties and Correlations with Responses. *PLOS ONE*, 7(6):e38392, 2012. doi: 10.1371/journal.pone.0038392.
- [263] G. R. Poe, S. Foote, O. Eschenko, et al. Locus coeruleus: A new look at the blue spot. *Nature Reviews Neuroscience*, 21(11):644–659, 2020. doi: 10.1038/s41583-020-0360-9.
- [264] G. Aston-Jones, J. Rajkowski, P. Kubiak, and T. Alexinsky. Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *Journal of Neuroscience*, 14(7):4467–4480, 1994. doi: 10.1523/JNEUROSCI.14-07-04467.1994.

- [265] N. K. Totalh, R. M. Neves, S. Panzeri, et al. The Locus Coeruleus Is a Complex and Differentiated Neuromodulatory System. *Neuron*, 99(5):1055–1068.e6, 2018. doi: 10.1016/j.neuron.2018.07.037.
- [266] D. Servan-Schreiber, H. Printz, and J. D. Cohen. A Network Model of Catecholamine Effects: Gain, Signal-to-Noise Ratio, and Behavior. *Science*, 249(4971):892–895, 1990. doi: 10.1126/science.2392679.
- [267] R. G. Abeysuriya, C. J. Rennie, and P. A. Robinson. Physiologically based arousal state estimation and dynamics. *Journal of Neuroscience Methods*, 253:55–69, 2015. doi: 10.1016/j.jneumeth.2015.06.002.
- [268] A. Schaefer, R. Kong, E. M. Gordon, et al. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28(9):3095–3114, 2018. doi: 10.1093/cercor/bhx179.
- [269] B. E. Jones. From waking to sleeping: Neuronal and chemical substrates. *Trends in Pharmacological Sciences*, 26(11):578–586, 2005. doi: 10.1016/j.tips.2005.09.009.
- [270] J.-P. Vilardaga. Recording Kinetics of Adrenergic Receptor Activation in Live Cells. In *Current Topics in Membranes*, volume 67, pages 101–111. Elsevier, 2011. doi: 10.1016/B978-0-12-384921-2.00005-7.
- [271] E. R. Samuels and E. Szabadi. Functional Neuroanatomy of the Noradrenergic Locus Coeruleus: Its Roles in the Regulation of Arousal and Autonomic Function Part I: Principles of Functional Organisation. *Current Neuropharmacology*, 6(3):235–253, 2008. doi: 10.2174/157015908785777229.
- [272] D. M. Devilbiss and B. D. Waterhouse. Phasic and Tonic Patterns of Locus Coeruleus Output Differentially Modulate Sensory Network Function in the Awake Rat. *Journal of Neurophysiology*, 105(1):69–87, 2011. doi: 10.1152/jn.00445.2010.
- [273] W. H. Press, editor. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, third edition, 2007. ISBN 978-0-521-88068-8.
- [274] G. Aston-Jones and J. D. Cohen. Adaptive gain and the role of the locus coeruleus–norepinephrine system in optimal performance. *Journal of Comparative Neurology*, 493(1):99–110, 2005. doi: 10.1002/cne.20723.
- [275] B. D. Waterhouse, H. C. Moises, and D. J. Woodward. Phasic activation of the locus coeruleus enhances responses of primary sensory cortical neurons to peripheral receptive field stimulation. *Brain Research*, 790(1):33–44, 1998. doi: 10.1016/S0006-8993(98)00117-6.
- [276] S. Ghosh and J. H. R. Maunsell. Locus coeruleus norepinephrine contributes to visual-spatial attention by selectively enhancing perceptual sensitivity. *Neuron*, 112(13):2231–2240.e5, 2024. doi: 10.1016/j.neuron.2024.04.001.
- [277] D. J. Chandler, W.-J. Gao, and B. D. Waterhouse. Heterogeneous organization of the locus coeruleus projections to prefrontal and motor cortices. *Proceedings of the National Academy of Sciences*, 111(18):6816–6821, 2014. doi: 10.1073/pnas.1320827111.
- [278] G. Aston-Jones and B. Waterhouse. Locus coeruleus: From global projection system to adaptive regulation of behavior. *Brain Research*, 1645:75–78, 2016. doi: 10.1016/j.brainres.2016.03.001.
- [279] R. Maran, E. J. Müller, and B. D. Fulcher. Modeling the influences of non-local connectomic projections on geometrically constrained cortical dynamics. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2506.19800.

- [280] C. Grimm, S. N. Duss, M. Privitera, et al. Tonic and burst-like locus coeruleus stimulation distinctly shift network activity across the cortical hierarchy. *Nature Neuroscience*, 27(11):2167–2177, 2024. doi: 10.1038/s41593-024-01755-8.
- [281] M. J. Redinbaugh, J. M. Phillips, N. A. Kambi, et al. Thalamus Modulates Consciousness via Layer-Specific Control of Cortex. *Neuron*, 106(1):66–75.e12, 2020. doi: 10.1016/j.neuron.2020.01.005.
- [282] R. V. Raut, Z. P. Rosenthal, X. Wang, et al. Arousal as a universal embedding for spatiotemporal brain dynamics. *Nature*, 647(8089):454–461, 2025. doi: 10.1038/s41586-025-09544-4.
- [283] S. Dehaene, C. Sergent, and J.-P. Changeux. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100(14):8520–8525, 2003. doi: 10.1073/pnas.1332574100.
- [284] S. Dehaene and J.-P. Changeux. Ongoing Spontaneous Activity Controls Access to Consciousness: A Neuronal Model for Inattentive Blindness. *PLOS Biology*, 3(5):e141, 2005. doi: 10.1371/journal.pbio.0030141.
- [285] A. Légaré, O. Ribordy, P. De Koninck, et al. Function aligns with geometry in locally connected neuronal networks. *bioRxiv preprint*, 2026. doi: 10.1101/2025.08.08.669348.
- [286] Y.-M. Wu, W.-C. Hung, Y. Chang, et al. Axon collateral pattern of a sparse locus coeruleus norepinephrine neuron in mouse cerebral cortex. *bioRxiv preprint*, 2026. doi: 10.1101/2025.02.04.636453.
- [287] B. E. Jones. Activity, modulation and role of basal forebrain cholinergic neurons innervating the cerebral cortex. In *Progress in Brain Research*, volume 145 of *Acetylcholine in the Cerebral Cortex*, pages 157–169. Elsevier, 2004. doi: 10.1016/S0079-6123(03)45011-5.
- [288] G. D. Andrews and A. Lavin. Methylphenidate Increases Cortical Excitability via Activation of Alpha-2 Noradrenergic Receptors. *Neuropsychopharmacology*, 31(3):594–601, 2006. doi: 10.1038/sj.npp.1300818.
- [289] E. J. Müller, B. Munn, L. J. Hearne, et al. Core and matrix thalamic sub-populations relate to spatio-temporal cortical connectivity gradients. *NeuroImage*, 222:117224, 2020. doi: 10.1016/j.neuroimage.2020.117224.
- [290] T. Babaie-Janvier and P. A. Robinson. Neural Field Theory of Corticothalamic Attention With Control System Analysis. *Frontiers in Neuroscience*, 13:1240, 2019. doi: 10.3389/fnins.2019.01240.
- [291] P. A. Robinson, N. C. Gabay, and T. Babaie-Janvier. Neural Field Theory of Evoked Response Sequences and Mismatch Negativity With Adaptation. *Frontiers in Human Neuroscience*, 15:655505, 2021. doi: 10.3389/fnhum.2021.655505.
- [292] C. O’Callaghan, I. C. Walpola, and J. M. Shine. Neuromodulation of the mind-wandering brain state: The interaction between neuromodulatory tone, sharp wave-ripples and spontaneous thought. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1817):20190699, 2021. doi: 10.1098/rstb.2019.0699.
- [293] R. A. España and C. W. Berridge. Organization of noradrenergic efferents to arousal-related basal forebrain structures. *Journal of Comparative Neurology*, 496(5):668–683, 2006. doi: 10.1002/cne.20946.

- [294] L. A. Briand, H. Gritton, W. M. Howe, et al. Modulators in concert for cognition: Modulator interactions in the prefrontal cortex. *Progress in Neurobiology*, 83(2):69–91, 2007. doi: 10.1016/j.pneurobio.2007.06.007.
- [295] E. S. Vizi. Interaction between adrenergic and cholinergic systems: Presynaptic inhibitory effect of noradrenaline on acetylcholine release. In B. Csillik and J. A. Kappers, editors, *Neurovegetative Transmission Mechanisms*, pages 61–78, Vienna, 1974. Springer Vienna. doi: 10.1007/978-3-7091-8341-0_4.
- [296] J. Tasserie, L. Uhrig, J. D. Sitt, et al. Deep brain stimulation of the thalamus restores signatures of consciousness in a nonhuman primate model. *Science Advances*, 8(11):eabl5547, 2022. doi: 10.1126/sciadv.abl5547.
- [297] S. J. Sara. The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience*, 10(3):211–223, 2009. doi: 10.1038/nrn2573.
- [298] V. Breton-Provencher and M. Sur. Active control of arousal by a locus coeruleus GABAergic circuit. *Nature Neuroscience*, 22(2):218–228, 2019. doi: 10.1038/s41593-018-0305-z.
- [299] E. Jodo and G. Aston-Jones. Activation of locus coeruleus by prefrontal cortex is mediated by excitatory amino acid inputs. *Brain Research*, 768(1):327–332, 1997. doi: 10.1016/S0006-8993(97)00703-8.
- [300] A. F. T. Arnsten and P. S. Goldman-Rakic. Selective prefrontal cortical projections to the region of the locus coeruleus and raphe nuclei in the rhesus monkey. *Brain Research*, 306(1):9–18, 1984. doi: 10.1016/0006-8993(84)90351-2.
- [301] E. S. Lein, M. J. Hawrylycz, N. Ao, et al. Genome-wide atlas of gene expression in the adult mouse brain. *Nature*, 445(7124):168–176, 2007. doi: 10.1038/nature05453.
- [302] K. M. Carley. On generating hypotheses using computer simulations. *Systems Engineering*, 2(2): 69–77, 1999. doi: 10.1002/(SICI)1520-6858(1999)2:2<69::AID-SYS3>3.0.CO;2-0.
- [303] T. O. Bergmann and G. Hartwigsen. Inferring Causality from Noninvasive Brain Stimulation in Cognitive Neuroscience. *Journal of Cognitive Neuroscience*, 33(2):195–225, 2021. doi: 10.1162/jocn_a_01591.
- [304] J. Gomez-Tames, I. Laakso, and A. Hirata. Review on biophysical modelling and simulation studies for transcranial magnetic stimulation. *Physics in Medicine & Biology*, 65(24):24TR03, 2020. doi: 10.1088/1361-6560/aba40d.
- [305] M. Dannhauer, L. J. Gomez, P. L. Robins, et al. Electric Field Modeling in Personalizing Transcranial Magnetic Stimulation Interventions. *Biological Psychiatry*, 95(6):494–501, 2024. doi: 10.1016/j.biopsych.2023.11.022.
- [306] S. Bai, C. Loo, and S. Dokos. A Review of Computational Models of Transcranial Electrical Stimulation. *Critical Reviews in Biomedical Engineering*, 41(1), 2013. doi: 10.1615/CritRevBiomedEng.2013007163.
- [307] C. C. McIntyre. Computational Modeling of Deep Brain Stimulation. In *Neuromodulation*, pages 171–178. Elsevier, 2009. doi: 10.1016/B978-0-12-374248-3.00017-3.

- [308] L. M. Li, K. Uehara, and T. Hanakawa. The contribution of interindividual factors to variability of response in transcranial direct current stimulation studies. *Frontiers in Cellular Neuroscience*, 9, 2015. doi: 10.3389/fncel.2015.00181.
- [309] V. López-Alonso, B. Cheeran, D. Río-Rodríguez, and M. Fernández-del-Olmo. Inter-individual Variability in Response to Non-invasive Brain Stimulation Paradigms. *Brain Stimulation*, 7(3): 372–380, 2014. doi: 10.1016/j.brs.2014.02.004.
- [310] U. Ziemann and H. R. Siebner. Inter-subject and Inter-session Variability of Plasticity Induction by Non-invasive Brain Stimulation: Boon or Bane? *Brain Stimulation*, 8(3):662–663, 2015. doi: 10.1016/j.brs.2015.01.409.
- [311] A. Guerra, V. López-Alonso, B. Cheeran, and A. Suppa. Variability in non-invasive brain stimulation studies: Reasons and results. *Neuroscience Letters*, 719:133330, 2020. doi: 10.1016/j.neulet.2017.12.058.
- [312] G. Rabuffo, M. Angiolelli, T. Fukai, et al. Pre-stimulus brain states predict and control variability in stimulation responses. *Brain Stimulation*, page 103118, 2026. doi: <https://doi.org/10.1016/j.brs.2026.103118>.
- [313] L. Sario, M. Nakai, C. Wang, and L. O. Chung. Laplace-Beltrami operator. In L. Sario, M. Nakai, C. Wang, and L. O. Chung, editors, *Classification Theory of Riemannian Manifolds: Harmonic, Quasiharmonic and Biharmonic Functions*, pages 12–25. Springer, 1977. doi: 10.1007/BFb0064419.
- [314] E. J. Allen, N. , S. J., and Z. and Zhang. Finite element and difference approximation of some linear stochastic partial differential equations. *Stochastics and Stochastic Reports*, 64(1-2):117–142, 1998. doi: 10.1080/17442509808834159.
- [315] L. R. Lines, R. Slawinski, and R. P. Bording. A recipe for stability of finite-difference wave-equation computations. *Geophysics*, 64(3):967–969, 1999. doi: 10.1190/1.1444605.

Appendix A

Supporting Information to Chapter 3




Figure A.1: **The short timescale influence of an FNP on stimulus-evoked dynamics is maximised when the stimulus is positioned at the source of the FNP, and drops off rapidly with increasing distance.** **A.** A plot of $C(t)$, the perturbation that an FNP from $\mathbf{p} = (3L/8, 3L/8)$ to $\mathbf{q} = (5L/8, 5L/8)$ induces on stimulus-evoked dynamics, for an ensemble of 200 impulse stimulus positions $\mathbf{r}_0$ sampled uniformly on Ω . The black curve is $C(t)$ when $\mathbf{r}_0 = \mathbf{p}$; i.e., when the stimulus is applied at the source of the FNP, which is identical to Fig. 3.3B. **B.** For the same ensemble of stimulus positions, **i.** $C_{\max}$, the maximum attained $C(t)$ over time, is plotted against $\|\mathbf{r}_0 - \mathbf{p}\|$, the distance between the stimulus and the source of the FNP, and **ii.** C_z , which quantifies the FNP-induced perturbation on the evoked BOLD response, is plotted against $\|\mathbf{r}_0 - \mathbf{p}\|$.




Figure A.2: A heatmap of $\gamma_{\mathbf{p}}^{\text{geo}}(\mathbf{r})$ —the spatial correlation function of the geometric model with reference point $\mathbf{p}$ —over Ω for four different spatial widths: **i.** $\sigma_n = 3L/80$, **ii.** $\sigma_n = L/8$, **iii.** $\sigma_n = L/4$, and **iv.** $\sigma_n = \infty$. The $\sigma_n = \infty$ case corresponds to a completely spatially unstructured input.




Figure A.3: **The influence of a connectome of $N = 100$ FNPs on spontaneous correlations is reduced by the incorporation of hub specificity or rich-club specificity.** **A.** For each model, we computed its correlation structure as the set of Pearson correlations between the activities of every pair of parcellated regions on Ω . We divided Ω into 10×10 parcels: $\Omega_{ij} = (i - 1, i]L/10 \times (j - 1, j]L/10$ for $i, j = 1, \dots, 10$. Parcel activities were obtained by integrating the pointwise activity $\phi(\mathbf{r}, t)$ over each parcel. For both **i.** the geometric model and **ii.** a hybrid model with a connectome of $N = 100$ randomly sampled FNPs, heatmaps of the correlations between the example parcel $\Omega_{5,5}$ and all other parcels are shown in **iii.** and **iv.**, respectively. Correlation structures for both models were computed using the same algorithm described in Sec. A.3. **B.** We quantified similarity between the spontaneous dynamics of the geometric and hybrid models using r , the Pearson correlation between the 100C2 elements of their parcellated correlation structures. Here, the correlation structures of the geometric and hybrid models from **Aii.** are compared. **C.** $1 - r$ quantifies the perturbation that the hybrid-model connectome induces on spontaneous dynamics. We compared the distribution of $1 - r$ over an ensemble of 20 connectomes containing $N = 100$ FNPs, for increasing hub specificity (λ_h) and rich-club specificity (λ_r). Each distribution is shown with error bars representing one standard deviation from the mean. The tolerance for sampled connectomes was increased to 10^{-5} to speed up computation.

Sec. A.1 Detailed derivations of model. In this appendix, we derive the equation of the mathematical model introduced and used in this work, describing the relevant variables and physiological mechanisms throughout. The goal of this model is to approximate the spatiotemporal evolution of macroscale cortical activity—i.e., mean-field population-level activity on the cortical surface of a given cerebral hemisphere at millimetre lengthscales and above—from

geometric propagation as travelling waves, concurrent with discrete connectomic propagation via fast-conducting non-local projections (FNPs) as rapid non-local interactions between specific populations on the surface. For simplicity, we derive the equation for the case of a single FNP, as this idealised case can be easily extended to multiple FNPs

The cortical surface (of either hemisphere) is mathematically represented as a closed continuous two-dimensional sheet denoted Ω . At each point $\mathbf{r}$ on Ω , we assign local geometric and non-local connectomic propagation to two separate populations. The locally propagating population, indexed $a = e$, has a spatially homogeneous and isotropic connectivity, and facilitates the geometric propagation. In other words, this population connects every pair of populations on Ω with a strength that is completely dependent on the distance between the two points. The non-local population, indexed $a = w$, is an FNP that connects population $\mathbf{a}$ to population $\mathbf{b}$, and facilitates rapid propagation between this pair only. Finally, the inhibitory population, indexed $a = i$, only connects with other neurons in the same position. At a given point $\mathbf{r}$ over time t , we track two mean population-level activities for each population-type a : the mean firing rate of generated pulses (spikes) outgoing from the somas positioned at $\mathbf{r}$, $Q_a(\mathbf{r}, t)$, and the mean rate of pulses terminating at the synapses positioned at $\mathbf{r}$, $\phi_a(\mathbf{r}, t)$. We assume that $\phi_a(\mathbf{r}, 0) = Q_a(\mathbf{r}, 0) = 0$ for all a , since the temporal dynamics of all variables are driven by external input rather than initial conditions. The output variable of the model is $\phi_e(\mathbf{r}, t)$, per the construction of neural field models [199].

The propagator mechanisms of each population, which determine the influence of Q_a on ϕ_a over space and time, is illustrated in Fig. A.4. Mathematically, the propagation of activity between two positions through the axons of population a is represented by the Volterra integral equation:

$$\phi_a(\mathbf{r}, t) = \int_{\Omega} d^2\mathbf{r}' \int_0^t dt' G_a(\mathbf{r}, t; \mathbf{r}', t') Q_a(\mathbf{r}', t') \text{ for all } (\mathbf{r}, t) \in \Omega \times [0, \infty), \quad a = e, i, w, \quad (\text{A.1})$$

where the propagator of type a , G_a , is the proportion of a pulse originating at a soma of type a at position $\mathbf{r}'$ and time t' , that terminates at a synapse of type a at position $\mathbf{r}$ and time t . Since all pulses terminate at a synapse, G_a has unit integral over space and time. Population e propagates activity between any two points on Ω , in accordance with a distance-dependent connectivity rule. Given that Ω is two-dimensional, we use the connectivity rule used in the two-dimensional neural field developed by Robinson et al. [136]:

$$G_e(\mathbf{r}, t; \mathbf{r}', t') = \frac{\gamma e^{-\gamma(t-t')}}{2\pi r} \frac{1}{\sqrt{r^2\gamma^2(t-t')^2 - |\mathbf{r} - \mathbf{r}'|^2}} \Theta(r\gamma(t-t') - |\mathbf{r} - \mathbf{r}'|), \quad (\text{A.2})$$

where r is the spatial lengthscale of the connectivity rule, and γ is the axonal propagation speed divided by r . In contrast, population w propagates activity strictly from point $\mathbf{a}$ to $\mathbf{b}$:

$$G_w(\mathbf{r}, t; \mathbf{r}', t') = \delta(\mathbf{r} - \mathbf{b}) \delta(\mathbf{r}' - \mathbf{a}) \delta(t - t' - \tau), \quad (\text{A.3})$$

where τ is the conduction time for propagation in population w . Finally, population i propagates

activity only to neurons in the same position:

$$G_i(\mathbf{r}, t; \mathbf{r}', t') = \delta(\mathbf{r} - \mathbf{r}') \delta(t - t'). \quad (\text{A.4})$$

Substituting the propagators of each type in Eqs (A.2), (A.3), (A.4) to Eq. (A.1) yields the equivalent equations:

$$\left(1 + \frac{2}{\gamma} \frac{\partial}{\partial t} + \frac{1}{\gamma^2} \frac{\partial^2}{\partial t^2} - r^2 \nabla^2\right) \phi_e(\mathbf{r}, t) = Q_e(\mathbf{r}, t); \quad (\text{A.5})$$

$$\phi_w(\mathbf{r}, t) = \delta(\mathbf{r} - \mathbf{b}) Q_w(\mathbf{a}, t - \tau); \quad (\text{A.6})$$

$$\phi_i(\mathbf{r}, t) = Q_i(\mathbf{r}, t), \quad (\text{A.7})$$

where ∇^2 is the Laplace–Beltrami operator acting on function $\phi_e(\cdot, t)$ [313]. We remark that Eq. (A.5), a damped wave equation, is exact under the assumption that Ω is unbounded [136], hence is an approximation since the physical cortical surface is bounded.




Figure A.4: **A schematic of the physiological model derived in this section.** At each point on the cortical sheet Ω , the model tracks the firing rate of the excitatory population (firing rate Q_e), and inhibitory population (firing rate Q_i). The synaptic activities of the populations at each point are then tracked by the variables ϕ_e, ϕ_i . Furthermore, in the presence of an FNP, the model tracks the firing rate of the FNP, Q_w , at its source $\mathbf{a}$, and the synaptic activity of the FNP, ϕ_w , at its target $\mathbf{b}$.

Complementing the propagator mechanisms of the model, are the synaptic mechanisms which map ϕ back to Q at each position $\mathbf{r}$ and time t , illustrated in Fig. A.4. The synaptic activities of population type e contribute to somatic activities of all population types ($\phi_e \rightarrow Q_e, Q_w, Q_i$), type i contributes to types e, i only ($\phi_i \rightarrow Q_e, Q_i$), and type w contributes to type e only ($\phi_w \rightarrow Q_e$). In addition, external input f contributes to the somatic activities of type e . The choices of these interactions is to allow the total system to be reduced to a single equation, while capturing the following synaptic contributions: $\phi_e \rightarrow Q_w$ and $\phi_w \rightarrow Q_e$ for non-local connectomic propagation, $\phi_e \rightarrow Q_e$ for local geometric propagation, and $f \rightarrow Q_e$ for external input to drive the system.

Mathematically, we used linear equations to represent the above synaptic mechanisms:

$$Q_e(\mathbf{r}, t) = G_{ee}(\mathbf{r})\phi_e(\mathbf{r}, t) + G_{ei}(\mathbf{r})\phi_i(\mathbf{r}, t) + G_{ew}(\mathbf{r})\phi_w(\mathbf{r}, t) + f(\mathbf{r}, t); \quad (\text{A.8})$$

$$Q_w(\mathbf{r}, t) = G_{we}(\mathbf{r})\phi_e(\mathbf{r}, t); \quad (\text{A.9})$$

$$Q_i(\mathbf{r}, t) = G_{ie}(\mathbf{r})\phi_e(\mathbf{r}, t) + G_{ii}(\mathbf{r})\phi_i(\mathbf{r}, t). \quad (\text{A.10})$$

In the equations above, each gain term $G_{ab}(\mathbf{r})$ is the change in Q_a at position $\mathbf{r}$, per unit increase in ϕ_b at the same position and time. G_{ab} is positive for $b = e, w$ and negative for $b = i$. Physiologically, this relationship arises from the pulses of the synapses of b being ‘transferred’ to the dendrites of a through neurotransmitter mechanisms, before arriving at the somas of a [137]. The equations neglect the temporal spreading that arises from conducting delays within the dendritic trees of individual neurons, as these delays do not differ between locally and non-locally propagating populations (e, w), thus out of the model’s scope. The equations also assume the transfer of pulses is linear, which is a reasonable approximation if the model’s variables are defined as first-order deviations about an assumed steady state [136]. The above Eqs (A.5), (A.6), (A.7), (A.8), (A.9), (A.10), when combined, reduces the system to a single non-local partial differential equation, with variable $\phi = \phi_e(\mathbf{r}, t)$, and zero initial conditions:

$$\left(1 - \tilde{G}_e(\mathbf{r}) + \frac{2}{\gamma} \frac{\partial}{\partial t} + \frac{1}{\gamma^2} \frac{\partial^2}{\partial t^2} - r^2 \nabla^2\right) \phi(\mathbf{r}, t) = \tilde{G}_w(\mathbf{r})\phi(\mathbf{a}, t - \tau) + f(\mathbf{r}, t), \text{ where} \quad (\text{A.11})$$

$$\tilde{G}_e(\mathbf{r}) = G_{ee}(\mathbf{r}) + \frac{G_{ei}(\mathbf{r})G_{ie}(\mathbf{r})}{1 - G_{ii}(\mathbf{r})}, \text{ and } \tilde{G}_w(\mathbf{r}) = G_{ew}(\mathbf{b})G_{we}(\mathbf{b})\delta(\mathbf{r} - \mathbf{b});$$

$$\phi(\mathbf{r}, 0) = \dot{\phi}(\mathbf{r}, 0) = 0. \quad (\text{A.12})$$

Finally we aim to calibrate the effective gain terms $\tilde{G}_e(\mathbf{r}), \tilde{G}_w(\mathbf{r})$ in space, to satisfy the property that, for any given input $f(\mathbf{r}, t)$, the total activity in the system integrated over space and time is preserved after the addition of population w . This preservation follows the approach through which neural field models are constructed, in which the total spacetime-integrated propagation of activity is preserved [234]. We let $\phi_0(\mathbf{r}, t)$ be the solution to the system in the absence of population w . Then, we have the constraint:

$$\int_0^\infty dt \int_\Omega d^2\mathbf{r} \phi(\mathbf{r}, t) = \int_0^\infty dt \int_\Omega d^2\mathbf{r} \phi_0(\mathbf{r}, t). \quad (\text{A.13})$$

In the case of the geometric model, where population w is absent, $\tilde{G}_e(\mathbf{r})$ is spatially uniform, and $\tilde{G}_w(\mathbf{r}) = 0$ for all $\mathbf{r}$. We let ν_0 be the constant value of $\tilde{G}_e(\mathbf{r})$ in the absence of population w . Thus ϕ_0 is the solution to the equation:

$$\left(1 - \nu_0 + \frac{2}{\gamma} \frac{\partial}{\partial t} + \frac{1}{\gamma^2} \frac{\partial^2}{\partial t^2} - r^2 \nabla^2\right) \phi_0(\mathbf{r}, t) = f(\mathbf{r}, t); \quad (\text{A.14})$$

$$\phi_0(\mathbf{r}, 0) = \dot{\phi}_0(\mathbf{r}, 0) = 0. \quad (\text{A.15})$$

The local synaptic mechanisms of population e in our model is identical at all points on Ω except $\mathbf{a}$, where synaptic activities of type e contribute to population type w . Accordingly, we

parameterise the effective gain term $\tilde{G}_e(\mathbf{r})$ to differ from ν_0 at point $\mathbf{a}$ only, i.e.:

$$\tilde{G}_e(\mathbf{r}) = \nu_0 - c\delta(\mathbf{r} - \mathbf{a}), \quad (\text{A.16})$$

for some constant $c > 0$. Under the initial conditions in Eqs (A.12), (A.15), and the fact that $\int_{\Omega} d^2\mathbf{r} \nabla^2 \phi(\mathbf{r}, t) = \int_{\Omega} d^2\mathbf{r} \nabla^2 \phi_0(\mathbf{r}, t) = 0$ for all t since Ω is a closed surface, integrating Eqs (A.11), (A.14) over space and time, and then applying constraint Eq. (A.13), yields equality $G_{ew}(\mathbf{b})G_{we}(\mathbf{b}) = c$, implying:

$$\tilde{G}_w(\mathbf{r}) = c\delta(\mathbf{r} - \mathbf{b}). \quad (\text{A.17})$$

Substituting Eqs (A.16) and (A.17) into Eq. (A.11) yields the simplified equation:

$$\begin{aligned} & \left(1 - \nu_0 + \frac{2}{\gamma} \frac{\partial}{\partial t} + \frac{1}{\gamma^2} \frac{\partial^2}{\partial t^2} - r^2 \nabla^2 \right) \phi(\mathbf{r}, t) \\ & = -c\delta(\mathbf{r} - \mathbf{a})\phi(\mathbf{a}, t) + c\delta(\mathbf{r} - \mathbf{b})\phi(\mathbf{a}, t - \tau) + f(\mathbf{r}, t), \\ & \phi(\mathbf{r}, 0) = \dot{\phi}(\mathbf{r}, 0) = 0. \end{aligned}$$

Finally, we generalise the equation to multiple FNPs, where the m th FNP has source and target $\mathbf{a}_m, \mathbf{b}_m$, and connectivity parameter c_m . Accordingly, we attain the equation:

$$\mathcal{D}_t \phi(\mathbf{r}, t) - r^2 \nabla^2 \phi(\mathbf{r}, t) - (C \circ \phi)(\mathbf{r}, t) = f(\mathbf{r}, t), \text{ where} \quad (\text{A.18})$$

$$\mathcal{D}_t \phi(\mathbf{r}, t) = \left(1 - \nu_0 + \frac{2}{\gamma} \frac{\partial}{\partial t} + \frac{1}{\gamma^2} \frac{\partial^2}{\partial t^2} \right) \phi(\mathbf{r}, t), \quad (\text{A.19})$$

$$(C \circ \phi)(\mathbf{r}, t) = \sum_m c_m [\delta(\mathbf{r} - \mathbf{b}_m)\phi(\mathbf{a}_m, t - \tau_m) - \delta(\mathbf{r} - \mathbf{a}_m)\phi(\mathbf{a}_m, t)]; \quad (\text{A.20})$$

$$\phi(\mathbf{r}, 0) = \dot{\phi}(\mathbf{r}, 0) = 0.$$

Sec. A.2 Details of numerical treatment of model. In this appendix, we detail the numerical methods used to solve the differential equation of the model used in this work. The equation can be found in Eq. (A.18), (A.19), (A.20). The methods used below also assume the particular model specifications used in this work: that the spatial domain is a square of length L with periodic boundary conditions ($\Omega = (0, L]^2$), and the conduction time of the m th FNP is zero ($\tau_m = 0$) for all m .

For accessibility, we first rewrite the differential equation to separate the temporal derivative terms from other terms:

$$\phi(\mathbf{r}, t) + \frac{2}{\gamma} \dot{\phi}(\mathbf{r}, t) + \frac{1}{\gamma^2} \ddot{\phi}(\mathbf{r}, t) = P(\mathbf{r}, t), \text{ for all } (\mathbf{r}, t) \in \Omega \times [0, \infty), \text{ where} \quad (\text{A.21})$$

$$P(\mathbf{r}, t) = \nu_0 \phi(\mathbf{r}, t) + r^2 \nabla^2 \phi(\mathbf{r}, t) + (C \circ \phi)(\mathbf{r}, t) + f(\mathbf{r}, t);$$

$$\phi(\mathbf{r}, 0) = \dot{\phi}(\mathbf{r}, 0) = 0.$$

From Ω , we uniformly partitioned a $N_x \times N_x$ computational grid of points with grid spacing

$\Delta x = L/N$:

$$\Omega^* = \{\mathbf{r}_{ij} : i, j = 1, \dots, N_x\} \subset \Omega, \text{ where } \mathbf{r}_{ij} = (i\Delta x, j\Delta x). \quad (\text{A.22})$$

For the sake of computer simulation, the time domain is truncated to a bounded interval $[0, T]$, from which we uniformly partitioned $1 + N_t$ time points $\{t_n = n\Delta t : n = 0, 1, \dots, N_t\}$ with a time step $\Delta t = T/N_t$.

We define $\phi_{ij}^n = \phi(\mathbf{r}_{ij}, t_n)$ and $P_{ij}^n = P(\mathbf{r}_{ij}, t_n)$ as the numerical solutions at each grid point and time step. At each time step n , the numerical scheme we outline below computes P_{ij}^n for i, j , and uses this value to compute ϕ_{ij}^{n+1} for all i, j .

P_{ij}^n is computed on the computational grid as the sum of: $r^2 \nabla^2(\mathbf{r}_{ij}, t_n)$, the local term; $(C \circ \phi)(\mathbf{r}_{ij}, t_n)$, the non-local term, and; $f(\mathbf{r}_{ij}, t_n)$, the external input. The first term, the local term, is approximated by the centred finite difference (five-point stencil), while incorporating the periodic boundary conditions of Ω :

$$r^2 \nabla^2 \phi(\mathbf{r}_{ij}, t_n) \approx B_{ij}^n = \frac{r^2}{(\Delta x)^2} \left(\phi_{i+1,j}^n + \phi_{i,j+1}^n + \phi_{i-1,j}^n + \phi_{i,j-1}^n - 4\phi_{ij}^n \right), \quad (\text{A.23})$$

$$\phi_{N+1,j}^n = \phi_{1,j}^n, \quad \phi_{0,j}^n = \phi_{N,j}^n, \quad \phi_{i,N+1}^n = \phi_{i,1}^n, \quad \phi_{i,0}^n = \phi_{i,N}^n.$$

The second term, the non-local term, is computed as:

$$(C \circ \phi)(\mathbf{r}_{ij}, t_n) = \sum_m c_m [\delta(\mathbf{r}_{ij} - \mathbf{b}_m) - \delta(\mathbf{r}_{ij} - \mathbf{a}_m)] \phi(\mathbf{a}_m, t_n).$$

since $\tau_m = 0$ for all m . Due to discontinuities of the delta terms $\delta(\mathbf{r}_{ij} - \mathbf{a}_m), \delta(\mathbf{r}_{ij} - \mathbf{b}_m)$ at $\mathbf{a}_m, \mathbf{b}_m$ respectively, and the possibility that the pointwise evaluation $\phi(\mathbf{a}_m, t_n)$ may not be attainable due to discretisation, $(C \circ \phi)(\mathbf{r}_{ij}, t_n)$ is numerically computed with mollifiers, which are smooth functions that approximate the delta function. We first rewrite the term as:

$$(C \circ \phi)(\mathbf{r}_{ij}, t_n) = \lim_{\epsilon \rightarrow 0} \sum_m c_m [k_\epsilon(\mathbf{r}_{ij} - \mathbf{b}_m) - k_\epsilon(\mathbf{r}_{ij} - \mathbf{a}_m)] \int_{\Omega} d^2 \mathbf{r}' k_\epsilon(\mathbf{r}' - \mathbf{a}_m) \phi(\mathbf{r}', t_n),$$

where k_ϵ is a mollifier of parameter ϵ with unit integral over Ω . For our experiments we approximate $(C \circ \phi)(\mathbf{r}_{ij}, t_n)$ by setting the mollifiers k_ϵ as Gaussian in space, and approximating the resultant integral by the Riemann sum expression:

$$(C \circ \phi)(\mathbf{r}_{ij}, t_n) \approx C_{ij}^n = \sum_{m=1} \frac{c_m}{(\Delta x)^2} [w_{ij}^{\mathbf{b}_m} - w_{ij}^{\mathbf{a}_m}] \sum_{i',j'} w_{i',j'}^{\mathbf{a}_m} \phi_{i',j'}^n, \text{ where} \quad (\text{A.24})$$

$$w_{ij}^{\mathbf{a}_m} \propto \exp\left(\frac{-\|\mathbf{r}_{ij} - \mathbf{a}_m\|^2}{2\epsilon^2}\right), \quad w_{ij}^{\mathbf{b}_m} \propto \exp\left(\frac{-\|\mathbf{r}_{ij} - \mathbf{b}_m\|^2}{2\epsilon^2}\right),$$

$$\sum_{i,j=1}^{N_x} w_{ij}^{\mathbf{a}_m} = \sum_{i,j=1}^{N_x} w_{ij}^{\mathbf{b}_m} = 1.$$

The third term, the external input term, is computed differently depending on whether the spatiotemporal profile of $f(\mathbf{r}, t)$ is an impulse-like input ($f^{\text{stim}}(\mathbf{r}, t)$ in Eq. (3.3)) used to simulate

stimulus-evoked responses, or a continuous space-time stochastic process ($f^{\text{noise}}(\mathbf{r}, t)$ in Eq. (3.4)) used to simulate noise-driven activity. In the first case, the input term is pointwise evaluated at each grid point and timestep:

$$f(\mathbf{r}_{ij}, t_n) \approx f_{ij}^{n, \text{stim}} \propto \exp\left(\frac{-\|\mathbf{r}_{ij} - \mathbf{r}_{os}\|^2}{2\sigma_x^2}\right) \exp\left(\frac{-(t_n - t_{os})^2}{2\sigma_t^2}\right), \quad (\text{A.25})$$

$$\sum_{n=0}^{1+N_t} \sum_{i,j=1}^{N_x} f_{ij}^{n, \text{stim}} = \frac{1}{(\Delta t)(\Delta x)^2}.$$

In the second case, the white noise process $\dot{W}$ is treated as a piecewise constant random process [314], whereas the deterministic function $g(\mathbf{r})$ is pointwise evaluated at each grid point:

$$f(\mathbf{r}_{ij}, t_n) \approx f_{ij}^{n, \text{noise}} = g_{ij} \frac{1}{\sqrt{\Delta t \Delta x}} \eta_{ij}, \quad (\text{A.26})$$

$$g_{ij} = g(\mathbf{r}_{ij}), \quad \eta_{ij} \sim \mathcal{N}(0, 1),$$

where $\mathcal{N}(0, 1)$ is a normal distribution with zero mean and unit standard deviation.

The above Eqs (A.23), (A.24), (A.25), (A.26) are used to compute P_{ij}^n . This value is then used to compute ϕ_{ij}^{n+1} from ϕ_{ij}^n and ϕ_{ij}^{n-1} . In this computation, the temporal derivatives in Eq. (A.21) are approximated by centred finite differences:

$$\ddot{\phi}(\mathbf{r}_{ij}, t_n) \approx \frac{\phi_{ij}^{n+1} - 2\phi_{ij}^n + \phi_{ij}^{n-1}}{(\Delta t)^2}, \quad \dot{\phi}(\mathbf{r}_{ij}, t_n) \approx \frac{\phi_{ij}^{n+1} - \phi_{ij}^{n-1}}{2(\Delta t)}.$$

In the case where $n = 0$, we substitute the initial conditions $\phi(\mathbf{r}, 0) = \dot{\phi}(\mathbf{r}, 0) = 0$ into the above finite difference approximations to obtain the approximations $\ddot{\phi}(\mathbf{r}_{ij}, t_0) = 2\phi_{ij}^1(\Delta t)^{-2}$, $\dot{\phi}(\mathbf{r}_{ij}, t_0) = 0$.

Put together, the numerical scheme used for the model is shown below:

$$\alpha = \frac{(\gamma \Delta t)^2}{2}; \quad \beta_1 = \frac{(\gamma \Delta t)^2}{\gamma \Delta t + 1}; \quad \beta_2 = \frac{2 - (\gamma \Delta t)^2}{\gamma \Delta t + 1}; \quad \beta_3 = \frac{\gamma \Delta t - 1}{\gamma \Delta t + 1}; \quad (\text{A.27})$$

$$\forall i, j = 1, \dots, N_x : \phi_{ij}^0 = 0;$$

$$\forall n = 0, 1, \dots, N_t - 1 :$$

$$\forall i, j = 1, \dots, N_x : P_{ij}^n = \nu_0 \phi_{ij}^n + B_{ij}^n + C_{ij}^n + f_{ij}^n; \quad (\text{A.23}), (\text{A.24}), (\text{A.25}), (\text{A.26})$$

$$\forall i, j = 1, \dots, N_x : \phi_{ij}^{n+1} = \begin{cases} \alpha P_{ij}^0, & n = 0, \\ \beta_1 P_{ij}^n + \beta_2 \phi_{ij}^n + \beta_3 \phi_{ij}^{n-1}, & n \geq 1. \end{cases}$$

The values of the quantities introduced above for numerical treatment, N_x, T, N_t, ϵ , used for our investigations are shown in Table A.1. N_x is fixed to the value 200 to finely resolve radial travelling waves at millimetre order spatial resolution ($\Delta x = 0.4 \text{ m}/200 = 2 \text{ mm}$). The default value of T was fixed to 70 ms to simulate the time-course of stimulus-evoked responses over short millisecond timescales. N_t was fixed to the value 988, as it satisfied the Courant–Friedrichs–Lewy condition of the two-dimensional wave equation in the absence of FNPs ($\Delta x/\Delta t > r\gamma\sqrt{2}$) [315],

and was found (from visual judgement) to not induce numerical oscillations in the presence of FNPs. ϵ was fixed to 0.002 m, which was small enough for the mollifier k_ϵ to mimic the delta functional without encountering spurious oscillations in the spontaneous experiments in Fig. 3.5.

Table A.1: **Quantities for numerical treatment**

Quantity	Value
N_x	200
T	0.07 s
N_t	988
ϵ	0.002 m

Sec. A.3 Details of computations in results.

Computation of BOLD responses. Given the model's evoked response $\phi(\mathbf{r}, t)$ to an impulse stimulus under a given structural connectivity, the BOLD response (BOLD signal observation of the stimulus-evoked response) is approximated by $z(\mathbf{r})$, the zero temporal frequency component of $\phi(\mathbf{r}, t)$:

$$z(\mathbf{r}) = \mathcal{F}[\phi](\mathbf{r}, \omega = 0) = \int_0^\infty dt \phi(\mathbf{r}, t). \quad (\text{A.28})$$

Here, $\mathcal{F}$ is the Fourier transform as an operator on ϕ as a function of t only. $z(\mathbf{r})$ is computed numerically at each grid point $\mathbf{r}_{ij}$ by computing the matrix $Z = (z(\mathbf{r}_{ij}))_{ij}$. Z is computed as the matrix series

$$Z = \Delta t \sum_{n=1}^{\infty} \Phi[n], \quad (\text{A.29})$$

where $\Phi[n] = (\phi_{ij}^n)_{ij}$. To compute the above series, we iteratively compute the partial sum

$$Y[m] = \Delta t \sum_{n=0}^{m \cdot N_t} \Phi[n] \text{ for } m = 1, 2, \dots$$

We stop the iteration when the cosine distance between the vectors $\text{vec}(Y[m]), \text{vec}(Y[m-1])$, the vectorisations of the matrices $Y[m], Y[m-1]$, is less than the threshold $\text{tol} = 10^{-5}$. The last computed $Y[m]$ is taken as an approximate scalar multiple of Z , since the cosine distance between $\text{vec}(Y[m])$ and $\text{vec}(Z)$ is close to zero. To compute Z from Y_m , we use the fact that, by integrating both sides of Eq. (A.18) over both space Ω and time $[0, \infty)$, one obtains

$$\begin{aligned} (1 - \nu_0) \left(\int_{\Omega} d^2\mathbf{r} z(\mathbf{r}) \right) &= \int_0^\infty dt \int_{\Omega} d^2\mathbf{r} \delta(\mathbf{r} - \mathbf{r}_0) \delta(t - t_0) = 1 \\ \implies \int_{\Omega} d^2\mathbf{r} z(\mathbf{r}) &= (1 - \nu_0)^{-1}. \end{aligned}$$

Accordingly, Z is computed by scalar multiplying $Y[m]$ so that the sum of all elements in Z is $(\Delta x)^{-2}(1 - \nu_0)^{-1}$.

Computation of Cosine Dissimilarity Metrics. The cosine dissimilarity between the stimulus-evoked dynamics of the geometric model and any hybrid model containing additional

FNPs, in response to the same impulse stimulus, at time $t > 0$, is defined by

$$C_\phi(t) = 1 - \frac{\int_\Omega d^2\mathbf{r} \phi(\mathbf{r}, t)\phi_0(\mathbf{r}, t)}{\sqrt{\int_\Omega d^2\mathbf{r} (\phi(\mathbf{r}, t))^2} \sqrt{\int_\Omega d^2\mathbf{r} (\phi_0(\mathbf{r}, t))^2}}, \quad (\text{A.30})$$

where ϕ_0 and ϕ are the geometric model's and hybrid model's activity respectively. Similarly, the cosine dissimilarity between the evoked BOLD response of the geometric model and any hybrid model containing additional FNPs, in response to the same impulse stimulus, is defined by

$$C_z = 1 - \frac{\int_\Omega d^2\mathbf{r} z(\mathbf{r})z_0(\mathbf{r})}{\sqrt{\int_\Omega d^2\mathbf{r} (z(\mathbf{r}))^2} \sqrt{\int_\Omega d^2\mathbf{r} (z_0(\mathbf{r}))^2}}, \quad (\text{A.31})$$

where z_0 and z are the geometric model's and hybrid model's BOLD responses respectively. $C_\phi(t)$ is used in our investigations to quantify the perturbation of an FNP, or a connectome of multiple FNPs on the model's stimulus-evoked dynamics over time, and C_z quantifies the perturbation on the evoked BOLD response. To compute $C_\phi(t_n)$, where $t_n = n\Delta t$ is a timestep in the simulation, we computed the cosine distance between the vectors $\text{vec}(\Phi[n])$ and $\text{vec}(\Phi_0[n])$, where $\Phi[n] = (\phi(\mathbf{r}_{ij}, t_n))_{ij}$, $\Phi_0[n] = (\phi_0(\mathbf{r}_{ij}, t_n))_{ij}$. Similarly, C_z is computed as the cosine distance between $\text{vec}(Z)$ and $\text{vec}(Z_0)$, where $Z = (z(\mathbf{r}_{ij}))_{ij}$, $Z_0 = (z_0(\mathbf{r}_{ij}))_{ij}$.

Computation of Spatial Correlation Function. The spatial correlation function with reference point $\mathbf{p}$, denoted $\gamma_{\mathbf{p}} : \Omega \rightarrow [-1, 1] : \mathbf{r} \rightarrow \text{corr}(\phi(\mathbf{p}, t), \phi(\mathbf{r}, t))$, returns the Pearson correlation between the model's noise-driven activity at point $\mathbf{p}$ and any input point $\mathbf{r}$ on Ω . In our investigations, we numerically compute $\gamma_{\mathbf{p}}$ evaluated at all grid points $\mathbf{r}_{ij}$ by computing the matrix $\Gamma = (\gamma_{\mathbf{p}}(\mathbf{r}_{ij}))_{ij}$. Each element of Γ is computed as

$$\gamma_{\mathbf{p}}(\mathbf{r}_{ij}) = \lim_{m \rightarrow \infty} r_{ij}[m],$$

$$r_{ij}[m] = \frac{\frac{\sum_{n=1}^{m \cdot N_t} \phi_{\mathbf{p}}^n \phi_{ij}^n}{m \cdot N_t} - \frac{\sum_{n=1}^{m \cdot N_t} \phi_{\mathbf{p}}^n}{m \cdot N_t} \cdot \frac{\sum_{n=1}^{m \cdot N_t} \phi_{ij}^n}{m \cdot N_t}}{\sqrt{\frac{\sum_{n=1}^{m \cdot N_t} (\phi_{\mathbf{p}}^n)^2}{m \cdot N_t} - \left(\frac{\sum_{n=1}^{m \cdot N_t} \phi_{\mathbf{p}}^n}{m \cdot N_t}\right)^2} \sqrt{\frac{\sum_{n=1}^{m \cdot N_t} (\phi_{ij}^n)^2}{m \cdot N_t} - \left(\frac{\sum_{n=1}^{m \cdot N_t} \phi_{ij}^n}{m \cdot N_t}\right)^2}},$$

where $\phi_{\mathbf{p}}^n$ is the simulated time series $\phi(\mathbf{r}, t_n)$ at point $\mathbf{p} = (0.15, 0.15)$. For grid size $N_x = 200$ and domain length $L = 0.4$, $\phi_{\mathbf{p}}^n$ is the time series $\phi_{75,75}^n$. Iteratively with $m = 1, 2, \dots$, we compute the matrix $R[m] = (r_{ij}[m])_{ij}$. We stop the iteration when the correlation between the vectors $\text{vec}(R[m])$, $\text{vec}(R[m-1])$ is less than the tolerance of 10^{-7} . The last attained $R[m]$ is then taken as the numerical computation of Γ .

Sampling of Connectomes with non-random properties Connectomes constrained by non-random properties (exponential distance rule, hub specificity and rich-club specificity) are sampled using a rejection-sampling algorithm. The source and target point of each FNP is sampled uniformly and independently on Ω , but is accepted only if the sampled source and target passes an acceptance-rejection test tailored to the non-random property under investigation; otherwise the source and target is sampled again. Once the sampled source and target is accepted, the next FNP is sampled in a similar manner. We emphasise, however, that each connectome we generate is constrained by one of the three non-random properties only.

Acceptance-rejection Test for Exponential Distance Rule constraint. To sample an FNP with parameter λ_e , a source $\mathbf{r}_1$ and target $\mathbf{r}_2$ is first sampled uniformly from Ω . A value u is then sampled from the Uniform $[0, 1]$ distribution. $\mathbf{r}_1$ and $\mathbf{r}_2$ are accepted if $u < \exp(-100\lambda_e l)$, where $\lambda_e \in [0, 1]$ is the selected decay parameter, and l is the Euclidean distance between the source and target (with periodic boundary conditions). Otherwise, $\mathbf{r}_1$, $\mathbf{r}_2$ and u are resampled. In this way, the probability that a FNP exists with source $\mathbf{r}_1$ and target $\mathbf{r}_2$, for any two given points $\mathbf{r}_1$ and $\mathbf{r}_2$, is

$$\mathbb{P}(\text{FNP } \mathbf{r}_1 \rightarrow \mathbf{r}_2 \mid \mathbf{r}_1, \mathbf{r}_2) = \exp(-100\lambda_e \|\mathbf{r}_1 - \mathbf{r}_2\|). \quad (\text{A.32})$$

Acceptance-rejection Test for Hub Specificity constraint. To sample an FNP with parameter λ_h , a value u is first sampled from the Uniform $[0, 1]$ distribution. If $u < \lambda_h$, the source and target points of the FNP are repetitively sampled uniformly from Ω and accepted when the FNP is hub-connecting, i.e., the FNP connects a point inside any one hub with a point outside that hub. Otherwise, if $u \geq \lambda_h$, the source and target points are sampled once and accepted. In this way, the probability that any given FNP is hub-connecting is

$$\mathbb{P}(\text{FNP is hub-connecting}) = (1 - \lambda_h)p_h + \lambda_h, \quad (\text{A.33})$$

where p_h is the probability that a uniformly sampled FNP is hub-connecting. As λ_h increases from 0 to 1, $\mathbb{P}(\text{FNP is hub-connecting})$ increases from p_h to 1. Theoretically,

$$p_h = 1 - n \left(\frac{|\Omega_h|}{|\Omega|} \right)^2 - \left(1 - n \frac{|\Omega_h|}{|\Omega|} \right)^2,$$

where $|\Omega_h|$ is the area of each hub, n is the number of hubs, and $|\Omega|$ is the area of Ω . For our investigations ($n = 4$, $|\Omega_h| = |\Omega|/34$), $p_h \approx 0.22$. u is then resampled for the next FNP.

Acceptance-rejection Test for Rich-Club Specificity constraint. To sample an FNP with parameter λ_r , a value u is first sampled from the Uniform $[0, 1]$ distribution. If $u < \lambda_r$, the source and target points of the FNP are repetitively sampled uniformly from Ω and accepted when the FNP is a rich-club connection, i.e., it connects a point inside any one hub with a point inside another hub. Otherwise, if $u \geq \lambda_r$, the source and target points are sampled once and accepted. In this way, the probability that any given FNP is hub-connecting is

$$\mathbb{P}(\text{FNP is a rich-club connection}) = (1 - \lambda_r)p_r + \lambda_r, \quad (\text{A.34})$$

where p_r is the probability that a uniformly sampled FNP is a rich-club connection. As λ_r increases from 0 to 1, $\mathbb{P}(\text{FNP is a rich-club connection})$ increases from p_r to 1. Theoretically,

$$p_r = n(n-1) \left(\frac{|\Omega_h|}{|\Omega|} \right)^2,$$

where $|\Omega_h|$ is the area of each hub, n is the number of hubs, and $|\Omega|$ is the area of Ω . For our investigations ($n = 4$, $|\Omega_h| = |\Omega|/34$), $p_r \approx 0.01$. u is then resampled for the next FNP.

Appendix B

Supporting Information to Chapter 4




Figure B.1: Effect of phasic modulation on stimulus-evoked responses under the Schaefer-400 parcellation. The experiments of Fig. 4.2A and B were repeated with regions positioned in three-dimensional space according to the Schaefer-400 left-hemisphere parcellation, rather than arranged on a two-dimensional ring. Internode structural connectivity was defined as $c_{ij} = \exp(-\lambda d_{ij})$, where d_{ij} is the 3D Euclidean distance between the centroids of regions i and j . The stimulus was applied to the most posterior region (the region with the most negative z -coordinate). Each node $j = 0, \dots, 199$ was indexed in decreasing order of time-to-peak activation, and the stimulus-evoked response $\phi^j(t)$ (s^{-1}) for each region was plotted. Snapshots of both responses (**A.** in the absence of modulation; **B.** with phasic modulation of strength $g = 0.01$) are shown on the cortical mesh at $t = 50, 100$, and 200 ms after onset.




Figure B.2: **Parcellated structural connectivity by distance.** The exponential distance rule (EDR) was fitted to the left hemisphere of the Schaefer-400 parcellated connectome. **A.** We obtained the high-resolution connectivity map from diffusion MRI tractography provided by Pang et al. [37] (*S255_high-resolution_group_average_connectome_cortex_nomedial-lh.mat*), which estimates structural connectivity between the 29,696 vertices of the fsaverage cortical surface (excluding the medial wall). We created a parcellated connectome by summing all vertex–vertex connections linking each pair of cortical regions. For our analyses, we used the Schaefer-400 parcellation of the left hemisphere [268], consisting of 200 cortical regions. Euclidean distance between region pairs was computed as the distance between their centroid coordinates, determined from the midthickness surface (*fs_LR.32k.L.midthickness.surf.gii*). An exponential function was fitted to the relationship between Euclidean distance and structural connectivity across region pairs. Fast-conducting non-local projections (FNPs) were defined as region pairs with log residuals more than 3 standard deviations above the exponential fit. **B.** Identified FNPs were visualised on the cortical mesh.




Figure B.3: **Effect of tonic LC firing rate on stimulus-evoked responses.** The model's stimulus-evoked responses over time are shown under tonic LC firing rates of **A.** $Q_l = 15$ s $^{-1}$ and **B.** $Q_l = 20$ s $^{-1}$. In both simulations, default geometric connectivity was used, node 100 was stimulated (red dot), and all nodes were modulated with strength $g = 0.01$.




Figure B.4: **Effect of modulation delay on stimulus-evoked responses.** The peak activation (maximum activity over time) of the stimulus-evoked response at each node (0–199) is shown across a finely sampled range of modulation delays (parameter τ_{lx}). In all simulations, default geometric connectivity was used, node 100 was stimulated, and all nodes were modulated with strength $g = 0.01$.




Figure B.5: **Threshold modulation for pairwise activation as a function of stimulus–target distance.** The peak activation of the model’s stimulus-evoked response at a target node was computed as a function of D , the distance between the stimulus and target node along the ring, and g , the modulation strength applied to the stimulus and target nodes (all other nodes unmodulated). Target activation was evaluated over parameter values $D = 1, 2, \dots, 100$ and $g = 0, 0.001, 0.002, \dots, 1$. Linear interpolation was then used to approximate contour plots of this function at target-activation levels of $\phi = 5, 10, 50$, and 100 s^{-1} .

Sec. B.1 Corticothalamic model.

We describe the corticothalamic dynamic gain model used in this work, describing the relevant variables and mechanisms throughout. This model approximates the spatially distributed dynamics of macroscale (millimetre lengthscales and above) cortical and thalamic activity, under the influence of modulation by the noradrenergic system.

Existing Model Description. The model is a network of N coupled nodes, each representing a distinct cortical region and its topographically aligned thalamic region. Each node contains four distinct populations: excitatory (e) and inhibitory (i) neurons in the cortex, and reticular (r) and relay (s) nuclei in the thalamus. For each population $a = e, i, r, s$ at each node j , the model tracks over time t : the (mean) firing rate of generated pulses (spikes) outgoing from the somas of population a , $Q_a^j(t)$ in units s^{-1} , and the (mean) membrane potential generated in the dendritic tree of population a , $V_a^j(t)$ in volts. Also, for each pair of populations a, b , the model tracks the (mean) rate of synaptic pulses at the axon terminals of b , that target the dendrites of population a , $\phi_{ab}^j(t)$ in units s^{-1} . The output variable of the model at each node is the self-excitatory propagator: $\phi^j(t) = \phi_{ee}^j(t)$, for all $j \in \{1, \dots, N\}$.

The firing rate of each population is determined by its membrane potential in accordance with the sigmoidal activation function:

$$Q_a^j(t) = S\left(V_a^j(t); Q_{\max}, \theta_a^j, \sigma_a^j\right) = \frac{Q_{\max}}{1 + \exp\left(-\frac{V_a^j(t) - \theta_a^j}{\sigma_a^j}\right)}. \quad (\text{B.1})$$

This relationship is derived from the assumption that any given neuron in a population fires

at its maximum firing rate $Q_{\max}$ if its membrane potential crosses a threshold [136], and the mean-field assumption that the threshold potentials across all neurons in the population follows a Gaussian-like distribution with mean θ_a^j and standard deviation $\sigma_a^j \pi / \sqrt{3}$ [74]. In the existing model, where modulation is not an explicit mechanism, θ_a^j, σ_a^j are fixed parameters $\bar{\theta}, \bar{\sigma}$ over time, and equal across all populations and nodes.

At each node j , populations i, r, s project to other populations in the same node in accordance with the equation:

$$\phi_{ab}^j(t) = Q_b^j(t - \tau_{ab}), \quad \forall a \in \{e, i, r, s\}, \quad b \in \{i, r, s\}, \quad j \in \{1, \dots, N\}, \quad (\text{B.2})$$

where τ_{ab} is the conduction time between the populations. If the projection is corticothalamic or thalamocortical, τ_{ab} is equal to the corticothalamic delay $t_0/2 > 0$, else equal to zero. Population e additionally projects to populations in other nodes in accordance to the more general differential equation

$$\left(1 + \frac{1}{\gamma} \frac{d}{dt}\right) \phi_{ae}^j(t) = \sum_{k=1}^N \tilde{c}_{jk} Q_e^k(t - \tau_{ae} - d_{jk}/v_{jk}), \quad \forall a \in \{e, i, r, s\}, \quad j \in \{1, \dots, N\}, \quad (\text{B.3})$$

where γ is the cortical damping rate, $\tilde{c}_{jk} = c_{jk} / \sum_j c_{jk}$ is the normalised connectivity strength of the corticocortical axonal projection from node k to j , and v_{jk} is its conduction speed. Equation (B.3) aims to capture the wave propagation dynamic of the original continuous-space neural field model [136], now on a discrete network topology. To justify Eq. (B.3), we equivalently express it as the (sum of) Volterra equations

$$\phi_{ae}^j(t) = \sum_{k=1}^N \int_{-\infty}^t dt' F_{jk}(t - t') Q_e^k(t' - \tau_{ae}), \quad \text{where} \quad (\text{B.4})$$

$$F_{jk}(t) = \tilde{c}_{jk} \gamma \exp(-\gamma(t - d_{jk}/v_{jk})) \Theta(t - d_{jk}/v_{jk}) \quad (\text{B.5})$$

is the response of node j to a impulse at node k ($Q_e^k(t' - \tau_{ae}) = \delta(t' - \tau_{ae})$), and Θ is the Heaviside function. In the default setting of the model, where the N nodes are arranged in a ring, connectivity strengths are determined by an EDR ($\tilde{c}_{jk} \propto \exp(-\lambda d_{jk})$), and all conduction velocities are equal ($v_{jk} \equiv v$), $F_{jk}(t)$ captures four important properties of $G(t, \mathbf{r})$, the impulse response (Green's function) of the wave equation in the original neural field model [136]:

- The impulse response is zero up until time d_{jk}/v , reflecting the finite velocity of axonal propagation from node k to j .
- The time-integrated impulse response of node j is $\int_0^\infty F_{jk}(t) dt = \tilde{c}_{jk} \propto \exp(-\lambda d_{jk})$, which abides by the EDR;
- The aggregate impulse response across all nodes at time t is

$$\sum_j F_{jk}(t) = \gamma \exp(-\gamma t) \frac{\sum_j \exp(-\lambda d_{jk}) \exp(\gamma d_{jk}/v) \Theta(t - d_{jk}/v)}{\sum_j \exp(-\lambda d_{jk})}, \quad (\text{B.6})$$

which is a function that is zero for time $t < 0$, then increases to a maximum before decaying back to zero at rate γ as $t \rightarrow \infty$. This property aligns with the space-integrated aggregate impulse response in the original neural field model [136],

$$\int_{\mathbb{R}^2} d^2\mathbf{r} G(t, \mathbf{r}) = \gamma^2 t \exp(-\gamma t) \Theta(t), \quad (\text{B.7})$$

and this property ensures that the model is linearly stable at its physiologically plausible fixed points (c.f. the Fourier transform of this response, $(1 - i\omega/\gamma)^2$, in Equation 11 of [267]);

- The time-integrated aggregate impulse response across all nodes is $\sum_j \tilde{c}_{jk} = 1$, which is required to ensure conservation of pulses.

By abiding these properties, we argue that Eq. (B.3) sufficiently recapitulates the travelling waves produced by the wave equation in the original neural field model.

Following propagation between populations, incoming pulses at a population of a given node are converted into postsynaptic membrane potentials in accordance with

$$\left(1 + \frac{1}{\alpha} \frac{d}{dt}\right) \left(1 + \frac{1}{\beta} \frac{d}{dt}\right) V_a^j(t) = \sum_{b \in \{e, i, r, s\}} \nu_{ab} \phi_{ab}^j(t),$$

$$\forall a \in \{e, i, r, s\}, j \in \{1, \dots, N\}, \quad (\text{B.8})$$

where α, β are the rates at which the potential rises and falls, and ν_{ab} is the synaptic strength. If population b is not synaptically coupled to a per the node architecture in Fig. 4.1B, then $\nu_{ab} = 0$. The membrane potential of the relay population s in each node is also driven by the stimulus propagator ϕ_{sx} with synaptic strength ν_{sx} .

The steady-state firing rates of each population, $\overline{Q_e}, \overline{Q_i} (= \overline{Q_e}), \overline{Q_r}, \overline{Q_s}$, are solved by setting all derivatives as zero, yielding the system of equations [267]:

$$\begin{cases} S^{-1}(\overline{Q_e}) &= (\nu_{ee} + \nu_{ei})\overline{Q_e} + \nu_{es}\overline{Q_s}, \\ S^{-1}(\overline{Q_r}) &= \nu_{rs}\overline{Q_s} + \nu_{re}\overline{Q_e}, \\ S^{-1}(\overline{Q_s}) &= \nu_{se}\overline{Q_e} + \nu_{sr}\overline{Q_r} + \nu_{sx}\overline{\phi_x}, \end{cases} \quad (\text{B.9})$$

where $S^{-1}(\cdot)$ is the inverse function of $S(\cdot; Q_{\max}, \bar{\theta}, \bar{\sigma})$, and $\overline{\phi_x}$ is the steady state component of $\phi_{sx}(t)$ across all nodes. After numerically solving this system, the steady-state membrane potentials and synaptic pulse rates are then computed as $\overline{V_a} = S^{-1}(\overline{Q_a})$, $\overline{\phi_{ab}} = \overline{Q_b}$. We then tracked variables $Q_a^j, V_a^j, \phi_{ab}^j$ (for all a, b, j) as perturbations from their steady states. From the equations above, only the nonlinear equation Eq. (B.1) is formulated differently under this redefinition:

$$Q_a^j(t) = S\left(\overline{V_a} + V_a^j(t); Q_{\max}, \theta_a^j, \sigma_a^j\right) - \overline{Q_a}, \quad \forall a \in \{e, i, r, s\}, j \in \{1, \dots, N\}. \quad (\text{B.10})$$

The model's stimulus-evoked response is treated as perturbations of the output, $\{\phi_{ee}^j(t)\}_j$, about its steady state, when the stimulus propagator at a single node j_0 is perturbed with a

Delta function:

$$\phi_{sx}^{j_0}(t) = \delta(t - t_{os}), \quad (\text{B.11})$$

where $t_{os} > 0$ is the stimulus onset time.

Modelling Noradrenergic Modulation of Stimulus-Evoked Responses. We introduce a single locus coeruleus (LC) population, indexed l , that is separate from the existing N nodes. The model tracks the LC's mean membrane potential ($V_l(t)$), mean firing rate ($Q_l(t)$), and mean synaptic pulse rate targeting the populations in each node, $\phi_{al}^j(t)$ for $a \in \{e, i, r, s\}$, $j \in \{1, \dots, N\}$, all as perturbations from their steady states.

The firing rate of the LC is a sum of a phasic and tonic component. The tonic component is a constant value ($Q_l^{\text{to}}(t) \equiv Q_l^{\text{to}}$). The phasic component is determined by the membrane potentials driven by the stimulus propagator ϕ_{sx} :

$$\left(1 + \frac{1}{\alpha} \frac{d}{dt}\right) \left(1 + \frac{1}{\beta} \frac{d}{dt}\right) V_l^{\text{ph}}(t) = \nu_{lx} \phi_{sx}^{j_0}(t - \tau_{lx}), \quad (\text{B.12})$$

where j_0 is the stimulated node and τ_{lx} is a time delay. By default, we assumed that LC activation is time-locked to the stimulus at the target node ($\tau_{lx} = 0$), however as part of our investigations we also consider a lagged relationship ($\tau_{lx} > 0$). The membrane potential is then related to the phasic firing rate by:

$$Q_l^{\text{ph}}(t) = \rho_l V_l^{\text{ph}}(t), \quad (\text{B.13})$$

where ρ_l is the gain of the LC. An important distinction to note is that Eq. (B.13) is a linearisation of the LC's nonlinear sigmoidal activation function about its steady state. The complete sigmoidal relationship in Eq. (B.10) is not used here, because computing the LC's firing rate from V_l^{ph} via the sigmoidal relationship requires its steady-state membrane potential ($\overline{V_l^{\text{ph}}}$) which, unlike the other populations, cannot be computed under our current model formulation.

The incoming pulse rate to the cortex or thalamus is computed from the firing rate in a similar manner to populations i, r, s :

$$\phi_{al}^j(t) = Q_l^{\text{to}} + Q_l^{\text{ph}}(t - \tau_{al}), \quad \forall a \in \{e, i, r, s\}, \quad j \in \{1, \dots, N\}, \quad (\text{B.14})$$

where τ_{al} is the conduction time from the LC to population a .

For each population a in node j , we define $\overline{\rho_a}$ as the steady-state gain of population a , and $\rho_a^j(t)$ the dynamic perturbation of this gain from steady state under the influence of modulation. We also define the modulation factor as the dynamic gain divided by its steady state, and $M_a^j(t)$, the perturbation in modulation factor from its steady-state value of 1. Incoming pulses then drive the modulation factor with timescale η and strength g_a^j :

$$\left(1 + \frac{1}{\eta} \frac{d}{dt}\right) M_a^j(t) = g_a^j \phi_{al}(t), \quad \forall a \in \{e, i, r, s\}, \quad j \in \{1, \dots, N\}. \quad (\text{B.15})$$

Changes in $M_a^j(t)$ then drive proportional changes in population gain, in accordance with the

definition $\rho_a^j(t) = \overline{\rho_a} M_a^j(t)$. Since the gain is not an explicit parameter of the sigmoidal activation function, it is adjusted indirectly by adjusting the mean firing threshold and threshold spread of each population in each node (θ_a^j, σ_a^j) . To do this, we note that the gain of population a is:

$$\overline{\rho_a} + \rho_a^j(t) = \left[\frac{dQ_a^j}{dV_a^j} \right]_{V=\overline{V_a}} = \frac{\overline{Q_a}}{\sigma_a^j} \left(1 - \frac{\overline{Q_a}}{Q_{\max}} \right). \quad (\text{B.16})$$

Under a fixed $Q_{\max}, \overline{Q_a}$, the gain is inversely proportional to σ_a^j . Hence to change the gain by $\overline{\rho_a} M_a^j(t)$, the total gain is multiplied by $1 + M_a^j(t)$, which is achieved by setting σ_a^j, θ_a^j as:

$$\sigma_a^j(t) = \frac{\overline{\sigma}}{1 + M_a^j(t)}; \quad \frac{\overline{V_a} - \theta_a^j(t)}{\sigma_a^j(t)} = \frac{\overline{V_a} - \overline{\theta}}{\overline{\sigma}}, \quad \forall a \in \{e, i, r, s\}, \quad j \in \{1, \dots, N\}. \quad (\text{B.17})$$

The second equation determines the value of θ_a^j that keeps the steady state $(\overline{V_a}, \overline{Q_a})$ unchanged.

To summarise, the dynamic gain corticothalamic model developed here is determined by Eqs (B.2), (B.3), (B.8), (B.10), (B.11), (B.12), (B.13), (B.14), (B.15), (B.17).

Sec. B.2 Numerical treatment of model. The equations of the model described in Sec. B.1 were numerically solved using a finite difference scheme. The time domain was truncated to a bounded interval $[0, T]$, from which we partitioned time points $\{t_n = n\Delta t : n = 0, \dots, N_t\}$ with time step $\Delta t = T/N_t$. We define $Q_a^j[n], V_a^j[n], \phi_{ab}^j[n], Q_l^{\text{ph}}[n], M_a^j[n], \theta_a^j[n], \sigma_a^j[n]$ as the numerical solutions at each node j and time step t_n .

We first computed the steady states $\overline{V_a}, \overline{Q_a}$ of the model. This was done by numerically solving the system of equations in Eq. (B.9) via the `fzero()` function in MATLAB.

We assumed the trivial initial condition that activity is initially at steady state and unmodulated — hence $Q_a^j[n] = V_a^j[n] = \phi_{ab}^j[n] = M_a^j[n] = 0$ for all a, b, j and time steps $n \leq 0$. Then, for $n \geq 1$, we numerically solved Eqs (B.2), (B.3), (B.8), (B.10), (B.12), (B.13), (B.14), (B.15), (B.17) using finite differences. We describe these computations below in the same order as the scheme:

$V_a^j[n]$ is computed from $V_a^j[n-1], V_a^j[n-2]$ and $\phi_{ab}^j[n-1]$ using the central difference approximation of Eq. (B.8):

$$\begin{aligned} V_a^j[n-1] + \left(\frac{1}{\alpha} + \frac{1}{\beta} \right) \frac{V_a^j[n] - V_a^j[n-2]}{2\Delta t} + \frac{1}{\alpha\beta} \frac{V_a^j[n] - 2V_a^j[n-1] + V_a^j[n-2]}{\Delta t^2} \\ = \sum_{b \in \{e, i, r, s, x\}} \nu_{ab} \phi_{ab}^j[n-1], \quad \forall a \in \{e, i, r, s\}. \end{aligned} \quad (\text{B.18})$$

$Q_l^{\text{ph}}[n]$ is computed from $Q_l^{\text{ph}}[n-1], Q_l^{\text{ph}}[n-2]$ and $\phi_{sx}^{j_0}[n-1]$ (j_0 = stimulated node) by combining Eqs (B.12) and (B.13) and applying a central difference approximation [273]:

$$\begin{aligned} Q_l^{\text{ph}}[n-1] + \left(\frac{1}{\alpha} + \frac{1}{\beta} \right) \frac{Q_l^{\text{ph}}[n] - Q_l^{\text{ph}}[n-2]}{2\Delta t} + \frac{1}{\alpha\beta} \frac{Q_l^{\text{ph}}[n] - 2Q_l^{\text{ph}}[n-1] + Q_l^{\text{ph}}[n-2]}{\Delta t^2} \\ = \rho_l \nu_{lx} \phi_{sx}^{j_0}[n-1 - \text{round}(\tau_{lx}/\Delta t)], \end{aligned} \quad (\text{B.19})$$

where $\text{round}(\tau_{lx}/\Delta t)$ is $\tau_{lx}/\Delta t$ rounded to the nearest integer.

$M_a^j[n]$ is computed from $M_a^j[n-1]$, Q_l^{to} and $Q_l^{\text{ph}}[n]$ by combining Eqs (B.14) and (B.15) and using a Backward Euler Scheme [273]:

$$M_a^j[n] + \frac{M_a^j[n] - M_a^j[n-1]}{\eta \Delta t} = g_a^j(Q_l^{\text{to}} + Q_l^{\text{ph}}[n - \text{round}(\tau_{al}/\Delta t)]). \quad (\text{B.20})$$

The firing threshold mean and spread $\theta_a^j[n]$, $\sigma_a^j[n]$ are computed from $M_a^j[n]$ using Eq. (B.17):

$$\sigma_a^j[n] = \frac{\bar{\sigma}}{1 + M_a^j[n]}, \quad \frac{\bar{V}_a - \theta_a^j[n]}{\sigma_a^j[n]} = \frac{\bar{V}_a - \bar{\theta}}{\bar{\sigma}}. \quad (\text{B.21})$$

$Q_a^j[n]$ is computed from V_a^j , $\theta_a^j[n]$ and $\sigma_a^j[n]$ using Eq. (B.10):

$$Q_a^j[n] = S\left(\bar{V}_a + V_a^j[n]; Q_{\max}, \theta_a^j[n], \sigma_a^j[n]\right) - \bar{Q}_a, \quad \forall a \in \{e, i, r, s\}. \quad (\text{B.22})$$

Since $\bar{V}_a, \bar{Q}_a$ are solved numerically from Eq. (B.9), their numerical errors will carry on to $Q_a^j[n]$ via Eq. (B.22). This means when $V_a^j[n]$ is zero, computing $Q_a^j[n]$ from Eq. (B.22) will return a non-zero value ϵ . This error will also be amplified by the modulation factor M_a^j , since $Q_a^j(t) = \bar{\rho}_a(1 + M_a^j)V_a^j(t) + o(V_a^j(t)^2)$. To correct this error, we set $Q_a^j[n] = 0$ if $|\bar{\rho}_a V_a^j[n]| \leq \epsilon$.

For populations $b = i, r, s$, $\phi_{ab}^j[n]$ is computed from $Q_b^j[n]$ using Eq. (B.2):

$$\phi_{ab}^j[n] = Q_b^j[n - \text{round}(\tau_{ab}/\Delta t)]. \quad (\text{B.23})$$

For the excitatory population e , $\phi_{ae}^j[n]$ is computed using Eq. (B.3). For accessibility, we first rewrite this equation as:

$$\left(1 + \frac{1}{\gamma} \frac{d}{dt}\right) \phi_{ae}^j(t) = P_e^j(t - \tau_{ae}), \quad \text{where } P_e^j(t) = \sum_{k=1}^N \widetilde{c}_{jk} Q_e^j(t - d_{jk}/v_{jk}). \quad (\text{B.24})$$

We then computed $P_e^j[n]$ and applied a Backward Euler Scheme:

$$P_e^j[n] = \sum_{k=1}^n \widetilde{c}_{jk} Q_e^j[n - \text{round}((d_{jk}/v_{jk})/\Delta t)]; \quad (\text{B.25})$$

$$\phi_{ae}^j[n] + \frac{\phi_{ae}^j[n] - \phi_{ae}^j[n-1]}{\gamma \Delta t} = P_e^j[n - \text{round}(\tau_{ae}/\Delta t)]. \quad (\text{B.26})$$

Finally, we computed the stimulus propagator in Eq. (B.11). This delta functional was numerically computed with a Gaussian mollifier for smoothness:

$$\phi_{sx}^{j_0}[n] \propto \exp\left(-\frac{(t_n - t_{os})^2}{\sigma_t^2}\right), \quad \text{such that } \sum_{n=1}^{N_t} \phi_{sx}^{j_0}[n] \Delta t = 1. \quad (\text{B.27})$$

For the investigations reported here, we set $\Delta t = 2^{-14}$ and $\sigma_t = \Delta t$.

Sec. B.3 Computation of cosine dissimilarity metrics.

The cosine dissimilarity between the model's stimulus-evoked dynamics under geometric

connectivity and the dynamics under hybrid connectivity (containing additional FNPs), in response to the same impulse stimulus, at time $t > 0$, is defined by:

$$C(t) = 1 - \frac{\sum_{j=1}^N (\phi_0^j(t) \phi_1^j(t))}{\sqrt{\sum_{j=1}^N (\phi_0^j(t))^2} \sqrt{\sum_{j=1}^N (\phi_1^j(t))^2}}, \quad (\text{B.28})$$

where ϕ_0^j and ϕ_1^j are the model's output at node j under geometric and hybrid connectivity respectively.

Similarly, the cosine dissimilarity between the zero-frequency components of the model's stimulus-evoked response under geometric and hybrid connectivity is defined by:

$$C_z = 1 - \frac{\sum_{j=1}^N (z_0^j z_1^j)}{\sqrt{\sum_{j=1}^N (z_0^j)^2} \sqrt{\sum_{j=1}^N (z_1^j)^2}}, \quad (\text{B.29})$$

where z_0^j and z_1^j are the zero-frequency components of $\phi_0^j(t)$ and $\phi_1^j(t)$, respectively.