ABSTRACT

Title of dissertation: ANALYZING THE DYNAMICS OF BIOLOGICAL AND ARTIFICIAL NEURAL NETWORKS WITH APPLICATIONS TO MACHINE LEARNING

Keshav Srinivasan, Doctor of Philosophy, 2024

Dissertation directed by: Professor Michelle Girvan
Department of Physics

Senior Investigator Dietmar Plenz
National Institutes of Health

The study of the brain has profoundly shaped the evolution of computational learning models and the history of neural networks. This journey began in the 1940s with Warren McCulloch and Walter Pitts' groundbreaking work on the first mathematical model of a neuron, laying the foundation for artificial neural networks. The 1950s and 60s witnessed a significant milestone with Frank Rosenblatt's development of the perceptron, showcasing the potential of neural networks for complex computational tasks. Since then, the field of neural networks has witnessed explosive growth, and terms like "Artificial Intelligence" and "Machine Learning" have become commonplace across diverse fields, including finance, medicine, and science.

This dissertation explores the symbiotic parallels between neuroscience and machine learning, focusing on the dynamics of biological and artificial neural networks. We begin by examining artificial neural networks, particularly in predicting the dynamics of large, complex networks—a paradigm where traditional machine learning algorithms often strug-

gle. To address this, we propose a novel approach utilizing a parallel architecture that mimics the network's structure, achieving scalable and accurate predictions.

Shifting our focus to biological neuronal networks, we delve into the theory of critical systems. This theory posits that the brain, when viewed as a complex dynamical system, operates near a critical point, a state ideal for efficient information processing. A key experimental observation of this type of criticality is neuronal avalanches—scale-free cascades of neuronal activity—which have been documented both in vitro (in neuronal cultures and acute brain slices) and in vivo (in the brains of awake animals). Recent advancements in experimental techniques, such as multi-photon imaging and genetically encoded fluorescent markers, allow for the measurement of activity in living organisms with unparalleled single-cell resolution. Despite these advances, significant challenges remain when only a fraction of neurons can be recorded with sufficient resolution, leading to inaccurate estimations of power-law relationships in size, duration, and scaling of neuronal avalanches. We demonstrate that by analyzing simulated critical neuronal networks alongside real 2-photon imaging data, temporal coarse-graining can recover the critical value of the mean size vs. duration scaling of neuronal avalanches, allowing for more accurate estimations of critical brain dynamics even from subsampled data.

Finally, we bridge the gap between machine learning and neuroscience by exploring the concept of excitatory-inhibitory balance, a crucial feature of neuronal networks in the brain, within the framework of reservoir computing. We emphasize the stabilizing role of inhibition in reservoir computers (RCs), mirroring its function in the brain. We propose a novel inhibitory adaptation mechanism that allows RCs to autonomously adjust inhibitory connections to achieve a specific firing rate target, motivated by the firing rate homeostasis observed in biological neurons.

Overall, this dissertation strives to deepen the ongoing collaboration between neuroscience and machine learning, fostering advancements that will benefit both fields.

ANALYZING THE DYNAMICS OF BIOLOGICAL AND ARTIFICIAL NEURAL NETWORKS WITH APPLICATIONS TO MACHINE LEARNING

by

Keshav Srinivasan

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2024

Advisory Committee:
Professor Michelle Girvan, Chair/Advisor
Senior Investigator Dietmar Plenz, Co-Advisor
Professor Brian Hunt (Dean's Representative)
Professor Edward Ott
Professor Wolfgang Losert

© Copyright by
Keshav Srinivasan
2024

Dedication

This thesis is dedicated to the memory of my grandfather, A. T. P. Sarathy. His discipline, kindness, and thirst for knowledge have been a constant source of inspiration throughout my life. His belief in me has fueled my determination to pursue this academic journey. While I deeply miss his presence, I know he would be immensely proud to witness the culmination of this work.

Acknowledgments

The research presented in this dissertation would not have been possible without the guidance of my two advisors, Dr. Michelle Girvan and Dr. Dietmar Plenz. Your constant support has shaped me not only as a scientist but also as a person. I am immensely grateful to have had such amazing advisors who have been supportive of everything I have done over these past six years.

I would also like to extend my thanks to all my colleagues and collaborators at the University of Maryland and the National Institutes of Health. In particular, I am grateful to Prof. Edward Ott, Prof. Brian Hunt, and Prof. Wolfgang Losert for not only being part of my thesis committee but also for being mentors and collaborators at various points during my PhD journey at UMD. The analyses I present would not have been possible without the efforts of all the experimentalists at NIH. Special thanks to Dr. Tiago Ribeiro, Dr. Patrick Kells, and Dr. Sinisa Pajevic for their constant guidance and support.

I also want to acknowledge Souad Nejjar and Jessica Crosby, program coordinators of the Biophysics program at UMD. Since day one, they have ensured that every student could put their best foot forward and were always available for a talk. On the NIH side, I would like to thank Craig Stewart and Latoya Evans for ensuring that everything in the lab functioned smoothly and without any hiccups.

I would like to highlight two programs at UMD: First, the TREND REU program, through which I had the opportunity to meet, collaborate with, and mentor two exceptional undergraduate students, Nolan Coble and Joy Hamlin. The work we did together formed a significant part of my first publication in grad school. Secondly, I am grateful for the COMBINE program, where I met many interesting people, formed academic and personal networks, and honed various communication skills. I was also honored to serve as the president of the COMBINE student body, through which I learned invaluable leadership and organizational skills.

As much as it is an academic accomplishment, the journey to a PhD has been a roller-coaster of emotions, and I couldn't have navigated it without my incredible wife, Tirtha Giri. She has been my constant source of strength, love, and inspiration. This dissertation is as much a testament to her unwavering support as it is to my academic pursuits. From late-night editing sessions to animating countless presentation slides, Tirtha has been my partner in every sense of the word. Her belief in me has carried me through the toughest challenges, and I am grateful for her love and companionship.

And speaking of emotional support, I would like to thank my family. First and foremost, I want to thank my parents, Radhika and A. T. Srinivasan. Dad, you have always supported me unconditionally, never trying to influence or change my decisions, but always encouraging me wholeheartedly and providing the best at every point. Mom, you have been my inspiration in academia. Seeing you complete your PhD before me (yes, you won the challenge) was a significant motivation for me to push on and work as hard as you do every day. Next, I would like to thank my sister, Keerthana, and my brother-in-law, Akash. They are probably the people I have spoken to the most (after Tirtha, of course), and they have always been my biggest cheerleaders. From philosophical discussions about life to technical advice on Adobe Illustrator, I appreciate and cherish everything you both have done and continue to do for me. I would also like to thank my grandmother, Bhooma Paati, for constantly showering me with love and blessings. Additionally, I extend my gratitude to all other members of the family who have always been so loving and supportive.

While I have been in the US, away from my family, I have found a new family at Blackfoot Nivas. This family has gone through multiple iterations, starting with Sanket, Spandan, Nikhil, and Sandesh, and later adding Akarsh, Vedanth, and Anirrudh. This place has been my home, and I cherish every memory we've made here. From cooking together and having debates that last over four hours to the numerous parties and shared experiences, Blackfoot Nivas has had my heart. Sanket, you were someone I knew from undergrad, but

despite being in the same program, we never really spoke much. I am thankful to UMD for bringing us together and allowing us to truly bond. All those nights in Erie, binging show after show, eating so much food, and having endless discussions – I am so happy it all happened. Nikhil, my permanent whiskey, food, and gym partner, I miss every part of our bond. Though it took some time for us to warm up to each other, I found in you a fierce friend who I can always rely on. I cherish each moment we shared and the countless experiences that brought us close. And finally, Spandan, my first new friend in the US and the one with whom I've shared every memory in College Park. You have been with me through every stage of my PhD, from our classes with Sukharev to all our "social elite" outings. From football to Bollywood, we didn't leave any stone unturned. You are truly the best friend anyone could ask for, and I am incredibly grateful to have had you as my partner in crime. Every memory of College Park that I fondly look back on will always include you, and I can't imagine any of it without you.

Last but certainly not least, I want to express my heartfelt gratitude to the entire Tiki-Taka squad. Football has been an invaluable escape during the challenging times of my PhD journey. The warmth and camaraderie I've experienced with each team member are truly special, and our two tournament victories will forever be cherished highlights of this chapter of my life. While I may not have explicitly mentioned countless colleagues, family, and friends (e.g. Magic Moments and the Scottishites) please know that your support means the world to me, and my gratitude is boundless.

Table of Contents

Dedication ii

Acknowledgements iii

Table of Contents vi

List of Tables x

List of Figures xi

List of Abbreviations and Symbols xiii

1 Introduction 1
- 1.1 A brief history of neural networks . . . 1
- 1.2 Introduction to Reservoir Computing . . . 3
- 1.3 Criticality in biological and artificial neural networks . . . 5
- 1.4 Introduction to neuronal avalanches . . . 8
- 1.5 Quantifying neuronal avalanches and methodological challenges . . . 9
- 1.6 Dissertation outline . . . 11

2 Parallel machine learning for predicting the dynamics of complex networks 13
- 2.1 Overview . . . 13
- 2.2 Introduction . . . 14
- 2.3 Methods . . . 15
  - 2.3.1 Test system: Frequency assoratative Kuramoto network . . . 15
  - 2.3.2 Background on non-parallel reservoir computing prediction . . . 17
  - 2.3.3 Parallel machine learning architecture . . . 20
  - 2.3.4 Generating the Reservoir Input Matrix, $\mathbf{W_{in}}$ . . . 20
  - 2.3.5 RC hyperparameters . . . 23
  - 2.3.6 Tests . . . 23
- 2.4 Results . . . 24
  - 2.4.1 Single non-parallel reservoir prediction . . . 24
  - 2.4.2 Parallel scheme with known links . . . 24
  - 2.4.3 Parallel scheme with unknown links . . . 26
  - 2.4.4 Node-level predictions . . . 27
  - 2.4.5 Effect of Hubs . . . 27

2.4.6 Dependence on size of the predicted network . . . 30
2.4.7 Coupled Lorenz System . . . 32
2.5 Conclusions . . . 34
2.6 Acknowledgements . . . 35

3 Overcoming sub-sampling bias in evaluating the role of neuronal avalanches in the brain 36
3.1 Overview . . . 36
3.2 Introduction . . . 37
3.3 Experimental Methods . . . 39
3.3.1 Overview of datasets analyzed . . . 39
3.3.2 Animal surgery . . . 39
3.3.3 Pre-processing pipeline . . . 40
3.4 Subsampled critical neuronal model . . . 41
3.4.1 Model Topology . . . 41
3.4.2 Model Dynamics: . . . 42
3.4.3 Model Simulations: . . . 43
3.5 Post-processing pipleine . . . 44
3.5.1 Continuous Epochs of Suprathreshold Population Activity: . . . 44
3.5.2 Scaling Curve Fit: . . . 45
3.5.3 Complexity Analysis: . . . 46
3.5.4 Temporal Profiles: . . . 47
3.6 Results . . . 49
3.6.1 Neuronal synchronization in the prefrontal cortex displays quadratic avalanche scaling . . . 49
3.6.2 Temporal coarse graining recovers avalanche scaling in a subsampled critical model . . . 51
3.6.3 Properly resolved avalanches maximize complexity and temporal correlations . . . 58
3.6.4 Scale-invariant inverted parabolic profile . . . 58
3.6.5 Summary . . . 63
3.7 Discussion . . . 64
3.8 Acknowledgements . . . 65

4 Limits in the recovery of parabolic avalanches in spatially subsampled neuronal networks at criticality 66
4.1 Overview . . . 66
4.2 Introduction . . . 67
4.3 Methods . . . 70
4.3.1 Correlation Analysis . . . 70
4.3.2 Deviation from criticality coefficient (DCC) . . . 71
4.3.3 Collapse Exponent of Threshold and Fractional Sampling . . . 71

4.4 Results . . . 72
4.4.1 High Coincident Firing Thresholds Underestimate χ in the Fully Sampled Model . . . 73
4.4.2 Temporal Coarse-Graining Rescues χ at High Thresholds in the Fully Sampled Model . . . 75
4.4.3 Rescue of $\chi = 2$ Is Only Seen in Critical Dynamics That Is Fractionally Sampled and Fails for Slightly Super- or Sub-critical Dynamics . . . 77
4.4.4 Scaling Collapse Demonstrates Equivalence of Thresholding and Subsampling in the Critically Balanced All-to-All E/I-Model . . . 79
4.4.5 The Rescue of $\chi = 2$ Indicates a Recovery of Mean Pair-Wise Correlations . . . 79
4.4.6 Subsampling and Spike-Resolution Introduce Unobservable Regions of Scaling Exponents . . . 80
4.4.7 Robust Recovery of $\chi = 2$ for Parabolic Avalanches in Frontal Cortex of Awake Mice . . . 82
4.5 Discussion . . . 83
4.5.1 Temporal coarse-graining when avalanche periods are unknown . . . 85
4.5.2 Limits in rescuing the scaling exponent under subsampling conditions 87
4.5.3 The DCC and crackling noise relationship fails to recover the scaling exponent under subsampling conditions . . . 87
4.5.4 The sensitivity of temporal coarse-graining approach to thresholds in avalanche size definition . . . 88
4.5.5 Synchronization, coincident firing and delayed neuronal correlations 89
4.6 Acknowledgements . . . 90

5 Heterogeneity and Local Excitatory-Inhibitory Plasticity Facilitate Robust Performance in Reservoir Computers 91
5.1 Overview . . . 91
5.2 Introduction . . . 92
5.3 Methods . . . 96
5.3.1 Excitatory-inhibitory reservoir computing . . . 96
5.3.2 Constructing the adjacency matrix, A_{ij} . . . 99
5.3.3 Neuron differential entropy . . . 100
5.3.4 Tasks for testing . . . 100
5.3.5 Inhibitory adaptation mechanism . . . 103
5.3.6 Heterogeneous distribution of firing rate targets . . . 104
5.3.7 Designing reservoirs with local E-I balance . . . 105
5.4 Results . . . 106
5.4.1 Robust performance in balanced and slightly over-inhibited regimes; brittleness with excess excitation . . . 106

5.4.2 Self-adapting E-I Reservoirs have locally balanced and higher entropy neurons . . . 112
5.4.3 Homogeneous and heterogeneous inhibitory adaptation significantly boosts RC performance . . . 115
5.4.4 Designing RCs for local balance yields similar results to adaptive mechanisms . . . 116
5.4.5 Memory-Nonlinearity trade-off in RCs . . . 117
5.5 Discussion . . . 120
5.6 Acknowledgements . . . 121

6 Summary and future directions 122

Bibliography 124

List of Tables

2.1 List of hyperparameter values obtained via a coarse grid search for single and parallel reservoir architectures 23

5.1 Default reservoir hyperparameters for Excitatory-Inhibitory RC framework 99
5.2 Optimal input scaling for various tasks in both non-adapting and inhibitory adaptive RCs 118

List of Figures

2.1 Frequency assortative network structure with 50 nodes 16
2.2 Largest Lyapunov Exponent as a function of the coupling constant, K 17
2.3 Schematic diagram for the parallel network ML architecture 20
2.4 Parallel prediction of Kuramoto system with known neighbors 22
2.5 Prediction of the order parameter in 3 different cases of ML implementations 25
2.6 Prediction of the dynamics of three representative individual oscillators . . 28
2.7 Predictive performance of the parallel method for different network architectures - Effect of hubs 29
2.8 Performance of the different Reservoir Computing methods as a function of the Kuramoto oscillator network size 31
2.9 Prediction of the dynamics of a coupled Lorenz network 33
2.10 Predictive performance of the parallel method for coupled Lorenz systems of different degrees 34

3.1 Neuronal synchronization in the prefrontal cortex displays quadratic avalanche scaling 48
3.2 Temporal coarse graining recovers avalanche scaling in a subsampled critical model 52
3.3 Examples of avalanche distributions as a function of threshold Θ and sampling fraction f for the critical E/I-model 53
3.4 Short simulation times provide robust estimates of χ_{sh} as a function of temporal coarse-graining k under fully-sampled and subsampled conditions . . 54
3.5 Impact of added irregular firing and subsampling as a function of threshold on the recovery of χ_{sh} 56
3.6 Avalanche distributions and scaling as a function of temporal coarse graining k for subcritical network dynamics 57
3.7 Properly resolved avalanches maximize complexity and temporal correlations 59
3.8 Neuronal assemblies follow the universal temporal profile of an inverted avalanche parabola with scaling collapse of $\chi = 2$ 60
3.9 Coarse-grained, few-generation avalanches with $\chi_{\mathrm{sh}} = 2$ exhibit inverted parabolic shape in cortical synchronization at high-resolution over many decades of generation 62
3.10 Summary sketch of parabolic avalanches represent properly recovered avalanche dynamics with maximal complexity in subsampled critical networks 63

4.1 Schematic of the balanced E/I-model and identification of suprathreshold sequences in population spiking activity with respect to minimal coincident spiking threshold and temporal coarse graining . . . 74
4.2 Increasing the threshold in the fully sampled model underestimates the scaling exponent, χ, which can be rescued by temporal coarse graining . . . 76
4.3 Temporal coarse graining at critical value of g: Rescue of scaling exponent, equivalence between threshold and fractional sampling, and increased mean pairwise neuron correlation . . . 78
4.4 Temporal coarse-graining identifies a robust rescue regime for $\chi = 2$ covering 0.1% - 100% fractional sampling at decreasing thresholding level, which is absent for the DCC. . . . 81
4.5 Consolidated view of the exponents α, β, and χ as a function of threshold and temporal coarse-graining for different sampling fractions . . . 82
4.6 Parabolic avalanches in ongoing activity of frontal cortex exhibit threshold and temporal coarse-graining rescue of scaling exponent $\chi = 2$ in line with critical model dynamics . . . 84

5.1 Schematic of a reservoir computer framework with tunable excitatory-inhibitory balance . . . 97
5.2 Schematic demonstrating the inhibitory adaptation mechanism . . . 105
5.3 The effects of varying the fraction of excitatory neurons in a network with $\mu_I = 4\mu_E$ . . . 107
5.4 The effects of varying the fraction of excitatory neurons in a network with $\mu_I = \mu_E$ . . . 108
5.5 Adjusting the relative inhibitory synaptic strength to modify the E-I balance while maintaining the proportion of inhibitory neurons at 20% (Memory Capacity) . . . 109
5.6 Adjusting the relative inhibitory synaptic strength to modify the E-I balance while maintaining the proportion of inhibitory neurons at 20% (NARMA-10, Mackey-Glass and Lorenz) . . . 110
5.7 Adaptation curves for inhibitory self-adaptive RCs with homogeneous and heterogeneous (beta-distributed) target firing rates . . . 112
5.8 The effects of the inhibitory adaptation mechanism on neuronal properties . 114
5.9 Consolidated view of the effect of the inhibitory adaptation mechanism for the Lorenz time-series prediction tasks highlighting input and non-input neurons . . . 115
5.10 Tuning curves for inhibitory self-adapting RCs with homogeneous firing rate targets across four different tests . . . 117
5.11 Histogram comparing the best performing globally tuned RC with two types of inhibitory adaptive RCs . . . 118
5.12 Memory-nonlinearity trade off in RCs- Tuning the input scaling parameter . 119

List of Abbreviations

2PI	2-photon imaging
AC	Autocorrelation
ACC	Anterior cingulate cortex
ACUC	Animal care and use committee
ANN	Artificial neural network
ARNI	Algorithm for Revealing Network Interactions
CNN	Convolutional neural network
DCC	Distance from criticality coefficient
DP	Directed percolation
E/I	Excitatory-Inhibitory
EEG	Electroencephalography
ESN	Echo state network
IF	Integrate-and-fire
IP	Intrinsic plasticity
LFP	Local field potential
LIF	Leaky integrate-and-fire
LSM	Liquid state machine
LSTM	Long-short term memory
MEG	Magnetoencephalography
MF-DP	Mean-field directed percolation
ML	Machine learning
mPFC	Medial prefrontal cortex
NA	Numerical aperture
NARMA	Nonlinear autoregressive moving average
NRMSE	Normalized root mean squared error
RC	Reservoir computer/ Reservoir computing
RNN	Recurrent neural network
RMS	Root mean squared
ROI	Region of interest
SD	Standard deviation
SOB	Self organized bistability
SOC	Self organized criticality
STDP	Spike timing dependent plasticity
TE	Transfer entropy
TREND	Training and Research Experiences in Nonlinear Dynamics
V1	Visual area 1 of the visual cortex
VPT	Valid prediction time

Chapter 1: Introduction

1.1 A brief history of neural networks

The quest to understand the human brain and replicate its incredible processing power has driven the development of artificial neural networks and machine learning for over eight decades. This journey showcases a fascinating interplay between neuroscience and machine learning, with each field inspiring and informing advancements in the other, leading to a rich symbiotic relationship.

The seeds of artificial neural networks were sown in the 1940s with Warren McCulloch and Walter Pitts' pioneering work. Their groundbreaking paper introduced a mathematical model of a neuron, laying the foundation for computational models of the brain [105]. Inspired by biological neurons, their model consisted of a simple summation unit with a binary threshold activation function. This rudimentary model paved the way for more complex network architectures.

The next significant leap came in the 1950s with Frank Rosenblatt's invention of the perceptron. This single-layer network, capable of learning simple linear patterns, marked a crucial milestone [150]. Rosenblatt's work sparked optimism about the potential of neural networks for various tasks. However, limitations in training algorithms and the inherent constraints of single-layer networks led to a period of decline in the 1960s and 70s.

During this lull, parallel developments were happening in neuroscience. Researchers like David Hubel and Torsten Wiesel made significant strides in understanding the visual

cortex [61]. Their work on the organization of visual processing in the brain, with its hierarchical architecture and local receptive fields, resonated with the idea of layered neural networks.

The tide turned for neural networks in the 1980s with the advent of new learning algorithms like backpropagation [151]. This powerful algorithm allowed for the training of multi-layer networks, overcoming the limitations of the perceptron. This resurgence coincided with the blossoming of cognitive neuroscience, with researchers like Eric Kandel studying the cellular basis of learning and memory [69]. The parallels between these investigations into biological learning mechanisms and the training algorithms used in neural networks became increasingly evident.

The 1990s witnessed a surge in neural network applications, particularly in areas like pattern recognition and image processing. Advances in computational power further fueled this growth. However, a new challenge emerged: the difficulty of training large, complex networks. Neuroscience again provided crucial insights. Studies on the brain's efficient information processing, its ability to learn with limited data, and its reliance on sparse connectivity inspired the development of techniques like sparsity and regularization in neural networks to improve efficiency and learning [82].

The 21st century has seen the rise of deep learning, a subfield of neural networks characterized by deep architectures with many layers. These networks have achieved remarkable success in tasks like image recognition, natural language processing, and machine translation. This progress reflects a deeper understanding of the brain's architecture and learning mechanisms, including the importance of unsupervised learning and hierarchical organization.

Today, the landscape of machine learning is dominated by large language models like ChatGPT and Gemini. These powerful tools, trained on massive datasets of text and code, can generate realistic and creative text formats, translate languages, write different kinds

of creative content, and answer questions in an informative way. This has opened doors in communication, education, and entertainment. However, the future of neural networks extends beyond these models. Advancements in neuromorphic computing, inspired by the brain's structure and efficiency, promise to create hardware that can process information in a similar way [63]. Additionally, research on the interpretability and explainability of neural networks is crucial for understanding information processing and building trust to ensure effective applications of this powerful technology.

All in all, the history of neural networks is a testament to the powerful synergy between neuroscience and machine learning. As we delve deeper into the intricacies of the brain, we can expect even more groundbreaking advancements in artificial intelligence, further blurring the lines between the biological and the artificial.

1.2 Introduction to Reservoir Computing

Within the context of artificial neural networks, a key focus of this dissertation will be reservoir computing. Reservoir computing is an innovative computational paradigm that utilizes the dynamic properties of recurrent neural networks (RNNs) to process temporal data. This approach simultaneously emerged from the principles of echo state networks (ESNs) and liquid state machines (LSMs) [64, 98]. It provides a robust framework for handling tasks that involve temporal dependencies, such as time series prediction, speech recognition, and control systems.

The core concept of reservoir computing revolves around mapping input data into a high-dimensional space through a fixed, randomly initialized recurrent neural network, known as the *reservoir*. This reservoir consists of a large number of interconnected neurons exhibiting complex, non-linear dynamics. One of the key advantages of this approach is that it decouples the complexity of the recurrent network dynamics from the training pro-

cess. Instead of training the entire network, only the output weights, which linearly combine the states of individual neurons in the reservoir, are adjusted during training [64, 94]. A typical reservoir computing system is composed of three main components:

- **Input Layer**: This layer is responsible for transforming the input data into a form that can be fed into the reservoir. It typically involves scaling and encoding the input signals.

- **Reservoir**: The reservoir is a large, fixed recurrent neural network with rich, nonlinear dynamics. The input signals are projected into this high-dimensional space, where the complex dynamics of the network create a diverse set of internal states.

- **Output Layer**: The output layer is a linear readout mechanism that maps the high-dimensional reservoir states to the desired output. This is the only part of the network that is trained, usually using linear regression techniques.

The reservoir's intricate dynamics enable it to capture and process temporal dependencies in the input data, making it a powerful tool for time-dependent tasks.

The training process in reservoir computing is significantly simplified compared to traditional RNN training. While traditional RNNs require the adjustment of all network weights through backpropagation through time, reservoir computing only involves training the output weights. The reservoir weights and input weights remain fixed after initialization. This simplification reduces the computational complexity and allows for faster training [179].

Reservoir computing has found applications in various fields due to its ability to handle temporal data effectively. Some of them include:

- **Time Series Prediction**: Reservoir computing has been successfully used for forecasting future values in temporal datasets. Its ability to capture complex temporal

patterns makes it well-suited for financial forecasting, weather prediction, and other time series applications [129, 183].

- **Speech Recognition**: By leveraging the temporal patterns in audio signals, reservoir computing has demonstrated promising results in speech recognition tasks. Its capacity to process continuous input streams makes it ideal for recognizing spoken language [48].

- **Control Systems**: In dynamic control systems, reservoir computing can implement effective control strategies. Its ability to adapt to temporal changes in the system's state makes it useful for robotics, automated driving, and other control applications [157].

Reservoir computing provides a powerful and efficient framework for processing temporal data by leveraging the dynamic properties of recurrent neural networks. This approach decouples the training process from the reservoir's complex dynamics, offering a practical method to utilize the computational power of recurrent networks. Its diverse applications, ranging from time series prediction to speech recognition and control systems, highlight its versatility and effectiveness in managing time-dependent tasks. Subsequent chapters will address detailed discussions on the specifics and implementation of reservoir computing.

1.3 Criticality in biological and artificial neural networks

Criticality, which will be another important focus of this dissertation, refers to the state of a system poised at the threshold between order and disorder, often exhibiting scale-free activity characterized by power-law distributions and long-range correlations. This concept, rooted in statistical physics, has found many applications in understanding complex

systems, including the brain and artificial neural networks.

The brain is a highly dynamic and complex system, displaying a diverse range of spatio-temporal activity patterns. Research suggests that the brain operates near a critical state, balancing between quiescence and chaos. This criticality is characterized by scale-invariance and power-law distributions in neural activity patterns [13]. Such a state is believed to optimize information processing capabilities, enhancing the brain's ability to respond to stimuli, process information efficiently, and store memories.

Early works by Packard [126], Langton [79], and Kauffman [71] independently suggested that the computational performance of a system would be optimized at the edge of chaos, at a critical point in a phase transition between two dynamical regimes. This hypothesis has been demonstrated experimentally in networks of cortical neurons [13], showing that they seem to be operating near a critical point.

Criticality has been shown to have significant implications in many areas of information processing, including:

- **Information transmission**: The scale-free nature of the dynamics at a critical point supports local information transmission and enables communication between neurons that are very far from each other [13, 73].

- **Memory**: When recurrent neuronal networks are at a critical point, the diversity in stable firing patterns is maximized, potentially leading to optimal information storage [18, 54].

- **Stability**: When a branching network model is tuned to the critical point, it produces stable trajectories that can be controlled by input [54].

- **Processing power**: Spiking neuronal network models show that computational power is maximized at the critical point between order and chaos. Completely ordered net-

works limit the range of mappings from input to output, while highly random networks have a larger range of mappings that are unreliable. Networks tuned between these extremes perform more effectively at computational tasks [98].

- **Dynamic range**: At criticality, the range of stimulus intensities that can be processed is maximized. This has been shown both experimentally in organotypic cultures and in models of coupled binary neurons [164].

Like the brain, artificial neural networks (ANNs) have also been investigated for signs of criticality. The hypothesis is that ANNs operating near criticality might exhibit improved learning dynamics, robustness, and generalization capabilities. In recent years, researchers have explored the conditions under which artificial networks might exhibit critical behavior.

Theoretical work has shown that deep neural networks can exhibit criticality by adjusting the balance between excitation and inhibition, much like biological neural networks [18]. For instance, criticality in these networks can be manipulated through weight initialization, network architecture, and training protocols. Networks initialized at the edge of chaos demonstrate enhanced training performance and better generalization, as they can explore a diverse range of network configurations efficiently.

Within the framework of reservoir computing, a recurring theme of this dissertation, criticality has been demonstrated to significantly impact performance. When the reservoir operates at the edge of chaos, it exhibits a balance between order and complexity that enhances its computational capabilities. Studies have shown that reservoirs poised at criticality can maximize memory capacity and improve the system's ability to handle nonlinear and temporal tasks effectively [83]. This is because criticality ensures that the network retains a rich repertoire of dynamical patterns, facilitating diverse and adaptable responses to inputs. Consequently, tuning reservoir networks to critical states can significantly boost their performance in tasks such as speech recognition, time series prediction, and adaptive

control.

The study of criticality in both biological and artificial neural networks reveals several parallels. Both systems benefit from operating near criticality, enhancing their ability to process information, adapt to new stimuli, and learn from experience. This convergence suggests that principles of criticality may be fundamental to understanding and designing intelligent systems.

Future research directions include further exploration of how criticality can be harnessed in artificial systems for improved performance and understanding the underlying mechanisms that drive criticality in the brain. Advances in neuromorphic computing, which aim to mimic the brain's architecture and dynamics, may offer new insights and applications by leveraging criticality for more efficient and powerful computation.

Overall, criticality provides a unifying framework for studying complex systems, offering valuable insights into the functioning of the brain and the optimization of artificial neural networks. By continuing to explore this frontier, we may unlock new paradigms for both neuroscience and artificial intelligence.

1.4 Introduction to neuronal avalanches

As alluded to in the previous section, the critical brain hypothesis posits that the brain operates near a critical point, a tipping point between two distinct states. This critical state is thought to be ideal for information processing, allowing the brain to efficiently respond to stimuli.

A key piece of evidence supporting this hypothesis is the presence of neuronal avalanches, which are scale-free cascades of neural activity. Neuronal avalanches were first identified in the spontaneous activity of organotypic cortex cultures and acute cortex slices in vitro [13]. Subsequent studies confirmed avalanche dynamics in various animal models,

including lightly anesthetized rodents [74, 162, 187], cats [53], awake rodents [15], and nonhuman primates [187, 188]. In humans, neuronal avalanches have been observed under various conditions, including the awake state [108], whereas avalanche parameters were systematically altered under anesthesia [43], and illness [5, 112].

Over the years, neuronal avalanches have significantly advanced our comprehension of information propagation and variability within synchronized cortical activity. They encapsulate both the rapid onset of synchronization and its subsequent dissemination across the cortex over brief intervals, offering a comprehensive perspective. Notably, avalanches exhibit scale invariance, meaning their size and duration follow predictable patterns described by power laws, regardless of the observational scale. The observed fractal pattern of neuronal avalanche sizes suggests a self-similar hierarchy. This allows us to look beyond the analysis of individual synchronization events and investigate the interwoven network of synchronization events with varying spatiotemporal scales within the cortex. Hence, this framework transcends the limitations of studying isolated events and facilitates a more comprehensive understanding of cortical dynamics. Furthermore, neuronal avalanches might offer quantitative measures for studying optimization in cortical networks, enabling us to objectively evaluate the efficiency of information processing within the brain.

1.5 Quantifying neuronal avalanches and methodological challenges

Neuronal avalanches, being scale-free cascades of neural activity, can be quantified by power-laws and their corresponding exponents. The magnitude or size, S, of an avalanche is defined as the number of activations occurring over the course of a single cascade, and the size probability distribution $P(S)$ is theoretically described by a power law [13] with critical exponent α,

$$P(S) \propto S^{-\alpha}. \tag{1.1}$$

Similarly, the duration of an avalanche follows a power law [13] with critical exponent β

$$P(T) \propto T^{-\beta}, \tag{1.2}$$

The average size as a function of duration is described by yet another power law,

$$\langle S \rangle(T) \propto T^{\chi}, \tag{1.3}$$

where the scaling exponent χ captures how the size of an avalanche scales with the temporal length of the neural cascade [24].

We also find that the precise profile of activations over the course of a cascade, $n(t,T)$, is a key measure. Critical systems theory predicts that avalanche shapes are scale-invariant [114] and thus obey the shape collapse function,

$$\langle n \rangle(t,T) = T^{(1-\chi_{\text{coll}})} \cdot \mathscr{F}\left(\frac{t}{T}\right) \tag{1.4}$$

For neuronal avalanches, the universality class is thought to be mean-field directed percolation (MF-DP). This universality class is associated with the following critical exponent values: $\alpha = \frac{3}{2}$, $\beta = 2$, and $\chi = 2$; as well as a parabolic shape profile with $\chi_{\text{coll}} = 2$.

However, accurately quantifying neuronal avalanches also presents several challenges. The limited spatial and temporal resolution of recording techniques can underestimate avalanche size and frequency due to incomplete capture of neuronal activity within a region. This issue is further compounded by subsampling, where in vivo recordings typically capture only a small fraction of the total neuronal population, leading to biased estimates of the avalanche distribution. We will discuss methods to address this key challenge throughout this dissertation. Additionally, distinguishing true neuronal activity from background

noise can be difficult, potentially introducing false positives. However, advancements in recording techniques and machine learning algorithms [81] have helped us optimize signal-to-noise ratios within the constraints of our experimental methods. Finally, finite-size effects can impact the analysis. Recordings often represent a limited portion of the brain, potentially affecting the observed scaling properties of avalanches. Furthermore, accurate estimation of critical exponents often requires large datasets, which can be computationally expensive to analyze. However, due to the comprehensive nature of the neuronal avalanche framework, overcoming these challenges in data acquisition and analysis is crucial. Unveiling the true dynamics of neuronal avalanches will provide a deeper understanding of information processing and communication within the brain, potentially unlocking new insights into brain function and dysfunction.

1.6 Dissertation outline

This dissertation begins with an exploration of efficient machine learning algorithms to address the problem of predicting the dynamics of large, complex networks. This chapter is based on the published work titled "Parallel Machine Learning for Forecasting the Dynamics of Complex Networks," which appeared in Physical Review Letters 128, 164101 (2022). Chapters 3 and 4 focus on criticality in both real and simulated biological neural networks, examining methods to overcome subsampling in such systems and the use of temporal coarse-graining to recover the scaling exponent. These chapters are based on the publications titled "Parabolic Avalanche Scaling in the Synchronization of Cortical Cell Assemblies," published in Nature Communications 14, 2555 (2023), and "The Recovery of Parabolic Avalanches in Spatially Subsampled Neuronal Networks at Criticality," which is currently under peer review. Chapter 5 integrates the fields of neuroscience and machine learning, exploring relative excitatory-inhibitory mechanisms and plasticity within

the context of reservoir computing. This chapter summarizes ongoing research and is being prepared for publication.

Chapter 2: Parallel machine learning for predicting the dynamics of complex networks

This chapter is adapted from the publication "Parallel Machine Learning for Forecasting the Dynamics of Complex Networks" by Keshav Srinivasan, Nolan Coble, Joy Hamlin, Thomas Antonsen, Edward Ott and Michelle Girvan published as Physical Review Letters 128, 164101 (2022).

2.1 Overview

This chapter marks the beginning of our exploration into the design of efficient machine learning algorithms. In particular, we focus on forecasting the dynamics of large, complex, and sparse networks using historical time series data. Accurate predictions of these network dynamics are crucial across various applications. In this chapter, we propose a machine learning approach that uses a parallel architecture to simulate the topology of the network in question. We demonstrate the effectiveness and scalability of our method, which employs reservoir computing to predict a chaotic network of oscillators.

We consider two scenarios with different levels of prior knowledge: (i) when the network links are known, and (ii) when the network links are unknown and need to be inferred through a data-driven approach, allowing for approximate optimization of prediction. This distinction offers insight into how machine learning can adapt to varying degrees of available information and still yield robust forecasting results.

2.2 Introduction

Machine learning (ML) has played a vital role in recent scientific advances in many disciplines. A key problem in these contexts is time series prediction of a dynamical system for which a first-principles, knowledge-based description is unavailable [66]. By using ML in combination with measured time-series data, one can hope to construct a faithful model of a system's dynamics and to then use this model to predict the future evolution of the system's state. Our aim in this chapter is to address this goal for *large systems of interacting components with complex connectivity and dynamics* - a system type of enormous technical and scientific interest in many fields, ranging, e.g., from neuroscience to power grids. However, straightforward application of the standard ML prediction schemes becomes problematic when applied to forecasting the dynamics of large networks. Leveraging the typical prevalence of sparse network topology in such systems, we introduce a parallel method for forecasting node dynamics. The key oft-occurring enabling system characteristic that our method requires is that the number of network nodes providing inputs to any node is small relative to the total number of nodes in the network. In our approach, we construct an ML architecture that mimics the topology of the network. Each node of the network to be predicted is assigned an individual small ML device and these individual ML devices are linked to each other based on the underlying connectivity of the network (either known *a priori* or inferred from the available time series data). We demonstrate and test this method by applying it to a network of Kuramoto oscillators [2, 78] constructed to exhibit chaotic dynamics. Our method is motivated in part by previous work on parallel ML prediction of large spatiotemporally chaotic systems [4, 129].

We consider two scenarios: (a) the connectivity of the network is known, and (b) the connectivity of the network is unknown *a priori*, yet may be approximately inferred from node time series data. Scenario (a) serves two purposes: first, as preparation for the more

challenging situation presented by scenario (b), and second, as a method applicable to cases where the connectivity is, in fact, known. The main conclusion of this work is that our proposed parallel ML scheme enables data-based network dynamics prediction in cases that would otherwise (i.e., without parallelization) be unattainable.

2.3 Methods

2.3.1 Test system: Frequency assoratative Kuramoto network

In order to demonstrate and test our approach, we consider the well-studied Kuramoto model of N network-coupled oscillators,

$$\dot{\theta}_i = \omega_i + K \sum_{j=1}^{N} A_{ij} \sin(\theta_j - \theta_i), \tag{2.1}$$

where θ_i is the phase angle of oscillator i, ω_i is the natural frequency of oscillator i when uncoupled, K is the coupling strength, and A_{ij} is the adjacency matrix that specifies the structure of the oscillator network ($A_{ij} = 1$, if there exists a network link from node j to node i with $i \neq j$, and $A_{ij} = 0$, otherwise). Here we consider an undirected ($A_{ij} = A_{ji}$), frequency assortative Kuramoto network [143]. By 'frequency assortative', we mean that two nodes are more likely to be linked if their natural oscillation frequencies are numerically close. The resulting frequency assortative system has chaotic dynamics for certain choices of parameters [168], hence serving as a good example of complex network dynamics whose evolution is challenging to predict. Each node is taken to have the same number of connections (this number is called the node's degree). The oscillator natural frequencies, ω_i, are drawn from a uniform random distribution from $-\pi/2$ to $\pi/2$. The frequency assortative network (i.e., the set of matrix elements A_{ij}) is constructed by starting with N_0 unlinked nodes, each with its assigned frequency (ω_i for node i) and then successively adding links,

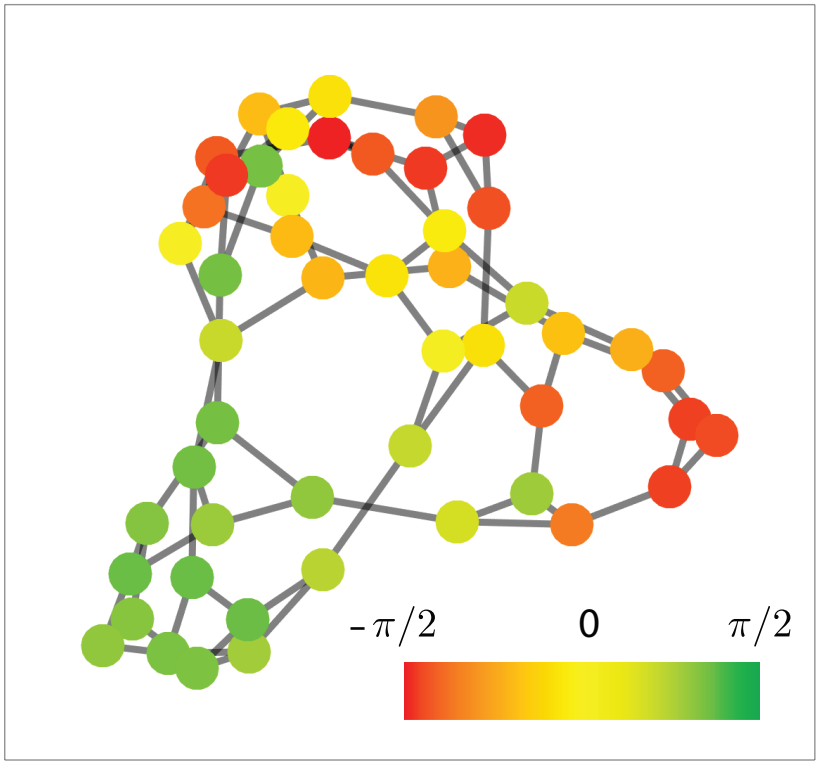

Figure 2.1: Frequency assortative network structure with 50 nodes displaying the natural frequencies of the individual oscillators, with nodes colored by frequency. Similar frequencies are connected with higher probability.

as follows. After randomly choosing a node i that still requires additional links, we next randomly pick another node j (not already connected to node i) which also still requires additional links, and then, with probability p_{ij}, we link nodes i and j, where

$$p_{ij} \propto \frac{\delta^\gamma}{\delta^\gamma + |\omega_i - \omega_j|^\gamma} \tag{2.2}$$

We continue in this way to make links until all nodes have the desired degree, d. Due to the form of p_{ij} (Eq. (2.2)), nodes with similar natural frequencies are connected with a higher probability (see Fig. 2.1). We use the global order parameter, R, as a metric to measure the

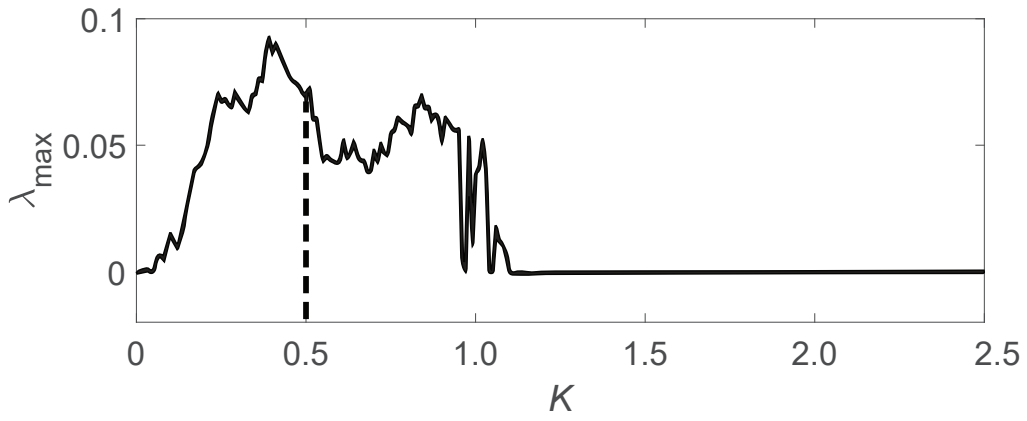

Figure 2.2: Largest Lyapunov Exponent as a function of the coupling constant, K. The dashed line represents the chosen value of K for our studies.

dynamics of the oscillator network, where

$$R(t) = \frac{\sum_{i,j=1}^{N} A_{ij} e^{i\theta_j}}{N\langle d \rangle} \tag{2.3}$$

For our frequency assortative network, with $N = 50$, a fixed nodal degree, d, of 3, $\delta = 0.8$, $\gamma = 5$ and $K = 0.5$ (our standard parameter set for most of out subsequent numerical experiments), we observe chaotic behavior, which is confirmed by the positive value of λ_{max}, the largest Lyapunov exponent of the system (Fig. 2.2).

2.3.2 Background on non-parallel reservoir computing prediction

In this paper we use Reservoir Computing (RC) [64, 98] as our ML scheme, because of its demonstrated utility for time series prediction [23, 33, 66]. We consider a reservoir computer constructed with an artificial high dimensional dynamical system, known as the reservoir, which is coupled to an input through an input layer, specified by a matrix $\mathbf{W_{in}}$

which maps the N_u dimensional input vector, $\mathbf{u}$, at discrete time t, to the reservoir state variables, which are collectively expressed as the scalar components of the reservoir state vector $\mathbf{r}$. In our RC implementation, the reservoir is a network (not to be confused with the network, e.g., Eq. (2.1), whose state we desire to predict), and the kth component of the vector $\mathbf{r}$ is the scalar state of reservoir node k. The RC network is directed, sparse, and random with N_r nodes having average input degree, $\kappa = 3$. The RC adjacency matrix is denoted B, with matrix elements $B_{kk} = 0$, and B_{kl} for $k \neq l$ chosen randomly and uniformly from $[-\beta, \beta]$ where β is chosen to yield a maximum eigenvalue of B denoted ρ (known as the spectral radius). Each input to the reservoir is sent to N_r/N_{in} reservoir nodes, where N_{in} is the number of inputs to the RC (Note: N_r is chosen to be an integer multiple of N_{in}). The input matrix, $\mathbf{W_{in}}$, is then a $N_r \times N_{in}$ dimensional matrix. The elements of $\mathbf{W_{in}}$ are chosen so that every node in the reservoir receives exactly one input from $\mathbf{u}(t)$ while each input in $\mathbf{u}(t)$ is connected to N_r/N_{in} nodes in the reservoir network (see Sec. 2.3.4 for further discussion). The non-zero elements are drawn from a uniform random distribution from $[-\sigma,\sigma]$, where σ is the input scaling. The reservoir state, $\mathbf{r}(t)$, is taken to evolve according to

$$\mathbf{r}(t+\Delta t) = \alpha \mathbf{r}(t) + (1-\alpha)\tanh[\mathbf{B}\mathbf{r}(t) + \mathbf{W_{in}}\mathbf{u}(t)], \tag{2.4}$$

where the tanh function is applied component-wise to its vector argument. Here α is the leak rate which controls the timescale of the reservoir nodes. The output of the system, $\tilde{\mathbf{u}}$, is obtained through the $N_r \times N_u$ dimensional output layer matrix $\mathbf{W_{out}}$

$$\tilde{\mathbf{u}}(t) = \mathbf{W_{out}}\mathbf{r}(t). \tag{2.5}$$

For the task of time-series prediction, the reservoir computer is used in two different modes: a training mode and a prediction mode. In the training mode, the reservoir computing

system, represented by Eqs. (2.4) and (2.5), is run for the time interval over which training data $u(t) = u(n\Delta t)$ $(n = -n_t, (1-n_t), (2-n_t), ..., 0)$ is available, $\mathbf{r}(n\Delta t)$ is computed, and the output matrix $\mathbf{W_{out}}$ is adjusted ('trained') so that the output of the reservoir computer $\tilde{\mathbf{u}}(t)$ best approximates $\mathbf{u}(t)$. This is done through a ridge regression procedure, wherein we minimize the error summed over the training times $t = n\Delta t$ for n running from $1-n_t$ to 0,

$$\min_{\mathbf{W}_{\text{out}}} \left\{ \sum \left[\|\mathbf{W}_{\text{out}}\mathbf{r}(t) - \mathbf{u}(t)\|^2 \right] + \beta \, \text{Tr} \left(\mathbf{W}_{\text{out}} \mathbf{W}_{\text{out}}^T \right) \right\} \tag{2.6}$$

Here β is the Tikhonov regularization parameter that is used to prevent over-fitting. The quantities (N_r, ρ, σ, α and β), referred to as hyperparameters of the reservoir computing setup, are collectively used to control the performance of system. In this paper, we chose the hyperparameters by a subsequent iterative process approximately maximizing the valid prediction time (See Eq. (2.8)) over the hyperparameters via a coarse grid search (See Table 2.1). In the prediction mode, the reservoir state now evolves autonomously in "closed-loop" mode; i.e., the output at time t, now serves as the input at time $t + \Delta t$,

$$\mathbf{r}(t + \Delta t) = \tanh[\mathbf{B}\mathbf{r}(t) + \mathbf{W_{in}}\mathbf{W_{out}}\mathbf{r}(t)]. \tag{2.7}$$

This procedure generates a predicted time series $\hat{\mathbf{u}}(n\Delta t) = \mathbf{W_{out}}\mathbf{r}(t)$ that is assumed to approximate the true future evolution of the state of the system, $\mathbf{u}(t)$ at a time $n\Delta t$ for $n > 0$ (we choose Δt small compared to the time scale for variation of u so that $u(n\Delta t)$ essentially specifies the continuous time function $\mathbf{u}(t)$).

2.3.3 Parallel machine learning architecture

In order to address the high computational complexity of predicting large networks, we introduce a parallel network RC architecture (see the schematic in Fig. 2.3; for code see Ref [87]). Each node, i, in the predicted network is assigned its own reservoir, R_i. The inputs to this reservoir are the signal of node i itself, as well as that of the nearest network neighbors of node i. The number of such neighbors is equal to the network degree. The reservoir R_i is then trained on these inputs to predict the signal of node i. Because each R_i predicts just one node, its size N_r can be relatively small. In addition, since our parallel scheme uses an interconnected network of independently trained reservoirs, we can efficiently parallelize our training process, making the system scalable to large networks.

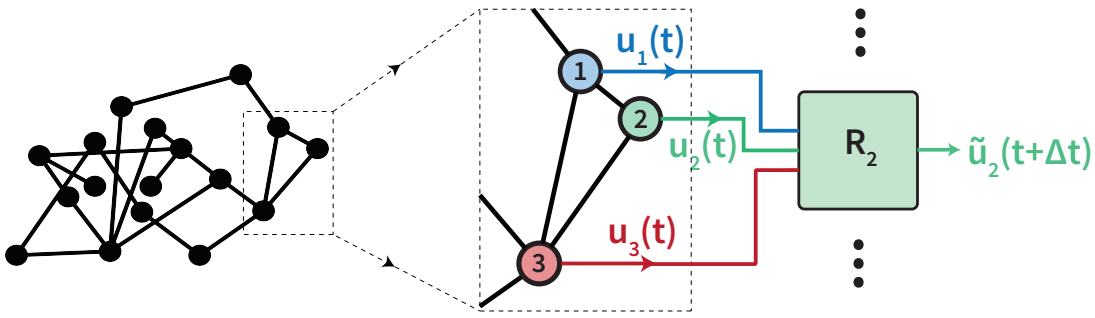

Figure 2.3: A schematic diagram for the parallel network ML architecture. Here we show Reservoir 2 (R_2), which receives input from its assigned node (node 2), plus inputs from nodes connected to node 2 (i.e., nodes 1 and 3). R_2 is then is then trained to predict its assigned node (node 2). This process is the same for each node in the network, such that the connectivity among the reservoirs mimics the network to be predicted.

2.3.4 Generating the Reservoir Input Matrix, $\mathbf{W_{in}}$

In this section, we provide more details on how we generated the input matrix, $\mathbf{W_{in}}$, that feeds into the reservoir for predicting our Kuramoto oscillator network. Here the state

of each oscillator, i, is specified as a tuple $[\sin\theta_i(t), \cos\theta_i(t)]$. The N_{in}^{osc} input oscillators to a particular reservoir are collapsed into a single input vector, $\mathbf{u}$, with N_{in}^{tot} ($N_{in}^{tot} = 2N_{in}^{osc}$) elements such that $\mathbf{u}_i(t) = \sin\theta_i(t)$, $\mathbf{u}_{(N_{in}^{osc}+i)}(t) = \cos\theta_i(t)$ for $1 \leq i \leq N_{in}^{osc}$. The elements of $\mathbf{W_{in}}$ are chosen so that each of the N_r nodes in the reservoir receives exactly one input from amongst the $2N_{in}^{osc}$ components of the vector $\mathbf{u}(t)$. In the parallel scheme with known neighbors, all the N_{in}^{tot} inputs to a reservoir are treated equally; i.e. each input in $\mathbf{u}(t)$ is connected to the same number (N_r/N_{in}^{tot}) of reservoir nodes (Note: N_r is chosen to be an integer multiple of N_{in}^{tot}). The non-zero elements of $\mathbf{W_{in}}$ are then randomly drawn from a uniform distribution in $[-\sigma,\sigma]$, where σ is the input scaling.

We compare the performance of an all-to-all input matrix to one that is generated as mentioned above, with the input scaling, σ, optimized separately for each case ($\sigma = 0.3$ for an all-to-all input matrix, for other hyperparameters see Table 2.1). Fig. 2.4 clearly highlights the superior performance of the latter method for a Kuramoto oscillator network with all nodes having degree 3, where each reservoir in the parallel scheme (with known neighbors) receives 8 inputs from exactly 4 oscillators: 1 "assigned" + 3 neighbors. The assigned oscillator is the oscillator the reservoir is trained to predict. As we increase the number of inputs to a reservoir, it might be beneficial for each input to be connected to a non-exclusive subset of the reservoir. This would mean each input is still connected to a fixed number of reservoir nodes, but we can relax the criteria that each node in the reservoir is connected to exactly one input.

For parallel prediction with unknown neighbors, we must slightly alter our scheme. In this case, each of the input oscillators to a reservoir are not treated equally. Of the N_{in}^{osc} input oscillators, we only know 1 "true" input oscillator- the assigned oscillator that reservoir is intended to predict. We cannot distinguish the true neighbors from the false positives in

the remaining $N_{in}^{osc} - 1$ input oscillators. We can account for this by "reserving" a larger part of the reservoir for the assigned oscillator and dividing the rest of the reservoir equally between all the inferred neighbors. This means, each of the 2 inputs from the assigned oscillator are connected to $N_{in}^{assign}/2$ reservoir nodes, while the remaining $N_{in}^{tot} - 2$ inputs are each connected to $(N_r - N_{in}^{assign})/(N_{in}^{tot} - 2)$ reservoir nodes. Here $N_r - N_{in}^{assign}$ is chosen to be an integer multiple of $N_{in}^{tot} - 2$ and hence the number of reservoir nodes used in this case can only be set approximately equal to one with known neighbors. For example, in our system with known neighbors we set $N_r = 200$. But for the case with unknown neighbors, where we set $N_{assign} = 50$, if we have 7 inferred neighbors for a particular oscillator, the N_r for the corresponding reservoir must be set to 204 to satisfy the criteria.

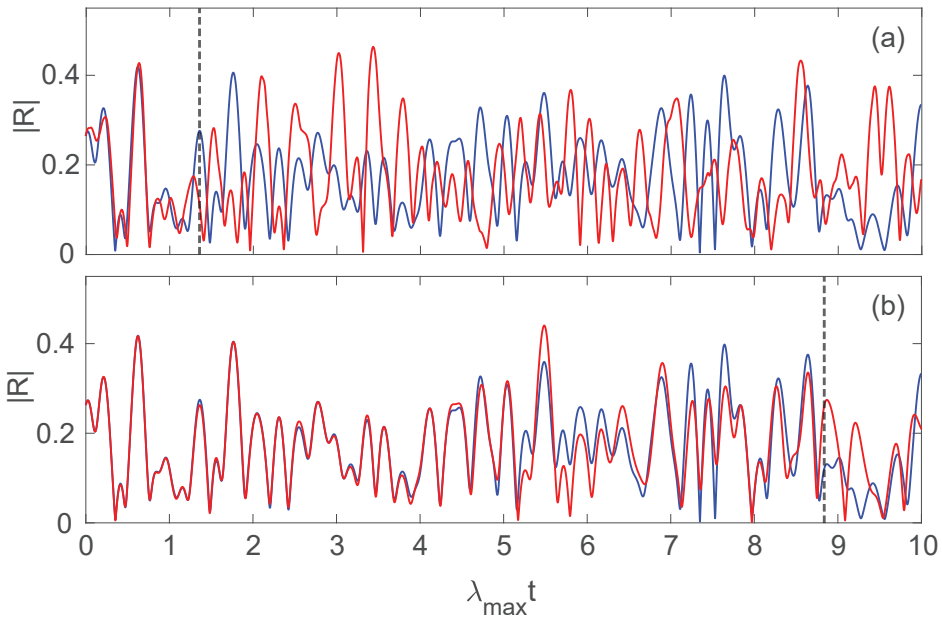

Figure 2.4: Parallel prediction of Kuramoto system with known neighbors. The blue curves are the data and red curves are predictions. The dotted lines in each plot denote the valid prediction time. (a) Using an all-to-all input matrix, i.e., all inputs map to all reservoir nodes. (b) Using an input matrix where each input is sent to a different disjoint set of reservoir nodes.

2.3.5 RC hyperparameters

We list the RC hyperparameters (Table 2.1) used to plot the results shown in this chapter. These hyperparameters were chosen by a subsequent iterative coarse grid search, which approximately maximized the valid prediction time.

	Res. nodes N_r	Spec. rad. ρ	Input scaling σ	Leak rate α	Regul. β
Single Reservoir	10000	0.9	0.1	0	10^{-7}
Parallel Scheme					
a) Known Neighbors	200	0.9	0.6	0.1	10^{-9}
b) Unknown Neighbors	≈ 200*	0.9	0.5	0.1	10^{-8}
Coupled Lorenz System	1000	0.9	0.2	0.1	10^{-7}

Table 2.1: Values of reservoir hyperparameters obtained via a coarse grid search for a single reservoir prediction as well as for the parallel reservoir scheme. *Reservoir size set approximately to 200; for details, see Section 2.3.4. **See Section 2.4.7

2.3.6 Tests

To compare the parallel RC scheme with the single RC approach, we use a $N_0 = 50$ node frequency assortative Kuramoto oscillator ($\delta = 0.8$, $\gamma = 5$, Eq. (2.2)) network with a coupling constant of $K = 0.5$ (See Section 2.4.7 for results using a coupled Lorenz system). We study the magnitude of the global order parameter $|R|$ (see Eq. (2.3)), which tells us about the network-level activity (see Fig. 2.5) and the prediction of the evolution of individual node states (see Fig. 2.6), both of which show the same main qualitative behavior. For the purpose of forming inputs to the reservoir, we specify the state of the oscillator i as $[\sin\theta_i(t), \cos\theta_i(t)]$. The input matrix is generated as described above and in Section 2.3.4. The degree of the test Kuramoto network is three links per node. This method works best

for networks with low in-degrees. Large in-degrees and/or the addition of in-hubs can reduce predictive performance (See Sections 2.4.5 and 2.4.7). These effects may be mitigated in part by increasing the number of nodes, N_r, in the parallel reservoirs.

2.4 Results

2.4.1 Single non-parallel reservoir prediction

The single reservoir computer prediction can fail as the size of the network we want to forecast increases. This is clearly demonstrated in Fig. 2.5a, where the prediction breaks down in a fraction of a Lyapunov time, $\lambda_{max}t$. We quantify the duration of an accurate prediction by a metric that we call the "valid prediction time". This metric is defined as the amount of time elapsed before the normalized root mean squared prediction error (NRMSE), $E(t)$, exceeds some chosen value f, $0 < f < 1$, for the the first time, where

$$E(t) = \frac{\|\mathbf{u}(t) - \widetilde{\mathbf{u}}(t)\|}{\langle \|\mathbf{u}(t)\|^2 \rangle^{1/2}}. \tag{2.8}$$

The valid prediction time for $f = 0.1$ is marked in Fig. 2.5 by vertical dotted lines. Even for the very large reservoir (N_r=10000), close to the limit of our computer resources, that is used in Fig. 2.5a, the system is still not able to predict past a fraction of a Lyapunov time.

2.4.2 Parallel scheme with known links

In cases where the network structure is known *a priori*, such as in certain social networks, we can construct our parallel reservoir architecture by using the known network links. In the case of our Kuramoto oscillator network, we demonstrate our results on the 50 node network by using 50 separate parallel reservoirs of relatively modest size ($N_r = 200$ as

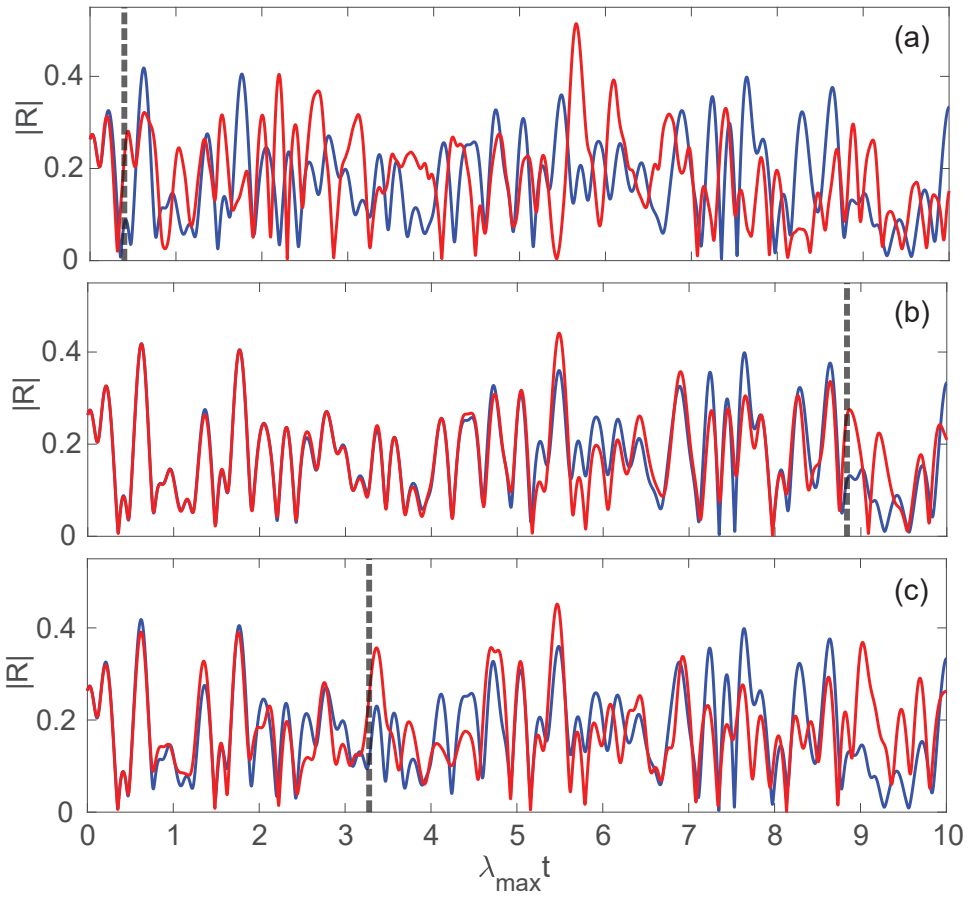

Figure 2.5: Prediction of the order parameter in 3 different cases of ML implementations. The blue curves are the data (the truth) and red curves are predictions. The dotted lines in each plot denote the valid prediction time. (a) Single, non-parallel RC prediction using a large reservoir ($N_r = 10000$), (b) Parallel prediction with known network links, using 50 separate reservoirs each having modest size ($N_r = 200$), (c) Parallel prediction with unknown network links, using 50 separate reservoirs each having modest size ($N_r = 200$). The network structure is estimated by using transfer entropy as a metric to draw network edges.

compared to $N_r = 10000$ for Fig. 2.5a), each having the same set of hyperparameters [see Fig. 2.5b]. The predictive performance of this architecture could potentially be enhanced by individually optimizing hyperparameters for each of the 50 reservoirs, but this would considerably increase both the time and computational resources required for this task. As seen from Fig. 2.5b, our parallel scheme does exceedingly well for multiples of the Lyapunov time, $\lambda_{max}t$. This is particularly clear from a comparison of valid prediction time (vertical dashed lines) in Fig. 2.5a versus those in Fig. 2.5b, the latter being $\gtrsim 10$ times larger, while at the same time being much less computationally demanding (mainly due to the difference in N_r, $N_r = 10000$ for the nonparallel case versus $N_r = 200$ in the parallel case).

2.4.3 Parallel scheme with unknown links

In many cases of interest, one may not have information about the underlying network structure. Using nodal time-series data for finding links in networks, such as metabolic [60] and gene-regulatory networks [11], is an active area of current research. Many heuristic-based [95] and statistics-based tools like conditional mutual information [174], and correlation [76], as well as recent nonlinear dynamics-based [28] and machine-learning-based techniques [10], have been used for link inference and might give useful approximations of the underlying network structure. These methods could then potentially be used in our parallel network scheme.

Link inference methods like the ones discussed above, typically use past nodal time series measurements to assign a score between each ordered pair of nodes in the network reflecting the likelihood of a directed link existing from the source to the target node. As we decrease this threshold, we draw more links and hence increase the average number of inferred neighbors for each node. Initially, decreasing the threshold, or in other words

increasing the number of inferred neighbors, increases the number of true positive links and improves the predictive performance of our reservoir scheme. However, if the threshold is decreased too much, the number of false positives increases substantially and significantly degrades the predictions. Since our goal is prediction, we view the link-inference threshold as an additional hyperparameter and choose it (along with the other hyperparameters) so as to optimize the valid prediction time. By this procedure, we effectively bootstrap our prediction process to determine the link threshold criterion. An example set of results for transfer entropy [159] for $N_r \approx 200$ (See Sec. 2.3.4) is shown in Fig. 2.5c. Again, in marked contrast with the results in Fig. 2.5a for a large single RC ($N_r = 10000$), we obtain good predictions, e.g., a valid time between 3 and 4 Lyapunov times for $|R|$.

2.4.4 Node-level predictions

In the previous section, we show the performance of our parallel RC scheme for predicting the macroscopic dynamics of the system. Here, we demonstrate that our approach also captures the node-level dynamics. Fig. 2.6 shows that a single large reservoir ($N_r = 10000$) is unable to capture the dynamics of individual oscillators in the network. However, our parallel schemes with known edges [Fig. 2.6, as well as unknown edges [Fig. 2.6], are able to capture the dynamics of the individual oscillators. The parallel schemes, also give good predictions of the system "climate" (i.e. the prediction still seems to resemble the original dynamics) even after the marked valid prediction time (vertical dotted lines in Fig. 2.6).

2.4.5 Effect of Hubs

Earlier, we mentioned that having hub nodes in the network can affect the performance of our parallel scheme. In this section, we analyze that effect in greater detail. For this section, we picked a node at random from the 50-node oscillator network used previously,

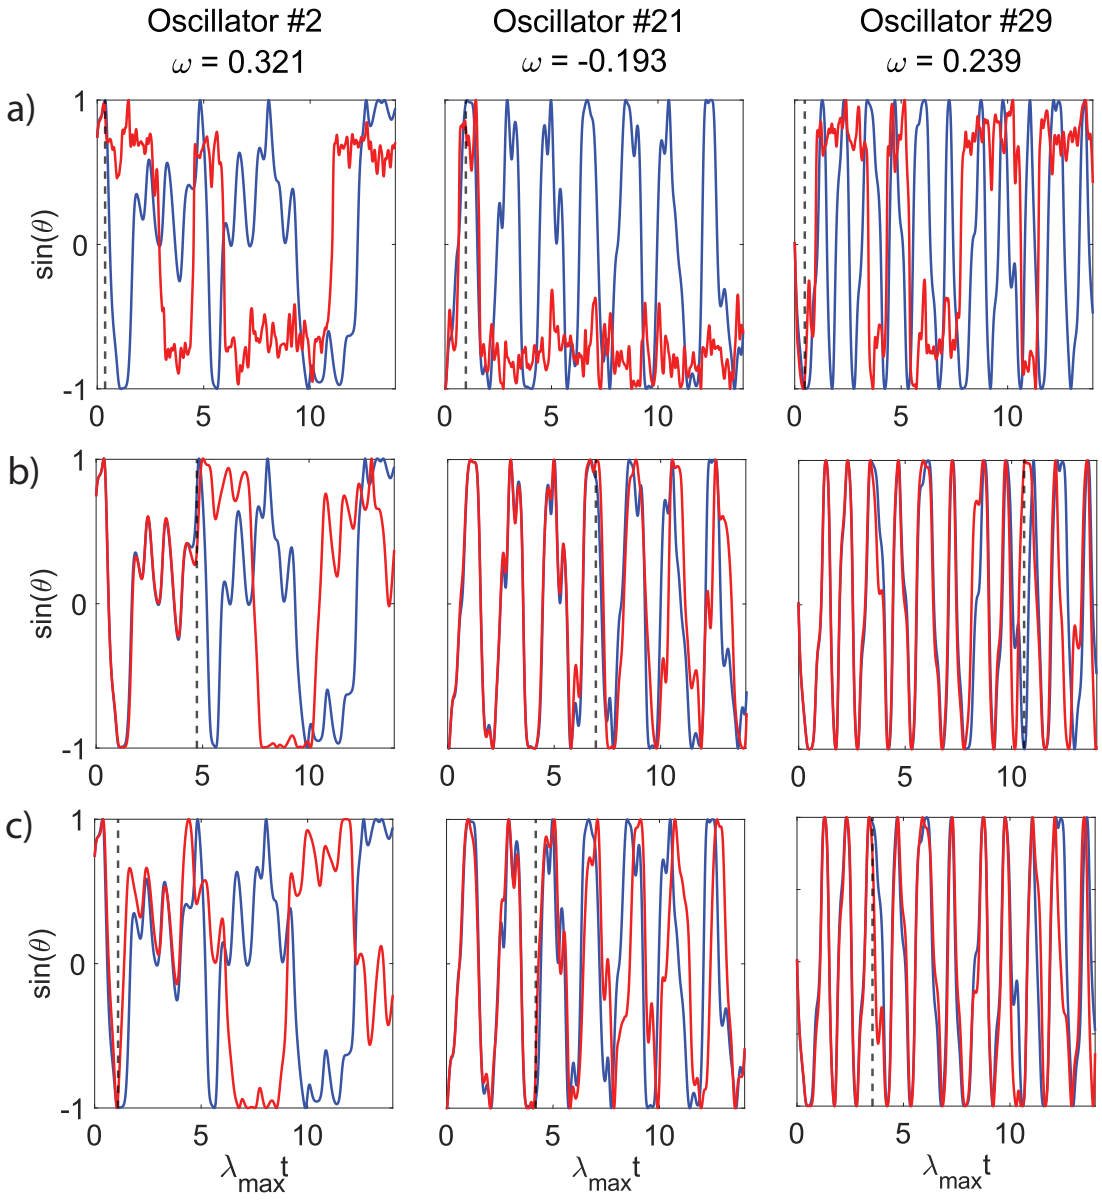

Figure 2.6: Prediction of the dynamics of three representative individual oscillators. The natural frequency of the oscillators is noted above. The blue curves are the data, and the red curves are predictions. The vertical dotted lines in each plot denote the valid prediction time. (a) Single, non-parallel RC prediction using a large reservoir ($N_r = 10000$), (b) Parallel prediction (with known network links) using 50 separate reservoirs each having modest size ($N_r = 200$), (c) Parallel prediction (with unknown network links) using 50 separate reservoirs each having modest size ($N_r \approx 200$). For row (c) the network structure is estimated by using transfer entropy as a metric for inferring network edges.

where all nodes had an in and out-degree of 3 and added additional in or out links to it until they reached the desired degree. In each case, we ran our parallel prediction scheme 50 times, with different random initializations of the reservoir and input weights, and reported the prediction time of the absolute value of the order parameter, $|R|$, as a measure of overall predictive performance. As we can see in the box plot below (Fig. 2.7), our parallel scheme is robust to the addition of an out-hub, which agrees with our expectations. However, adding an in-hub did reduce the predictive performance. Due to the architecture of the parallel method, its performance is affected more strongly by the in-degree of a node than its respective out-degree.

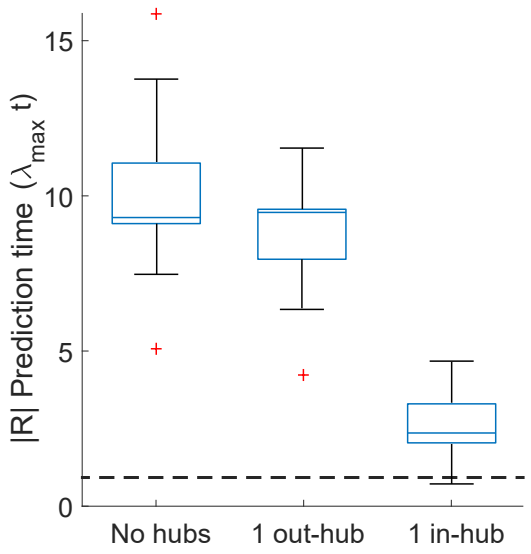

Figure 2.7: Predictive performance of the parallel method for different network architectures. a) A 50-node assortative Kuramoto oscillator network (degree = 3). b) Network with one out-hub (out-degree = 20) c) Network with one in-hub (in-degree = 20). The dashed line shows the prediction time of the 50-node network with no hubs using a single large reservoir ($N_r = 10000$). The blue line represents the median value of the prediction time, and the blue box represents the inter-quartile range. The black bars mark the lower and upper adjacent values. The red points are potential outliers.

2.4.6 Dependence on size of the predicted network

Considering the case where the links are unknown, we test the performance of our parallel approach as a function of network size using two different link inference methods: first, we use the frequently employed method based on transfer entropy [159], implemented via a MATLAB package [133]; and second, we use a method called Algorithm for Revealing Network Interactions (ARNI) [28] using code available online [32]. Fig. 2.8 shows a plot of the valid prediction time as a function of the oscillator network size, N_0. As we increase the size of the oscillator network, the prediction using a single reservoir ($N_r = 10000$) quickly degrades and becomes unable to capture the network dynamics at all. Since the parallel method assigns a reservoir to each oscillator in the network, for the case with known links, as expected, it maintains constant performance within the estimated uncertainty of the valid times. In the case of link inference using transfer entropy (TE), the performance is far better than for the single reservoir, but not nearly as strong as for the case of link inference using ARNI. For a small oscillator network size of $N_0 = 10$, TE has a true positive rate of 100%, and a false discovery rate of 37.5%. ARNI, on the other hand, has a true positive rate of 100% and a false discovery rate of 3.3%. When we increase the network size to $N_0 = 500$, TE has a true positive rate of 96% and a false discovery rate of 69.6%. ARNI gives us improved performance with a true positive rate of 96% and a false discovery rate of only 7.3%. The parallel approach can tolerate imperfect link inference with a significant false discovery rate as long as the true positive rate is high enough. The false negative inference of a link to node i deprives reservoir R_i of vital information needed for prediction of the state of node i. In contrast, reservoir R_i can compensate for a false positive link from node j to node i by learning, through its training, to ignore its time series input from node j.

We emphasize that we have not tried to determine an optimal link inference method, and

other methods may yield longer prediction times. We use transfer entropy as a relatively standard link inference method with which to demonstrate our network parallel approach and ARNI as an appealing link inference choice for further investigation in other contexts.

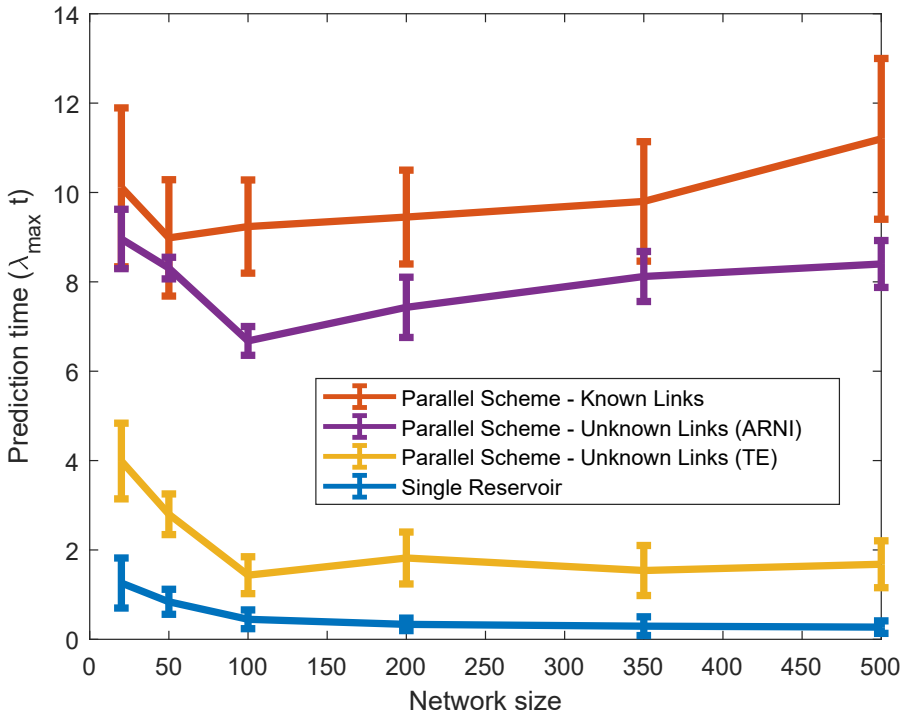

Figure 2.8: Performance of the different Reservoir Computing methods as a function of the Kuramoto oscillator network size

2.4.7 Coupled Lorenz System

In this section, we demonstrate the utility of our parallel method in a network of weakly coupled Lorenz systems [10] governed by the following equations:

$$dx_k/dt = -10\left[x_k - y_k + c\sum_{l=1}^{N} a_{kl}\,(y_l - y_k)\right] \tag{2.9}$$

$$dy_k/dt = 28x_k - y_k - x_k z_k \tag{2.10}$$

$$dz_k/dt = -(8/3)z_k + x_k y_k \tag{2.11}$$

Here a_{kl} is the network adjacency matrix which is chosen to be a random regular graph where each node has the same fixed in-degree, and c is the coupling constant which was set to 0.09 (degree 3) or 0.02 (degree 10). The coupling constant was chosen to be strong enough so that the node dynamics were significantly different from the traditional Lorenz ($c = 0$) system yet weak enough to avoid global synchrony.

Since a node state in this network is 3-dimensional, we use a larger reservoir for each node ($N_r = 1000$). As we can see in the left panels of Fig. 2.9, a single large reservoir ($N_r = 10000$) is unable to predict the node dynamics for even a very short time. In the middle panels, we see that the parallel method is able to predict the node dynamics for ≈ 9 seconds (in Lorenz time units) for a representative node. On the right panels, we observe that increasing the in-degree to 10 reduces the prediction time to ≈ 3 seconds for a representative node. To get more details on the prediction time of nodal dynamics across the network, we show a box plot of the 50 individual nodal prediction times for networks with different in-degrees. (Fig. 2.10). Here, we see that increasing the average in-degree reduces predictive performance, but the parallel method still performs better than the traditional method, which is unable to predict the nodal dynamics at all.

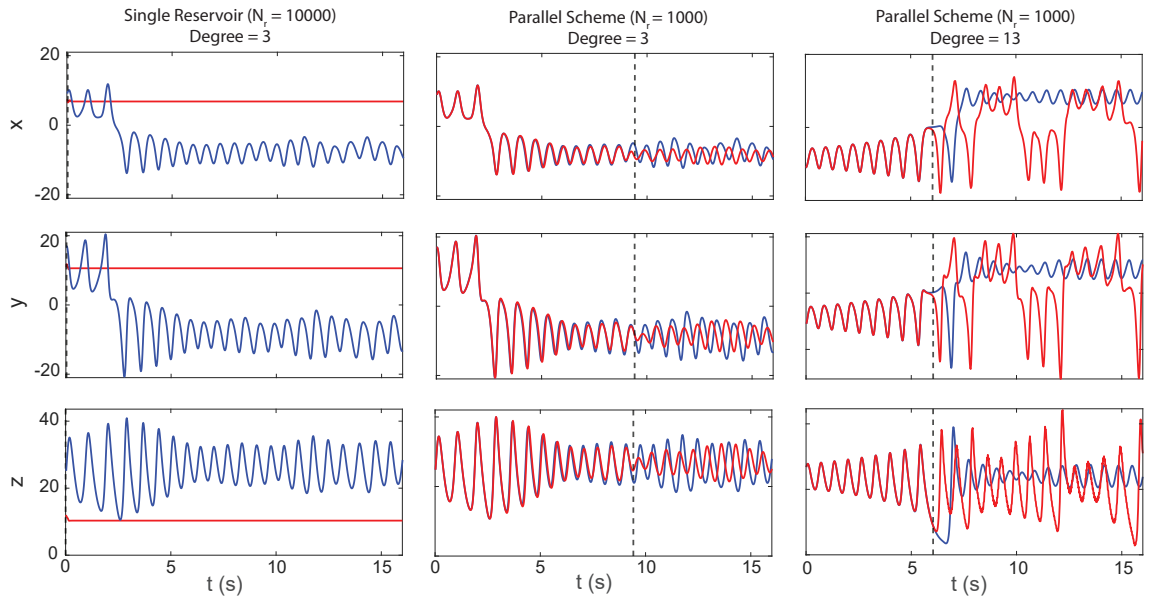

Figure 2.9: Prediction of the dynamics of a representative node from the coupled Lorenz network. The left panels are non-parallel RC predictions using a large reservoir ($N_r = 10000$) for a network of in-degree 3. The middle panels are parallel predictions using 50 separate reservoirs ($N_r = 1000$) for a network of in-degree 3. The right panels are parallel predictions using 50 separate reservoirs ($N_r = 1000$) for a network of in-degree 10. The blue curves are the data, and the red curves are predictions. The vertical dotted lines in each plot denote the valid prediction time.

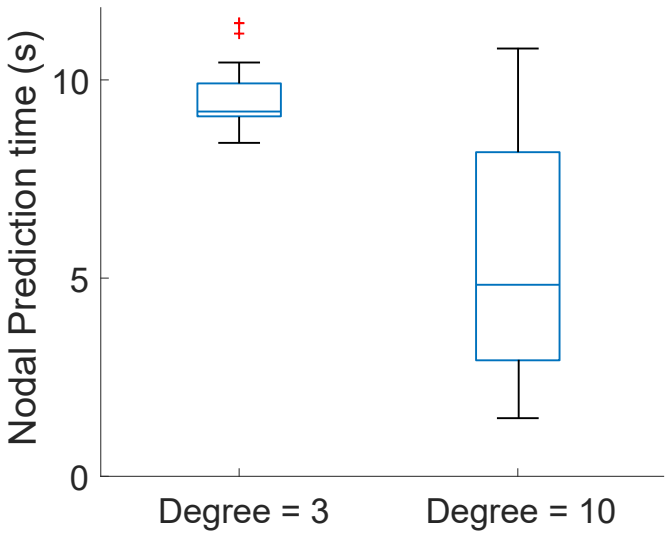

Figure 2.10: Predictive performance of the parallel method for the nodal dynamics of coupled Lorenz systems of different degrees. a) Fixed in-degree = 3. b) Fixed in-degree = 10. The blue line represents the median value of the prediction time, and the blue box represents the inter-quartile range. The black bars mark the lower and upper adjacent values. The red points are potential outliers.

2.5 Conclusions

We are able to construct accurate, data-driven forecasts for the dynamics of large complex networks using a parallel ML architecture that reflects the topology of the network to be predicted. We demonstrate our approach using a chaotic network of oscillators, but we believe it should be widely applicable in a variety of contexts. In cases for which a non-parallel approach with comparable resources fails, our scheme is successful when the network links are either known or unknown a priori. The parallel nature makes our approach scalable for extremely large networks, creating potential applications for many fields.

2.6 Acknowledgements

I would like to extend my heartfelt thanks to TREND program undergraduate students Nolan Coble and Joy Hamlin for their dedicated work on this project. I am also deeply grateful to the project advisors, Prof. Edward Ott, Prof. Thomas Antonsen, and Prof. Michelle Girvan, for their invaluable guidance and support. The project team collectively acknowledges the financial support from the National Science Foundation under Grant Nos. PHY-1461089, DGE-1632976, and DMS-1813027..

Chapter 3: Overcoming sub-sampling bias in evaluating the role of neuronal avalanches in the brain

This chapter is adapted from the publication "Parabolic avalanche scaling in the synchronization of cortical cell assemblies" by Elliott Capek, Tiago L. Ribeiro, Patrick Kells, Keshav Srinivasan, Stephanie R. Miller, Elias Geist, Mitchell Victor, Ali Vakili, Sinisa Pajevic, Dante R. Chialvo & Dietmar Plenz published as Nature Communications 14, 2555 (2023). Refer to Section 3.8 Acknowledgements for a detailed list of contributors responsible for collecting and analyzing the experimental datasets.

3.1 Overview

This chapter explores the phenomenon of neuronal avalanches in 2-photon imaging data from neurons in the superficial cortex of awake mice, as well as in a subsampled critical branching model. Our analysis reveals that synchronized cell assemblies organize into scale-invariant avalanches, where their size increases quadratically with their duration. This specific quadratic scaling pattern was observed only among correlated neurons and required temporal coarse-graining to compensate for the spatial subsampling inherent in imaging cortical tissue.

Our results suggest that cortical dynamics are critical, which we demonstrate through simulations of balanced excitatory/inhibitory (E/I) networks. These simulations and experimental data show that the time course of these avalanches follows an inverted parabola,

describing the synchronized firing of neurons for up to 5 seconds.

These parabolic avalanches were found to maximize temporal complexity as well as pairwise correlations. Our findings indicate the presence of a scale-invariant spatiotemporal order in the synchronization of cortical cell assemblies, suggesting that parabolic avalanches may play a fundamental role in cortical processing.

3.2 Introduction

Neuronal synchronization is a key concept in many theories about the cerebral cortex. Cortical neurons can integrate activity from neighboring cells, as well as from distant inputs [21], resulting in two primary synchronization patterns: oscillations [167] and cascades. The latter can take the form of waves [121, 139], synfire chains [40, 57], or neuronal avalanches[13, 15]. Cortical cascades occur when a group of neurons that fire synchronously facilitates synchronization in downstream neurons through select, converging connections. This mechanism can support robust yet flexible information processing in the cortex [13, 40, 57]. However, simulations suggest that both the growth of cascades and their overall temporal profile—the number of neurons or spikes encountered over time—can vary significantly, posing a challenge to the reliable transmission of information within the cortical network [1, 13, 40, 52, 135].

This variability issue is especially apparent in neuronal avalanches [13], which are highly diverse, scale-invariant cascades of neuronal activity primarily found in the superficial layers of the cortex. Avalanches can be detected in local field potentials [134] and can selectively engage single neurons [16, 136], with a high information capacity [14, 165]. However, it's uncertain whether avalanches measured at the single-cell level in vivo [15, 20, 70, 96] exhibit robust neuronal synchronization with a consistent and predictive unfolding pattern.

Traditionally, the defining characteristic of neuronal avalanches has been their scale-invariant spatiotemporal statistics, typically exhibiting power laws in both size and duration [13] with exponents $\alpha \approx 3/2$ and $\beta \approx 2$, respectively. More recent studies have suggested that the temporal profile, which describes how the size of an avalanche develops over time, can differentiate between models of avalanche generation. This scaling behavior, including the mean size relative to duration and a universal, duration-invariant temporal profile [77, 128, 163], can be encapsulated by a single scaling exponent χ.

However, models of avalanches that lack interactions, are driven by noise, or reflect certain 2PI data patterns often show noise-induced shifts in the recovered value of χ_{sh} to levels below 2, without a change in the key scaling factor (Fig. 3.5a). Additionally, better tracking of spiking activity can shift χ_{sh} toward smaller scaling values, as found in cases of deep interpolation (Fig. 3.5b). Our research found that when sampling fewer neurons (low f) and using less sensitive thresholds (high θ), greater temporal coarse-graining was needed to recover $\chi_{\mathrm{sh}} \approx 2$ (Fig. 3.2e). Our model also indicated that the scaling range ϕ for which $\chi_{\mathrm{sh}} \approx 2$ reflects a finite-size effect and can be consistently recovered within the limits of temporal coarse-graining (Fig. 3.2f).

Flat, many-generation avalanches showed $\chi_{\mathrm{lg}} \ll 2$, similar to those in our data, which appear in the cut-off of their size and duration distribution (Fig. 3.2c, d). Crucially, our simulations revealed that achieving $\chi_{\mathrm{sh}} \approx 2$ under spatial subsampling conditions is only feasible when the network operates with critical dynamics, whereas subcritical dynamics result in $\chi_{\mathrm{sh}} \approx 1$, irrespective of temporal coarse-graining (Fig. 3.2g; Fig. 3.6). These findings suggest that our 2PI data represent spatially subsampled activity that reflects critical dynamics in the cortex.

3.3 Experimental Methods

3.3.1 Overview of datasets analyzed

In the paper [24], six different datasets were analyzed and four distinct probes were used to monitor ongoing activity in the superficial layers of the frontal cortex (ACC/mPFC) and evoked activity in the superficial layers of the primary visual cortex. However, in this dissertation, to maintain focus and clarity, we concentrate on just one of these datasets that recorded ongoing activity in the contralateral ACC/mPFC, monitored using the calcium indicator jRGECO1a.

The study involved C57BL/6 mice (Jackson Laboratories), aged over six weeks, that were injected with a viral construct to express jRGECO1a in cortical neurons using the Syn promoter. Chronic 2-photon imaging (2PI) began after more than two weeks in the contralateral ACC/mPFC at an estimated depth of approximately 150–300 μm, using a microprism. Recordings were obtained from five mice (three males, two females; aged 8–20 weeks) with jRGECO1a expression, resulting in a total of 17 recordings, each lasting 30 minutes. These recordings were conducted over several days and analyzed separately for each mouse.

3.3.2 Animal surgery

All experimental procedures were approved by the NIH Animal Care and Use Committee (ACUC), and the experiments complied with the NIH Guide for the Care and Use of Laboratory Animals. Mice were obtained from Jackson Labs, with a C57BL/6 background, and bred in-house under a reversed 12:12-hour light/dark cycle. Chronic 2-photon imaging (2PI) was conducted on adult mice (older than 6 weeks), using a head bar in combination with a cranial window.

The cranial implants were constructed in one of two ways: (1) a microprism/coverslip assembly, used for imaging through the medial wall of the contralateral ACC/mPFC, as described in [92], or (2) a stack of circular glass cover slips, following established protocolsˆ67. For ACC/mPFC recordings, 2-3 injections of viral vectors (100–400 nL per injection, with a total volume under 1 μL, at 10^{13} vg/mL) were administered into the contralateral hemisphere relative to the prism implant, with injection sites at +0.75–1.25 mm anterior-posterior, 0.1–0.3 mm lateral, and approximately 250 μm below the pial surface. For V1 activity monitoring, the cranial window was placed centered about 2.5 mm from the midline in the right hemisphere, roughly 1 mm rostral to the lambdoid suture.

3.3.3 Pre-processing pipeline

Standard 2-photon imaging (2PI) was performed using a scanning microscope (Bergamo II series, B248, Thorlabs Inc.) coupled with a pulsed femtosecond Ti:Sapphire laser with dispersion compensation (Chameleon Discovery NX, Coherent Inc.). The microscope was controlled by ThorImageLS and ThorSync software (Thorlabs Inc.), and the laser wavelength was set to 1120 nm to excite jRGECO1a. The signals were collected using a 16 x 0.8 NA microscope objective (Nikon) and were separated through 607/70 nm band-pass filters. GaAsP photomultiplier tubes were used to detect the emitted photons.

The imaging field of view was approximately 450 x 450 μm, with frames of 512 x 512 pixels acquired at a frequency of 45.527 Hz using bidirectional scanning via a 12 kHz Galvo-resonant scanner. Beam turnarounds at the image edges were blanked with a Pockels cell, ensuring clean transitions. The average power used for imaging was below 70 mW, as measured at the sample. The resulting tif movie files in uint16 format underwent rigid motion correction using the Python-based software suite2p [125]. Following this, the registered images were further denoised using a machine-learning-based technique called deep

interpolation [81] (Deep-IP), and then processed by suite2p for region-of-interest (ROI) selection and fluorescence signal extraction.

For each labeled neuron, the raw soma and neuropil fluorescence signals (red for jRGECO1a) were extracted from their corresponding ROIs over time. Spiking probabilities were calculated from neuropil-corrected fluorescence traces ($F_{corrected} = F_{ROI} - 0.7 * F_{neuropil}$) using the MLspike library (https://github.com/MLspike), which includes an auto-calibration feature to determine unitary spike event amplitude, decay time, and channel noise for each ROI.

Deep interpolation [81](https://github.com/AllenInstitute/deepinterpolation) is a noise reduction technique that uses local spatiotemporal data from a noisy image stack to predict, or interpolate, pixel intensity values in a withheld central frame. This deep neural network uses a UNet-inspired encoder-decoder architecture with 2D convolutional layers and doesn't require ground truth data for training and validation. After rigid motion correction, individual denoised frames were generated by streaming a registered 60-frame stack (N˙pre = N˙post = 30 frames) through the Ai-93 pretrained model [81] to interpolate each frame. Given the imaging rate of approximately 45 Hz, these 60 frames covered about 1.3 seconds of data surrounding each frame to be interpolated. To assess the impact of Deep-IP, this step was omitted in our analysis of ongoing activity from the jRGECO1a recordings.

3.4 Subsampled critical neuronal model

3.4.1 Model Topology

We adapted the model developed by Girardi-Schappo and colleagues [49], an excitatory/inhibitory (E/I) balanced system consisting of integrate-and-fire (IF) neurons. This model exhibits four distinct firing domains: low-rate irregular, high-rate, quiescent, and

intermittent, along with a second-order continuous phase transition. At this phase transition, the model demonstrates the dynamics of a critical branching process with a slope of $\alpha = 3/2$ in the avalanche size distribution and $\beta = 2$ for the avalanche duration distribution, consistent with experimental observations of neuronal avalanches [13].

Our standard network consisted of $N = 10^6$ non-leaky, probabilistic integrate-and-fire neurons, with an E/I ratio of 4:1 to approximate the prevalence of excitatory over inhibitory neurons in the cortex. The neurons were connected in an all-to-all topology. The elements of the connectivity matrix W were initialized as $W_{EE} = W_{IE} = J$, $W_{II} = W_{EI} = -gJ$. Connections from excitatory presynaptic neurons had synaptic strength J, while those from inhibitory presynaptic neurons had synaptic strength $-gJ$. The synaptic balance parameter g was used to tune the model to an E/I-balanced state with critical dynamics, with a critical value $g_c = 3.5$ for our network parameters [49]. Values of g above g_c lead to inhibition-dominated, subcritical dynamics, while values below g_c result in excitation-dominated, supercritical dynamics.

3.4.2 Model Dynamics:

The state of each neuron in the model was described by two variables: a Boolean variable X, indicating whether a neuron is firing ($X(t) = 1$) or quiescent ($X(t) = 0$), and a membrane potential V, which controls the probability of firing. The membrane potential evolves as follows:

$$V_i(t+1) = \left(\mu V_i(t) + \sum_{j=1}^{N} W_{ij} X_j(t) \right) \cdot (1 - X_i(t)), \tag{3.1}$$

where μ is a leakage parameter, set to 0 in our simulations. The term $1 - X_i(t)$ creates an absolute refractory period by resetting the membrane potential after a spike, effectively

preventing the neuron from firing again in the next time step. The probability of a neuron firing is proportional to its membrane potential, with:

$$P(X = 1 \mid V) = \begin{cases} \Gamma V, & \text{if } \Gamma V < 1, \\ 1, & \text{if } \Gamma V \geq 1, \end{cases} \tag{3.2}$$

where Γ is the neuronal gain, set to 1 in our simulations, allowing for probabilistic firing near a second-order phase transition [49]. This network is non-conservative, with energy dissipating through inhibition and spike collisions.

3.4.3 Model Simulations:

The mean-field activity of the neurons in our model exhibits hallmarks of neuronal avalanches, with $\alpha = 3/2$, $\beta = 2$, and $\chi = 2$ (as shown in Fig. 3.3). This was confirmed in our finite-size, fully sampled network. External activity in the network is triggered by independent Poisson processes to each neuron at a low rate, activating approximately 20 neurons (about 0.002% of the network) per time step. This external drive induces intermittent, reverberating activity in the network, leading to a neuronal firing rate of approximately 2.5 spikes per 1,000 time units for the critical state. A single time unit in our model corresponds to about 2.5 ms in real time, reflecting the 1 Hz average neuronal firing rate observed in our experimental data. However, further refinement of the model to match in vivo data is challenging due to the unknown sampling fraction f, actual spike rates, neuronal topology, and the corresponding threshold dependence of k.

For our analysis, we applied a threshold to population activity above the expected average number of spikes from the external drive to focus on actual cascading activity in the network. Unless stated otherwise, the network was simulated for 10^8 time-steps, and the

resulting time series were analyzed. Epochs were extracted as described for the experimental data, using the same size and duration definitions. A random sampling fraction f of the total neurons in the network was considered, systematically varying from 0.01% (100 neurons) to the full 10^6-neuron network. The corresponding changes in threshold Θ were used to study the effects of subsampling on the network activity. All scaling curves for χ_{sh} and χ_{lg} from the simulated data employed the fitting function described in the previous section.

3.5 Post-processing pipleine

This section dissects the analytical approach, detailing the steps and methods applied to both the experimental and model data sets.

3.5.1 Continuous Epochs of Suprathreshold Population Activity:

Continuous periods of population activity were identified by applying a threshold (Θ) to the population activity, $p(t)$, which is the sum of the spike densities from all neuronal ROIs at a given time t. This thresholding process is described mathematically as follows:

$$p_\Theta(t) = \begin{cases} p(t) & \text{if } p(t) > \Theta \\ 0 & \text{if } p(t) \leq \Theta \end{cases}. \tag{3.3}$$

This method, known as "hard-thresholding," was used throughout the analysis unless stated otherwise. Given a specific recording, $p(t)$, and a coarse-graining value k, the relationship between the number of continuous epochs (N) and the threshold (Θ) was analyzed over a range of thresholds, $\Theta \in [\Theta_1, \Theta_2]$. The threshold Θ_1 was set low enough to retain all population activity, while Θ_2 was high enough to exclude all population activity. The function $N(\Theta)$ typically follows a log-normal distribution, with shape parameters μ and σ

derived from fitting. For our analysis, we selected the threshold as $\Theta = \mu - 2\sigma$, estimating it for each coarse-graining factor k.

Hard-thresholding of the population activity obtained from a set of spiking neurons was first applied by Poil et al. [137] in neuronal simulations to study neuronal avalanches. This method is similar to the traditional approach of thresholding the local field potential (LFP) used in population-based analyses of neuronal avalanches (see, for example, ref.[136]). It has been shown that the amplitude of the local negative deflection in the LFP increases with the number of neurons firing within the vicinity of a microelectrode (e.g., ref.[134]). Importantly, we chose a low threshold to minimize potential errors in estimating the scaling exponent χ [180], thus reducing the number of epochs for our scaling analysis.

3.5.2 Scaling Curve Fit:

To precisely evaluate the scaling exponents χ_{sh} and χ_{lg}, we employed the following fitting function:

$$S(d) = \frac{C d^{\chi_{\text{sh}}}}{(1 + (d/\Phi)^{\gamma})^{\left(\chi_{\text{sh}} + \chi_{\text{lg}}\right)/\gamma}} \tag{3.4}$$

where $S(d)$ represents a double power law with an initial slope χ_{sh}, which transitions to a second slope χ_{lg} around the point $d = \Phi$. The parameter γ controls the steepness of this transition and was fixed at 4 for all curves presented. The other parameters were allowed to vary to fit the data, and all fitting was conducted in log-space, where the log of S was fitted against the log of d (involving the logarithms of both the average sizes and the durations).

3.5.3 Complexity Analysis:

Complexity analysis was performed as described in [91], examining patterns of depths D in the range of 4–7. The complexity measure C was calculated from population activity as a function of the coarse-graining factor k and the thresholding method, similarly to how epochs were computed for obtaining scaling. Subthreshold activity was not considered in this analysis.

For each segment of length D, the pattern p is defined as the rank order of the time series. For example, a monotonically increasing trace over $D = 4$ time points would result in the pattern 0123. From all possible segments of length D, a probability distribution $P = \{p_j \mid j = 1, 2, \ldots, N\}$ is derived, where N is the number of possible states. For example, when $D = 3$, the possible states are 012, 021, 102, 120, 201, and 210.

Shannon's logarithmic information is then calculated for the distribution P as follows:

$$S[P] = -\sum_{j=1}^{N} p_j \ln(p_j), \tag{3.5}$$

and the normalized Shannon entropy is defined as $H[P] = S[P]/S[Pe]$, where Pe represents the uniform distribution (the entropy is maximized for the uniform distribution, resulting in $0 \leq H \leq 1$).

Finally, the complexity measure is defined as:

$$C[P] = Q_J[P, P_e] \times H[P], \tag{3.6}$$

where $Q_J[P, P_e]$ is defined as:

$$Q_J[P, P_e] = Q_0 \left(S\left[\frac{P + P_e}{2}\right] - \frac{S[P]}{2} - \frac{S[Pe]}{2} \right), \tag{3.7}$$

with Q_0 being a normalization constant such that $0 \leq Q_J \leq 1$. Q_0 is the inverse of the maximum possible value of $J[P, P_e]$.

3.5.4 Temporal Profiles:

Temporal profiles were calculated from individual recordings after applying thresholding and temporal coarse-graining. Avalanches were identified as described in the first paragraph of this section, and then grouped by their duration. The population activity during all avalanches of a particular duration was averaged to generate the temporal profile for that specific duration.

To determine the profile collapse exponent, χ^{coll}, the temporal profiles were first rescaled along the x-axis to fit between 0 and 1, then linearly interpolated to 500 points. Each interpolated temporal profile within the target duration range was then rescaled along the y-axis by $d^{\chi^{\text{coll}}-1}$. The value of χ^{coll} was selected to minimize the root mean square (RMS) error among all collapsed profiles within the specified duration range.

To compute temporal profiles at a specific coarse-graining level k using the original experimental frame rate, avalanche epochs were calculated at the coarse-grained resolution, as previously described. These epochs were then mapped back onto the raster at the original frame rate, and the temporal profiles were calculated from these remapped epochs. The standard deviations from the mean rescaled profiles were visualized as shaded regions unless stated otherwise.

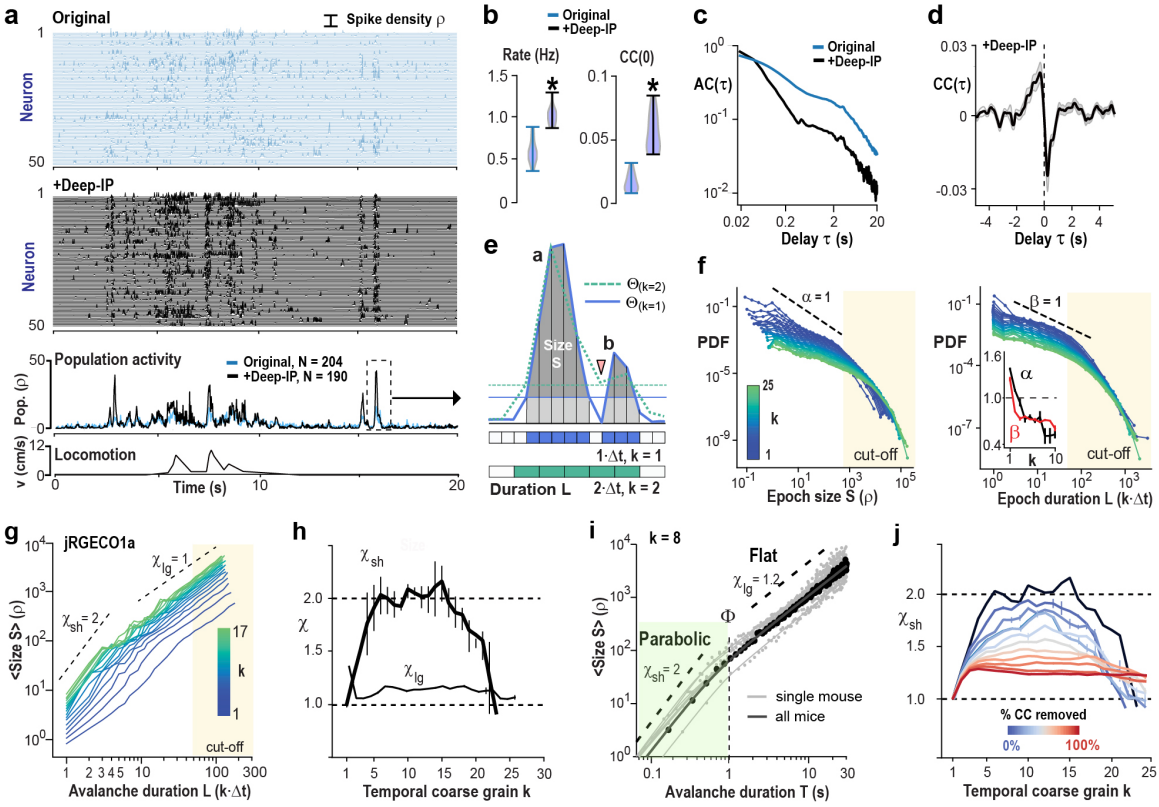

Figure 3.1: (a) Ongoing neuronal activity in mPFC/ACC of awake mice exhibits epochs of increased population activity (columnar organization; original) enhanced after deep-interpolation (+Deep-IP; 2PI, jRGECO1a). Spike density raster of 50 randomly selected neurons each (single mouse). Middle: Full population activity (Pop). Bottom: Locomotion velocity (v). (b, c, d) Deep-IP increases neuronal firing (Rate), cross-correlation (CC) and sharpens population activity transients visible in a faster decay of the autocorrelation (AC). Locomotion onset is preceded by population activity increases (CC). Corresponding experiment in (a). (e) Sketch of continuous epoch in population synchrony per Δt above 'hard' threshold $\Theta(k=1)$ (blue line). Corresponding size S (gray) and duration L in multiples of time bins (blue). A 'soft' threshold discards the subthreshold area (light gray; $k=1$). Note that the two suprathreshold epochs a and b, separated at $1\cdot\Delta t$ (arrowhead), will merge upon temporal coarse-graining ($2\cdot\Delta t$, $k=2$; green dotted line), despite an increase in synchrony requirement per Δt ($\Theta(k=2)$; broken green line). (f) Suprathreshold epochs fulfill the criteria of neuronal avalanches. Epoch size S (left) and duration L (right) exhibit power laws robust to temporal coarse-graining ($k=1-25$; color code). Note cut-offs for $S>1{,}000$ and $L>50$ (beige areas). Corresponding decreasing slopes $\alpha(k)$ and $\beta(k)$ cross the value of 1.

(g) Change in mean size $\langle S \rangle$ vs. duration L for $k = 1 - 17$. Slope estimates for $L = 1 - 4$ (χ_{sh}) and $L \geq 10$ (χ_{lg}). (h) Summary of χ_{sh} and χ_{lg} as a function of k. SD from $n = 10$ bootstrapping repeats. (i) For 'parabolic' avalanches with durations of $T < 1$ s, $\chi_{sh} = 2$ at $k = 8$, whereas 'flat' avalanches exhibit $\chi_{lg} \approx 1.3$ and $T > 2$ s. Avalanches combined for single mice and for all mice. Solid line: fit. (j) $\chi_{sh} = 2$ reflects inter-neuronal correlations and $\chi_{sh} \ll 2$ when neurons are made independent by random temporal shifts (black: see h; 10 surrogates/condition). (f–i) Data pooled from $n = 5$ mice and 17 experiments; jRGECO1a. Values presented as mean or mean ± SD.

3.6 Results

3.6.1 Neuronal synchronization in the prefrontal cortex displays quadratic avalanche scaling

We examined neuronal synchronization in the prefrontal cortex of awake mice during resting and spontaneous locomotion using two-photon imaging (2PI) (Fig. 3.1). With a chronically implanted prism [92], we imaged 200 to 300 neurons within a 450 x 450 μm window, crossing the fissure in the contralateral, intact superficial layers of the anterior cingulate and medial prefrontal cortex (ACC/mPFC). The imaging had a temporal resolution of Δt = 22 ms (45.5 Hz frame rate; Fig. 3.1a). Using a machine-learning-based deep interpolation technique [81], we extracted spike densities from each neuron, significantly enhancing the average firing rate and pairwise correlation among neurons (n = 5 mice; Fig. 3.1b). Summing the activity across all neurons generated the time course of population activity, which exhibited rapid transients of coincident neuronal firing.

We analyzed these epochs of synchronization by applying two criteria: (1) a threshold (Θ) to the summed spike density within a Δt time bin, and (2) concatenation of successive suprathreshold population events (Fig. 3.2e). This approach (e.g., ref. [137]) is sensitive to spatial subsampling and noise. Spatial subsampling can prematurely terminate contiguous periods of activity by missing neurons' firing, while noise can lead to misidentification

of suprathreshold periods. To compensate, we systematically relaxed our concatenation criterion and increased Θ to select higher coincident neuronal activity, indicating stronger synchronization.

Suprathreshold population events identified at Δt were concatenated at a temporally coarse-grained resolution of $k \cdot \Delta t$, where $k = 1, \ldots, k_{\max} < 40$. The threshold Θ(k) was set to -2 standard deviations (SD) of the z-scored distribution in terms of the number of epochs, increasing the absolute value of Θ with k (Fig. 3.1e; see Post-processing pipeline). This analysis showed that population synchrony exhibited the hallmarks of neuronal avalanches [13, 15, 134, 136], with size S (summed suprathreshold spike densities) and duration L (number of time bins or generations per epoch) following power laws (Fig. 3.1f). The corresponding power-law exponents, α and β, were both smaller than 2 and crossed the value of 1 with temporal coarse-graining, indicating broad distributions despite robust size and duration cut-offs of $S > 10^3$ and $L > 50$ (Fig. 3.1f; inset).

To determine the scaling exponent, χ, we plotted mean avalanche size against duration, following a power-law relationship. Short-lasting avalanches ($L < 5$) had χ_{sh} values that increased with temporal coarse-graining, reaching maximal values around 2 (Fig. 3.1g, h). Long-lasting avalanches ($L > 10 - 30$) exhibited χ_{lg} values close to 1, regardless of the temporal scale at which epochs were observed (Fig. 3.1g, h). The transition between these two regimes, characterized by scaling range Φ, demonstrated that shorter avalanches (absolute duration $T = L \cdot k \cdot \Delta t$) shorter than 1 second, with $\Phi = 0.52 \pm 0.21$), exhibited quadratic, rapid growth in size, while longer avalanches (lasting more than 2 seconds) showed a flatter growth pattern.

Parabolic avalanches, indicating rapid growth, were strongly dependent on spatial correlations. These avalanches disappeared when neuronal time series were randomly shifted in time.

The robustness of our approach to subsampling was confirmed, as a 50% reduction in

the neuronal sampling fraction still allowed us to recover $\chi_{\text{sh}} \approx 2$. This was because of a corresponding decrease in the minimal requirement for coincident activity Θ. The identification of $\chi_{\text{sh}} \approx 2$ was also consistent across different forms of thresholding, indicating its reliability even with systematic errors introduced by the thresholding process. Despite these errors, the results remained consistent across a wide range of thresholds.

Our findings suggest that parabolic and flat avalanches represent distinct phenomena, with parabolic avalanches exhibiting rapid, quadratic growth, while flat avalanches show slower, steady patterns. Differences in spatial sampling and experimental methodology could account for these variations.

3.6.2 Temporal coarse graining recovers avalanche scaling in a subsampled critical model

In this section, we use simulations to test whether temporal coarse-graining can recover parabolic avalanche synchronization in critical networks that are spatially subsampled. We designed a network with $N = 10^6$ binary, probabilistic, integrate-and-fire neurons, comprising 80% excitatory and 20% inhibitory neurons, with an all-to-all connectivity pattern (Fig. 3.2a). The excitatory and inhibitory connections were represented by constant matrices, with $W_{EE} = W_{IE} = J$ for excitatory connections, and $W_{II} = W_{EI} = -g \cdot J$ for inhibitory connections. We set $g = 3.5$ to achieve an E/I balance that supports avalanche dynamics in the fully sampled system, yielding critical exponents $\alpha = 3/2$, $\beta = 2$, and $\chi = 2$, approximating a critical branching process triggered by external Poisson inputs (3.3).

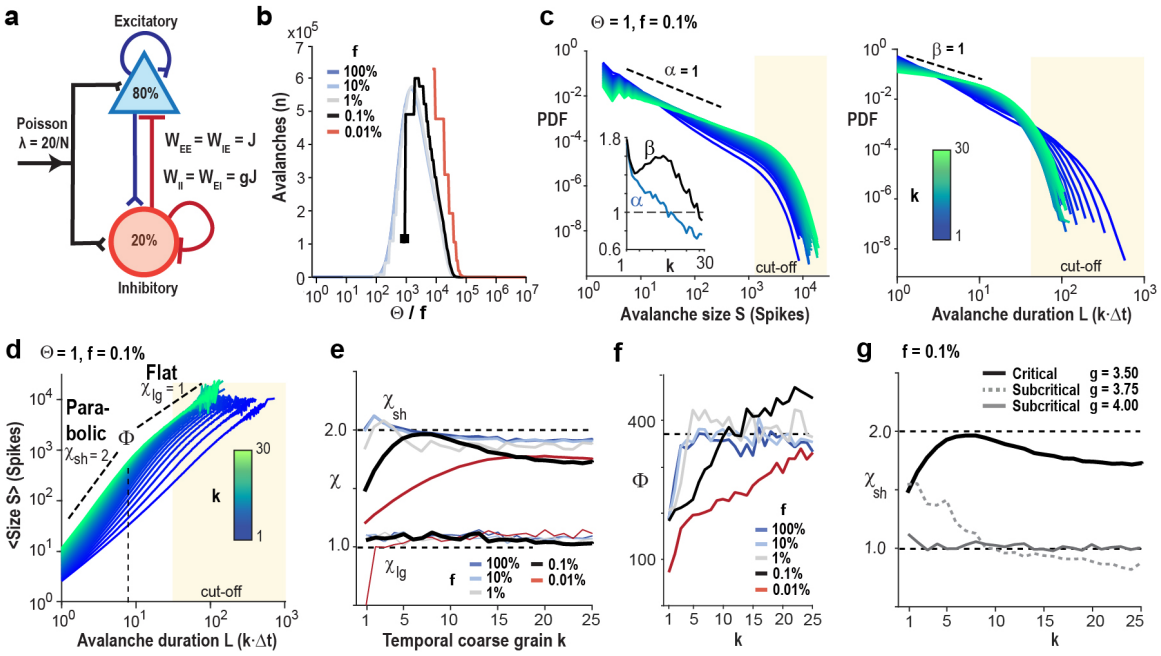

Figure 3.2: (a) Sketch of the neuronal network with $N = 10^6$ neurons (80% excitatory, E; 20% inhibitory, I) and external Poisson drive of rate $\lambda = 20/N$ per time step. The E/I balance is controlled by the scalar g, which scales inhibitory weight matrices $W_{\mathrm{II}} = W_{\mathrm{EI}}$ as a function of excitatory weight matrices $W_{\mathrm{EE}} = W_{\mathrm{IE}} = J$. (b) Change in number of avalanches as a function of threshold Θ normalized by sampling fraction f. (c) Power law in size (left) and duration (right) distributions for avalanches become shallower with temporal coarse-graining k. Note cut-off regimes for $S \gtrsim 10^3$ and duration $L \gtrsim 50$. Inset: Corresponding slopes $\alpha(k)$ and $\beta(k)$. (d) Temporal coarse-graining uncovers $\chi_{\mathrm{sh}} = 2$ for short-duration avalanches ($L = 1 - 10$), whereas χ_{lg} remains $\sim 1 - 1.2$ for long-duration avalanches ($L > 10$). (e) Summary of change in χ_{sh} and χ_{lg} with k for 5 sampling fractions f. A decrease in f requires a higher k to recover $\chi_{\mathrm{sh}} = 2$. Note failure of recovery for very low f. χ_{lg} does not depend on k. (f) Temporal coarse-graining recovers avalanches up to the finite-size cut-off of $\Phi \approx 400$ time steps. Note plateau in maximal scaling range Φ for $\chi_{\mathrm{sh}} = 2$ plotted in simulation time steps as a function of k. (g) Temporal coarse-graining recovers $\chi_{\mathrm{sh}} = 2$ for critical network dynamics but fails for subcritical dynamics. Note that weak subcritical ($g = 3.75$) and critical conditions ($g = 3.5$) can exhibit similar χ_{sh} at the original temporal resolution yet diverge with temporal coarse-graining. Broken, black lines: visual guide to the eye. Results obtained from $T = 10^8$ simulation time steps.

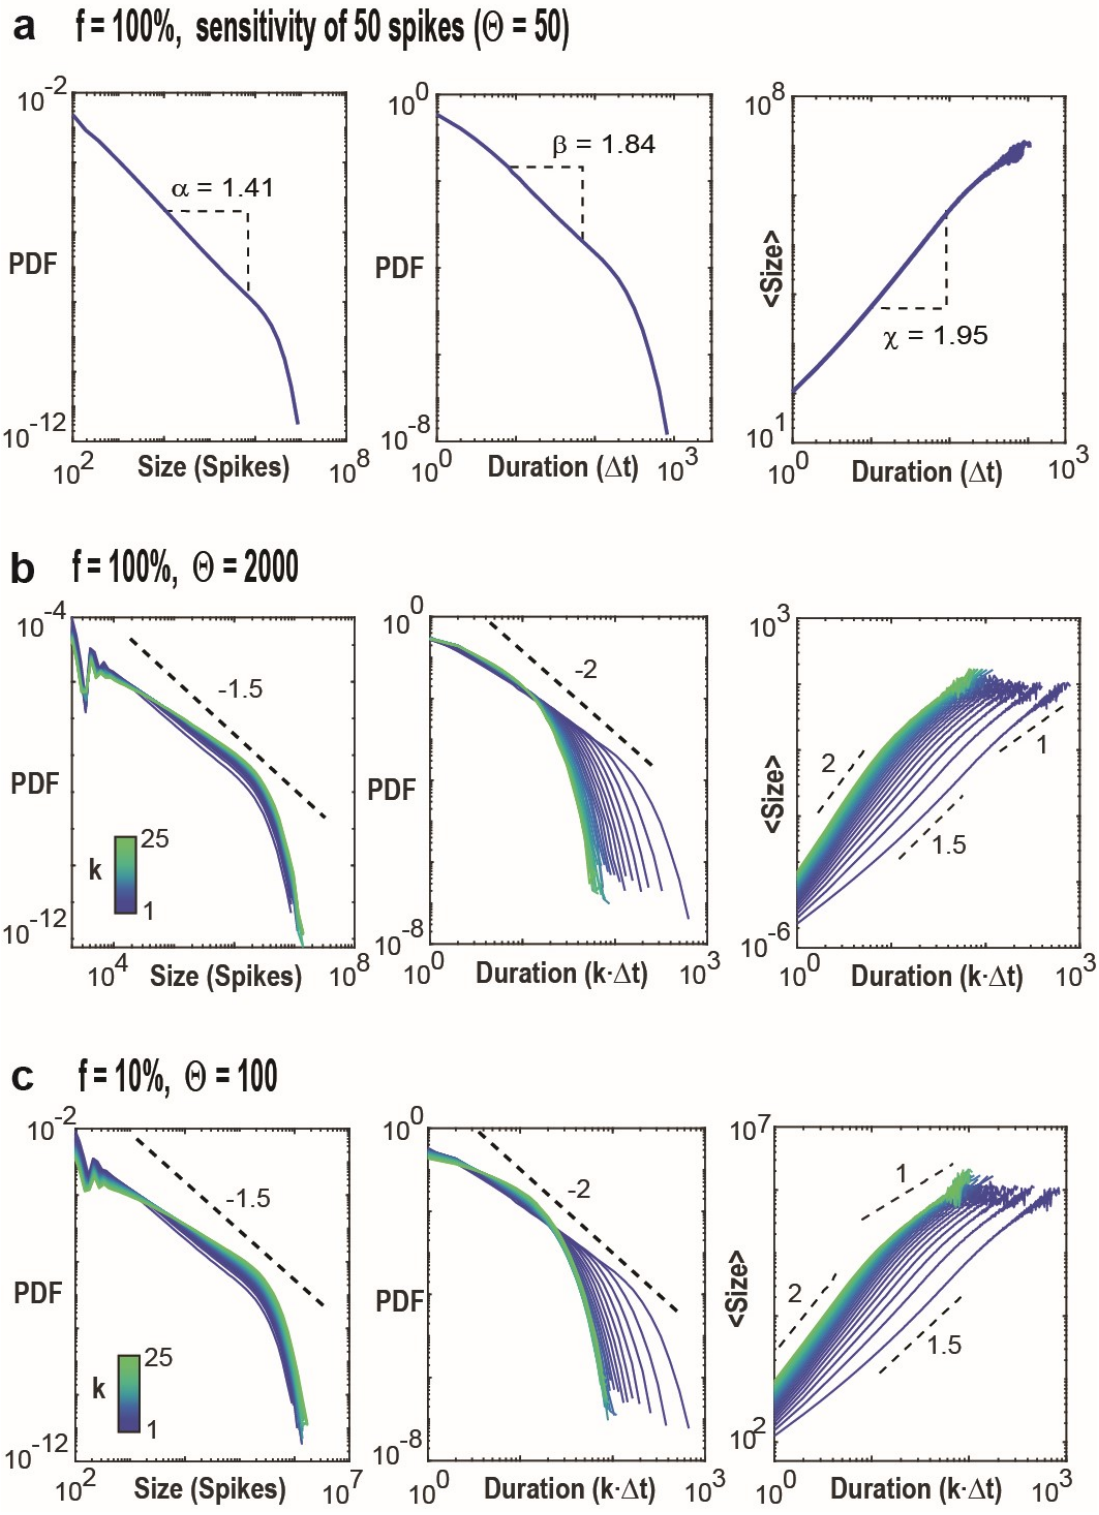

Figure 3.3: Examples of avalanche distributions as a function of threshold Θ and

sampling fraction, f for the critical E/I-model. **(a)** Distributions for the fully sampled model ($f = 100\%$; $N = 10^6$ neurons, $g = 3.5$; 10^8 time steps) at population spike sensitivity of 50 spikes/time step ($\Theta = 50$). *Left*: size. *Middle:* duration. *Right:* mean size vs. duration scaling with $\chi = 1.95$. Note that the Kinouchi model requires a minimum of about 100 spikes in $N = 10^6$ to initiate successful propagation. The fully sampled model exhibits critical exponents close to $\alpha = 3/2$, $\beta = 2$ that fulfill the analytical prediction of $\chi = (\beta - 1)/(\alpha - 1) = 2$. **(b)** Temporal coarse graining recovers $\chi \approx 2$ at low sensitivity, i.e., high Θ. Example distributions and scaling for the fully sampled model ($f = 100\%$) and high threshold $\Theta = 2000$. **(c)** As in (b) with 10 times reduced sampling ($f = 10\%$) and 20 times reduced threshold ($\Theta = 100$). For corresponding $\chi(k)$ plots see Fig. 3.2. Note that in this model, uncorrelated propagated activity yields long-duration avalanches with $\chi_{lg} = 1.5$. $T = 10^8$ time steps. *Broken, black lines:* slopes as a visual guide to the eye

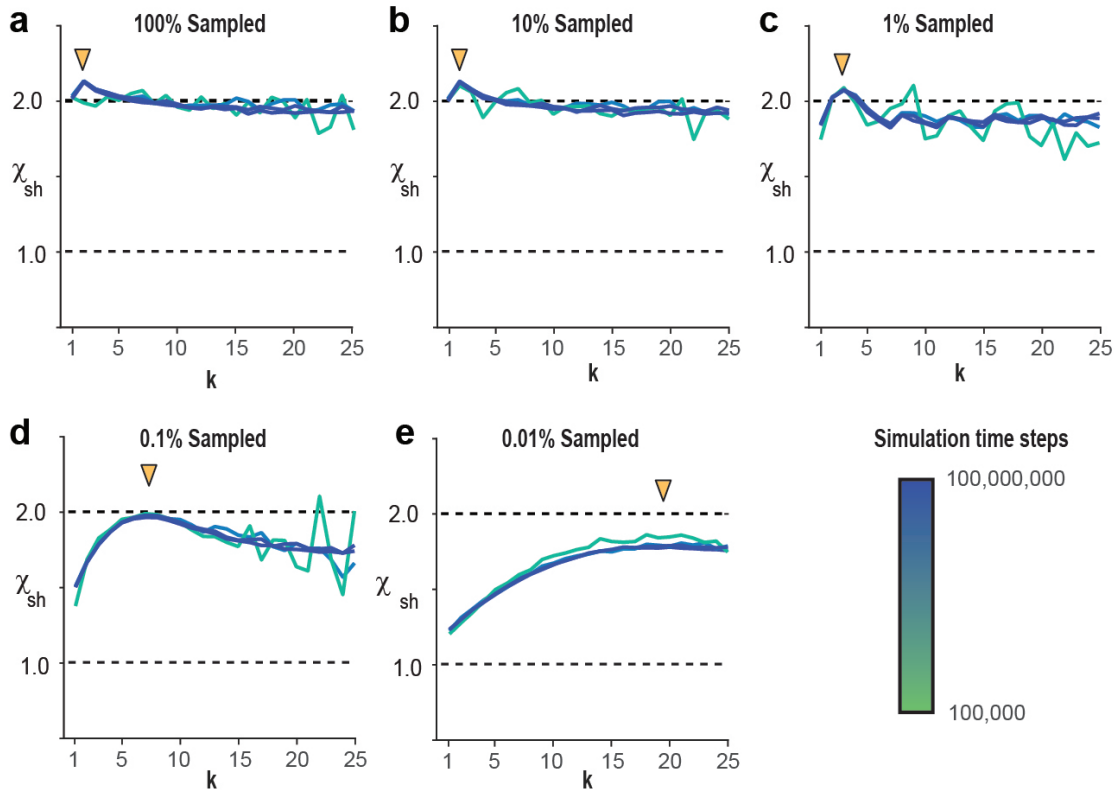

Figure 3.4: Short simulation times provide robust estimates of χ_{sh} as a function of temporal coarse-graining k under fully-sampled and subsampled conditions. **(a)** In the fully sampled model ($f = 100\%$), reasonable estimates of $\chi_{sh}(k)$ are obtained with $100,000$ simulation time steps that only marginally improve even for $1,000$ times longer simulations. **(b) – (e)** Same as in (a) but for decreasing sampling fractions f. When lowering the sampling ($f = 10\% - 0.1\%$), the recovery of $\chi_{sh} = 2$ requires a corresponding increase in temporal coarse-graining factor k until recovery fails for very low sampling fractions ($f = 0.01\%$; *arrowheads*).

We found that temporal coarse-graining could recover $\chi_{sh} \approx 2$ for few-generation avalanches across a wide range of subsampling conditions. These conditions were defined by the neuronal sampling fraction, f, which ranged from 0.01% to 100%, and the population activity threshold, Θ, which ranged from 0 to below 20,000 spikes per time unit. To demonstrate these results, we used a sampling fraction of $f = 0.1$ and a threshold $\Theta = 1$, which required at least two spikes per time unit (Fig. 3.2b–f). As observed in our experimental data, temporal coarse-graining of subsampled avalanches reduced the power-law slopes for both size and duration, confirming that the mean size-to-duration power-law relationship yielded $\chi_{sh} \approx 2$ for parabolic avalanches (Fig. 3.2e).

It's important to note that successive avalanches in the model are uncorrelated, leading to a reduction in χ_{sh} for very high coarse-graining values (large k), regardless of the total simulation time (Fig. 3.4). As observed in our 2PI data, noise can reduce the recovered maximum value of $\chi_{sh} \leq 2$ without changing the corresponding coarse-graining factor k (Fig. 3.5a), while improved tracking of spiking activity tends to shift the recovery of χ_{sh} towards lower values of k, as seen with deep interpolation (Fig. 3.5b). Generally, we found that when sampling fewer neurons (low f) and using less sensitive thresholds (high Θ), more temporal coarse-graining was required to recover $\chi_{sh} \approx 2$ (Fig. 3.2e). The model also demonstrated that the scaling range, Φ, where $\chi_{sh} \approx 2$, reflects a finite-size effect that can be reliably recovered with accurate temporal coarse-graining (Fig. 3.2f). As expected, $\chi_{lg} \ll 2$ for many-generation, or flat, avalanches, which typically occur near the cut-off in the corresponding size and duration distribution (Fig. 3.2c, d).

Crucially, our simulations demonstrated that achieving $\chi_{sh} \approx 2$ under spatial subsampling conditions is only possible when the network exhibits critical dynamics. In contrast, for subcritical dynamics, χ_{sh} remains around 1, regardless of temporal coarse-graining (Fig. 3.2g; Fig. 3.6). These results suggest that our 2PI data represent spatially subsampled activity from critical dynamics in the cortex.

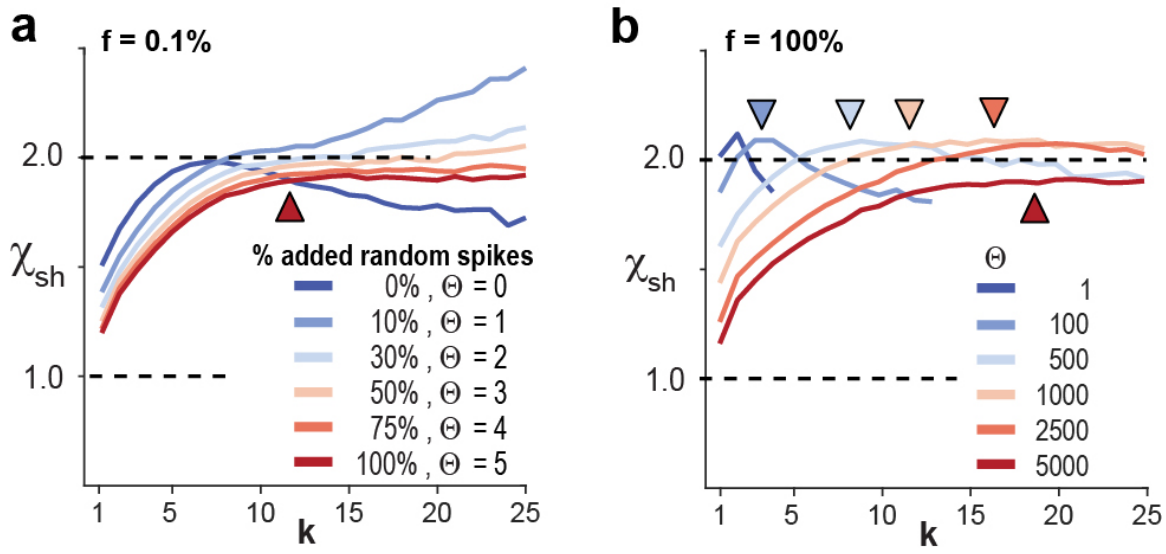

Figure 3.5: Impact of added irregular firing and subsampling as a function of threshold on the recovery of χ_{sh}. **(a)** Increasing percentage of irregular spikes in the subsampled model ($f = 0.1\%$) necessitates a decrease in sensitivity, i.e., higher Θ, to offset the effect of increased noise. This increase in Θ shifts the recovery of $\chi(k) \approx 2$ towards larger k and this recovery begins to fail ($\chi_{sh}(k) < 2$) at higher percentages of irregular firing (*arrowhead*). **(b)** Increased sensitivity, i.e., lower Θ, in the fully sampled model ($f = 0.1\%$) shifts the recovery of $\chi(k) \approx 2$ towards smaller k (arrow heads) as observed with Deep-IP. Note failure of recovery when Θ is too high, for which $\chi_{sh}(k)$ remains below 2 ($f = 100\%$). $N = 10^6$ neurons, $T = 10^8$ time steps.

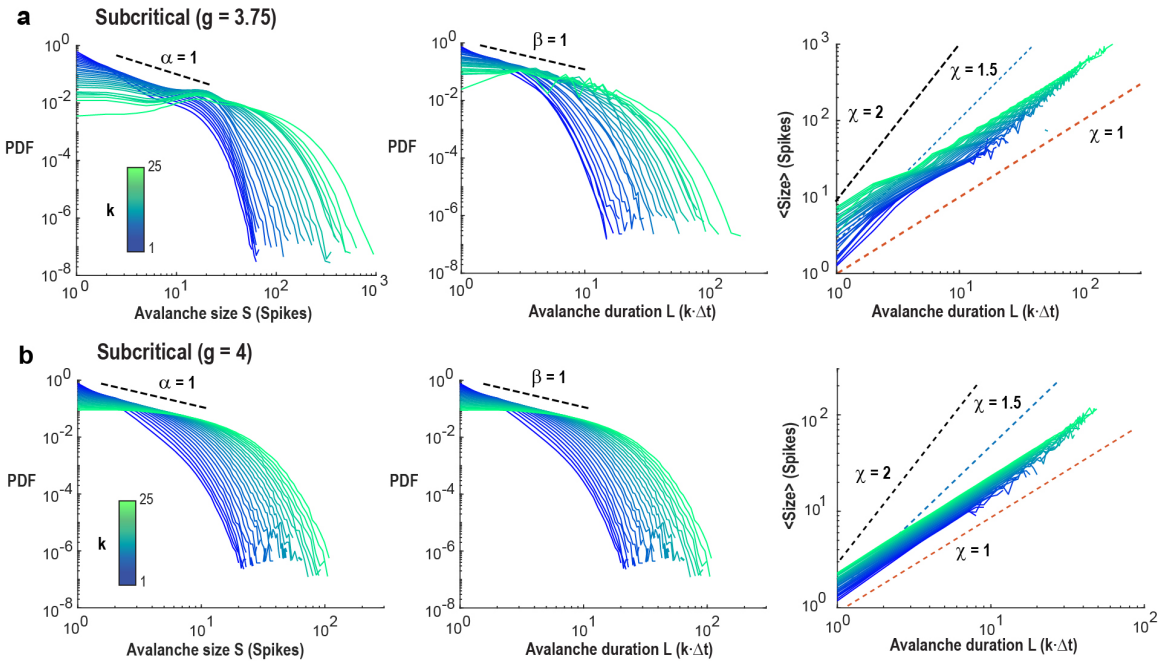

Figure 3.6: Avalanche distributions and scaling as a function of temporal coarse graining k for subcritical network dynamics. **(a)** Distributions and scaling for subcritical dynamics (increased inhibition; $g = 3.75$). *Left:* size. *Middle:* duration. *Right:* mean size vs. duration scaling with $\chi = 1$, 1.5 and 2 highlighted. **(b)** Distributions and scaling for stronger subcritical dynamics (increased inhibition; $g = 4$). Note failure of coarse graining to reach $\chi = 2$ for short or long duration avalanches. Instead, χ remains close to 1 for all avalanches. For corresponding $\chi(k)$ plots. For all simulations: $f = 0.1\%$; $N = 10^6$ neurons, $T = 10^8$ time steps. *Broken, black lines:* an eye guide for the corresponding slope values.

3.6.3 Properly resolved avalanches maximize complexity and temporal correlations

To understand how the initial rise and subsequent fall of χ_{sh} with temporal coarse-graining relates to the temporal organization of neuronal population activity, we quantified the temporal complexity in population activity. We divided the thresholded population activity time course into successive temporal bins of duration $k \cdot \Delta t$, with depths $D = 4, \ldots, 7$. For each pattern depth D and coarse-graining factor k, we calculated pattern complexity C_x, which measures the number of distinct temporal sequences of length D (ref. 39; see Post-processing pipeline).

Our results showed that C_x peaked near the value of k where $\chi_{sh} = 2$, for both ongoing and evoked activity, regardless of the pattern depth (Fig. 3.7). This observation was confirmed in our E/I-balanced model but only for critical dynamics (Fig. 3.7a, bottom); subcritical dynamics did not exhibit peak complexity with temporal coarse-graining. Our simulations also demonstrated that lower sampling fractions required higher temporal coarse-graining to recover peak complexity, which is consistent with our findings on χ_{sh} in the data (cf. Fig. 3.2e). This increase in temporal complexity aligned with our observation that the delayed autocorrelation $AC(1)$ peaked for parabolic avalanches at the value of k where $\chi_{sh} = 2$, a pattern not seen for flat avalanches or when all avalanches were analyzed (Fig. 3.7b).

3.6.4 Scale-invariant inverted parabolic profile

Theory and experiments suggest that a slope of $\chi_{sh} \approx 2$, as consistently found in our experiments, predicts scale-invariant, inverted parabolic profiles for avalanches that can be collapsed with an exponent $\chi_{coll} \approx 2$ (refs.[77, 128, 163]; see Post-processing pipeline). In

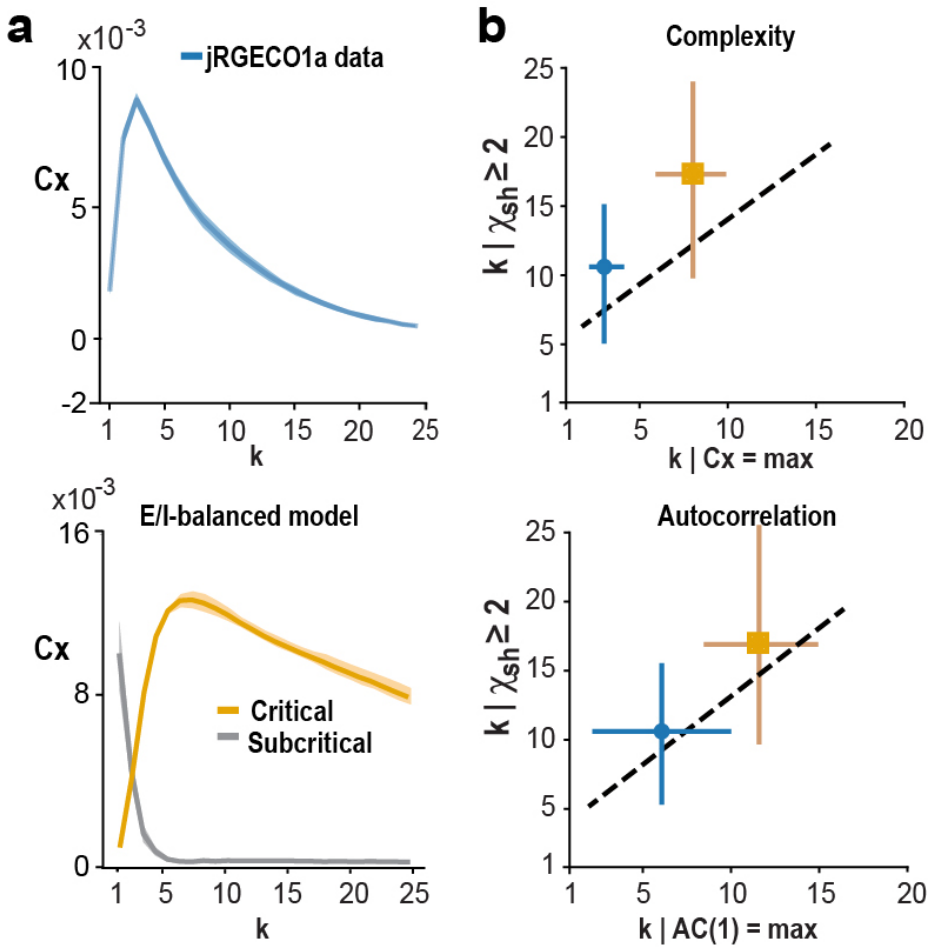

Figure 3.7: (a) Peak in pattern complexity (Cx) with temporal coarse-graining (pattern depth $D = 5$, thresholded population activity). *Top:* Each experimental condition. *Bottom:* E/I-model. With temporal coarse-graining, the pattern complexity peaks in subsampled, critical dynamics ($g = 3.5$; $f = 0.1\%$, $\Theta = 1$), but monotonically decreases for subsampled, subcritical dynamics ($g = 3.75$); pattern depth $D = 5$. For each condition and model, mean complexity is presented. Error bars are the SD across 10 equal sections of each experiment/simulation. (b) Temporal coarse-graining at which $\chi_{sh} \geq 2$ maximizes pattern complexity *(top)* and temporal correlations *(bottom)*. Averages over all experimental conditions and model. Error bars indicate range of k for which $\chi_{sh} \geq 2$ and 90 – 100% of maximal $AC(1)$ or Cx, at baseline of $k = 1$ (Pearson's $r = 0.53$ for Cx vs. k; Pearson's $r = 0.85$ for $AC(1)$ vs. k; $n = 6$). *Broken line:* slope of 1. For color code see (a).

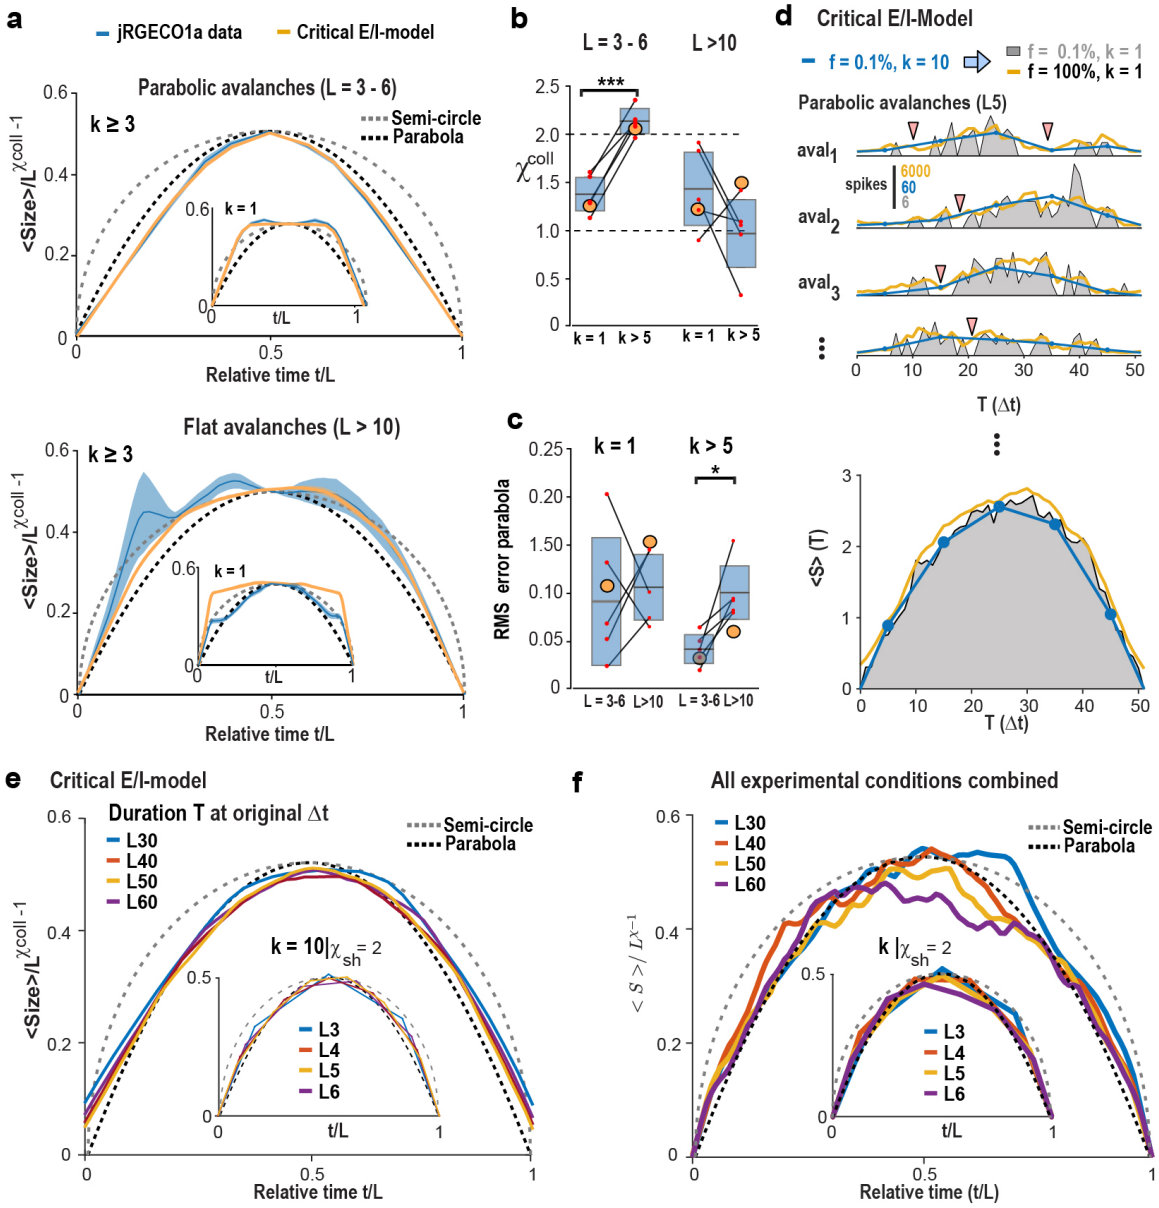

Figure 3.8: (a) Top: For few-generation avalanches ($L = 3-6$; $k > 5$), inverted parabolas are recovered from flattened profiles present at $k = 1$ *(inset)*. *Bottom:* Many-generation avalanches ($L = 10-20$) exhibit non-parabolic profiles at $k = 1$ *(inset)* and after temporal coarse-graining. Average profile $\pm$ SD. Avalanches were pooled from all recordings and mice per condition *(color)*. Mean activity within an avalanche normalized by the number of generations L to the power of $\chi_{\text{coll}} - 1$ plotted for each relative time step, t/L. Note the match for the critical, subsampled E/I-model ($f = 0.1\%$, $\Theta = 1$), in line with the recovery of the critical exponent, and inverted-parabola collapse for few-generation avalanches.

(b) For few-generation avalanches, the scaling exponent for profile collapse, χ_{coll}, is significantly higher and close to 2 after coarse-graining compared to $k = 1$, whereas no difference is found for many-generation avalanches ($t = -11.94$, DF = 4, $p = 0.00028$ vs. $t = 1.6$, DF = 4, $p = 0.185$, two-sided paired t-test). Summary statistics for all experimental conditions. Orange circles: Critical E/I-model (SD $<$ symbol size). (c) After temporal coarse-graining, short-lasting avalanches exhibit profiles closer to a parabola compared to long-lasting avalanches, but not at $k = 1$ ($t = -3.7$, DF = 4, $p = 0.021$) vs. $t = -0.25$, DF = 4, $p = 0.81$, two-sided paired t-test). (d) In the critical model, population activity sequences t_i defined by $L = 5$ at $k = 10$ *(blue)* exhibit interruption of contiguous suprathreshold epochs due to subsampling *(gray)*, not found in the fully sampled model *(yellow)*. Bottom: Parabolic profiles of $L = 5$, $k = 10$ epochs *(blue)* match corresponding time averages from subsampled and fully sampled conditions. (e) Summary for L_{30}, L_{40}, L_{50}, and L_{60} at Δt and corresponding avalanche profiles at $k = 10$ and $\chi_{\text{coll}} = 2$. (f) Corresponding analysis pooled over all experiments. (d–f) Averages are shown for clarity. *Box plots:* mean $\pm$ SD.

Fig. 3.8, we illustrate the collapsed avalanche profiles for both few-generation and many-generation avalanches across all experimental conditions and the critical model. An inverted parabolic shape was found only at the coarse-graining factor k that maximized χ_{sh} (Fig. 3.8a, top), which applied to short-lasting, few-generation avalanches, but not to long-lasting, many-generation avalanches (Fig. 3.8a, bottom). The corresponding value of χ_{coll} increased to approximately 2 after temporal coarse-graining for few-generation avalanches, while many-generation avalanches maintained a flatter profile that deviated from a parabola (Fig. 3.8c).

Interestingly, parabolic avalanches ($L = 3-6$, $\chi_{\text{sh}} \geq 2$) did not occur with regular periodicity, which distinguishes them from oscillatory activity. By using the temporally coarse-grained segments of parabolic avalanches, we could recover their corresponding inverted parabolic shape at the original temporal resolution Δt for both the model and our data (Fig. 3.8d–f; Fig. 3.9a, b). In the model, this approach helped identify segments of parabolic and flat avalanches in a subsampled system at the original temporal resolution (Fig. 3.9c).

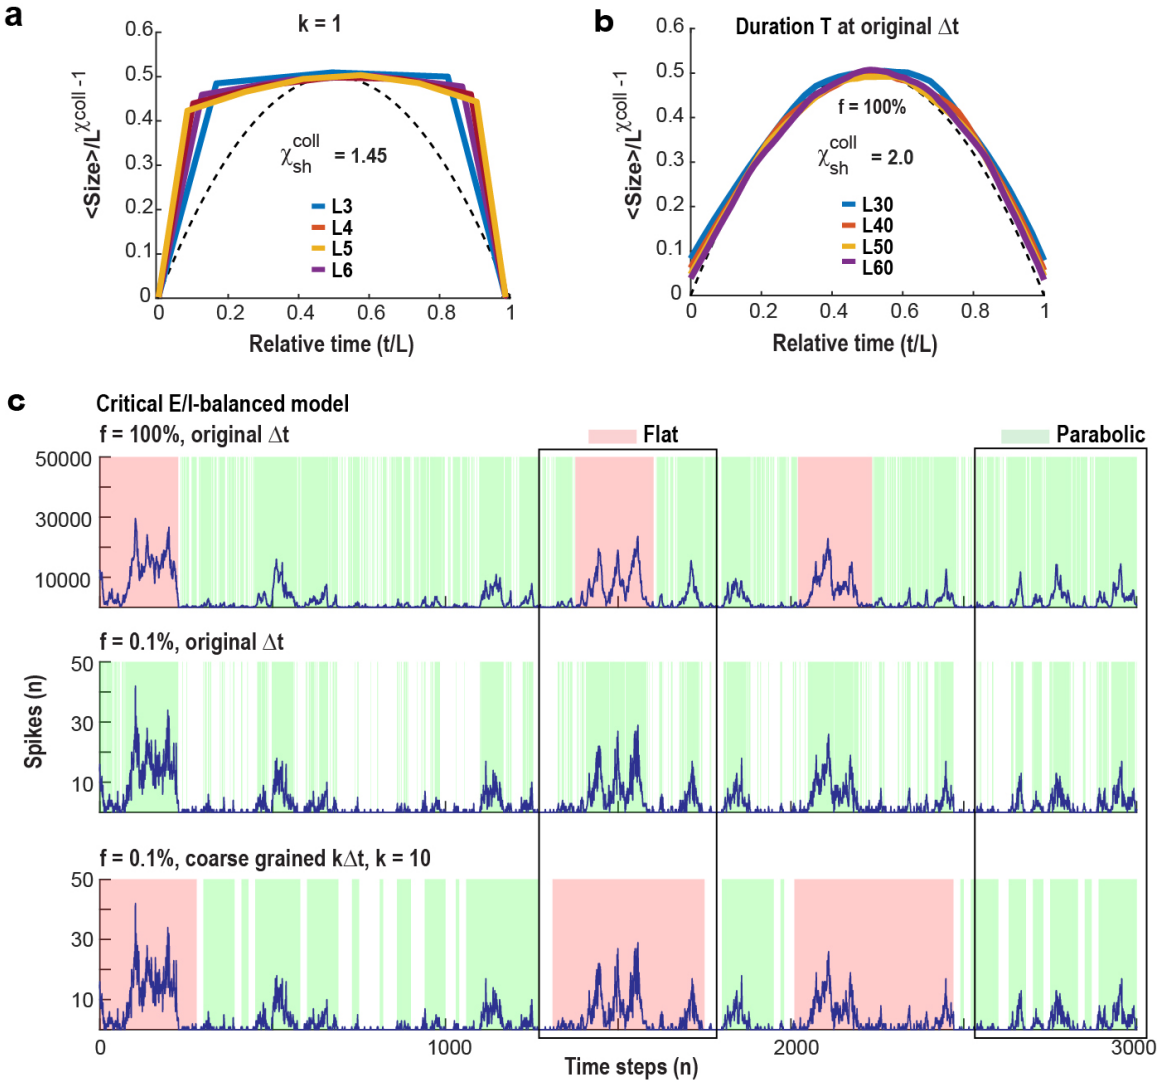

Figure 3.9: Coarse-grained, few-generation avalanches with $\chi_{\text{sh}} = 2$ exhibit inverted parabolic shape in cortical synchronization at high-resolution over many decades of generation. **(a)** Critical E/I-model at $k = 1$ and corresponding avalanche collapse for $L = 3 - 6$; see also Fig. 3.8 **(b)** Critical E/I-model. As in Fig. 5e for the fully sampled model, i.e., $f = 100\%$. Same as in (a) for $L = 30, 40, 50, 60$ generation avalanches at high temporal resolution ($k = 1$). Note $\chi_{\text{sh}}^{\text{coll}} = 2.0$ for all generations. **(c)** Recovery of flat and parabolic avalanches at $f = 0.1\%$ and temporal coarse graining of $k = 10$. Color bars indicate duration of flat and parabolic avalanches, respectively. Rectangles provide two periods of recovery as visual guides. Example for 3,000 consecutive time steps.

3.6.5 Summary

A summary of our findings is presented in Fig. 3.10. In short, parabolic avalanches, unlike flat avalanches, can be accurately recovered in spatially subsampled critical networks through a combination of temporal coarse-graining and thresholding. We conclude that parabolic avalanches represent spontaneous and evoked synchronized cell assemblies in the cerebral cortex, exhibiting high temporal complexity.

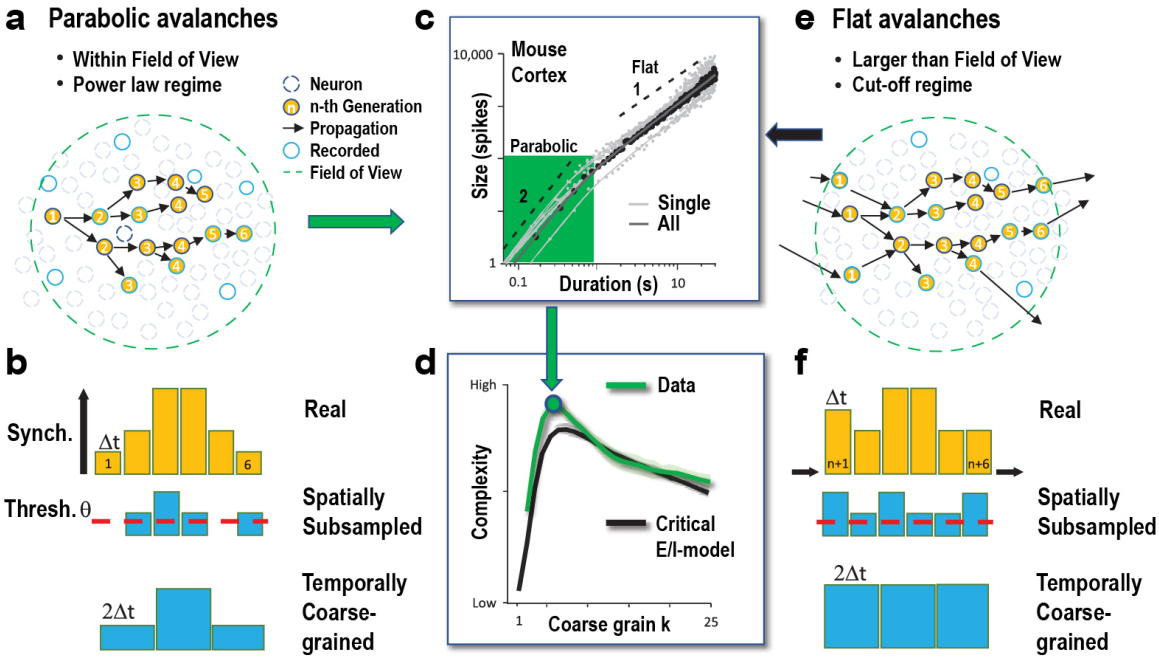

Figure 3.10: (a, b) Parabolic avalanches are constrained within the field of view and reside within the power law regime of the measured avalanche size and duration distributions. Despite the constraint of spatial subsampling, their temporal profile of synchronized neuronal activity can be recovered using thresholding and temporal coarse-graining. (c) Parabolic avalanches grow quadratically in size with duration. (d) The sequence of parabolic avalanches in the network exhibits maximal temporal complexity. (e, f) Flat avalanches are spatially larger than the field of view, preventing a proper recovery of their time course using coarse-graining. This results in a flat average temporal profile and corresponding near-linear growth with duration (see (c), 1). Flat avalanches exhibit size and duration that lie beyond the cut-off of the size and duration distribution. For details, see also Figs. 3.1, 3.7.

3.7 Discussion

In this chapter, we found that synchronization of cell assemblies in the prefrontal cortex and primary sensory cortex can exhibit an inverted parabolic profile. This symmetrical pattern shows scale-invariance across hundreds of milliseconds to several seconds, supporting the concept of parabolic avalanches. The temporal appearance of these parabolic avalanches indicates high complexity in neuronal synchronization during resting states, spontaneous locomotion, and visually evoked responses under various recording conditions.

The synchronized groups identified here at the cellular level in superficial layer 2/3 align with the framework of neuronal avalanches, characterized by power laws in group size and duration within the defined field of view. These power laws, when interpreted within the context of self-organized criticality (SOC; ref. [9]), suggest a critical state in the cortex that supports the propagation of synchronized activity across wide regions [14, 16, 148]. This interpretation has faced challenges from other models of avalanche generation that focus on balanced noise without neuronal interactions [103, 175], or from noise-induced fractures near a discontinuous first-order phase transition [154], summarized as self-organized bistability (SOB; refs. [22, 39]).

Our finding of a scaling exponent of $\chi = 2$, which is sensitive to spatial correlations, does not support these alternative models, which typically have χ values ranging from 1 to 1.5 and exhibit non-parabolic shapes resembling semi-circles or sawtooth patterns [39, 175, 180]. This discrepancy helps explain recent experimental results on neuronal avalanches in the cortex, which showed a scaling exponent $\chi \approx 1.3$ [44, 70, 96]. By introducing temporal coarse-graining along with deep-interpolation-based denoising, we overcame experimental limitations from spatial subsampling and demonstrated that neuronal avalanches in the cortex indeed exhibit a scaling exponent of $\chi = 2$, consistent with

a critical branching process near a second-order phase transition [13, 136, 148].

Our simulations suggest that achieving $\chi = 2$ under subsampling conditions is only possible at criticality, regardless of the level of coarse-graining or threshold applied. While complete 'separation of time scales' can help compensate for subsampling in critical models [84], this assumption is not applicable to experimental data where the exact beginning and end of avalanches are unknown. However, our simulations demonstrate that temporal coarse-graining can recover $\chi = 2$ even when 'separation of time scales' cannot be assumed.

The key takeaway here is that parabolic avalanches represent the spiking activity of synchronized neuronal groups in the cortex. This inverted parabolic profile stands out as a robust pattern of synchronization, suggesting a distinct mechanism for how neurons coordinate their activity over time.

3.8 Acknowledgements

I would like to express my gratitude to Elliott Capek and Dr. Tiago L. Ribeiro, who were the primary authors of this project, leading much of the data analysis. I also want to thank Dr. Patrick Kells and Dr. Tiago L. Ribeiro for conducting the experiments discussed in this chapter, with Dr. Kells also taking charge of the deep-interpolation analysis. Special thanks go to Dr. Stephanie R. Miller, Mitchell Victor, Elias Geist, Dr. Sinisa Pajevic, Dr. Ali Vakili, and Prof. Dante Chialvo for their contributions to various aspects of the data collection and analysis. Finally, I want to thank my co-advisor, Dr. Dietmar Plenz, for his role in planning and coordinating this study, as well as for his guidance in shaping every aspect of the research and analysis.

Chapter 4: Limits in the recovery of parabolic avalanches in spatially subsampled neuronal networks at criticality

This chapter is adapted from an article currently under revision in Scientific Reports titled "The recovery of parabolic avalanches in spatially subsampled neuronal networks at criticality," authored by Keshav Srinivasan, Tiago L. Ribeiro, Patrick Kells, and Dietmar Plenz. You can access the current preprint on bioRxiv.

4.1 Overview

In the previous chapter, we introduced the concept of parabolic avalanches and showed how temporal coarse-graining can mitigate the impact of subsampling in both experimental and critical model contexts. This chapter explores the results from an all-to-all balanced network of excitatory and inhibitory neurons with critical branching process dynamics. We demonstrate that by applying temporal coarse-graining and setting a threshold for the minimum level of coincident firing in the network, it is possible to recover $\chi \approx 2$, even when sampling as little as 0.1% of the neurons. By contrast, the commonly proposed 'crackling noise' approach does not achieve this under similar conditions.

In this chapter, we delve into the specifics of the strengths and limitations in recovering parabolic avalanches in severely subsampled critical models. We aim to identify the scenarios where temporal coarse-graining can successfully recover critical scaling behavior, and where it may fall short due to high levels of subsampling or other factors.

4.2 Introduction

Complex systems, composed of many local components or agents that interact weakly, often exhibit event cascades that cover a wide range of scales in both space and time. Such scale-invariant cascades, typically identified by power-law distributions in their duration and size – defined as the total activity observed during its lifetime, have been observed in many real systems, including solar flares [55], earthquakes [8, 123], superconductors [12], sandpiles [56], forest fires [41], and in the brain in the form of neuronal avalanches [13, 15, 24]. The scaling of such system cascades, specifically how their mean size grows with their duration, has been particularly informative as to the potential underlying dynamics of cascade generation. For example, if local system components activate independently at a constant rate, the mean size, $\langle S \rangle$, of observed cascades will scale linearly with duration, T, exhibiting a scaling exponent of $\chi \approx 1$, i.e., temporally contiguous system events reflect a constant, yet random probability of their continuation. Values of $\chi \approx 1$ are often found near a 1st order phase transition [86, 152, 153]. Conversely, if cascades exhibit $\chi = 2$, their mean growth is rapid and parabolic as can be found at the second-order phase transition of a high-dimensional critical branching process, supporting local system interactions [180]. In general, values of $\chi \gg 1$ are of particular interest as they suggest the ability of a system to bind spatially distributed events rapidly and selectively into diverse internal system configurations which is generally considered beneficial for a complex adaptive system such as the brain (e.g.,[59, 117]).

Here, we focus on neuronal avalanches, which were first shown to spontaneously emerge in isolated brain networks, i.e., acute cortex slices and long-term cortex cultures [13]. By measuring locally synchronized neuronal activity using the local field potential (LFP) and tracking the spatiotemporal spread of LFP cascades with microelectrode arrays, power-law distributions for both size, S, and duration, T, and were revealed, with exponents

$\alpha \approx 3/2$ and $\beta \approx 2$ respectively, in addition to a critical branching parameter of $\sigma \approx 1$, and significant correlation of neuronal activity within and between [88, 89] neuronal cascades (for review see [136]). These hallmarks of fast, neuronal synchronization in the form of avalanches were subsequently confirmed in the normal brain for ongoing and evoked LFP cascades in the cortex of nonhuman primates [114, 134, 189] and in humans using magnetoencephalography [51, 127, 166]. Their sensitivity to even slight pharmacological alterations in the balance of excitation and inhibition [13, 16, 51, 136] and changes during pathological conditions such as epilepsy [104, 107, 111, 112] or changes in brain states such as sleep [144, 145] and sleep deprivation [108–110, 113], firmly indicate that these avalanches represent neuronal synchronization of cell assemblies in the normal brain.

To understand avalanche synchronization mechanistically, the contribution of single neuron firing to avalanches needs to be properly identified. This task is extremely challenging as studies in nonhuman primates and isolated slice preparations show that neurons participate sparsely and highly selectively even in spatially extended, large LFP avalanches [16]. Accurate reconstruction of avalanche synchronization from single neuron activity, therefore, requires the appropriate sampling of a significant portion of the brain network, which poses an enormous technical challenge in neuroscience. Conversely, spatial subsampling, that is, sampling of only a fraction of neurons, is limited in identifying neuronal synchronization largely by failing to identify avalanche continuation in non-sampled neurons [147]. This underestimates avalanche duration and size, biasing outcomes towards expectations for decorrelated, random processes [24, 38, 175]. Accordingly, simulations of avalanche generating critical branching processes and fractional sampling of neurons reduces the expected growth in mean avalanche size with avalanche duration, χ, from the expected value of $\chi = 2$, to a value of χ closer to $1 - 1.3$ [27, 44, 84]. Not surprisingly, estimates of $\chi \approx 1.3$ were reported for fractionally sampled neuronal spikes of potential avalanche activity in vivo [44, 70, 96, 102, 184] and in vitro [45].

If the exact start and end times of avalanches are known, for example in simulations under the assumption of 'separation of time scales', spatially subsampled avalanches can be easily corrected for by temporal integration thereby recovering the correct critical exponents in size and duration [84] (for a review, see also [85]). Avalanche times, however, are largely unattainable in experimental data for which only a small fraction of neurons is observed and temporal integration eventually will degenerate avalanche scaling by indiscriminately concatenating successive avalanches [24]. Recently, Capek & Ribeiro et al. [24] demonstrated that, even when the exact times of avalanches are not known, $\chi = 2$ can be recovered for subsampled excitation-inhibition (E-I) balanced networks of integrate-and-fire neurons which are known to belong to the mean-field directed percolation (MF-DP) universality class [49, 50]. This rescue requires (1) reducing uncorrelated neuronal action potential firing and (2) using temporal coarse-graining with an increase in minimal synchronization requirements to identify synchronized population activity in the network. Under these conditions, the authors could show $\chi = 2$ in vivo 2-photon imaging recordings of cellular activity with high signal-to-noise ratio for sensory evoked and spontaneous spike avalanches in frontal and sensory cortices of awake mice in line with their simulations [24]. This is in line with a recent EEG study [155] in humans that demonstrated $\chi = 2$.

In the current work, we extend our previous investigations of rescuing χ in subsampled, critical networks. We systematically study the rescue of the scaling exponent χ to its critical value of 2 in a continuously driven, subsampled all-to-all DP-model [49] for which we demonstrate the equivalence of fractional sampling and thresholding. Importantly, we show that the rescue requires the network to be very close to criticality and fails in the near sub- and supercritical regimes. We demonstrate the robust rescue of χ in the critical model over a domain covering 4 orders of magnitude in fractional sampling within the temporal coarse graining vs. thresholding parameter regime. Using this approach, we demonstrate the robust recovery of $\chi \approx 2$ in cellular 2-photon imaging data of the frontal cortex of awake

mice over a large range of varying thresholds and recording conditions. We further show that, in contrast, the 'crackling noise' relationship and related measures that can predict χ [67, 77, 106] fail under conditions of fractional sampling. Our findings demonstrate the pitfalls of assessing avalanche synchronization when only parts of a critical network are observed and demonstrate the regime of potential scaling rescue for synchronization in the form of parabolic avalanches under experimental conditions of fractional sampling.

4.3 Methods

We will focus on the same model used in the previous chapter. For more details, see Section 3.4.

4.3.1 Correlation Analysis

To assess the mean pairwise correlation within our simulated network, we selected a smaller temporal segment consisting of 10,000 time steps, allowing us to manage the large datasets generated from neuron raster analyses. For each pair of neurons (represented as an ordered pair $[i, j]$), we calculated the time-lagged correlation between the activity of neuron i and neuron j, shifted by one time step, $k \cdot \Delta t$. The mean correlation across the network was then derived by averaging the correlation coefficients for all these pairs.

This analysis considered temporal bins ranging from 0 to 100 and spatial sampling fractions from 0.01% to 1%. Due to the exponential growth in both duration and memory requirements with increased system size, higher sampling rates made it challenging to store the raster data and perform pairwise analyses. Consequently, we focused on these smaller sampling fractions to maintain computational feasibility while exploring the correlation structure within the network.

4.3.2 Deviation from criticality coefficient (DCC)

The scaling exponent χ, discussed in detail in the previous chapter, determines how the average avalanche size $\langle S \rangle$ grows with its duration T. Mathematically, this relationship can be expressed as $\langle S \rangle \sim T^{\chi}$. To estimate χ, we computed the slope of the plot showing average avalanche size $\langle S \rangle$ against duration T.

In critical systems, the scaling exponent χ has been related to the critical exponents α' and β' via the formula:

$$\chi_{\text{cn}} = \frac{\beta' - 1}{\alpha' - 1}, \tag{4.1}$$

where χ_{cn} is often called the "crackling noise" scaling exponent [50, 106, 155] and depends on slope estimates of the critical exponents α and β. This relationship has been used to gauge how close a system is to criticality, with a derived measure $DCC = \chi - \chi_{\text{cn}}$. A *DCC* approaching zero has been interpreted to suggest that network dynamics are near criticality [29]. However, as we'll demonstrate in this chapter, this measure can be highly sensitive to operations like thresholding and temporal coarse-graining, indicating the need for a more cautious approach when using it to assess criticality.

4.3.3 Collapse Exponent of Threshold and Fractional Sampling

In our analysis of the effects of temporal coarse graining on thresholding and fractional sampling, we collapsed the χ vs. k curves for k within the range 1 – 100. In the fully (100%) sampled model, we examined various thresholds (θ_{100}) ranging from 10^2 to 10^5. In the case of fractionally sampled models, the threshold was adjusted proportionally to an exponent, ξ, of the sampling fraction, f, according to the equation. For each ξ value

spanning from 0 to 2 in steps of 0.1, we computed the collapse error of the five curves representing sampling rates from 0.01% to 100% of χ vs. k. The collapse exponent ξ is determined by the minimum collapse error and was found to be close to 1.

4.4 Results

Neuronal avalanches were studied using a critically balanced excitatory-inhibitory (E-I) model with all-to-all connectivity consisting of one million neurons driven externally using independent, stationary Poisson processes [49] (Fig. 4.1a). The external driving triggered an average of 20 neurons to fire per time step, Δt, which intermittently induced spreading of neuronal activity in the network. Simulations spanned across 10^8 time steps. An avalanche was defined as a series of contiguous time periods, T, of coincident firing within Δt surpassing a predetermined threshold, q, over a specific number of time steps, k (Fig. 4.1b, e). Thus, our model-based approach was characterized by two basic parameters: $k \times \Delta t$ and q, which were explored in the context of the subsampling fraction of observed neurons, f, ranging from 1 out of every 10,000 neurons observed ($f = 0.01\%$) to all neurons observed ($f = 100\%$). We identified the statistical distribution of sequence size, S, and duration, T, as a function of q, f, and k and estimated the corresponding size exponent, α, and duration exponent, b, within the constraints of the lower and upper cut-offs (Fig. 4.1c; see Methods; [29, 101]). In our critical model, these distributions take on the form of power laws at $k = 1$, $f = 100\%$, and $q = 100$, the defining characteristics of avalanches [13]. The scaling exponent, χ, governs the growth in average avalanche size, $\langle S \rangle$, for a given duration, T,

$$\langle S \rangle \propto T^{\chi}, \tag{4.2}$$

and accordingly, we estimated the scaling factor χ from the slope of the average avalanche size, $\langle S \rangle$, for a given avalanche duration, T, outside the minimal and maximal cut-off in

logarithmic scales (Fig. 4.1d). This scaling factor was previously referred to as χ_{sh} as it captures the scaling of short-duration avalanches that occur within the power law regime of the size and duration distribution (see also [24]). The fitting procedure (see Methods) also identifies a second slope, χ_{lg}, for long-duration avalanches that occupy the cut-off regime of the corresponding size and duration distributions. As demonstrated previously and in line with expectations for uncorrelated activity [24], χ_{lg} remains close to 1 independent of thresholding and coarse-graining and this trivial scaling factor will not be considered further.

4.4.1 High Coincident Firing Thresholds Underestimate χ in the Fully Sampled Model

In systems characterized by a clear separation of timescales and an absence of noise, the necessity for a coincident firing threshold diminishes. However, within real systems lacking such temporal separation and affected by significant intrinsic and measurement noise, the statistical analysis becomes intricately tied to our selection of threshold values. Setting the threshold too low poses the risk of incorporating substantial noise from the system, while setting it too high leads to data sparsity, resulting in the exclusion of a significant portion of the system's activity. To emulate the noisy environment and absence of timescale separation between avalanche generation and propagation inherent in real systems, we introduce a minimal external Poisson drive in our simulations. This Poisson drive establishes a noise floor within the spiking activity, serving as the origin of avalanches within the system. This is clearly visible in the size distribution of avalanches, which exhibits a power law only for avalanche sizes of $S_0 > 50 - 100$ spikes up to the finite size of the network [49] (cf. Fig. 4.1c left). To isolate avalanches above the noise floor (i.e., $S > S_0$), we first studied the dependence of α, β, χ on θ and k in the fully sampled network ($f = 100\%$; Fig. 4.2). In

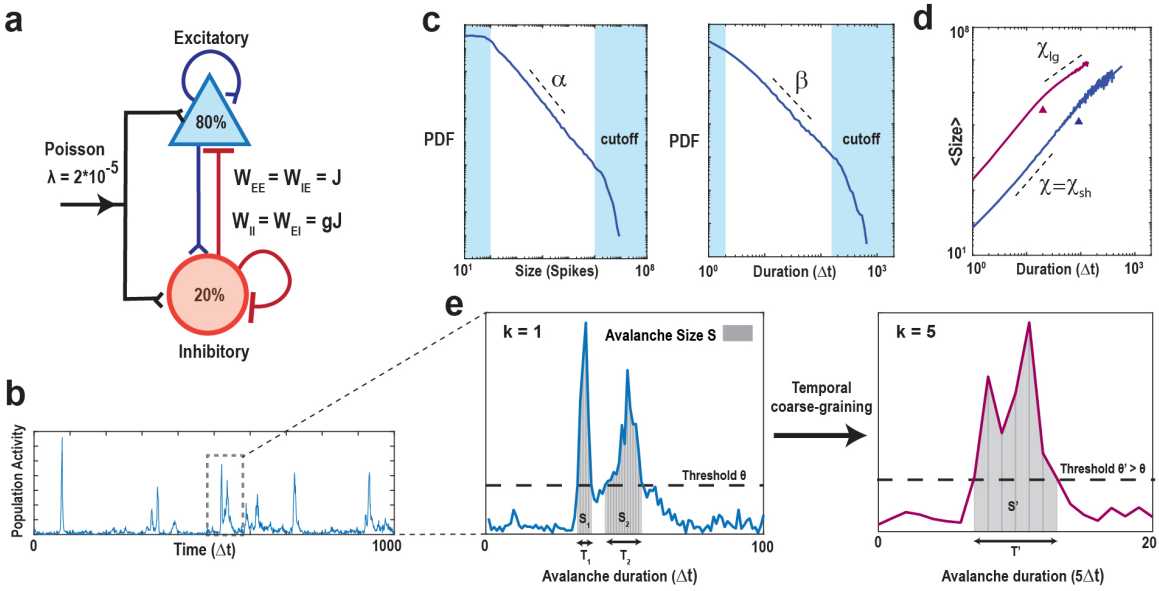

Figure 4.1: Schematics of the balanced E/I-model and identification of suprathreshold sequences in population spiking activity with respect to minimal coincident spiking threshold and temporal coarse graining. (a) Schematic of the neuronal network consisting of $N = 10^6$ neurons (80% excitatory and 20% inhibitory) with a low external Poisson drive of rate $\lambda = 2 \times 10^{-5}$ per time step per neuron. The E/I balance is controlled by g, which scales the inhibitory weight matrices (W_{II}, W_{EI}) as a function of the excitatory weight matrices (W_{EE}, W_{IE}). If not otherwise stated, we set $g = 3.5$ to obtain critical dynamics (see Methods). (b) High variability in intermittent population activity characterizes critical dynamics. Snapshot of the summed neuronal spiking activity as a function of time. (c) Size and duration of avalanches in the critical model follow power-laws with corresponding exponent estimates α and β, respectively. Note that the external Poisson drive and the finite size of the network introduce lower and higher cut-offs, respectively (shaded areas). (d) Scaling of mean avalanche size as a function of duration also follows a power-law. χ is the estimate of the scaling exponent for short avalanches (below cut-off point, $\triangle$). Avalanches within the upper cut-off (above cut-off point, $\triangle$) exhibit a trivial scaling exponent close to 1, denoted as χ_{lg}, that is largely independent of threshold and temporal coarse-graining and will not be considered further. Purple: Corresponding scaling relation and cut-off for temporal coarse-graining with $k = 5$ (see below). (e) Zoomed population activity from B. At the original temporal resolution Δt and given the coincident spiking threshold, θ (left), we can identify two sequences of suprathreshold activity, S_1 and S_2 with durations T_1 and T_2, respectively. Temporally coarse-graining the population activity (right; $k = 5$, binning the data into new bins of 5 time points) and increasing the threshold to $\theta' > \theta$, absorbs S_1 and S_2 into a new suprathreshold activity period T' with size S' with the corresponding change in scaling (see also d).

Fig. 4.2a, we show an overview in the form of color maps for all three parameters when increasing θ. At the highest temporal resolution ($k = 1$), we observed a rapid decrease in χ below its expected value of 2 as θ increases. This change largely follows the dependence of β while α remains close to 1.5, the latter being in line with findings from LFP avalanches in non-human primates [134] (Fig. 4.2b). At high thresholds, χ tends to approach 1.3, a value commonly reported in experiments [27, 44, 70, 84, 102, 184], which is likely a consequence of subsampling biases inherent in these experimental settings. We conclude that high thresholds in weakly driven, noisy networks bias towards longer suprathreshold periods, significantly underestimating the scaling exponent in critical systems.

A similar underestimation holds true for the χ_{cn} value. Inserting the corresponding values of exponents extracted from our fully sampled system for small thresholds, we obtain $\chi_{\mathrm{cn}} \approx 2$ and DCC ≈ 0, in line with expectations for the mean-field directed percolation (MF-DP) universality class [120]. When θ increases, χ_{cn} drops significantly below 2, which is similar to our observation for χ and accordingly the DCC remains close to zero (see Fig. 4.2a). We attribute the small, yet consistently observed underestimation of χ with respect to χ_{cn} (Fig. 4.2b) to the error-bounds of our fitting procedure. We conclude that avalanche scaling will be underestimated even in fully sampled networks if there are technical constraints on minimal coincident spike counts that require the use of high thresholds.

4.4.2 Temporal Coarse-Graining Rescues χ at High Thresholds in the Fully Sampled Model

Next, we show that temporal coarse graining rescues χ in the presence of a high threshold in fully sampled networks. This is shown in Fig. 4.2c, in which χ approaches 2 for $k > 8$ at $\theta = 1000$ (see also corresponding projections in Fig. 4.2a). In contrast, the size and duration distributions exhibit a gradual reduction in steepness when k increases, due to

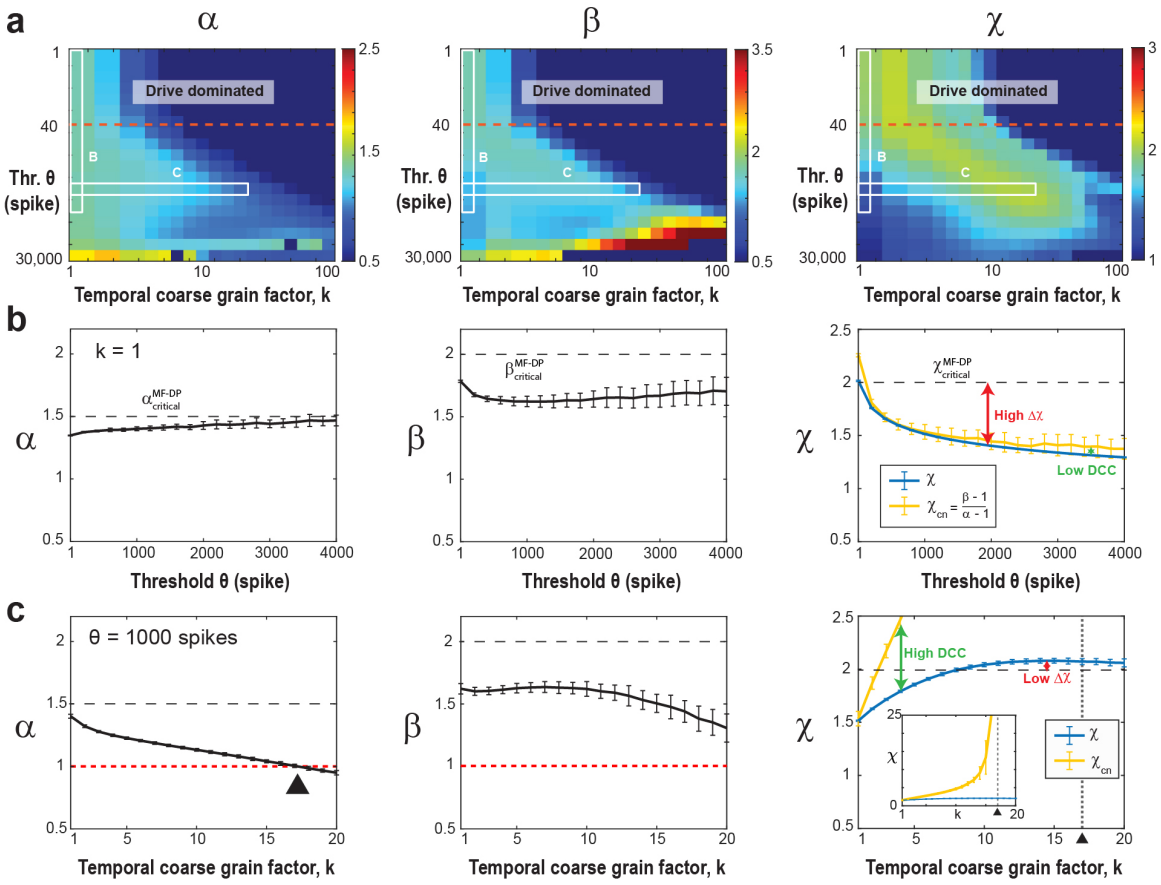

Figure 4.2: Increasing the threshold in the fully sampled model underestimates the scaling exponent, χ, which can be rescued by temporal coarse graining. (a) Consolidated view of the avalanche size (α), duration (β), and scaling (χ) exponents as a function of threshold θ and temporal coarse-graining factor k ($f = 100\%$, fully sampled model). Drive dominated: Low threshold regime (above the red broken line) dominated by external Poisson drive. White frames: Parameter regions displayed in b and c. (b) At the highest temporal resolution ($k = 1$), high thresholds underestimate the scaling exponent, χ, as well as χ_{cn}, with DCC remaining low. Size exponent, α (left), duration exponent, β (middle), scaling exponent, χ, and expected crackling noise relation, χ_{cn}, (right) as a function of θ. (c) At high threshold ($\theta = 1000$), temporal coarse graining recovers χ, but the χ_{cn} exhibits a singularity leading to high DCC. Size exponent, α (left), duration exponent, β (middle), as a function of coarse graining factor, k, ($\theta = 1000$). α passes through 1 (red broken line; black triangle), which causes a singularity in χ_{cn}. (right). Note that the scaling exponent, χ, stabilizes to a value of 2 (right, black dashed line), whereas χ_{cn} grows until it passes through a singularity at the temporal coarse graining value of $k = 17$ (vertical broken line; triangle; see inset).

concatenating subsequent cascades, in line with experimental findings [13, 84, 134, 140]. In the model, this eventually causes α to cross the value of 1 (Fig. 4.2c), which forces χ_{cn} to undergo a singularity (Equation 4.1; Fig. 4.2c, inset). Accordingly, we find a rapid breakdown in the agreement between χ and χ_{cn} with an increase in k, with diverging DCC (Fig. 4.2c). We conclude that while temporal coarse-graining can mitigate the bias introduced by thresholding in the fully sampled model, the DCC remains an unreliable estimate for avalanche scaling even in fully sampled networks under high thresholding conditions.

4.4.3 Rescue of $\chi = 2$ Is Only Seen in Critical Dynamics That Is Fractionally Sampled and Fails for Slightly Super- or Sub-critical Dynamics

We recently demonstrated that the rescue of the critical scaling exponent, χ, was exclusive to the critical model ($g_c = 3.5$) and failed for subcritical conditions [24]. Here, we explored this dependency more systematically and show that small deviations from criticality by changing the E/I-balance parameter, g (see Methods), significantly impede rescue to the expected critical (MF-DP) value of 2. In Fig. 4.3a, we illustrate the maximum χ across various temporal coarse graining factors k, plotted against g for subsampled models ($f = 0.1\%$) utilizing a 1-spike threshold ($\theta = 1$). In subcritical models ($g > 3.5$), the restored scaling exponent swiftly deviates from its critical value of 2 towards 1. Conversely, in slightly supercritical models ($3.45 < g < 3.5$), χ also rapidly degrades to below 2. We note that for $g < 3.45$, persistent activity (shaded gray regions in Fig. 4.3a) obstructs the identification of avalanches at the specified threshold. This demonstration carried out for a critical branching model suggests that a successful rescue of the critical value of χ implies the dynamics to be at an E/I-balance of the order of 0.05 or 1.4% close to criticality. While further analytical studies are warranted to establish the rescue of χ as a key hallmark of critical dynamics, our findings suggest temporal coarse-graining combined with thresholding

to recover the expected value χ within a particular universality class might be a sensitive approach when confronted with fractionally sampled critical dynamics.

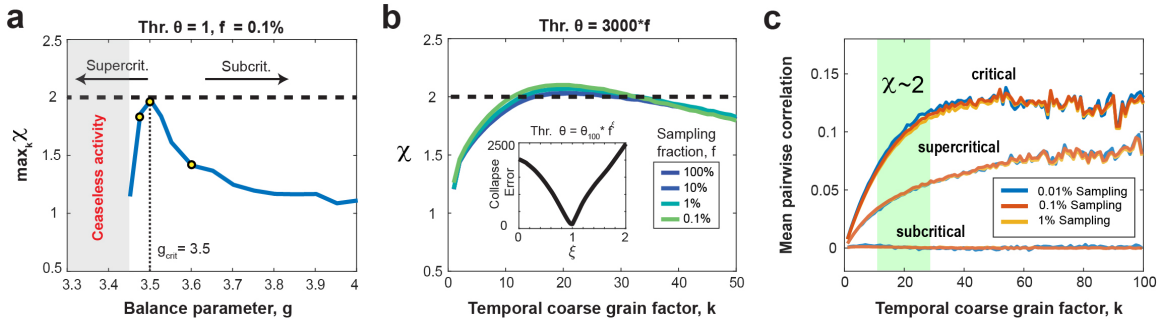

Figure 4.3: Temporal coarse graining at critical value of g: Rescue of scaling exponent, equivalence between threshold and fractional sampling, and increased mean pairwise neuron correlation. (a) The maximum scaling exponent across different temporal binning values as a function of the balance parameter, g, for a subsampled model ($f = 0.1\%$) employing a 1-spike threshold. Notably, the retrieval of $\chi = 2$ is exclusively observed in critical models ($g_c = 3.5$), but promptly diminishes in slightly subcritical or supercritical models. In supercritical models, we obtain ceaseless activity (shaded gray region), precluding the identification of avalanches at the specified threshold. Yellow circles mark g values used in c. (b) Collapsed scaling exponent curves for different values of the sampling fraction, f, ranging from 100% to 0.1% for a threshold that is scaled by the sampling fraction as $\theta = 3{,}000 * f$ (i.e., a threshold of 3,000 spikes for the fully sampled model translates to a threshold of 3 spikes for the 0.1% sampled model). Inset: The total collapse error as a function of collapse exponent, ξ. The error is taken over multiple different curves with θ ranging from 100 to 100,000. The error shows a clear minimum at $\xi = 1$, indicating the threshold can be scaled proportionally with the sampling fraction to obtain the best collapse. (c) The rescue of $\chi = 2$ correlates with an increase of the mean delayed pair-wise correlation among neurons, which is independent of the fraction of neurons sampled. However, the temporal coarse graining regime of k for which $\chi = 2$ (green area) is not predicted by this pair-wise correlation. For a supercritical ($g = 3.475$) network, the mean pairwise correlation is lower, but still increases with temporal coarse graining. For a subcritical network ($g = 3.6$), the mean delayed pair-wise correlation is close to 0.

4.4.4 Scaling Collapse Demonstrates Equivalence of Thresholding and Subsampling in the Critically Balanced All-to-All E/I-Model

Both the increase in coincident spiking threshold θ as well as a decrease in sampling fraction, f, underestimate the scaling exponent χ due to prematurely terminating avalanches by unobserved network activity. On the other hand, temporal coarse graining allows for concatenating sequential activity periods potentially recovering the proper duration for a subset of avalanches. We therefore examined the impact of temporal coarse graining on thresholding and fractional sampling. To achieve this, we first collapsed the χ vs. k curves over a wide range of θ (10^2 – 10^5) by scaling θ with an exponent of f, ξ (see Methods). We found that the corresponding functions, obtained for different values of f, exhibited a minimum total collapse error for $\xi = 1$, demonstrating a linear relationship between θ and f in our all-to-all model. This implies that a fully sampled network (10^6 neurons) with $\theta = 3,000$ spikes is qualitatively equivalent to a 0.1% sampled model (1,000 neurons) with $\theta = 3$ spikes (Fig. 4.3b). This equivalence between θ and f allows us to rescale the rescue of χ for different f in the corresponding (k, θ)-plane for ease of comparison (see Fig. 4.4).

4.4.5 The Rescue of $\chi = 2$ Indicates a Recovery of Mean Pair-Wise Correlations

Subsampling as well as decorrelation of spiking activity have been shown to reduce the expected growth in mean avalanche size with avalanche duration from $\chi = 2$, to a value of χ closer to 1 [27, 44, 84]. We therefore studied how the mean delayed pair-wise correlation, $\langle C \rangle$, changes with temporal coarse graining, k, when recovering χ (see Methods). Temporal coarse graining increases $\langle C \rangle$ in the critical balanced E/I-network and saturates at high k,

independent of the amount of fractional sampling (Fig. 4.4c). The increase in correlation did not predict the temporal coarse graining value of k at which $\chi = 2$ and was not found in strongly subcritical networks, for which the average pairwise correlation remains near zero and independent of k (Fig. 4.3c; $g = 3.75$). In slightly supercritical models, the recovery of $\langle C \rangle$ was significant but did not support $\chi = 2$ for the critical model. Taken together, these findings demonstrate that $\chi = 2$ identifies the expansion of coincidental firing among groups of neurons in the critical model and is not adequately captured by a particular value of $\langle C \rangle$, a measure limited to studying pairs of neurons.

4.4.6 Subsampling and Spike-Resolution Introduce Unobservable Regions of Scaling Exponents

By taking advantage of the trade-off between thresholding and sub-sampling (see Subsection above), we can now rescale our k vs. θ parameter space by $f \cdot \theta$ to account for different fractional sampling. In the consolidated analysis (Fig. 4.4), we demonstrate reasonable agreement with the rescue of $\chi = 2$ from 0.1% - 100% fractional sampling. When we increase the threshold (or equivalently, decrease f), we observe a deviation from $\chi = 2$, yet the exponent can be reliably restored at higher values of k (Fig. 4.4, white broken lines). On the other hand, reducing f eventually introduces unobservable (k, θ)-regions which arise from the 1 spike resolution of the model excluding fractional thresholds. We note that even when sampling only 0.1% of the neurons, one can still recover $\chi = 2$. However, this approach fails to simultaneously restore all exponents while maintaining the crackling-noise relationship (see Fig. 4.2 and 4.5) because the point of concurrent rescue lies within the unobservable regions of parameter space. Decreasing sampling further to $f = 0.01\%$ (100 neurons out of 1 million neurons) prevents rescue of $\chi = 2$.

Temporal coarse graining does not rescue the 'crackling noise' relationship of χ_{cn}

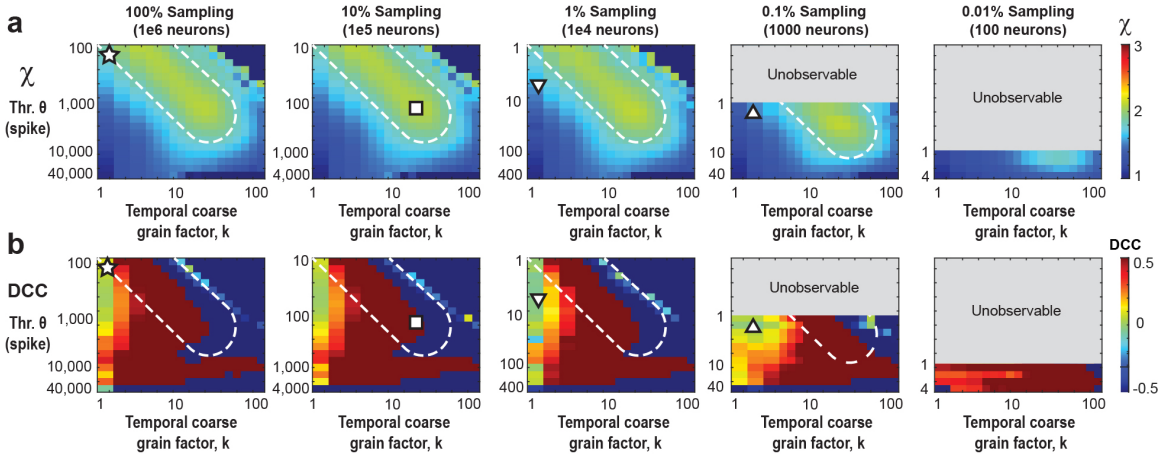

Figure 4.4: Temporal coarse-graining identifies a robust rescue regime for $\chi = 2$ covering 0.1% - 100% fractional sampling at decreasing thresholding level, which is absent for the DCC. (a) Consolidated view of χ as a function of θ and k for different values of f. For $f = 100\%$, $\chi \approx 2$ for low θ and k (star), but as we make the data sparser by increasing the threshold (or reducing the sampling), we need a higher coarse-graining factor, k, to compensate and rescue χ back to 2 (square). White dotted region: a visual guide for χ close to 2 for increasing θ and k. The grey parts of the plots for $f = 0.1\%$ and 0.01%, respectively, are unobservable parameter regions since they would require fractional thresholds below the 1 spike minimum resolution of the model. At $f = 0.1\%$ sampling, χ can be rescued to a value of 2, but not for more severe subsampling, $f = 0.01\%$. For triangles, see b. (b) The DCC remains close to 0 only for low k, and quickly breaks down at higher k values. For the fully sampled model, there exists a region with the correct scaling exponent $\chi \approx 2$ as well as low DCC (star in A and B). However, as we move to lower sampling, this region of correspondence becomes more and more difficult to maintain. For $f = 1\%$ and 0.1%, regions at low k, and moderate threshold (or equivalently lower sampling) remain at low DCC, but underestimate the true scaling exponent, χ (up and down triangles, respectively).

(Equation 4.1). The DCC, introduced recently as the difference between the data-based value of χ and χ_{cn}, should measure close to 0 if both estimates agree. As can be seen in Figure 4.4, co-regions with DCC = 0 and $\chi = 2$ exist in the fully sampled system (Fig. 4.4; star). In contrast, no such co-alignment is found for subsampled systems, and temporal coarse graining fails to establish a DCC = 0 regime for most values of k and θ (Fig. 4.4; square, triangles), and accordingly, DCC = 0 does not reflect a proper rescue of the true $\chi = 2$ under these conditions.

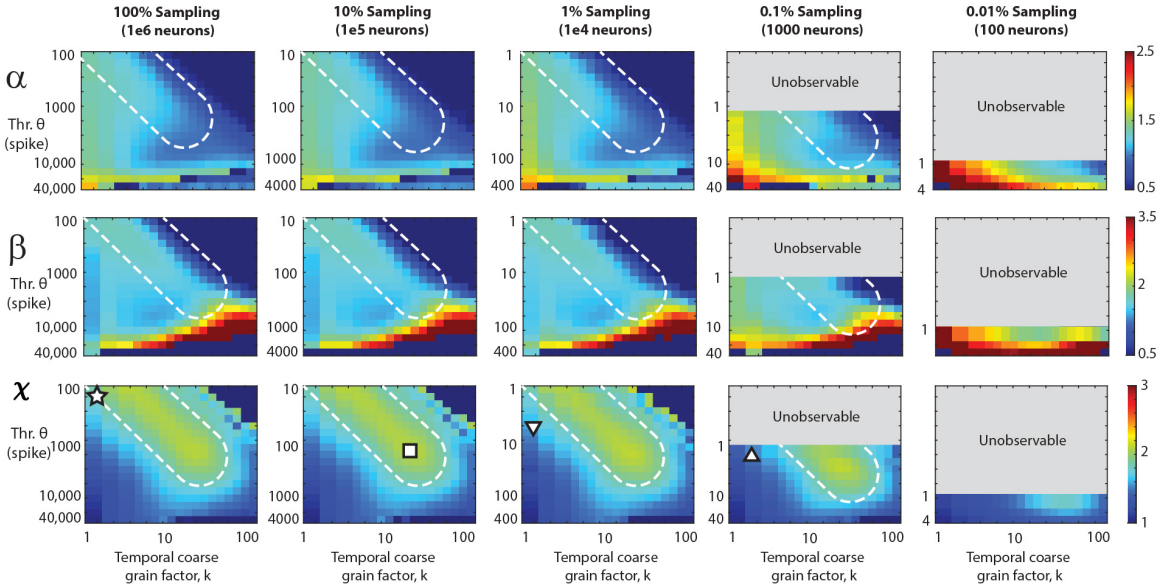

Figure 4.5: Consolidated view of the exponents α, β, and χ as a function of threshold and temporal coarse-graining for different sampling fractions. Consolidated view of α, β, and χ as a function of θ and k for different values of f. For $f = 100\%$, $\chi \approx 2$ for low θ and k (star), but as we make the data sparser by increasing the threshold (or reducing the sampling), we need a higher coarse-graining factor, k, to compensate and rescue χ back to 2 (square). Bottom row is replotted for χ from Figure 4 for ease of comparison with the corresponding size slopes, α, and duration slopes, β. White dotted region: visual guide for χ close to 2 for increasing θ and k. The grey parts of the plots for $f = 0.1\%$ and 0.01% respectively are unobservable parameter regions since they would require fractional thresholds below the 1 spike minimum resolution of the model (cf. Figure 4.4).

4.4.7 Robust Recovery of $\chi = 2$ for Parabolic Avalanches in Frontal Cortex of Awake Mice

In the previous section, we demonstrated that temporal coarse-graining and thresholding effectively identify a robust subregion where the scaling exponent χ is rescued in the model. This rescue can be collapsed across different levels of fractional sampling. We now present similar findings from cellular 2-photon imaging of ongoing neuronal activity in the anterior cingulate cortex of awake mice ($n = 17$ experiments, each 30 minutes long; $n = 5$ mice; see Methods for details). The avalanche number distribution, obtained for each

recording, follows a lognormal pattern as a function of the threshold (Fig. 4.6a, d; see also [24]) and guides our analysis. In addition, this lognormal distribution varies with temporal coarse-graining and is thus a function of k (Fig. 4.6a, d).

In a first approach, we pooled all data sets for a given absolute spike threshold (Fig. 4.6a), which revealed a consistent subregion of $\chi \approx 2$ for different spike density threshold levels (Fig. 4.6b). Specifically, we observe that the recovery of $\chi \approx 2$ shows a θ vs. k dependency, consistent with our simulations, whereas the DCC remains variable and noisy (Fig. 4.6c).

In a second approach, we z-scored the avalanche number distribution for each data set in order to account for experimental variability as well as for variable shifts induced by temporal coarse graining (Fig. 4.6d). Figure 4.6e shows that, regardless of the exact choice of z-scored threshold within a wide range (-2 to $+1.5$), χ is robustly recovered to approximately 2. The range of k values that recover χ decreases as the threshold increases. This robust recovery supports previous studies that adopted an avalanche-maximizing threshold (z-score = 0) [15] or lower thresholds (z-score = -2) [24]. We note that at very low thresholds, noise becomes problematic, and at very high thresholds, data scarcity leads to poor statistics. In contrast, the DCC measure remains highly variable and unreliable even when the threshold is individually selected for each experiment and coarse-graining value (Fig. 4.6f).

4.5 Discussion

The study of criticality in neuronal systems, particularly neuronal avalanches, has been extensively researched and debated [30, 31, 58, 67, 119, 136]. This study addresses the intricacies and challenges of estimating scaling relationships in the brain, especially under conditions of fractional subsampling. A significant finding of our research is the ability

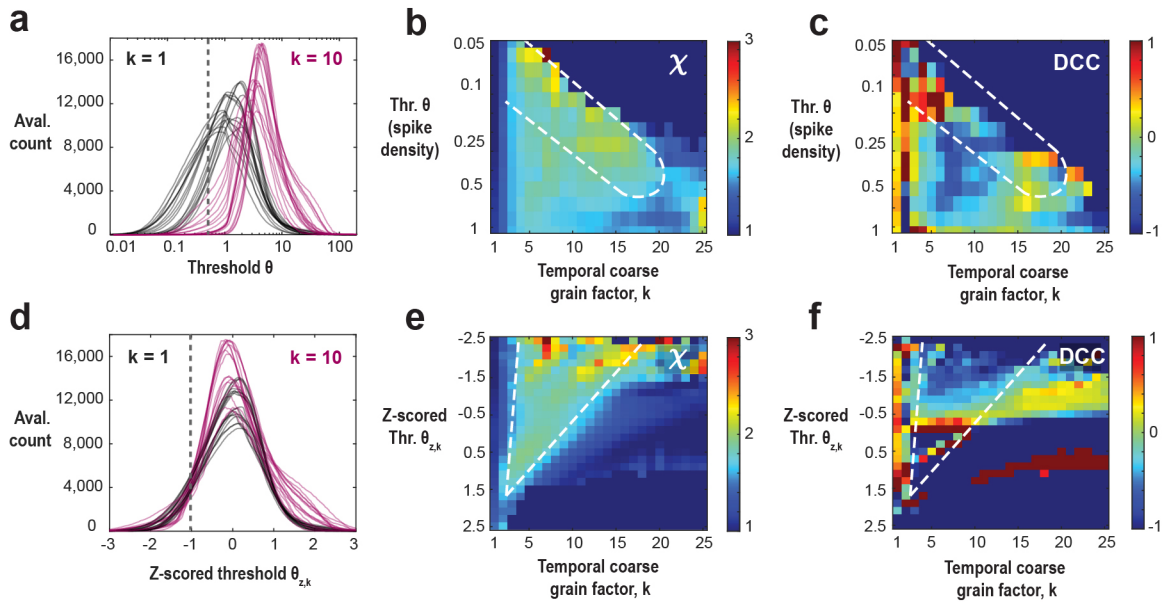

Figure 4.6: Parabolic avalanches in ongoing activity of frontal cortex exhibit threshold and temporal coarse-graining rescue of scaling exponent $\chi = 2$ in line with critical model dynamics. (a) Lognormal distribution of avalanche number as a function of threshold for 17 different recordings ($n = 5$ mice; 2-photon imaging of ongoing activity in frontal cortex) at the original temporal resolution ($k = 1$, black) and after temporal coarse graining ($k = 10$, purple). Note the experimental variability at a given k, as well as the systematic shift in the distribution with change in k. A sample threshold (θ, grey dotted) shows how a given fixed threshold relates to these distributions. (b) Scaling exponent, χ, and (c) crackling noise deviation, DCC, as a function of temporal coarse grain factor, k, and the spike density thresholds. The subregion in which $\chi \approx 2$ (broken line) is similar to the one identified in the critically balanced E-I model. At low thresholds, the scaling is rescued at smaller values of temporal coarse graining factor, k and at high thresholds, it is rescued at a larger value of k. Over this large range of parameters, the DCC is very noisy, and even a small change in the parameters can cause a large change in the DCC. (d) Z-scoring procedure for the lognormal distributions corrects for experimental variability as well as temporal coarse graining. A sample z-scored threshold ($\theta_{z,k}$, grey dotted) shows how a given z-scored threshold relates to these distributions. (e) Scaling exponent, χ, and (f) crackling noise deviation, DCC, as a function of temporal coarse grain factor, k, and the z-scored threshold at the given k, $\theta_{z,k}$. Figure d shows a robust rescue of $\chi \approx 2$ (broken line) for a large range of z-scored threshold values (-1.5 to 1.5). The range of k values over which $\chi \approx 2$ becomes smaller with an increase in the z-scored threshold. Like b, Figure e shows that the DCC remains unreliable and noisy even when the thresholds are z-scored and tuned for each k.

to retrieve scaling exponents with fractional subsampling by employing temporal coarse graining and threshold adjustments. The sensitivity of this approach to deviations from criticality supports our method to study potential universality classes of criticality in real-world systems such as the brain. We demonstrate that temporal coarse-graining, combined with stricter criteria for coincident neuronal firing, allows for accurate estimation of the critical scaling exponent ($\chi = 2$), even when observing as little as 0.1% of the network's neurons. This suggests that it's possible to accurately determine critical scaling characteristics despite heavy subsampling limitations. In contrast, previous studies on similar directed percolation (DP) models found a reduction in the scaling exponent to 1.3 with spatial subsampling of 10% and below [27]. Additionally, subsampling has been observed to falsely suggest population independence of correlations, impacting markers of criticality such as specific heat [122]. It's important to note that our model, due to its all-to-all connectivity, does not allow for the exploration of spatial correlation dependencies.

4.5.1 Temporal coarse-graining when avalanche periods are unknown

Temporal coarse-graining has previously been suggested as a means to preserve the size exponents of avalanches in subsampled systems, assuming an infinite separation of time scales [84]. This assumption would theoretically allow researchers to distinguish sequentially occurring avalanches, a scenario that may not hold in natural systems. In contrast, our analysis shows that temporal coarse-graining can effectively maintain the scaling exponent in models lacking a clear separation of time scales, even with fractional sampling as low as 0.1%. This suggests its utility in analyzing highly subsampled experimental data within the directed percolation (DP) class. Our approach introduces a threshold above the Poisson noise floor to approximate the separation of successive avalanches, revealing a trade-off between subsampling intensity and the threshold level. In our all-to-all model, increasing

the threshold of the fully sampled network mimics the effects of subsampling combined with a lower threshold, illustrating the interaction between these factors and their collective impact on criticality evaluations. Leveraging this trade-off, we apply our methodology to cellular-resolution 2-photon imaging data from the anterior cingulate cortex of awake mice, consistently recovering the scaling exponent to a value of 2, in line with expectations for a critical branching process. Additionally, new avalanche detection methods that utilize continuous time models have been developed, though their effectiveness compared to traditional methods under experimental and subsampling conditions remains to be fully evaluated [132]. Our recovery of the scaling exponent $\chi \approx 2$ in ongoing neuronal spike synchronization of the awake mouse frontal cortex provides compelling evidence for the critical role of $\chi \approx 2$ in capturing essential features of neuronal avalanches (see also [24, 114]). For instance, models with slow latent variables can generate heavy-tailed or even power-law distributions of neuronal activity size and duration, but they typically show $\chi \approx 1$, suggesting a lack of correlated activity [3, 135, 141]. While the critical branching process used here achieves $\chi \approx 2$, it does not account for temporal correlations among successive avalanches (e.g., [88, 89, 135, 136]), limiting its ability to fully model neuronal avalanches in the brain.

It's important to distinguish temporal coarse-graining from renormalization procedures. While renormalization in neuronal avalanches can preserve estimates of maximal mutual information during spatial coarse-graining under complex LFP sampling conditions [165], our focus here is on addressing stochastic subsampling bias through temporal coarse-graining, as detailed previously [24].

4.5.2 Limits in rescuing the scaling exponent under subsampling conditions

Our simulations also clearly show that subsampling in discrete models renders certain parts of the phase space inaccessible (see Fig. 4.4), preventing the recovery of the scaling exponent under heavy subsampling ($f \ll 0.1\%$). A similar limitation exists for temporal subsampling, especially when the required precision exceeds the system's resolution [27]. Properly addressing these spatial and temporal constraints is crucial for the analysis of experimental data with limitations in both resolution domains. Further limitations arise with extensive subsampling of network topologies, particularly when estimating heavy-tailed or scale-free distributions [19, 161, 173]. While under-sampled critical branching processes have been explored in extracellular spike recordings with microelectrode arrays in the anesthetized rat [147], these limitations shouldn't overshadow our findings' potential. Careful consideration of experimental parameters can ensure feasible and informative criticality assessments, providing valuable insights into the core mechanisms of neural synchronization.

4.5.3 The DCC and crackling noise relationship fails to recover the scaling exponent under subsampling conditions

The Deviation from Criticality Coefficient (DCC) approach [44, 96, 97] has gained popularity for quantifying criticality based on the "crackling noise relationship" [67, 77, 106]. However, concerns have been raised about its ability to identify critical dynamics [38, 118] given that random processes can exhibit a DCC close to 0, thus limiting general claims of this approach to identify critical dynamics even in fully sampled systems. For instance, an unbiased random walk will produce excursions above the mean that are power-law distributed in sizes, with exponent 1.5, durations, with exponent 1.33, and scaling with expo-

nent 1.5, which result in the random walk obeying the crackling noise relationship (DCC = 0), despite not being a critical system. Our study demonstrates several limitations of the DCC approach, particularly under subsampling conditions: (1) DCC estimates become unreliable with limited data, (2) the DCC is highly sensitive to temporal coarsening, leading to noisy results, (3) the DCC metric has a singularity under fractional sampling, making it a questionable marker for criticality. Furthermore, the DCC approach assumes power-law slopes obtained from size and duration distributions directly translate to critical exponents. This is flawed, as analysis parameters like temporal resolution and spatial sampling distance systematically influence these slopes [13, 134]. These factors must be considered when interpreting criticality from experimental data using the DCC approach.

4.5.4 The sensitivity of temporal coarse-graining approach to thresholds in avalanche size definition

Our analysis emphasizes the scaling exponent's sensitivity to synchronization thresholds. As the threshold increases, we observe a noticeable deviation of the scaling exponent and the crackling noise ratio (χ_{cn}) from the expected value of 2 (see also [180]). This sensitivity highlights the importance of careful consideration when choosing thresholding strategies in experimental data. Thresholding of time-series and ignoring the subthreshold contribution during suprathreshold periods can lead to a change in the avalanche size exponent itself as demonstrated in simulations of critical dynamics [180]. Capek, Ribeiro et al. [24], in their supplemental figure 6, experimentally demonstrated that including or excluding the subthreshold regime for size estimates provides narrow upper and lower bounds for χ that center around 2. Ideally, threshold selection should strike a balance between preserving genuine synchronization events and minimizing noise, especially when accurate scaling exponent estimates are crucial.

4.5.5 Synchronization, coincident firing and delayed neuronal correlations

While traditional studies of spatiotemporal synchronization in neural systems have focused on phase-locked oscillations, our system operates in a unique scenario—one devoid of a dominant frequency. This lack of a dominant frequency becomes particularly important when considering 'noise correlations,' which encapsulate the intrinsic non-stimulus-induced correlation structure among neurons. These correlations exert a notable influence on the decoding capabilities of neuronal populations, constituting a pivotal aspect of neural information processing [6, 47, 146].

In this study, we investigate scale-free neuronal avalanches as a context for exploring neuronal synchrony and criticality. Avalanches offer a powerful tool for understanding neuronal synchronization, from local clusters to the entire brain. As we've shown, analyzing avalanches depends crucially on choosing thresholds and temporal resolution. This dependence is similar to the analysis of spatial and temporal neuronal correlations, which are heavily influenced by their respective coarse-graining [160, 169]. Interestingly, as shown in Figure 4.3, restoring the scaling exponent χ to its critical value of 2 doesn't directly correspond to the recovery in pairwise correlations. This emphasizes that χ at 2 signifies a critical system that describes the coordination among many neurons, not just an average estimate of correlation among pairs of neurons.

Our findings have implications for both neuroscience and complex systems research. Neuronal avalanches are valuable markers of synchronization in the brain. Precise scaling exponent estimation, even with limited data, is crucial for understanding mechanisms of neural synchronization in healthy and diseased brains. This supports the idea that the brain operates within a critical regime, aligning with the concept of self-organized criticality observed in many complex systems.

4.6 Acknowledgements

The research in this chapter was supported by the Division of the Intramural Research Program (DIRP) of the National Institute of Mental Health (NIMH), USA, ZIAMH002797, ZIAMH002971, and the BRAIN initiative Grant U19 NS107464-01. This research utilized the supercomputing resources of the National Institutes of Health (NIH, USA; Biowulf, http://hpc.nih.gov) and the University of Maryland (UMD College Park, USA, https://hpcc.umd.edu/hpcc/dt2.html).

Chapter 5: Heterogeneity and Local Excitatory-Inhibitory Plasticity Facilitate Robust Performance in Reservoir Computers

This chapter focuses on ongoing research exploring the role of relative excitatory-inhibitory balance and adaptive mechanisms on the dynamics and performance of reservoir computers and is currently being prepared for publication.

5.1 Overview

This chapter marks the culmination of this dissertation and aims to bring together the two themes of neuroscience and machine learning. While deep learning with backpropagation dominates machine learning, reservoir computers (RCs) offer a biologically plausible and efficient alternative. By leveraging the complexity of random networks and training only the output layer, RCs avoid the costly process of training the internal network weights. However, their performance is sensitive to the hyperparameters involved in constructing the random internal network, or 'reservoir'. A key yet under-explored feature of the reservoir is its relative balance of excitation and inhibition (E-I balance).

In this chapter, we study the role of E-I balance in RC performance and introduce mechanisms that leverage local E-I balance for improved performance. We first demonstrate that RCs with balanced to slightly over-inhibited dynamics offer better performance and greater robustness compared with even slightly over-excitatory dynamics. We then introduce an unsupervised local inhibitory plasticity mechanism that dynamically adjusts

individual neurons to a target firing rate. We also explore the utility of introducing heterogeneity in neuronal target firing rates. We find that networks with tuned local E-I balance and networks with untuned but heterogeneous local balance perform better than networks within which only the global balance is optimized. Our study highlights how biologically-inspired heterogeneity and local excitatory-inhibitory balance can improve robustness in information processing.

5.2 Introduction

In recent years, reservoir computing has provided a computationally inexpensive method to implement recurrent neural networks (RNNs) [65, 99]. Unlike deep learning models such as convolutional neural networks (CNNs), long-short term memory networks (LSTMs), and transformers, reservoir computers (RCs) offer a more biologically plausible paradigm for information processing [42, 170, 185]. Despite their notable success, deep learning methodologies frequently necessitate the training of extensive neural networks, which require substantial amounts of data and computational resources. This reliance on large-scale models poses significant challenges, underscoring the need for smaller, more efficient models capable of rapid inference and adaptation. Drawing inspiration from the brain's efficient information processing and adaptability with minimal energy expenditure, there is a growing interest in developing biologically inspired architectures that may provide a viable solution.

Reservoir computing (RC) leverages a three-layer architecture loosely inspired by the brain's intricate connectivity. The core element is a high-dimensional "reservoir," typically a randomly fixed recurrent neural network. This reservoir acts as a powerful computational unit, processing incoming data fed from the input layer. Finally, the output layer extracts relevant information by combining reservoir signals with task-specific weights learned dur-

ing training. The key advantage of RCs lies in their ease of training. By keeping the reservoir fixed and only adjusting the final layer weights, RCs achieve remarkable speed and computational efficiency.

However, the reliance on randomly generated reservoirs also presents a key hurdle. Performance can exhibit significant variability across different reservoir configurations, and more importantly, there is limited guidance for optimal construction of these random networks. Addressing these challenges, is crucial to unlocking the full potential of reservoir computing and its integration into broader machine-learning applications.

Several variations have been explored to enhance the performance of traditional reservoir computer (RC) systems. These include the development of hybrid systems that integrate knowledge-based models [131], the creation of parallel RC architectures optimized for predicting the dynamics of large spatiotemporal [130] and networked [171] systems, and the design of more intricate, layered RC structures aimed at improving memory and predictive capabilities [46]. While these approaches target performance enhancement across diverse practical scenarios, they do not directly address the intrinsic variability and dynamics of the reservoirs. Since the internal links of the reservoir are fixed, the performance of RC models critically depends on the reservoir's internal architecture. Recent theoretical studies have demonstrated that the network properties of the reservoir are connected to performance [25], with suggestions that RCs may perform best when the network is poised at the edge of chaos [26]. Additionally, incorporating features of biological neural networks, such as modularity [116, 149], small-world connectivity [37, 72], and scale-free degree distribution [37], has been shown to be beneficial for RC models.

Yet, a critical factor in biological neural networks, the relative balance between excitation and inhibition (E-I balance), is not explicitly addressed in most current reservoir computing (RC) approaches. E-I balance is known to play a crucial role in computational functions in cortical circuits, and pharmacological perturbations of this balance can induce

deficits across various cognitive functions [186]. Although RCs typically lack dedicated excitatory and inhibitory neurons, individual neurons can exert both influences due to the random distribution of outgoing positive and negative connections. However, the systematic impact of relative E-I balance on RC performance remains an under-explored area of research. Unlike traditional RCs, biological networks have distinct excitatory and inhibitory populations, often adhering to Dale's principle (where neurons have exclusively excitatory or inhibitory outputs). In addition to the anatomical ratio of excitatory to inhibitory neurons, the functional E-I balance in biological networks depends critically on the dynamical activity patterns in the network.

Here, we explicitly study the effects of relative E-I balance. Our results demonstrate the brittleness of excitation-dominated networks and highlight the critical role of inhibition in enhancing the robustness of the RC system across various parameters. By constructing RCs that are robust across a range of hyper-parameters, we aim to side-step the costly process of hyper-parameter tuning as much as possible.

Typically, hyperparameter tuning reservoir computers for specific applications involves heuristic methods, stochastic optimization, or, in some cases, exhaustive parameter grid searches. However, these approaches can be computationally expensive, with optimization often focused on achieving enhanced performance for a specific task rather than identifying trends in improved performance across tasks. Additionally, the inherent variability of randomly generated reservoirs necessitates multiple performance evaluations to obtain a representative picture of the RC capabilities.

Inspired by the brain's plasticity, researchers have developed adaptive reservoirs as a promising approach to address the challenge of inherent reservoir variability. This approach leverages unsupervised learning, reminiscent of biological adaptation, to adjust internal reservoir properties before training the output weights. In a 2007 study, Babinec et al. [7] included Hebbian plasticity to modify reservoir weights. Subsequently, other groups have

incorporated additional biologically-inspired mechanisms, such as Spike-Timing-Dependent Plasticity (STDP) [36, 80] and non-synaptic plasticity rules like Intrinsic Plasticity (IP) [158, 176]. This exploration has resulted in a rich landscape of plasticity models [115, 172, 181, 190, 191] successfully integrated into RC networks, yielding encouraging results. These adaptive learning rules aim to enhance reservoir robustness by enabling the reservoir to autonomously adjust its internal dynamics in response to the external input and its own internal dynamics. Typically, these adaptations occur independently of various hyperparameters, such as initial weight settings or network topology, facilitating the transition to an optimal dynamical regime. In this chapter, we address the challenge of local variability in global hyperparameter optimization by proposing inhibitory plasticity - a novel, unsupervised adaptive mechanism inspired by biological neurons.

Similar to how biological neurons are thought to adjust to maintain a stable average firing rate [34, 138, 142, 177, 178], our mechanism allows a neuron's incoming inhibitory weights to adapt autonomously. By tuning these weights to achieve an appropriate firing rate target, our model achieves both local and global balance, in contrast to typical implementations, which generally have only a global balance of positive and negative input weights. Additionally, we incorporate heterogeneity in target neuron firing rates to eliminate the need to fine-tune the target firing rate for the task at hand. We also introduce a streamlined method for designing networks with local E-I balance (i.e., one-step adaptation) for scenarios where complete knowledge of network connectivity is available.

We characterize the dynamics of the reservoirs using the average firing rate and a balance parameter, β, which indicates the relative balance of excitation and inhibition. Additionally, we measure the entropy of the neuron states, a metric previously correlated with reservoir performance [25, 26, 62, 124]. To evaluate performance, we train RCs for 4 different taks. First, we train for the memory capacity task [65, 149] to assess RC memory. Next, we train RCs for the Nonlinear Autoregressive Moving Average (NARMA) task, a discrete

temporal task with a 10th-order lag requiring both memory and nonlinearity, commonly used to evaluate reservoir computing (RC) systems. Lastly, we train RCs on two chaotic time-series prediction tasks, a known strength of RCs [130], using the Mackey-Glass and Lorenz systems.

Overall, this research sits at the intersection of neuroscience and artificial intelligence, leveraging observed brain properties to develop improvements in machine learning architectures. It builds on a rich body of prior work in this area [42, 156, 170, 185]. Recent findings suggest that biological neural networks exhibit properties similar to reservoir computers [42, 68], presenting an opportunity to glean insights from our results and enhance our understanding of the role of inhibition in biological neural networks.

5.3 Methods

5.3.1 Excitatory-inhibitory reservoir computing

A reservoir computer is a type of recurrent neural network (RNN) in which an input signal is fed into a network of artificial neurons known as the "reservoir." As mentioned earlier, the reservoir computing framework is structured into three layers: the input layer, the reservoir, and the output layer (See Fig. 5.1A). Training occurs exclusively in the output layer, leaving the internal properties of the reservoir unchanged during operation.

Traditional reservoir computers (RCs) have randomly distributed outgoing weights and lack distinct neuronal populations; however, our framework uses a reservoir with tunable excitatory-inhibitory balance, where neurons adhere to Dale's law. According to this law, inhibitory neurons can only send inhibitory signals, represented by negative link weights, while excitatory neurons can only send excitatory signals, represented by positive link weights (See Fig. 5.1A). Our E-I reservoir consists of rate-encoding neurons where each neuron's state at any given time is defined by two primary terms:

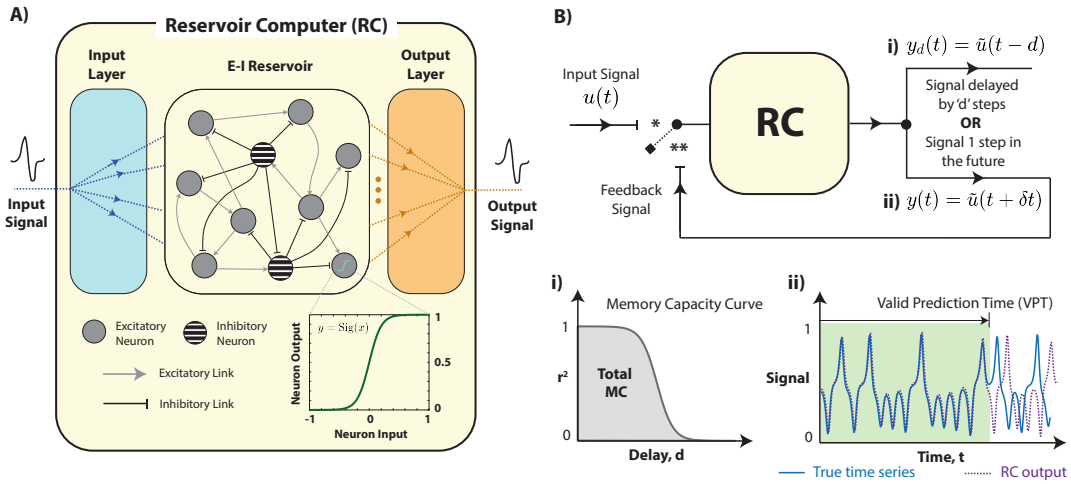

Figure 5.1: A) Schematic of our reservoir computer framework with tunable excitatory-inhibitory balance, highlighting the three main components: input layer, reservoir, and output layer. The inset illustrates the activation function for each neuron. B) The two key types of tests used to benchmark the RC setup: i) Memory Capacity Test, where the total memory capacity (MC) is calculated as the area under the memory capacity curve; ii) Time-Series Prediction, where the reservoir receives a time-series input (such as Lorenz, Mackey-Glass, or NARMA-10) and is trained to predict the signal one step ahead. Performance is measured using either root mean square error (RMSE) or valid prediction time (VPT, see Methods)

• Membrane potential (V): This internal variable describes the state of a neuron based on all its influencing factors.

• Firing variable (r): This external variable represents the observable firing rate state of a neuron, dependent on the membrane potential and the neuron's activation threshold. Values close to 0 represent no firing, and values close to 1 represent saturated firing.

The following equations govern the dynamics of these two variables:

$$V_i(t+1) = \lambda_i V_i(t) + \sum_{j=1}^{N} (A_{ij}^E - A_{ij}^I) \cdot r_j(t) + W_i^{in} x(t)$$
$$r_i(t) = \mathrm{Sig}\Big(V_i(t) - \theta_i(t)\Big) \tag{5.1}$$

Here, λ_i is the leakage rate of neuron i, which indicates how slowly a neuron's state decays

over time. The sparse connectivity matrix, A, defines how neurons in the network are connected. A fraction of the neurons are labeled 'excitatory,' with their outgoing synapses making up the excitatory submatrix A^E, and the remaining neurons are labeled 'inhibitory,' with their outgoing synapses making up the inhibitory submatrix A^I (See Sec. 5.3.2 for details on the construction of the adjacency matrix). The global relative E-I balance is controlled by a global balance parameter, β, which represents the relative strength of the excitatory and inhibitory subpopulations. This parameter is defined as follows:

$$\beta = \left\langle \beta_i = \sum_j A^E_{ij} - A^I_{ij} \right\rangle, \tag{5.2}$$

where $A^E_{ij}, A^I_{ij} \geq 0$. and β_i is the local balance of neuron i. In our studies, we vary the E-I balance by varying the strength of inhibitory links, keeping other parameters fixed.

The firing variable, r, is determined by the threshold, θ_i, and the activation function of each neuron (as illustrated in the inset of Fig. 5.1A). In our setup, the activation function is a sigmoid function, given by $\mathrm{Sig}(x) = 1/(1+e^{-cx})$, with c set to 10. This model approximates a spiking network as $c \to \infty$.

The input is fed into the reservoir via the input layer, which links the input signal to a fraction, f_{in}, of the reservoir neurons. The non-zero values in the input matrix are drawn from a uniformly distributed random variable within the interval $[-\sigma/2, \sigma/2]$, where σ is the input strength.

The output layer functions as a readout, with weights connecting to each neuron in the reservoir. These weights are trained using ridge regression method with a regularization parameter, η. Unless stated otherwise, the default values for the parameters used in this study are provided in Table 5.1.

Description	Symbol	Value
Number of neurons	N	500
Average network degree	k	50
Leakage rate	λ_i	0
Excitatory fraction	f_E	0.8
Mean exc. synaptic strength	μ_E	$\frac{1}{k \cdot f_E} = \frac{1}{40}$
Scaling factor	α	1
Synaptic spread (Exc. & Inh.)	$\sigma_I = \sigma_E$	$0.2\mu_E$
Balance parameter	β	0*
Threshold	θ	0*
Sigmoid steepness	c	10
Input fraction	f_{in}	0.3
Input strength	σ	Table. 5.2
Regularization parameter	η	10^{-7}

Table 5.1: Default reservoir hyperparameters for Excitatory-Inhibitory RC framework

5.3.2 Constructing the adjacency matrix, A_{ij}

In the previous section, we describe the construction of a recurrent circuit (RC) with a tunable excitatory-inhibitory (E-I) balance. The adjacency matrix, A, is divided into two submatrices: A_E and A_I, corresponding to excitatory and inhibitory links, respectively. Neurons in the network are classified as either excitatory or inhibitory, with all outgoing links of a neuron being of the same type, in accordance with Dale's Law. The non-zero entries in the excitatory submatrix are drawn from a normal distribution with a mean μ_E given by $1/(d \cdot f_E)$, where d is the average degree and f_E is the fraction of excitatory neurons. This is done so that the average incoming excitatory current to a neuron is 1. The standard deviation of this distribution is $\sigma_E = 0.2\mu_E$, or 20%. The non-zero entries in the inhibitory submatrix are drawn from a different normal distribution with a mean μ_I, which is varied to control the relative E-I balance in the system. The mean μ_I is defined by the

following expression:

$$\mu_I = \frac{f_E \cdot \mu_E - \beta/k}{1 - f_E} \tag{5.3}$$

Here, β is the global balance parameter, with $\beta = 0$ indicating global balance. The standard deviation of the inhibitory link distribution, σ_I, is the same as that of the excitatory link distribution, σ_E.

5.3.3 Neuron differential entropy

In order to characterize the internal dynamics of the reservoir, we study the entropy of the neuronal states, which has been previously suggested to correlate with the reservoir's performance [25, 26, 62, 124]. To avoid the problem of binning our continuous variable, r, we use a non-parametric estimator of differential entropy [17] based on the nearest neighbors of a particular value, $r_i(t)$. This estimator, first introduced by Kozachenko and Leonenko [75], is given by the following equation:

$$H(r_i) = \psi(T) - \psi(1) + \log(2) + \frac{1}{T}\sum_{t=1}^{T} \log \varepsilon_i(t) \tag{5.4}$$

Here, $H(r_i)$ is the entropy estimate of the state of the ith neuron. ψ is the digamma function, T is the total number of time steps used to estimate the entropy, which in our case was 10^4, and $\varepsilon_i(t)$ is the nearest neighbor distance of $r_i(t)$. Unlike Shannon entropy, commonly used for discrete values, differential entropy can be negative.

5.3.4 Tasks for testing

We assess the performance of our reservoir computer by conducting a set of well-known tasks. These tasks are designed to test various computationally relevant aspects. The tasks

we focus on include:

• **Memory Capacity Test:** A widely used memory benchmark task, introduced by Jaeger (2002), involves feeding the reservoir a random signal, $u(t)$ with values uniformly sampled from the interval [0,1]. The goal of this task is to train the reservoir to remember and recall a value it encountered d steps earlier. For each delay, d, the accuracy of the recall is measured using a coefficient of determination, r^2, given by

$$\frac{\text{cov}(u(t-d), y_d(t))^2}{\text{var}(u(t)) \cdot \text{var}(y_d(t))} \tag{5.5}$$

Here $u(t)$ represents the input and y_d denotes the reservoir output for a given delay, d. A memory capacity curve (Fig. 5.1B, i), shows the plot of r^2 as a function of delay. The overall memory capacity is determined by calculating the total area under this curve, approximated by summing the values of r^2. As a practical constraint, this summation is capped at $d = 70$

• **NARMA-10:** Another commonly used benchmark for evaluating a system's computational performance is the Nonlinear Autoregressive Moving Average (NARMA) test [182], which is a discrete temporal task with a 10th-order lag. The formulation for the NARMA-10 time series is:

$$\begin{aligned} y(t) = py(t-1) + qy(t-1)\sum_{i=1}^{10} y(t-i) \\ + g \cdot u(t-10)u(t-1) + \Delta \end{aligned} \tag{5.6}$$

where the parameters are defined as $p = 0.3$, $q = 0.05$, $g = 1.5$, $\Delta = 0.1$. The input $u(t)$ is sampled from a uniform distribution over the range [0, 0.5]. This test is designed to chal-

lenge the system's ability to handle nonlinearities as well as its capacity to retain memory across large time lags. Performance was evaluated by calculating the root mean squared error (RMSE) between the value given by Equation 5.6 and the trained reservoir output.

• **Mackey-Glass System:** Chaotic time-series prediction is a field where reservoir computers (RCs) are particularly effective. As dynamical systems themselves, RCs can model the underlying behavior of chaotic systems after training, allowing them to predict system behavior for extended periods before deviations occur due to the accumulation of small errors, a characteristic of chaotic dynamics. In this context, the Mackey-Glass system is a widely used chaotic model [100] frequently employed to evaluate the performance of RCs. It is described by the following one-dimensional delay differential equation:

$$\frac{dx}{dt} = \xi \frac{x(t-\tau)}{1+x^n(t-\tau)} - \gamma x \tag{5.7}$$

The specific parameter values used to generate chaotic behavior are $\xi = 0.2$, $\gamma = 0.1$, $\tau = 17$, and $n = 10$. Performance is quantified as the valid prediction time (VPT). This is defined as the first time point at which the normalized root mean squared error (NRMSE) between the reservoir output and truth values (using Eqn. 5.7) exceeds 0.4.

• **Lorenz System:** Another popular paradigm in RCs for chaotic time series prediction is the Lorenz system. The Lorenz system is a set of three ordinary differential equations representing a simplified model for atmospheric convection [90] and is known for its chaotic

behavior under specific parameter values. The equations defining the Lorenz system are:

$$\begin{aligned} \dot{x} &= a(y-x) \\ \dot{y} &= -xz + bx - y \\ \dot{z} &= xy - cz \end{aligned} \tag{5.8}$$

Where x, y, and z represent the internal state variables, each evolving according to the given set of differential equations. The parameter values $a = 10$, $b = 28$, and $c = 8/3$ are known to generate chaotic dynamics. For our study, we focus on one of the system variables, x, and keep the other 2 hidden from the reservoir. This approach simulates scenarios where not all system variables are available and is particularly relevant for real-world applications where data may be incomplete or partially observed. Additionally, since the range of the Mackey-Glass system and the x values in the Lorenz system can vary widely, they are normalized to values between 0 and 1 to facilitate comparisons.

5.3.5 Inhibitory adaptation mechanism

A novel technique presented in this chapter is the inhibitory adaptation mechanism, which plays a vital role in guiding neurons toward a locally balanced state (Fig. 5.2). The mechanism operates on a simple principle: neurons with activity levels exceeding their target firing rate receive a slight increase in the strength of their inhibitory synapses, while neurons with activity below their target experience a slight decrease in their incoming inhibitory connections (see Fig. 5.2). The following equation defines the adaptation of inhibitory connections:

$$A^I_{ij}(t+1) = A^I_{ij}(t) + \delta(r_i(t) - \rho_i) \tag{5.9}$$

In this context, $A^I_{ij} \geq 0$ denotes strength of the inhibitory link from j to neuron i, δ is the learning rate which in our case is set to 10^{-3}, r_i represents the state variable of neuron i, and ρ_i is the target firing rate of neuron i. Neurons can have a common (homogeneous) target firing rate, or they can have varying targets. A target rate of 0.5 corresponds locally balanced excitation and inhibition, while values above (below) 0.5 correspond to over-excitation (-inhibition). This parameter is crucial to our adaptation process.

To eliminate the necessity for fine-tuning the target firing rate (Fig. 5.10a), we explored brain-inspired heterogeneous firing rate targets (using the beta distribution, See Section 5.3.6 for details).

5.3.6 Heterogeneous distribution of firing rate targets

In section 5.3.5 and the results, we refer to the heterogeneous distribution of firing rate targets. This section provides detailed information on the subject. Figure 5.8a illustrates that neurons with extreme firing rates (~ 0 or ~ 1) exhibit low entropy and contribute less to decision-making processes. These neurons are effectively "wasted." In designing a heterogeneous distribution, we aimed to satisfy two criteria: i) limit the range of neuronal firing rates to avoid extreme values, and ii) maintain a reasonable fraction of neurons near 0.5 (higher memory) while preserving a spread of firing rates to exploit their nonlinearity.

The β distribution proved suitable for this purpose. Since our firing rate values are confined between 0 and 1, the β distribution is a natural fit. For the parameters of the distribution ($\beta(a,b)$), we set $a = b$ to ensure a symmetric distribution with a mean of 0.5. We chose the parameter value of 9 to achieve a range that avoids excessively high or low values (see the top histogram in Figure 5.8c). Consequently, for all analyses involving a heterogeneous distribution of firing rate targets, we utilize the $\beta(9,9)$ distribution. However, other

distributions that meet the mentioned criteria should yield comparable results.

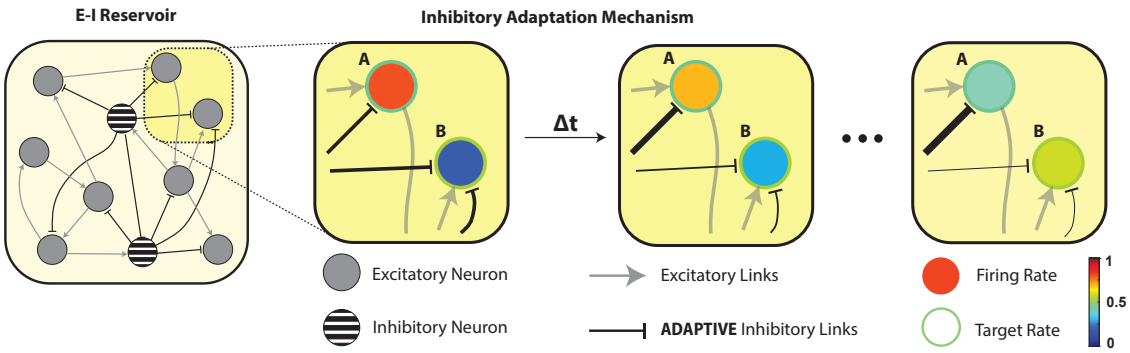

Figure 5.2: Schematic demonstrating the inhibitory adaptation mechanism. We examine two specific neurons (A and B) within the E-I reservoir. The target firing rate of each neuron is represented by its border color, while the solid fill color depicts its current value. Neuron A exhibits higher activity than its target, prompting the strengthening of inhibitory links to bring it closer to the desired level. Conversely, Neuron B displays lower activity than its target, resulting in the weakening of inhibitory links. Following the adaptation process, we observe that both the highlighted neurons converge toward the target firing rate

5.3.7 Designing reservoirs with local E-I balance

When we have complete information about the network's connectivity and the strengths of its links, we can start with a randomly configured network and compute the adjustments needed to bring each neuron to a state with desired relative local balance. For each neuron, we can determine a multiplicative factor, Ω_i, that should be applied to the inhibitory synapses targeting that neuron. This factor is calculated as follows:

$$\Omega_i = \frac{\text{Sig}^{-1}(\rho_i) + \theta_i - W_i^{in}\langle x \rangle - \sum A_{ij}^{E}\rho_j}{\sum A_{ij}^{I}\rho_j} \tag{5.10}$$

Sig^{-1} function is an inverse sigmoid function, $\langle x \rangle$ is the mean input and ρ_j is the target firing rate of neuron j.

5.4 Results

5.4.1 Robust performance in balanced and slightly over-inhibited regimes; brittleness with excess excitation

The impact of inhibition, and consequently the relative balance between excitation and inhibition (E-I balance), is a key area of interest in both neuroscience and information processing. While the significance of E-I balance has been extensively examined in the brain, its application to machine learning systems, particularly reservoir computers, is a less explored area. Recent studies suggest that E-I balance plays a crucial role in performance, with some emphasizing its importance for optimal results while others highlight how inhibition contributes to a stable and effective operating regime [35]. Here, we find that reservoir computer (RC) performance is relatively strong in E-I balanced networks and often improves in slightly over-inhibited regimes. By contrast, we find that performance degrades rapidly even with slight over-excitation. We investigate how this effect depends on both neuronal (node) and synaptic (link) properties In Figure 5.3, we demonstrate the effect of varying the fraction of inhibitory nodes in a bio-inspired configuration where the relative strength of inhibitory synapses is four times that of the excitatory synapses. In a purely excitatory network (0% inhibition), the transition (Fig. 5.3a, *dashed lines*) from quiescence (no neurons firing) to saturation (all neurons firing) is abrupt (Fig. 5.3a, *top*). The average neuronal entropy (Fig. 5.3a, *middle*) is highest close to the transition and is a reasonable predictor of performance (Fig. 5.3a, *bottom*). Notably, performance is low and highly sensitive to parameter choices near the transitions. However, when we increase the proportion of inhibitory neurons to 20%, we achieve a balanced state since the ratio of E to I neurons (4:1) balances the strength disparity (1:4). This results in a significant shift in the phase diagram, where the transition between saturation and quiescence becomes much smoother,

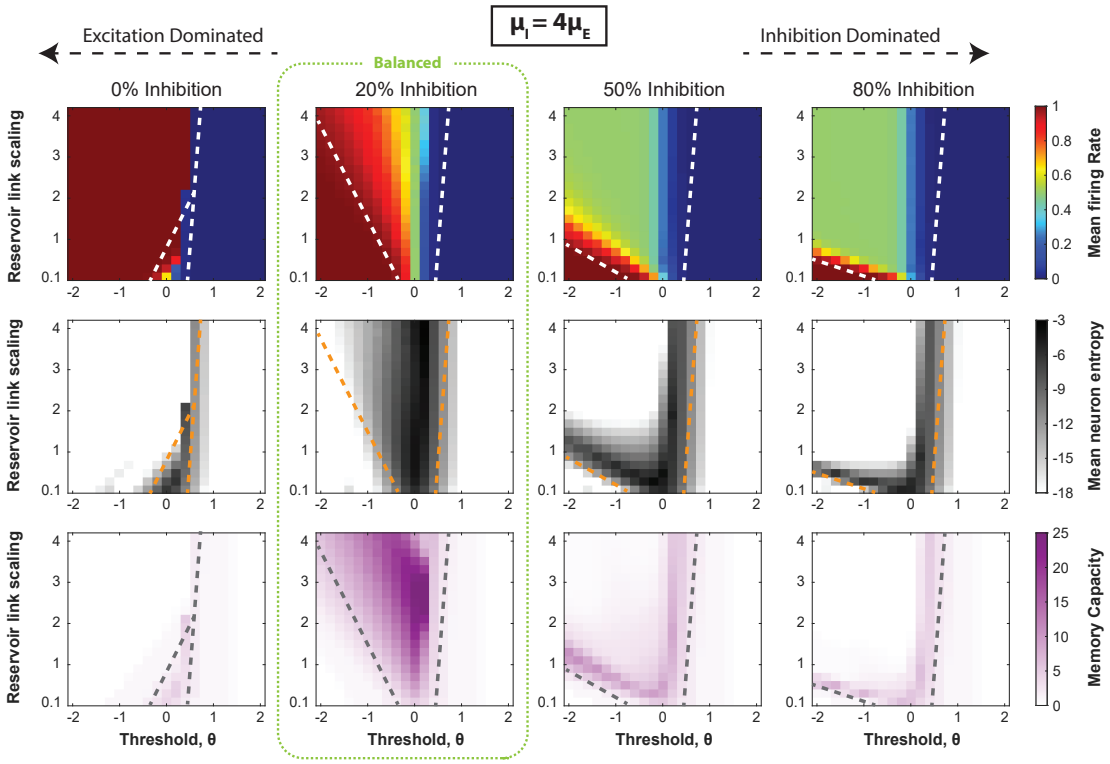

Figure 5.3: The effects of varying the fraction of excitatory neurons. Here, we maintain a fixed, brain-inspired four-fold ratio between the average strength of inhibitory neurons and excitatory neurons ($\mu_I = 4\mu_E$). Each column has a different fraction of inhibitory neurons, which increases from left to right. The top row displays the mean firing rate of the network neurons. The middle shows the mean differential entropy of neurons in the network, and the bottom row illustrates the total memory capacity of the network. Each plot is presented as a heatmap, with the neuronal threshold on the horizontal axis and the link scaling factor for synapses in the reservoir on the vertical axis. Dashed lines highlight the transition away from a quiescent state (firing rate ≈ 0) and the transition to a saturated state (firing rate ≈ 1). We observe low and brittle performance in purely excitatory networks *(left column)*, while balanced networks *(second column, highlighted in green)* exhibit high and robust memory capacity. Although over-inhibition leads to a decline in performance, it remains elevated and more robust compared to the over-excited counterpart

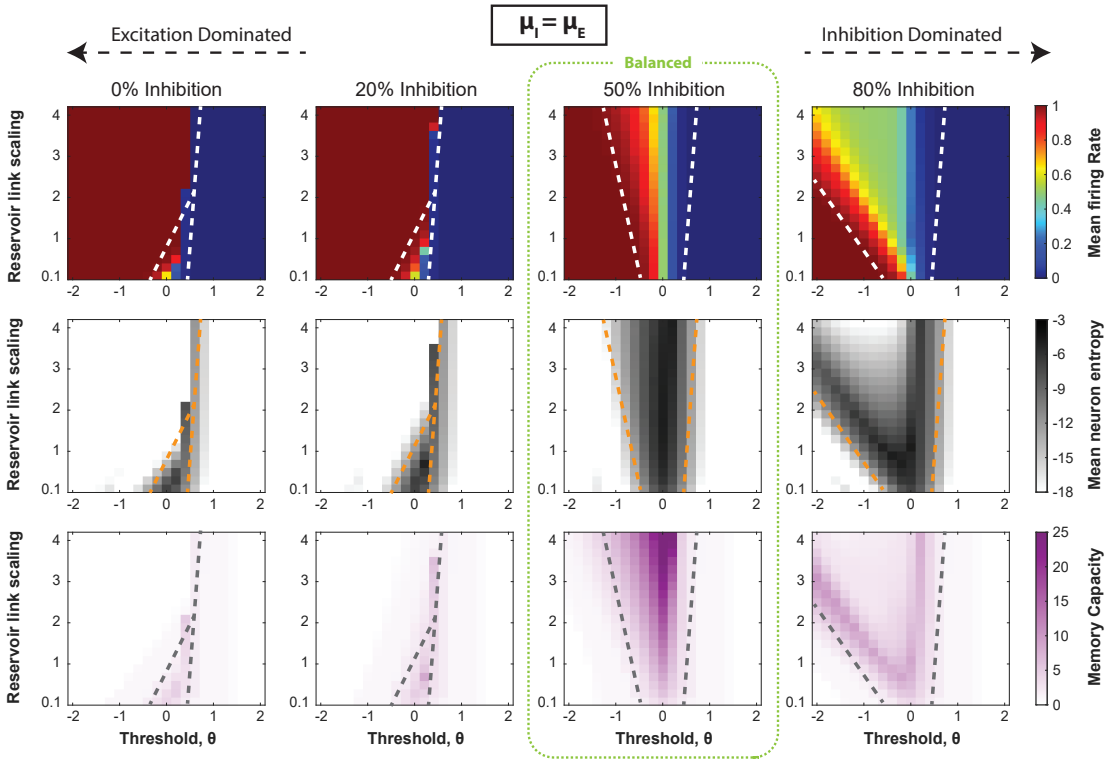

Figure 5.4: The effects of varying the relative fraction of excitatory and inhibitory neurons with $\mu_I = \mu_E$. Each column has a different fraction of inhibitory neurons, which increases from left to right. The top row displays the mean firing rate of the network neurons. The middle shows the mean differential entropy of neurons in the network, and the bottom row illustrates the total memory capacity of the network. Each plot is presented as a heatmap, with the horizontal axis representing the neuronal threshold and the y-axis denoting the link scaling factor for synapses in the reservoir. Dashed lines denote the transition away from a quiescent state (firing rate = 0) and the transition to a saturated state (firing rate = 1).

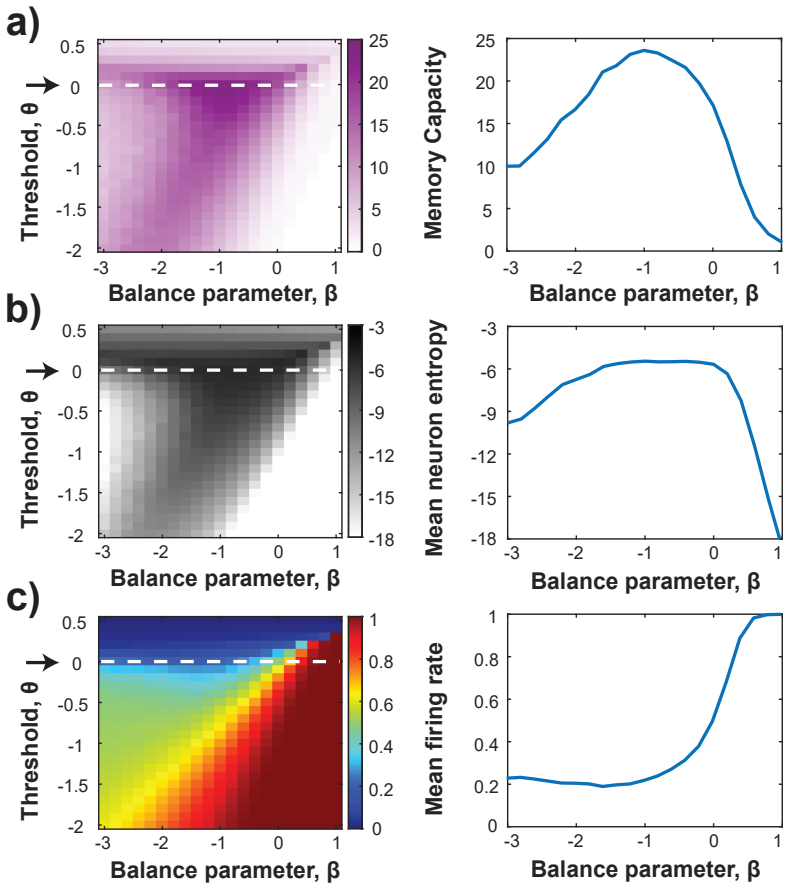

Figure 5.5: Adjusting the relative inhibitory synaptic strength to modify the E-I balance while maintaining the proportion of inhibitory neurons at 20%. The net E-I balance is now quantified by a balance parameter, β (See Sec. 5.3.2) The heatmaps depict this balance parameter on the horizontal axis and neuronal thresholds on the vertical axis. The plots on the right display a slice (indicated by an arrow) of the heatmap.

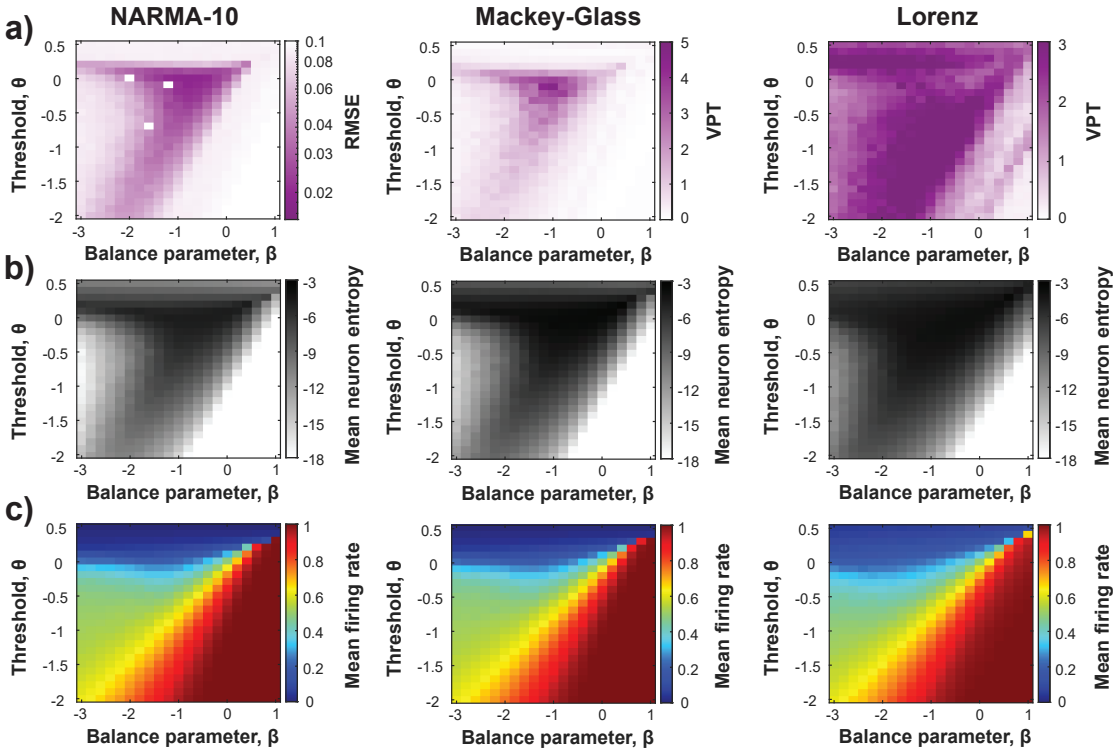

Figure 5.6: Adjusting the relative inhibitory synaptic strength to modify the E-I balance while maintaining the proportion of inhibitory neurons at 20% for 3 different tasks: NARMA-10, Mackey-Glass and Lorenz time series prediction. The net E-I balance is now quantified by a balance parameter, β (See Sec. 5.3.1 and 5.3.2 for details). The heatmaps depict this balance parameter on the x-axis and neuronal thresholds on the y-axis. Note the robust and elevated performance in the balanced and slightly over-inhibitory regimes, and low, brittle performance in the over-excitatory regime.

leading to a broader range of intermediate firing rates (Fig. 5.3b, *top*). This broadening is also reflected in both the neuronal entropy (Fig. 5.3b, *middle*) and performance (Fig. 5.3b, *bottom*), which shows a larger region of higher memory. As the ratio of inhibitory neurons rises to 50% (Fig. 5.3c), the network operates in a slightly over-inhibited regime, but significant regions of good performance and elevated entropy persist. When inhibition is further increased to 80% (Fig. 5.3d), performance declines and becomes very sensitive to hyperparameters once more. We observe a similar pattern in systems with a different value of mean inhibitory strength ($\mu_I = \mu_E$, see Fig. 5.4). However, one must note that neuronal

entropy alone is not a perfect predictor for performance. We can have low-performing internal representations that are mimic random noise (high entropy) while being decoupled from the input. However, when the input scaling is tuned for a particular task (see section 5.4.5), we find that the neuronal entropy does track the performance reasonably well.

In Figure 5.5, we demonstrate the effects of varying the balance of excitation and inhibition in an Excitatory-Inhibitory (E-I) reservoir while keeping the ratio of inhibitory cells fixed at 20%. This balance is represented by the variable β, where $\beta = 0$ indicates balance, negative values indicate an inhibition-dominated regime, and positive values indicate an excitation-dominated regime.

The heatmaps in Figure 5.5 show that when $\beta > 0$ (over-excited networks), the transition from quiescence to saturation is abrupt (Fig. 5.5c). However, as inhibition increases, the transition is more gradual. Each heatmap has β on the horizontal axis and the neuronal firing threshold θ on the vertical axis. The data indicate strong performance (Fig. 5.5a) over a wide range θ when the network is balanced to slightly over-inhibited (-2 $\lessapprox \beta \leq 0$). We see that the average neuronal entropy (Fig. 5.5b) is also elevated in the regions where good performance is observed. (Similar heatmaps for the other tests can be found in Fig. 5.6) A slice through the heatmap at threshold $\theta = 0$ (marked by a solid arrow) shows that memory capacity is increased and neuronal entropy is elevated when the network is balanced to slightly over-inhibited but drops off quickly in the over-excited regime. This trend is consistent with the observation that slightly over-excited networks quickly reach saturation, while over-inhibited networks gradually become less active as inhibition is increased.

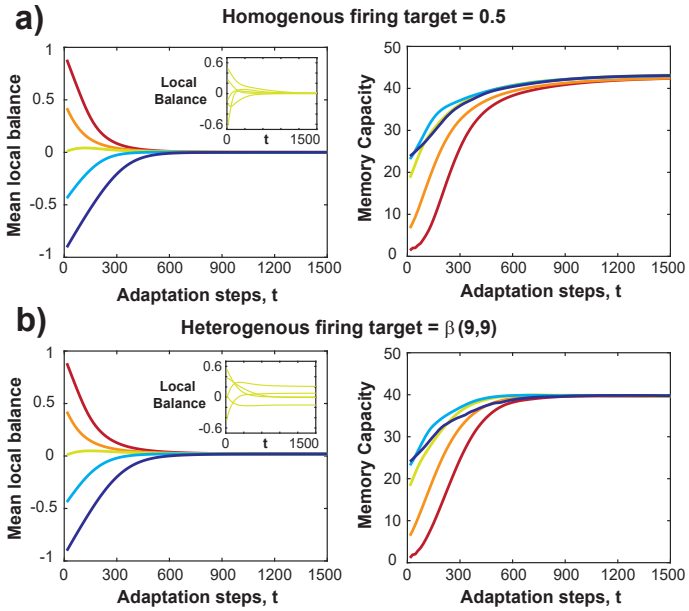

Figure 5.7: Adaptation curves for inhibitory self-adaptive RCs with (a) homogeneous target firing rates and (b) heterogeneous (beta-distributed) target firing rates. Each plot includes (*left*) the mean local balance as a function of adaptation steps for five different initializations of global relative E-I balance. (*inset*) Adaptation curves for five representative neurons for an RC initialized with global balance (*lime*). The plot on the (*right*) displays the memory capacity as a function of adaptation steps for the same five initializations.

5.4.2 Self-adapting E-I Reservoirs have locally balanced and higher entropy neurons

The primary method introduced in this chapter is the inhibitory self-adaptation mechanism. This approach adjusts the inhibitory connections to each neuron based on a predefined firing rate target (see Figure 5.2 and Section 5.3.5 for details). In panel a) of Figure 5.7, we observe that irrespective of the initial relative E-I balance of the reservoir, with a homogeneous firing target of 0.5, all the different starting conditions converge to a mean local balance of 0. Even the initially globally balanced reservoir (*lime curve*) contains locally neurons with small local imbalances, which can be seen in the individual neuron traces (*inset*, n=5 neurons). However, when we apply our inhibitory adaptation mechanism

with fixed targets for all neurons (Figure 5.7b), we achieve precise control over the firing rate and maintain a relative local balance at the neuronal level. By using a heterogeneous distribution of firing rates targets — which is inspired by observations of biological neural networks — we achieve a useful range of firing rates (*inset*, n=5 neurons) while keeping the network relatively balanced on the local and global scales. This heterogeneity enriches the RC's dynamics and reduces the need for meticulous fine-tuning of the firing rate target. Figure 5.8 presents a comprehensive view of the inhibitory adaptive mechanism's effect in both homogeneous and heterogeneous firing target scenarios. In the non-adapting RC (Fig. 5.8a, *top*), the scatter plot of firing rates and the local balance parameter, β_i, reflects the sigmoid activation function of each neuron. The histogram indicates a wide range of firing rates and local balance.

In Fig. 5.8a, *middle*, neurons with intermediate firing rates exhibit higher differential entropy compared to those firing very little (~ 0) or excessively (~ 1). The bottom panel shows that neurons with extreme firing rates and low entropy have low absolute output weights, indicating a systematically lower contribution to the reservoir's performance. Conversely, in Fig. 5.8b, the adaptation mechanism successfully guides neurons to a specific firing rate target and a fixed local balance. Here, neuron entropy is elevated offering a more useful population of neurons from which to construct the output weights.

Finally, Fig. 5.8c illustrates the effect of imposing a heterogeneous (beta-distributed, see Sect. 5.3.6) firing rate target. With a range of firing rates that suppresses quiescence and saturation, the local balance remains relatively narrow. The distribution of neuronal differential entropy is similar to the homogeneous case (Fig. 5.8b) and higher than the non-adapted case (Fig. 5.8a which contained a significant fraction of low entropy neurons. Additionally, no systematic trend of neurons with low output weights is observed. This heterogeneity leverages variation to eliminate the need for fine-tuning the firing rate target. See Figure 5.9 for analysis of the Lorenz time-series prediction task highlighting input and

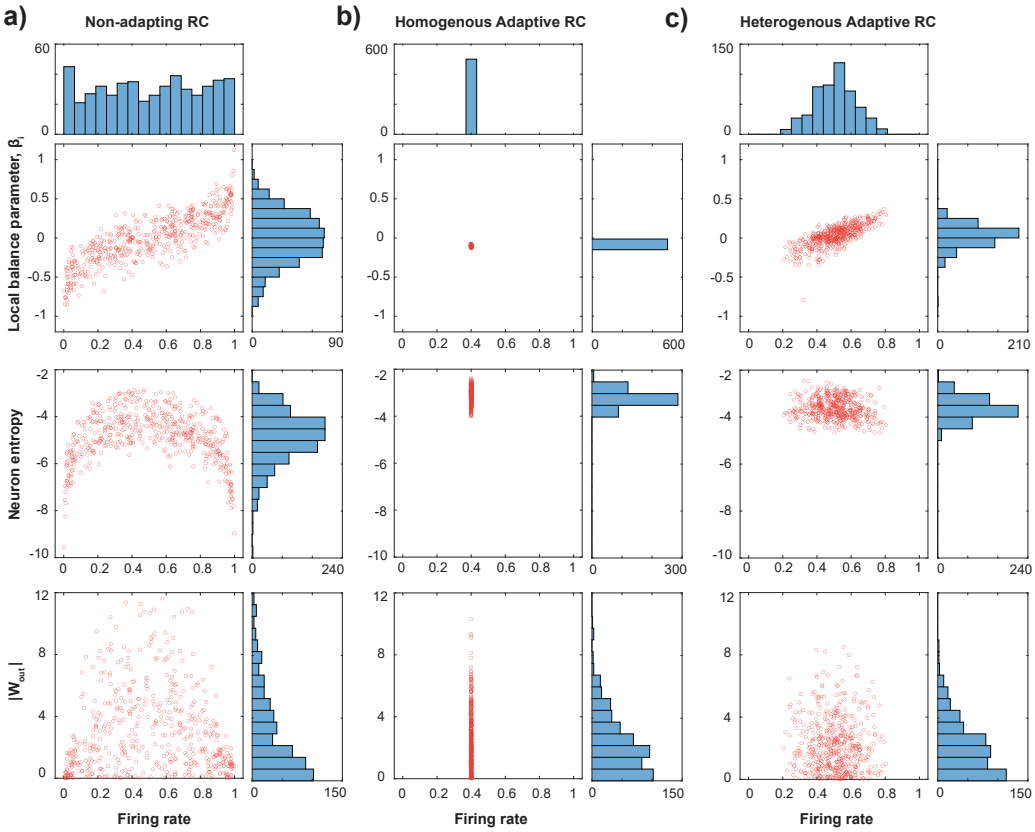

Figure 5.8: The effects of the inhibitory adaptation mechanism on neuronal properties: (a) RC without local adaptation, (b) Inhibitory adaptive mechanism with a homogeneous firing rate target of 0.4 (slightly over-inhibited), (c) Inhibitory adaptive mechanism with a heterogeneous (beta-distributed) firing rate target. Each panel includes scatter plots and histograms of the (*top*) local balance parameter, (*middle*) neuronal differential entropy, and (*bottom*) absolute value of the output weight in a 5-step memory task. Each scatter plot is plotted against the individual firing rate of the neurons. Even though the non-adapting RC is globally balanced, it exhibits a wide range of local balance parameters and, consequently, firing rates, with neurons close to quiescence (firing rate 0) or saturation (1) having low entropy and low output weights. The adaptive mechanism (homogeneous and heterogeneous) effectively brings the whole population of neurons to more useful local balances and entropies.

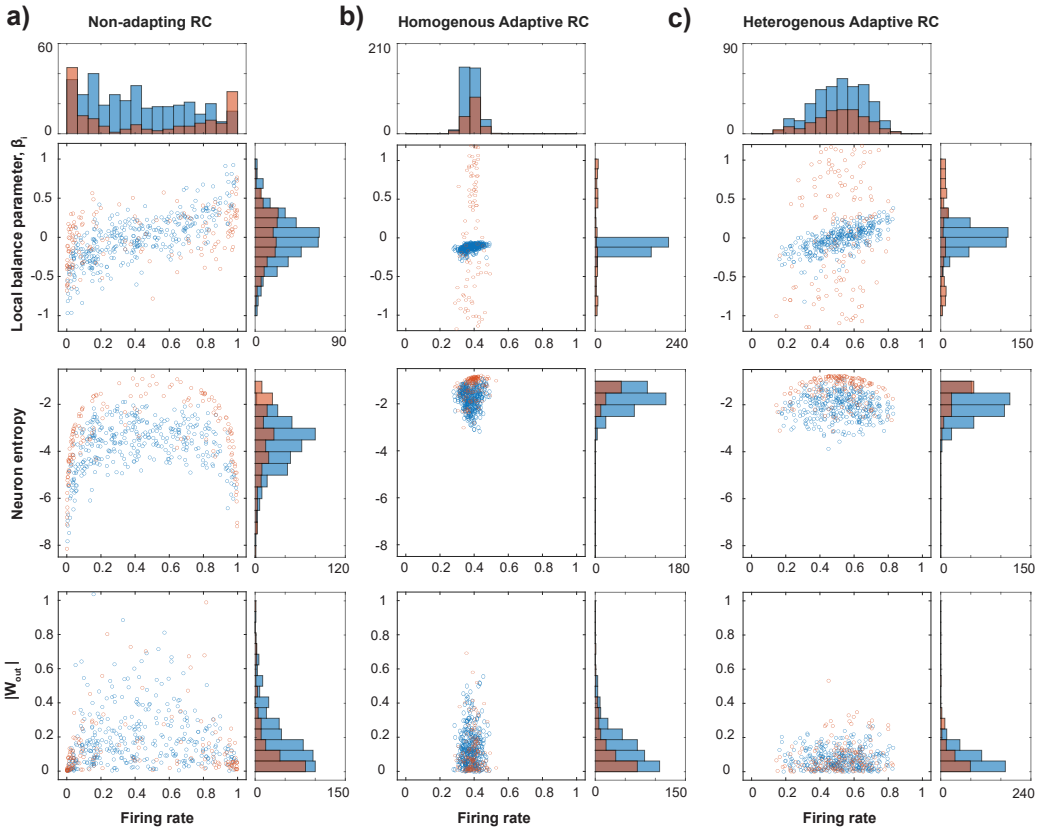

Figure 5.9: Consolidated view of the effect of the inhibitory adaptation mechanism for the Lorenz time-series prediction tasks highlighting input and non-input neurons for: (a) Non-adapting RC, (b) Inhibitory adaptive mechanism with a homogeneous firing rate target of 0.4, (c) Inhibitory adaptive mechanism with a heterogeneous (beta-distributed) firing rate target.

non-input neurons.

5.4.3 Homogeneous and heterogeneous inhibitory adaptation significantly boosts RC performance

Figures 5.7 and 5.8 demonstrate that our inhibitory adaptation mechanism consistently meets the desired firing rate target, while Figure 5.10 reveals the significant improvements in RC performance that result from this approach. In Figure 5.10a, we plot performance as

a function of a uniform firing rate target across four different tasks: Memory, NARMA-10, Mackey-Glass, and Lorenz time series prediction. Notably, the Memory task performs best with a target firing rate of 0.5, corresponding to the linear region of the sigmoid curve, which is known to support improved memory capacity. In contrast, the other tasks achieve optimal results at firing rate targets slightly removed from 0.5, where they also benefit from the nonlinearity. None of the tasks performs well in the saturated (≈ 1) or quiescent (≈ 0) states.

In Figure 5.10b, we compare the performance of the best non-adaptive RC (*blue*) with the top-performing homogeneous adaptive RC (*solid green*) and a heterogeneous adaptive RC (*striped green*). Clearly, the adaptive approaches offer significant performance gains over the non-adaptive method. Although the heterogeneous method generally exhibits slightly reduced performance across most tasks (except for NARMA-10, as shown in Fig. 5.10b.ii), it enhances the robustness of the RC and eliminates the need for precise tuning of the target (Fig. 5.10a).

5.4.4 Designing RCs for local balance yields similar results to adaptive mechanisms

The adaptation rule in RCs was designed with two main objectives. First, it aims to guide the RC toward a state of local balance, achieving the desired distribution of firing rates. Second, it introduces a local feedback mechanism that leverages the reservoir's current state to adjust synapses for future behavior. Biological systems inspire this second objective and assume only local information availability. However, with complete knowledge of the system and its synaptic weights, we can design a reservoir with varying levels of local E-I balance by simultaneously adjusting all inhibitory synapses to attain the desired target firing-rate distribution (see Section 5.3.7 for details).

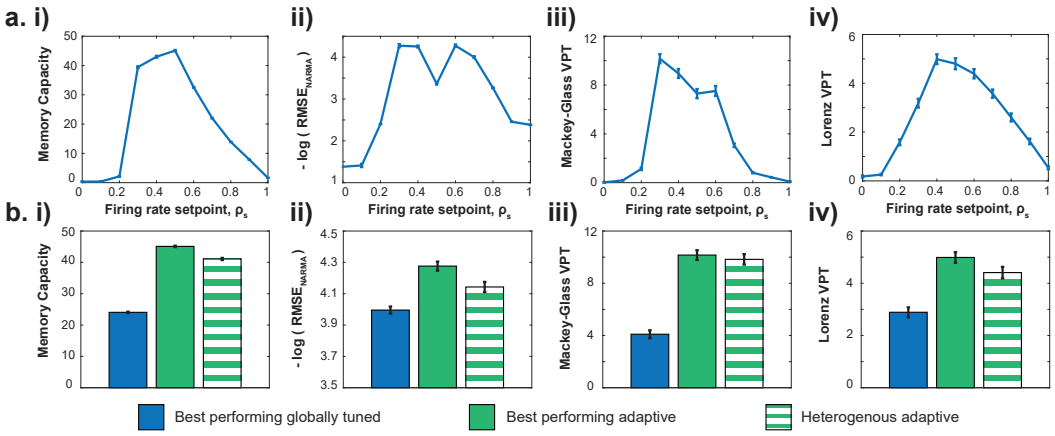

Figure 5.10: a) Tuning curves for inhibitory self-adapting RCs with homogeneous firing rate targets across four different tests. b) Histograms showing the performance of the best performing globally tuned RC (*blue*), the best inhibitory adaptive RC from panel (a) (*green*), and an adaptive RC with heterogeneous firing rate targets(*green striped*).

Similar to the local *adaptation mechanisms*, the reservoir can be *designed* in two ways: i) using homogeneous firing rate targets and ii) using heterogeneous (beta-distributed) targets. In Figure 5.11, the performance of the designed reservoirs (*orange bars*) matches that of the inhibitory adaptation rule (*green bars*). This performance level significantly exceeds that of the best-performing globally tuned RC (*blue bar*). Thus, this design rule can be a valuable approach to designing better, more locally balanced RCs, offering a clear improvement over a globally tuned random initialization.

5.4.5 Memory-Nonlinearity trade-off in RCs

A crucial factor in the design of machine learning systems, especially reservoir computers (RCs), is balancing the trade-off between memory and nonlinearity. A common parameter to adjust in this context is the input scaling. When the input scaling parameter, σ, is low, the transfer curve of the neurons is mostly linear; as σ increases, the neuronal responses become more nonlinear. In the linear regime (small σ), RCs tend to exhibit

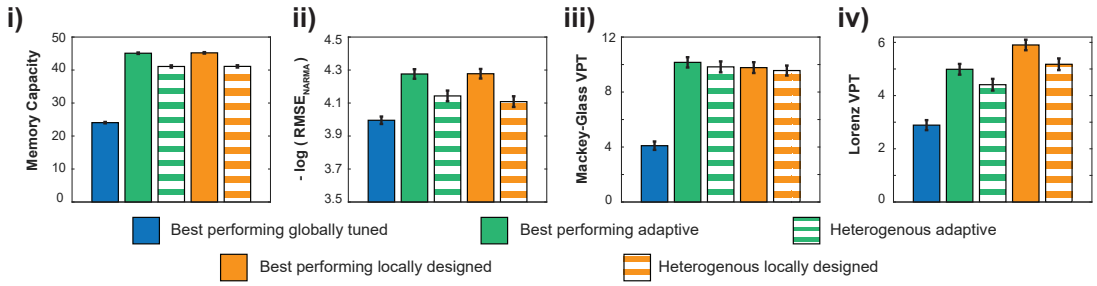

Figure 5.11: Histogram comparing the best performing globally tuned RC (*blue*) with two types of inhibitory adaptive RCs: homogeneous (*solid green*) and heterogeneous (*striped green*). It also compares these with two types of locally designed RCs: homogeneous (*solid orange*) and heterogeneous (*striped orange*).

strong memory capacity. However, many complex real-world tasks require both memory and nonlinear mapping, necessitating careful tuning of the input scaling.

Task	Memory Capacity	NARMA 10	Mackey Glass	Lorenz
Non-adapting RC	0.04	0.16	0.63	3.98
Inhibitory adaptive RC	0.02	0.16	0.63	2.51

Table 5.2: Optimal input scaling for various tasks in both non-adapting and inhibitory adaptive RCs

In Figure 5.12, we arrange a series of tasks from those requiring the most memory (linear) to those requiring the most nonlinearity, moving from left to right (Fig. 5.12a-d) based on their optimal input scaling. The top row in Figure 5.12 shows a heatmap of the performance of a non-adapting RC as a function of input scaling and the balance parameter across the 4 tasks. Notably, we observe that optimal input scaling for a given task tends to be consistent across different values of the balance parameter, suggesting that the ideal

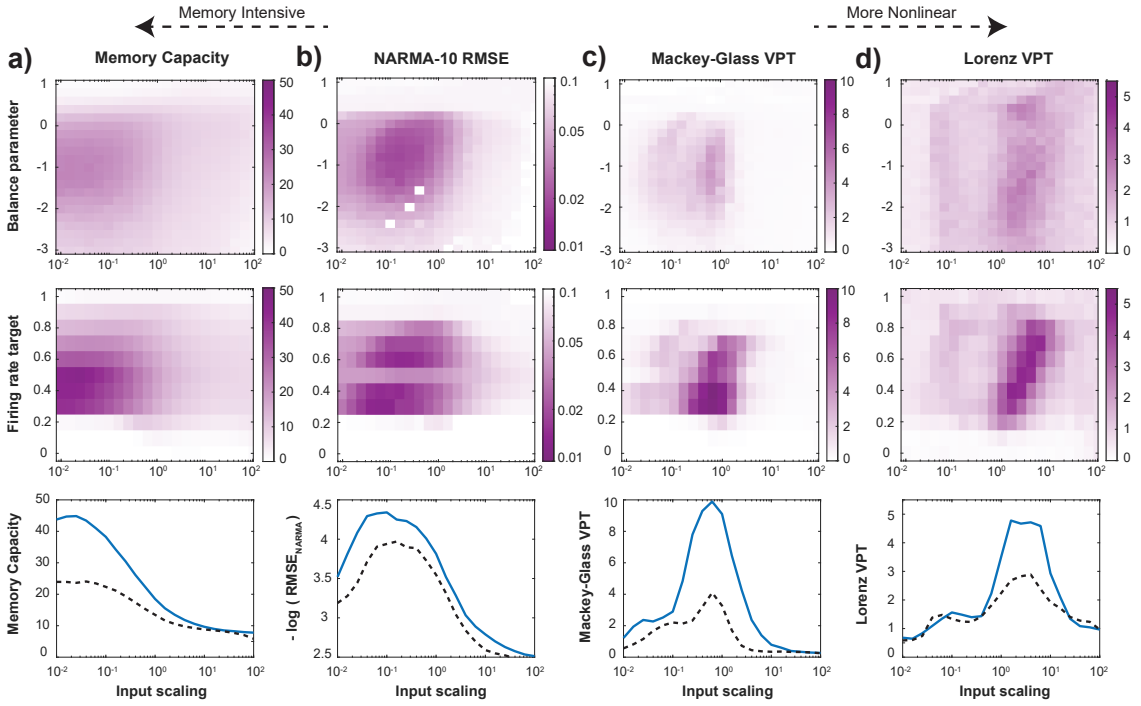

Figure 5.12: (*top*) Heatmaps for four tasks using a non-adaptive reservoir, with the input scaling on the horizontal axis and balance parameter β on the vertical axis. (*middle*) Heatmaps for four tasks using an adaptive reservoir, with the same horizontal axis and the homogeneous firing rate target on the vertical axis. (*bottom*) The dashed line shows the best-performing non-adaptive RC as a function of input scaling, while the solid blue line represents the best-performing inhibitory self-adapting system with homogeneous targets.

input scaling is intrinsic to the problem being solved and is largely independent of other parameters in the reservoir. In the middle row, we plot similar heatmaps with an adaptive RC. The vertical axis of the heatmaps is now the homogeneous target firing rate. We observe elevated performance across a wide range of input scaling and firing rate targets. In the bottom row, we indicate the best-performing non-adaptive RC for each level of input scaling with a dashed line and plot the performance of the best performing inhibitory adaptive RC with homogeneous targets in blue. The optimal input scaling is tabulated in Table 5.2. The results show that the adaptation mechanism generally leads to better performance across a wide range of input scaling levels. However, the most significant improvements

are observed when the RC is operating close to the optimal input scaling for the given task.

5.5 Discussion

In this work, we have explored how tuning the relative excitation-inhibition (E-I) balance affects the dynamics and performance of reservoir computers. Our results echo the role of inhibition in the brain, which allows for a robust range of performance. This result is consistent with phenomena studied by Lu et al. [93] in models of neuronal dynamics, which demonstrate that inhibition can create a robust region of critical-like behavior in neuronal models, a regime known to have information processing benefits.

Building on these findings, we tackled the problems associated with random realization of reservoir networks and hyperparameter optimization by introducing a novel inhibitory adaptation mechanism. This mechanism guides the network to a particular firing rate target, similar to homeostatic mechanisms in the brain.

We measured the internal dynamics of the network using two metrics: the average firing rate and differential entropy of each neuron. These metrics are known to correlate with performance, with neurons that have extremal values of mean firing rate showing low entropy and consequently low performance. Our results show an improvement in performance across four different tasks: Memory, NARMA-10, Mackey-Glass prediction, and Lorenz time series prediction. Each of these tasks has different requirements for memory and nonlinearity and hence is optimized at different input scaling. Compared with RCs in which the global balance is optimized, our local inhibitory adaptation mechanism is able to achieve elevated performance across all these tasks.

However, a potential drawback of homogeneous adaptive mechanism is the need to tune a firing rate target. We addressed this issue by leveraging heterogeneity. By using a beta-distribution for firing rate targets (suppressing quiescent and saturated firing rates),

we achieved elevated performance without the need for careful tuning of firing rate targets. Additionally, in cases where we have explicit knowledge of all the synapses in the network, we introduced a simple design principle to achieve local balance and a particular firing rate distribution by applying an inhibitory link rescaling that corrects the local imbalances of each neuron. These designed reservoirs were able to replicate the performance gains seen in our adaptive systems.

This work adds to the growing body of research drawing parallels between machine learning and neuroscience. The inclusion of more brain-inspired features like neuronal criticality, spike-timing-dependent plasticity, and refractory periods to name a few, will only enhance the machine learning landscape and also deepen our understanding of the role these mechanisms play in information processing.

5.6 Acknowledgements

I would like to thank my advisors Dr. Michelle Girvan and Dr. Dietmar Plenz for their constant guidance on this project. KS and DP's contributions were supported by the Division of the Intramural Research Program (DIRP) of the National Institute of Mental Health (NIMH), USA, ZIAMH002797, ZIAMH002971, and the BRAIN initiative Grant U19 NS107464-01. MG's contributions were supported by ONR award number N000142212656. The research in this chapter utilized the supercomputing resources of the University of Maryland (UMD College Park, USA, https://hpcc.umd.edu/hpcc/zaratan.html).

Chapter 6: Summary and future directions

This thesis explores the intersection of artificial and biological neural networks, aiming to illuminate the dynamics of both systems.

Chapter 2 addresses the challenge of predicting behaviors in large, complex networks where traditional machine learning methods often fail. We propose a novel architecture that emulates the network's topology, enabling highly accurate and scalable predictions. This method proves effective not only when the network links are known but also when they are unknown and inferred from past data.

Chapters 3 and 4 delve into neuronal avalanches, specifically tackling the challenge of accurately identifying scaling relationships between the size and duration of these avalanches in subsampled networks. Subsampling, a bias introduced by limited measurement capabilities, can skew avalanche analysis. We introduce temporal coarse-graining as a method to recover the true scaling exponent in both experimental and simulated subsampled data, allowing for accurate observation of neuronal avalanches despite data collection limitations.

Chapter 5 integrates neuroscience and machine learning to investigate the role of excitatory-inhibitory balance in the brain. We demonstrate the stabilizing effect of inhibition and propose a novel inhibitory adaptation mechanism that directs the network towards a desired firing rate distribution, resulting in robust and improved performance.

Several exciting research areas remain unexplored in the field of neuronal avalanches. One key direction involves delving deeper into the neuronal contribution to avalanche dynamics. Current analysis often focuses on population-level activity, neglecting the specific

participation (or non-participation) of individual neurons within an avalanche. By investigating avalanches from a single-neuron perspective, we could gain valuable insights into how information is optimally transferred between neurons, ultimately shedding light on brain function and information processing.

Another promising area lies in incorporating the concept of avalanche criticality into the field of reservoir computing. Here, we could analyze the information processing capabilities of subcritical, critical, and supercritical networks. This analysis has the potential to solidify the role of critical brain networks in information processing and computation, bridging the gap between biological and artificial systems.

Finally, self-adapting systems hold immense potential for generating scalable machine learning architectures. By leveraging principles gleaned from the analysis of both artificial and biological neural networks, we could design systems capable of rapid adaptation and learning of diverse patterns and dynamics. This could lead to the development of more efficient and data-cost-effective machine learning systems.

This thesis underscores the value of exploring the parallels between neuroscience and machine learning. I hope the findings in this thesis will inspire further research and collaboration between these two disciplines.

Bibliography

[1] M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1 edition, February 1991.

[2] Juan A. Acebrón, L. L. Bonilla, Conrad J. Pérez Vicente, Félix Ritort, and Renato Spigler. The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1):137–185, April 2005.

[3] Laurence Aitchison, Nicola Corradi, and Peter E. Latham. Zipf's Law Arises Naturally When There Are Underlying, Unobserved Variables. *PLOS Computational Biology*, 12(12):e1005110, December 2016.

[4] Troy Arcomano, Istvan Szunyogh, Jaideep Pathak, Alexander Wikner, Brian R. Hunt, and Edward Ott. A Machine Learning-Based Global Atmospheric Forecast Model. *Geophysical Research Letters*, 47(9), May 2020.

[5] Oshrit Arviv, Mordekhay Medvedovsky, Liron Sheintuch, Abraham Goldstein, and Oren Shriki. Deviations from Critical Dynamics in Interictal Epileptiform Activity. *The Journal of Neuroscience*, 36(48):12276–12292, November 2016.

[6] Bruno B. Averbeck, Peter E. Latham, and Alexandre Pouget. Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5):358–366, May 2006.

[7] Stefan Babinec and Jiri Pospichal. Improving the prediction accuracy of echo state neural networks by Anti-Oja's learning. In *Proceedings of the 17th international conference on Artificial neural networks*, ICANN'07, pages 19–28, September 2007.

[8] Per Bak and Chao Tang. Earthquakes as a self-organized critical phenomenon. *Journal of Geophysical Research: Solid Earth*, 94(B11):15635–15637, November 1989.

[9] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Physical Review Letters*, 59(4):381–384, July 1987.

[10] Amitava Banerjee, Jaideep Pathak, Rajarshi Roy, Juan G. Restrepo, and Edward Ott. Using machine learning to assess short term causal dependence and infer network links. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(12):121104, December 2019.

[11] Michael Banf and Seung Y. Rhee. Enhancing gene regulatory network inference through data integration with markov random fields. *Scientific Reports*, 7(1):41174, February 2017.

[12] Kevin E. Bassler and Maya Paczuski. Cellular Model of Superconducting Vortex Dynamics. In A. T. Skjeltorp and T. Vicsek, editors, *Complexity from Microscopic to Macroscopic Scales: Coherence and Large Deviations*, pages 215–227. Springer Netherlands, Dordrecht, 2002.

[13] John M. Beggs and Dietmar Plenz. Neuronal Avalanches in Neocortical Circuits. *The Journal of Neuroscience*, 23(35):11167–11177, December 2003.

[14] John M. Beggs and Dietmar Plenz. Neuronal Avalanches Are Diverse and Precise Activity Patterns That Are Stable for Many Hours in Cortical Slice Cultures. *The Journal of Neuroscience*, 24(22):5216–5229, June 2004.

[15] Timothy Bellay, Andreas Klaus, Saurav Seshadri, and Dietmar Plenz. Irregular spiking of pyramidal neurons organizes as scale-invariant neuronal avalanches in the awake state. *eLife*, 4:e07224, July 2015.

[16] Timothy Bellay, Woodrow L. Shew, Shan Yu, Jessica J. Falco-Walter, and Dietmar Plenz. Selective Participation of Single Cortical Neurons in Neuronal Avalanches. *Frontiers in Neural Circuits*, 14:620052, January 2021.

[17] Thomas B. Berrett, Richard J. Samworth, and Ming Yuan. Efficient multivariate entropy estimation via k-nearest neighbour distances. *The Annals of Statistics*, 47(1), February 2019.

[18] Nils Bertschinger and Thomas Natschläger. Real-Time Computation at the Edge of Chaos in Recurrent Neural Networks. *Neural Computation*, 16(7):1413–1436, July 2004.

[19] Mattia Bonzanni, Kimberly M. Bockley, and David L. Kaplan. On the effect of neuronal spatial subsampling in small-world networks. *European Journal of Neuroscience*, 53(2):485–498, January 2021.

[20] Zac Bowen, Daniel E. Winkowski, Saurav Seshadri, Dietmar Plenz, and Patrick O. Kanold. Neuronal Avalanches in Input and Associative Layers of Auditory Cortex. *Frontiers in Systems Neuroscience*, 13:45, September 2019.

[21] Valentino Braitenberg and Almut Schüz. *Anatomy of the Cortex*, volume 18 of *Studies of Brain Function*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1991.

[22] Victor Buendía, Serena Di Santo, Pablo Villegas, Raffaella Burioni, and Miguel A. Muñoz. Self-organized bistability and its possible relevance for brain dynamics. *Physical Review Research*, 2(1):013318, March 2020.

[23] Daniel Canaday, Aaron Griffith, and Daniel J. Gauthier. Rapid time series prediction with a hardware-based reservoir computer. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(12):123119, December 2018.

[24] Elliott Capek, Tiago L. Ribeiro, Patrick Kells, Keshav Srinivasan, Stephanie R. Miller, Elias Geist, Mitchell Victor, Ali Vakili, Sinisa Pajevic, Dante R. Chialvo, and Dietmar Plenz. Parabolic avalanche scaling in the synchronization of cortical cell assemblies. *Nature Communications*, 14(1):2555, May 2023.

[25] T. L. Carroll. Do reservoir computers work best at the edge of chaos? *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(12):121109, December 2020.

[26] T. L. Carroll and L. M. Pecora. Network structure effects in reservoir computers. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(8):083130, August 2019.

[27] Tawan T. A. Carvalho, Antonio J. Fontenele, Mauricio Girardi-Schappo, Thaís Feliciano, Leandro A. A. Aguiar, Thais P. L. Silva, Nivaldo A. P. De Vasconcelos, Pedro V. Carelli, and Mauro Copelli. Subsampled Directed-Percolation Models Explain Scaling Relations Experimentally Observed in the Brain. *Frontiers in Neural Circuits*, 14:576727, January 2021.

[28] Jose Casadiego, Mor Nitzan, Sarah Hallerberg, and Marc Timme. Model-free inference of direct network interactions from nonlinear collective dynamics. *Nature Communications*, 8(1):2192, December 2017.

[29] Yan-Jiun Chen, Stefanos Papanikolaou, James P. Sethna, Stefano Zapperi, and Gianfranco Durin. Avalanche spatial structure and multivariable scaling functions: Sizes, heights, widths, and views through windows. *Physical Review E*, 84(6):061103, December 2011.

[30] Dante R. Chialvo. Emergent complex neural dynamics. *Nature Physics*, 6(10):744–750, October 2010.

[31] Luca Cocchi, Leonardo L. Gollo, Andrew Zalesky, and Michael Breakspear. Criticality in the brain: A synthesis of neurobiology, models and cognition. *Progress in Neurobiology*, 158:132–152, November 2017.

[32] ARNI code:. https://github.com/networkinference/ARNI.

[33] Joni Dambre, David Verstraeten, Benjamin Schrauwen, and Serge Massar. Information Processing Capacity of Dynamical Systems. *Scientific Reports*, 2(1):514, July 2012.

[34] Graeme W. Davis. Homeostatic control of neural activity: From phenomenology to molecular design. *Annual Review of Neuroscience*, 29(1):307–323, July 2006.

[35] Myriam Lauren De Graaf, Luis Mochizuki, Frederik Thies, Heiko Wagner, and Charlotte Le Mouel. Motor pattern generation is robust to neural network anatomical imbalance favoring inhibition but not excitation, April 2022.

[36] Bruno Del Papa, Viola Priesemann, and Jochen Triesch. Criticality meets learning: Criticality signatures in a self-organizing recurrent neural network. *PLOS ONE*, 12(5):e0178683, May 2017.

[37] Zhidong Deng and Yi Zhang. Collective Behavior of a Small-World Recurrent Neural System With Scale-Free Distribution. *IEEE Transactions on Neural Networks*, 18(5):1364–1375, September 2007.

[38] Alain Destexhe and Jonathan D. Touboul. Is There Sufficient Evidence for Criticality in Cortical Systems? *eneuro*, 8(2):ENEURO.0551–20.2021, March 2021.

[39] Serena Di Santo, Raffaella Burioni, Alessandro Vezzani, and Miguel A. Muñoz. Self-Organized Bistability Associated with First-Order Phase Transitions. *Physical Review Letters*, 116(24):240601, June 2016.

[40] Markus Diesmann, Marc-Oliver Gewaltig, and Ad Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533, December 1999.

[41] B. Drossel and F. Schwabl. Self-organized critical forest-fire model. *Physical Review Letters*, 69(11):1629–1632, September 1992.

[42] Pierre Enel, Emmanuel Procyk, René Quilodran, and Peter Ford Dominey. Reservoir Computing Properties of Neural Dynamics in Prefrontal Cortex. *PLOS Computational Biology*, 12(6):e1004967, June 2016.

[43] Tomer Fekete, David B. Omer, Kazunori O'Hashi, Amiram Grinvald, Cees Van Leeuwen, and Oren Shriki. Critical dynamics, anesthesia and information integration: Lessons from multi-scale criticality analysis of voltage imaging data. *NeuroImage*, 183:919–933, December 2018.

[44] Antonio J. Fontenele, Nivaldo A.P. De Vasconcelos, Thaís Feliciano, Leandro A.A. Aguiar, Carina Soares-Cunha, Bárbara Coimbra, Leonardo Dalla Porta, Sidarta Ribeiro, Ana João Rodrigues, Nuno Sousa, Pedro V. Carelli, and Mauro Copelli. Criticality between Cortical States. *Physical Review Letters*, 122(20):208101, May 2019.

[45] Nir Friedman, Shinya Ito, Braden A. W. Brinkman, Masanori Shimono, R. E. Lee DeVille, Karin A. Dahmen, John M. Beggs, and Thomas C. Butler. Universal Critical Dynamics in High Resolution Neuronal Avalanche Data. *Physical Review Letters*, 108(20):208102, May 2012.

[46] Claudio Gallicchio, Alessio Micheli, and Luca Pedrelli. Deep reservoir computing: A critical experimental analysis. *Neurocomputing*, 268:87–99, December 2017.

[47] Timothy J. Gawne, Troels W. Kjaer, John A. Hertz, and Barry J. Richmond. Adjacent Visual Cortical Complex Cells Share About 20% of Their Stimulus-Related Information. *Cerebral Cortex*, 6(3):482–489, 1996.

[48] Arfan Ghani, T. Martin McGinnity, Liam P. Maguire, and Jim Harkin. Neuro-inspired Speech Recognition with Recurrent Spiking Neurons. In Véra Kůrková, Roman Neruda, and Jan Koutník, editors, *Artificial Neural Networks - ICANN 2008*, volume 5163, pages 513–522. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISSN: 0302-9743, 1611-3349 Series Title: Lecture Notes in Computer Science.

[49] Mauricio Girardi-Schappo, Ludmila Brochini, Ariadne A. Costa, Tawan T. A. Carvalho, and Osame Kinouchi. Synaptic balance due to homeostatically self-organized quasicritical dynamics. *Physical Review Research*, 2(1):012042, February 2020.

[50] Mauricio Girardi-Schappo, Emilio F Galera, Tawan T A Carvalho, Ludmila Brochini, Nilton L Kamiji, Antonio C Roque, and Osame Kinouchi. A unified theory of E/I synaptic balance, quasicritical neuronal avalanches and asynchronous irregular spiking. *Journal of Physics: Complexity*, 2(4):045001, December 2021.

[51] Elakkat D. Gireesh and Dietmar Plenz. Neuronal avalanches organize as nested theta- and beta/gamma-oscillations during development of cortical layer 2/3. *Proceedings of the National Academy of Sciences*, 105(21):7576–7581, May 2008.

[52] James P. Gleeson and Rick Durrett. Temporal profiles of avalanches on networks. *Nature Communications*, 8(1):1227, October 2017.

[53] Gerald Hahn, Thomas Petermann, Martha N. Havenith, Shan Yu, Wolf Singer, Dietmar Plenz, and Danko Nikolić. Neuronal Avalanches in Spontaneous Activity In Vivo. *Journal of Neurophysiology*, 104(6):3312–3322, December 2010.

[54] Clayton Haldeman and John M. Beggs. Critical Branching Captures Activity in Living Neural Networks and Maximizes the Number of Metastable States. *Physical Review Letters*, 94(5):058101, February 2005.

[55] D. Hamon, M. Nicodemi, and H. J. Jensen. Continuously driven OFC: A simple model of solar flare statistics. *Astronomy & Astrophysics*, 387(1):326–334, May 2002.

[56] G. A. Held, D. H. Solina, H. Solina, D. T. Keane, W. J. Haag, P. M. Horn, and G. Grinstein. Experimental study of critical-mass fluctuations in an evolving sandpile. *Physical Review Letters*, 65(9):1120–1123, August 1990.

[57] Mike Hemberger, Mark Shein-Idelson, Lorenz Pammer, and Gilles Laurent. Reliable Sequential Activation of Neural Assemblies by Single Pyramidal Cells in a Three-Layered Cortex. *Neuron*, 104(2):353–369.e5, October 2019.

[58] Janina Hesse and Thilo Gross. Self-organized criticality as a fundamental property of neural systems. *Frontiers in Systems Neuroscience*, 8, September 2014.

[59] John H. Holland. *Signals and Boundaries: Building Blocks for Complex Adaptive Systems*. The MIT Press, 2012.

[60] Petter Holme and Mikael Huss. Role-similarity based functional prediction in networked systems: application to the yeast proteome. *Journal of the Royal Society Interface*, 2(4):327–333, September 2005.

[61] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1):106–154, January 1962.

[62] Lauren A. Hurley, Juan G. Restrepo, and Sean E. Shaheen. Tuning the activation function to optimize the forecast horizon of a reservoir computer, 2023. Version Number: 1.

[63] Giacomo Indiveri, Bernabé Linares-Barranco, Robert Legenstein, George Deligeorgis, and Themistoklis Prodromakis. Integration of nanoscale memristor synapses in neuromorphic computing architectures. *Nanotechnology*, 24(38):384010, September 2013.

[64] Herbert Jaeger. The 'echo state' approach to analyzing and training recurrent neural networks. GMD Report 148, German National Research Center for Information Technology, 2001.

[65] Herbert Jaeger. The "echo state" approach to analysing and training recurrent neural networks-with an erratum note. *German National Research Center for Information Technology GMD Technical Report*, 148(34):13, 2001.

[66] Herbert Jaeger and Harald Haas. Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication. *Science*, 304(5667):78–80, April 2004.

[67] Henrik Jeldtoft Jensen. *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems*. Cambridge University Press, 1 edition, January 1998.

[68] Han Ju, Mark R. Dranias, Gokulakrishna Banumurthy, and Antonius M.J. VanDongen. Spatiotemporal Memory Is an Intrinsic Property of Networks of Dissociated Cortical Neurons. *The Journal of Neuroscience*, 35(9):4040–4051, March 2015.

[69] Eric R. Kandel. *Psychiatry, psychoanalysis, and the new biology of mind*. American Psychiatric Pub, Washington, DC, 1st ed edition, 2005.

[70] Yahya Karimipanah, Zhengyu Ma, Jae-eun Kang Miller, Rafael Yuste, and Ralf Wessel. Neocortical activity is stimulus- and scale-invariant. *PLOS ONE*, 12(5):e0177396, May 2017.

[71] Stuart A. Kauffman. *The origins of order: self-organization and selection in evolution*. Oxford University Press, New York, 1993.

[72] Yuji Kawai, Jihoon Park, and Minoru Asada. A small-world topology enhances the echo state property and signal propagation in reservoir computing. *Neural Networks*, 112:15–23, April 2019.

[73] Osame Kinouchi and Mauro Copelli. Optimal dynamical range of excitable networks at criticality. *Nature Physics*, 2(5):348–351, May 2006.

[74] Andreas Klaus, Shan Yu, and Dietmar Plenz. Statistical Analyses Support Power Law Distributions Found in Neuronal Avalanches. *PLoS ONE*, 6(5):e19779, May 2011.

[75] L. F. Kozachenko and N. N. Leonenko. Sample estimate of the entropy of a random vector. *Problems of Information Transmission*, 23(1-2):95–101, 1987.

[76] Sunil Kumar and Nivedita Deo. Correlation and network analysis of global financial indices. *Physical Review E*, 86(2):026101, August 2012.

[77] Matthew C. Kuntz and James P. Sethna. Noise in disordered systems: The power spectrum and dynamic exponents in avalanche models. *Physical Review B*, 62(17):11699–11708, November 2000.

[78] Yoshiki Kuramoto. Self-entrainment of a population of coupled non-linear oscillators. In *International Symposium on Mathematical Problems in Theoretical Physics*, Lecture Notes in Physics, pages 420–422, Berlin, Heidelberg, 1975. Springer.

[79] Chris G. Langton. Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D: Nonlinear Phenomena*, 42(1-3):12–37, June 1990.

[80] Andreea Lazar. SORN: a Self-organizing Recurrent Neural Network. *Frontiers in Computational Neuroscience*, 3, 2009.

[81] Jérôme Lecoq, Michael Oliver, Joshua H. Siegle, Natalia Orlova, Peter Ledochowitsch, and Christof Koch. Removing independent noise in systems neuroscience data using DeepInterpolation. *Nature Methods*, 18(11):1401–1408, November 2021.

[82] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.

[83] Robert Legenstein and Wolfgang Maass. Edge of chaos and prediction of computational performance for neural circuit models. *Neural Networks*, 20(3):323–334, April 2007.

[84] A. Levina and V. Priesemann. Subsampling scaling. *Nature Communications*, 8(1):15140, May 2017.

[85] Anna Levina, Viola Priesemann, and Johannes Zierenberg. Tackling the subsampling problem to infer collective properties from limited data. *Nature Reviews Physics*, 4(12):770–784, November 2022.

[86] Junhao Liang and Changsong Zhou. Criticality enhances the multilevel reliability of stimulus responses in cortical neural networks. *PLOS Computational Biology*, 18(1):e1009848, January 2022.

[87] Github link:. https://github.com/keshav224/NetworkReservoir.

[88] F. Lombardi, H. J. Herrmann, D. Plenz, and L. De Arcangelis. Temporal correlations in neuronal avalanche occurrence. *Scientific Reports*, 6(1):24690, April 2016.

[89] Fabrizio Lombardi, Hans J. Herrmann, Dietmar Plenz, and Lucilla De Arcangelis. On the temporal organization of neuronal avalanches. *Frontiers in Systems Neuroscience*, 8, October 2014.

[90] Edward N. Lorenz. Deterministic Nonperiodic Flow. *Journal of Atmospheric Sciences*, 20(2):130 – 141, 1963. Place: Boston MA, USA Publisher: American Meteorological Society.

[91] Nastaran Lotfi, Thaís Feliciano, Leandro A. A. Aguiar, Thais Priscila Lima Silva, Tawan T. A. Carvalho, Osvaldo A. Rosso, Mauro Copelli, Fernanda S. Matias, and Pedro V. Carelli. Statistical complexity is maximized close to criticality in cortical dynamics. *Physical Review E*, 103(1):012415, January 2021.

[92] Ryan J. Low, Yi Gu, and David W. Tank. Cellular resolution optical access to brain regions in fissures: Imaging medial prefrontal cortex and grid cells in entorhinal cortex. *Proceedings of the National Academy of Sciences*, 111(52):18739–18744, December 2014.

[93] Zhixin Lu, Shane Squires, Edward Ott, and Michelle Girvan. Inhibitory neurons promote robust critical firing dynamics in networks of integrate-and-fire neurons. *Phys. Rev. E*, 94:062309, Dec 2016.

[94] Mantas Lukoševičius and Herbert Jaeger. Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3):127–149, August 2009.

[95] Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications*, 390(6):1150–1170, March 2011.

[96] Zhengyu Ma, Haixin Liu, Takaki Komiyama, and Ralf Wessel. Stability of motor cortex network states during learning-associated neural reorganizations. *Journal of Neurophysiology*, 124(5):1327–1342, November 2020.

[97] Zhengyu Ma, Gina G. Turrigiano, Ralf Wessel, and Keith B. Hengen. Cortical Circuit Dynamics Are Homeostatically Tuned to Criticality In Vivo. *Neuron*, 104(4):655–664.e4, November 2019.

[98] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, November 2002.

[99] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*, 14(11):2531–2560, November 2002.

[100] Michael C. Mackey and Leon Glass. Oscillation and Chaos in Physiological Control Systems. *Science*, 197(4300):287–289, July 1977.

[101] Alessandro Magni, Gianfranco Durin, Stefano Zapperi, and James P Sethna. Visualization of avalanches in magnetic thin films: temporal processing. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(01):P01020, January 2009.

[102] Benedetta Mariani, Giorgio Nicoletti, Marta Bisio, Marta Maschietto, Roberto Oboe, Alessandro Leparulo, Samir Suweis, and Stefano Vassanelli. Neuronal Avalanches Across the Rat Somatosensory Barrel Cortex and the Effect of Single Whisker Stimulation. *Frontiers in Systems Neuroscience*, 15:709677, August 2021.

[103] Matteo Martinello, Jorge Hidalgo, Amos Maritan, Serena Di Santo, Dietmar Plenz, and Miguel A. Muñoz. Neutral Theory and Scale-Free Neural Dynamics. *Physical Review X*, 7(4):041071, December 2017.

[104] Matias I. Maturana, Christian Meisel, Katrina Dell, Philippa J. Karoly, Wendyl D'Souza, David B. Grayden, Anthony N. Burkitt, Premysl Jiruska, Jan Kudlacek, Jaroslav Hlinka, Mark J. Cook, Levin Kuhlmann, and Dean R. Freestone. Critical slowing down as a biomarker for seizure susceptibility. *Nature Communications*, 11(1):2172, May 2020.

[105] Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133, December 1943.

[106] Amit P. Mehta, Andrea C. Mills, Karin A. Dahmen, and James P. Sethna. Universal pulse shape scaling function and exponents: Critical test for avalanche models applied to Barkhausen noise. *Physical Review E*, 65(4):046139, April 2002.

[107] Christian Meisel. Antiepileptic drugs induce subcritical dynamics in human cortical networks. *Proceedings of the National Academy of Sciences*, 117(20):11118–11125, May 2020.

[108] Christian Meisel, Kimberlyn Bailey, Peter Achermann, and Dietmar Plenz. Decline of long-range temporal correlations in the human brain during sustained wakefulness. *Scientific Reports*, 7(1):11825, September 2017.

[109] Christian Meisel, Andreas Klaus, Vladyslav V. Vyazovskiy, and Dietmar Plenz. The Interplay between Long- and Short-Range Temporal Correlations Shapes Cortex Dynamics across Vigilance States. *The Journal of Neuroscience*, 37(42):10114–10124, October 2017.

[110] Christian Meisel, Eckehard Olbrich, Oren Shriki, and Peter Achermann. Fading Signatures of Critical Brain Dynamics during Sustained Wakefulness in Humans. *The Journal of Neuroscience*, 33(44):17363–17372, October 2013.

[111] Christian Meisel, Dietmar Plenz, Andreas Schulze-Bonhage, and Heinz Reichmann. Quantifying antiepileptic drug effects using intrinsic excitability measures. *Epilepsia*, 57(11), November 2016.

[112] Christian Meisel, Andreas Schulze-Bonhage, Dean Freestone, Mark James Cook, Peter Achermann, and Dietmar Plenz. Intrinsic excitability measures track antiepileptic drug action and uncover increasing/decreasing excitability over the wake/sleep cycle. *Proceedings of the National Academy of Sciences*, 112(47):14694–14699, November 2015.

[113] Christian Meisel, Alexander Storch, Susanne Hallmeyer-Elgner, Ed Bullmore, and Thilo Gross. Failure of Adaptive Self-Organized Criticality during Epileptic Seizure Attacks. *PLoS Computational Biology*, 8(1):e1002312, January 2012.

[114] Stephanie R. Miller, Shan Yu, and Dietmar Plenz. The scale-invariant, temporal profile of neuronal avalanches in relation to cortical γ–oscillations. *Scientific Reports*, 9(1):16403, November 2019.

[115] Guillermo B. Morales, Claudio R. Mirasso, and Miguel C. Soriano. Unveiling the role of plasticity rules in reservoir computing. *Neurocomputing*, 461:705–715, October 2021.

[116] Satoshi Moriya, Hideaki Yamamoto, Hisanao Akima, Ayumi Hirano-Iwata, Shigeru Kubota, and Shigeo Sato. Mean-field analysis of directed modular networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(1):013142, January 2019.

[117] Harold J. Morowitz. *The Mind, The Brain, and Complex Adaptive Systems*. Routledge, 1 edition, March 2018.

[118] Mia C Morrell, Ilya Nemenman, and Audrey Sederberg. Neural criticality from effective latent variables. *eLife*, 12:RP89337, March 2024.

[119] Miguel A. Muñoz. *Colloquium* : Criticality and dynamical scaling in living systems. *Reviews of Modern Physics*, 90(3):031001, July 2018.

[120] Miguel A. Muñoz, Ronald Dickman, Alessandro Vespignani, and Stefano Zapperi. Avalanche and spreading exponents in systems with absorbing states. *Physical Review E*, 59(5):6175–6179, May 1999.

[121] Ian Nauhaus, Laura Busse, Dario L. Ringach, and Matteo Carandini. Robustness of Traveling Waves in Ongoing Activity of Visual Cortex. *The Journal of Neuroscience*, 32(9):3088–3094, February 2012.

[122] Marcel Nonnenmacher, Christian Behrens, Philipp Berens, Matthias Bethge, and Jakob H. Macke. Signatures of criticality arise from random subsampling in simple population models. *PLOS Computational Biology*, 13(10):e1005718, October 2017.

[123] Zeev Olami, Hans Jacob S. Feder, and Kim Christensen. Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes. *Physical Review Letters*, 68(8):1244–1247, February 1992.

[124] Mustafa C. Ozturk, Dongming Xu, and José C. Príncipe. Analysis and Design of Echo State Networks. *Neural Computation*, 19(1):111–138, January 2007.

[125] Marius Pachitariu, Carsen Stringer, Mario Dipoppa, Sylvia Schröder, L. Federico Rossi, Henry Dalgleish, Matteo Carandini, and Kenneth D. Harris. Suite2p: beyond 10,000 neurons with standard two-photon microscopy, June 2016.

[126] N.H. Packard. *Adaptation Toward the Edge of Chaos*. University of Illinois at Urbana-Champaign, Center for Complex Systems Research, 1988.

[127] J. Matias Palva, Alexander Zhigalov, Jonni Hirvonen, Onerva Korhonen, Klaus Linkenkaer-Hansen, and Satu Palva. Neuronal long-range temporal correlations and avalanche dynamics are correlated with behavioral scaling laws. *Proceedings of the National Academy of Sciences*, 110(9):3585–3590, February 2013.

[128] Stefanos Papanikolaou, Felipe Bohn, Rubem Luis Sommer, Gianfranco Durin, Stefano Zapperi, and James P. Sethna. Universality beyond power laws and the average avalanche shape. *Nature Physics*, 7(4):316–320, April 2011.

[129] Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach. *Physical Review Letters*, 120(2):024102, January 2018.

[130] Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach. *Physical Review Letters*, 120(2):024102, January 2018.

[131] Jaideep Pathak, Alexander Wikner, Rebeckah Fussell, Sarthak Chandra, Brian R. Hunt, Michelle Girvan, and Edward Ott. Hybrid forecasting of chaotic processes: Using machine learning in conjunction with a knowledge-based model. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(4):041101, April 2018.

[132] Johannes Pausch. From neuronal spikes to avalanches: Effects and circumvention of time binning. *Physical Review Research*, 4(2):023212, June 2022.

[133] Hanchuan Peng, Fuhui Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238, August 2005.

[134] Thomas Petermann, Tara C. Thiagarajan, Mikhail A. Lebedev, Miguel A. L. Nicolelis, Dante R. Chialvo, and Dietmar Plenz. Spontaneous cortical activity in awake monkeys composed of neuronal avalanches. *Proceedings of the National Academy of Sciences*, 106(37):15921–15926, September 2009.

[135] D. Plenz. Neuronal avalanches and coherence potentials. *The European Physical Journal Special Topics*, 205(1):259–301, May 2012.

[136] Dietmar Plenz, Tiago L. Ribeiro, Stephanie R. Miller, Patrick A. Kells, Ali Vakili, and Elliott L. Capek. Self-Organized Criticality in the Brain. *Frontiers in Physics*, 9:639389, July 2021.

[137] Simon-Shlomo Poil, Arjen Van Ooyen, and Klaus Linkenkaer-Hansen. Avalanche dynamics of human brain oscillations: Relation to critical branching processes and temporal correlations. *Human Brain Mapping*, 29(7):770–777, July 2008.

[138] Karine Pozo and Yukiko Goda. Unraveling Mechanisms of Homeostatic Synaptic Plasticity. *Neuron*, 66(3):337–351, May 2010.

[139] J. C. Prechtl, L. B. Cohen, B. Pesaran, P. P. Mitra, and D. Kleinfeld. Visual stimuli induce waves of electrical activity in turtle cortex. *Proceedings of the National Academy of Sciences*, 94(14):7621–7626, July 1997.

[140] Viola Priesemann, Matthias Hj Munk, and Michael Wibral. Subsampling effects in neuronal avalanche distributions recorded in vivo. *BMC Neuroscience*, 10(1):40, December 2009.

[141] Viola Priesemann and Oren Shriki. Can a time varying external drive give rise to apparent criticality in neural systems? *PLOS Computational Biology*, 14(5):e1006081, May 2018.

[142] Astrid A Prinz, Dirk Bucher, and Eve Marder. Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7(12):1345–1352, December 2004.

[143] Juan G. Restrepo and Edward Ott. Mean-field theory of assortative networks of phase oscillators. *EPL (Europhysics Letters)*, 107(6):60006, September 2014.

[144] Sidarta Ribeiro, Damien Gervasoni, Ernesto S Soares, Yi Zhou, Shih-Chieh Lin, Janaina Pantoja, Michael Lavine, and Miguel A. L Nicolelis. Long-Lasting Novelty-Induced Neuronal Reverberation during Slow-Wave Sleep in Multiple Forebrain Areas. *PLoS Biology*, 2(1):e24, January 2004.

[145] Sidarta Ribeiro and Miguel A.L. Nicolelis. Reverberation, storage, and postsynaptic propagation of memories during sleep. *Learning & Memory*, 11(6):686–696, November 2004.

[146] Tiago L. Ribeiro, Peter Jendrichovsky, Shan Yu, Daniel A. Martin, Patrick O. Kanold, Dante R. Chialvo, and Dietmar Plenz. Trial-by-trial variability in cortical responses exhibits scaling of spatial correlations predicted from critical dynamics. *Cell Reports*, 43(2):113762, February 2024.

[147] Tiago L. Ribeiro, Sidarta Ribeiro, Hindiael Belchior, Fábio Caixeta, and Mauro Copelli. Undersampled Critical Branching Processes on Small-World and Random Networks Fail to Reproduce the Statistics of Spike Avalanches. *PLoS ONE*, 9(4):e94992, April 2014.

[148] Tiago L. Ribeiro, Sidarta Ribeiro, and Mauro Copelli. Repertoires of Spike Avalanches Are Modulated by Behavior and Novelty. *Frontiers in Neural Circuits*, 10, March 2016.

[149] Nathaniel Rodriguez, Eduardo Izquierdo, and Yong-Yeol Ahn. Optimal modularity and memory capacity of neural reservoirs. *Network Neuroscience*, 3(2):551–566, January 2019.

[150] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.

[151] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, October 1986.

[152] Silvia Scarpetta. Critical Behavior and Memory Function in a Model of Spiking Neurons with a Reservoir of Spatio-Temporal Patterns. In Nergis Tomen, J. Michael

Herrmann, and Udo Ernst, editors, *The Functional Role of Critical Dynamics in Neural Systems*, volume 11, pages 179–197. Springer International Publishing, Cham, 2019. Series Title: Springer Series on Bio- and Neurosystems.

[153] Silvia Scarpetta, Ilenia Apicella, Ludovico Minati, and Antonio De Candia. Hysteresis, neural avalanches, and critical behavior near a first-order transition of a spiking neural network. *Physical Review E*, 97(6):062305, June 2018.

[154] Silvia Scarpetta and Antonio De Candia. Alternation of up and down states at a dynamical phase-transition of a neural network with spatiotemporal attractors. *Frontiers in Systems Neuroscience*, 8, May 2014.

[155] Silvia Scarpetta, Niccolò Morisi, Carlotta Mutti, Nicoletta Azzi, Irene Trippi, Rosario Ciliento, Ilenia Apicella, Giovanni Messuti, Marianna Angiolelli, Fabrizio Lombardi, Liborio Parrino, and Anna Elisabetta Vaudano. Criticality of neuronal avalanches in human sleep and their relationship with sleep macro- and microarchitecture. *iScience*, 26(10):107840, October 2023.

[156] Felix Johannes Schmitt, Vahid Rostami, and Martin Paul Nawrot. Efficient parameter calibration and real-time simulation of large-scale spiking neural networks with GeNN and NEST. *Frontiers in Neuroinformatics*, 17:941696, February 2023.

[157] Benjamin Schrauwen, David Verstraeten, and Jan Van Campenhout. An overview of reservoir computing: theory, applications and implementations. In *Proceedings of the 15th european symposium on artificial neural networks. p. 471-482 2007*, pages 471–482, 2007.

[158] Benjamin Schrauwen, Marion Wardermann, David Verstraeten, Jochen J. Steil, and Dirk Stroobandt. Improving reservoirs using intrinsic plasticity. *Neurocomputing*, 71(7-9):1159–1171, March 2008.

[159] Thomas Schreiber. Measuring Information Transfer. *Physical Review Letters*, 85(2):461–464, July 2000.

[160] David P. A. Schulz, Maneesh Sahani, and Matteo Carandini. Five key factors determining pairwise correlations in visual cortex. *Journal of Neurophysiology*, 114(2):1022–1033, August 2015.

[161] Matteo Serafino, Giulio Cimini, Amos Maritan, Andrea Rinaldo, Samir Suweis, Jayanth R. Banavar, and Guido Caldarelli. True scale-free networks hidden by finite size effects. *Proceedings of the National Academy of Sciences*, 118(2):e2013825118, January 2021.

[162] Saurav Seshadri, Andreas Klaus, Daniel E. Winkowski, Patrick O. Kanold, and Dietmar Plenz. Altered avalanche dynamics in a developmental NMDAR hypofunction model of cognitive impairment. *Translational Psychiatry*, 8(1):3, January 2018.

[163] James P. Sethna, Karin A. Dahmen, and Christopher R. Myers. Crackling noise. *Nature*, 410(6825):242–250, March 2001.

[164] Woodrow L. Shew, Hongdian Yang, Thomas Petermann, Rajarshi Roy, and Dietmar Plenz. Neuronal Avalanches Imply Maximum Dynamic Range in Cortical Networks at Criticality. *The Journal of Neuroscience*, 29(49):15595–15600, December 2009.

[165] Woodrow L. Shew, Hongdian Yang, Shan Yu, Rajarshi Roy, and Dietmar Plenz. Information Capacity and Transmission Are Maximized in Balanced Cortical Networks with Neuronal Avalanches. *The Journal of Neuroscience*, 31(1):55–63, January 2011.

[166] Oren Shriki and Dietmar Plenz. Neuronal Avalanches in the Human Brain. In Dietmar Plenz and Ernst Niebur, editors, *Criticality in Neural Systems*, pages 177–190. Wiley, 1 edition, April 2014.

[167] Wolf Singer. Cortical dynamics revisited. *Trends in Cognitive Sciences*, 17(12):616–626, December 2013.

[168] Per Sebastian Skardal, Juan G. Restrepo, and Edward Ott. Frequency assortativity can induce chaos in oscillator networks. *Physical Review E*, 91(6):060902, June 2015.

[169] Matthew A. Smith and Adam Kohn. Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. *The Journal of Neuroscience*, 28(48):12591–12603, November 2008.

[170] Miguel C. Soriano, Daniel Brunner, Miguel Escalona-MorÃ¡n, Claudio R. Mirasso, and Ingo Fischer. Minimal approach to neuro-inspired information processing. *Frontiers in Computational Neuroscience*, 9, June 2015.

[171] Keshav Srinivasan, Nolan Coble, Joy Hamlin, Thomas Antonsen, Edward Ott, and Michelle Girvan. Parallel Machine Learning for Forecasting the Dynamics of Complex Networks. *Physical Review Letters*, 128(16):164101, April 2022.

[172] Jochen J. Steil. Online reservoir adaptation by intrinsic plasticity for backpropagation–decorrelation and echo state learning. *Neural Networks*, 20(3):353–364, April 2007.

[173] Michael P. H. Stumpf, Carsten Wiuf, and Robert M. May. Subnets of scale-free networks are not scale-free: Sampling properties of networks. *Proceedings of the National Academy of Sciences*, 102(12):4221–4224, March 2005.

[174] Fei Tan, Yongxiang Xia, and Boyao Zhu. Link Prediction in Complex Networks: A Mutual Information Perspective. *PLOS ONE*, 9(9):e107056, September 2014.

[175] Jonathan Touboul and Alain Destexhe. Power-law statistics and universal scaling in the absence of criticality. *Physical Review E*, 95(1):012413, January 2017.

[176] Jochen Triesch. Synergies Between Intrinsic and Synaptic Plasticity Mechanisms. *Neural Computation*, 19(4):885–909, April 2007.

[177] Gina G. Turrigiano. The Self-Tuning Neuron: Synaptic Scaling of Excitatory Synapses. *Cell*, 135(3):422–435, October 2008.

[178] Gina G. Turrigiano and Sacha B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(2):97–107, February 2004.

[179] D. Verstraeten, B. Schrauwen, M. D'Haene, and D. Stroobandt. An experimental unification of reservoir computing methods. *Neural Networks*, 20(3):391–403, April 2007.

[180] Pablo Villegas, Serena Di Santo, Raffaella Burioni, and Miguel A. Muñoz. Time-series thresholding and the definition of avalanche size. *Physical Review E*, 100(1):012133, July 2019.

[181] Xinjie Wang, Yaochu Jin, and Kuangrong Hao. Echo state networks regulated by local intrinsic plasticity rules for regression. *Neurocomputing*, 351:111–122, July 2019.

[182] Zhanshan Wang, Xianshuang Yao, Zhanjun Huang, and Lei Liu. Deep Echo State Network With Multiple Adaptive Reservoirs for Time Series Prediction. *IEEE Transactions on Cognitive and Developmental Systems*, 13(3):693–704, September 2021.

[183] Francis Wyffels, Benjamin Schrauwen, and Dirk Stroobandt. Stable Output Feedback in Reservoir Computing Using Ridge Regression. In Véra Kůrková, Roman Neruda, and Jan Koutník, editors, *Artificial Neural Networks - ICANN 2008*, volume 5163, pages 808–817. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISSN: 0302-9743, 1611-3349 Series Title: Lecture Notes in Computer Science.

[184] Yifan Xu, Aidan Schneider, Ralf Wessel, and Keith B. Hengen. Sleep restores an optimal computational regime in cortical networks. *Nature Neuroscience*, 27(2):328–338, February 2024.

[185] Tadashi Yamazaki and Shigeru Tanaka. A spiking network model for passage-of-time representation in the cerebellum. *European Journal of Neuroscience*, 26(8):2279–2292, October 2007.

[186] Ofer Yizhar, Lief E. Fenno, Matthias Prigge, Franziska Schneider, Thomas J. Davidson, Daniel J. O'Shea, Vikaas S. Sohal, Inbal Goshen, Joel Finkelstein, Jeanne T. Paz, Katja Stehfest, Roman Fudim, Charu Ramakrishnan, John R. Huguenard, Peter

Hegemann, and Karl Deisseroth. Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature*, 477(7363):171–178, September 2011.

[187] Shan Yu, Andreas Klaus, Hongdian Yang, and Dietmar Plenz. Scale-Invariant Neuronal Avalanche Dynamics and the Cut-Off in Size Distributions. *PLoS ONE*, 9(6):e99761, June 2014.

[188] Shan Yu, Tiago L Ribeiro, Christian Meisel, Samantha Chou, Andrew Mitz, Richard Saunders, and Dietmar Plenz. Maintained avalanche dynamics during task-induced changes of neuronal activity in nonhuman primates. *eLife*, 6:e27119, November 2017.

[189] Shan Yu, Hongdian Yang, Hiroyuki Nakahara, Gustavo S. Santos, Danko Nikolić, and Dietmar Plenz. Higher-Order Interactions Characterized in Cortical Activity. *The Journal of Neuroscience*, 31(48):17514–17526, November 2011.

[190] Mohd-Hanif Yusoff, Joseph Chrol-Cannon, and Yaochu Jin. Modeling neural plasticity in echo state networks for classification and regression. *Information Sciences*, 364-365:184–196, October 2016.

[191] Anguo Zhang, Wei Zhu, and Ming Liu. Self-organizing reservoir computing based on spiking-timing dependent plasticity and intrinsic plasticity mechanisms. In *2017 Chinese Automation Congress (CAC)*, pages 6189–6193, Jinan, October 2017. IEEE.