2025

Decoding computational neuroscience: accessible frameworks for sophisticated neural analyses

BOSTON UNIVERSITY

COLLEGE OF ENGINEERING

Dissertation

DECODING COMPUTATIONAL NEUROSCIENCE:
ACCESSIBLE FRAMEWORKS FOR SOPHISTICATED NEURAL ANALYSES

by

DANIEL CARBONERO

B.S., University of Miami, 2019
M.S., Boston University, 2022

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2025

Approved by

First Reader ____

John A. White, Ph.D.
Professor and Chair of Biomedical Engineering
Professor of Pharmacology and Experimental Neuroscience
Professor of Neuroscience

Second Reader ____

Mark A. Kramer, Ph.D.
Professor of Mathematics and Statistics

Third Reader ____

Kamal Sen, Ph.D.
Associate Professor of Biomedical Engineering

Fourth Reader ____

Laura D. Lewis, Ph.D.
Athinoula A. Martinos Associate Professor of IMES and EECS
Massachusetts Institute of Technology

Fifth Reader ____

Christopher V. Gabel, Ph.D.
Associate Professor of Pharmacology
Associate Professor of Physiology and Biophysics

Through action, calm.

DEDICATION

I dedicate this work to my grandparents, for risking future for opportunity; my father, for always leading, and by example; my mother, for her unconditional love and unwavering support; and my sister, for being my best friend.

ACKNOWLEDGMENTS

I would like to thank my advisor John A. White for having me in his research group, and for providing an environment for his students to be curious, and independently pursue interesting research. He has never been anything other than supportive, and he always consistently ensures his trainees are in a position to succeed.

I would also like my co-advisor Mark A. Kramer. I was always made to feel welcome in his group, and he always provided excellent mentorship and affirmation. Our three-month data-augmentation project was proper intellectual fun.

I would like to thank Fernando R. Fernandez for the day-to-day iterating towards the beginning of my research at Boston University. He provided critical support and was instrumental in developing my scientific thought process.

To Jacob Norman, and Jad Noueihed, I am very appreciative for their contemporary support, and mentorship as I learned to navigate science on my own.

To the rest of my committee, I thank them for their guidance and support in the completion of this work.

To the rest of the NDL Lab, Kevin Ghaemi, Bahar Rahsepar, Akemi Ito, Audrey Teasley, and Irene Zhou, I am thankful for their efforts in creating such a supportive and collegial lab atmosphere that made long days short, and easy days long. I would especially like to thank Brandon Williams for being my everyday "partner-in-crime." I am fortunate to have worked every day with such an excellent group of people.

I would like to thank Emily Schlafly for including me in our fast-paced collaboration in the Kramer group, and for her openness to my ideas, and her mentorship in a different field. She was instrumental in my completion of this work.

To the rest of my colleagues and friends in the city of Boston, I am thankful they have helped me create a home away from home. I am truly grateful for their support and companionship through difficult times, and the laughs and celebrations through the good. I need to thank Ashutosh Agarwal, without whom I would not have chosen to pursue this career trajectory. He was my first scientific advisor at the University of Miami, and he remains a very trusted colleague, mentor, and friend.

I would like to thank my cousin Paul for always including me in his pursuits. He has always provided an example to live by and has challenged me to always be intellectually curious. To my godfather, Tutti, I am thankful for him always being a guiding voice in my life, and a steady and grounding presence. I would like to thank the rest of my extended family — my cousins, my uncles, and aunts — for being the village it takes to pursue one's goals. I would also like to thank my Cuban cousins, Ish and Carlos, for being a part of my journey, and having me be a part of theirs.

I would like to thank my grandparents, for risking their future in coming to the United States to give their children and grandchildren a better life. It created opportunities for me to pursue whatever I wanted.

Finally, I want to thank my family. To my father, for helping instill in me my values and way of being, I am eternally grateful. To my mother, for the long days and nights of raising me, and beyond, I cannot think of anyone who has been more instrumental in my success. To my sister, for being the best companion anyone could ask for, I am thankful to have her by my side.

DECODING COMPUTATIONAL NEUROSCIENCE:

ACCESSIBLE FRAMEWORKS FOR SOPHISTICATED NEURAL ANALYSES

DANIEL CARBONERO

Boston University College of Engineering, 2025

Major Professor: John A. White, Ph.D. Professor and Chair of Biomedical Engineering, Professor of Pharmacology and Experimental Neuroscience, Professor of Neuroscience

ABSTRACT

Advancements in neural recording technologies have led to an explosion in the scale and complexity of neuroscientific datasets, demanding analytical frameworks that are not only computationally powerful, but also interpretable and accessible. This dissertation introduces a cohesive suite of computational methodologies, validations, and tools to address these demands, focusing on the application of both unsupervised and supervised machine learning techniques to extract meaningful neural dynamics and predict emergent network behavior, or clinically significant events, from a high volume of recordings.

We first focus on the application and validation of dimensionality reduction (DR), rigorously evaluating Nonnegative Matrix Factorization (NMF) for analyzing state-dependent neuronal network dynamics in calcium imaging data. We demonstrate its superiority over commonly used DR methods, specifically PCA, ICA, and UMAP, in recovering biologically plausible, interpretable subnetwork activity.

Building on this validation, we present CaNetiCs, an open-source software toolbox that integrates standardized implementations of NMF and complementary DR

methods, alongside modules for geometric low-dimensional space modeling, and synthetic neuronal simulations. We apply this framework to datasets from C. elegans and murine recordings under graded anesthesia, recapitulating known physiological trends while uncovering novel insights into neural population dynamics and state transitions.

Finally, we extend this analytical ideology to the supervised learning domain by introducing a data augmentation strategy for enhancing epileptic spike ripple detection in human electrophysiological recordings. Employing LSTM-based neural networks, we show that training on a combination of empirically grounded synthetic and in vivo data improves classifier performance, offering a scalable solution to training-data scarcity in clinical neurodiagnostics.

Together, these contributions provide accessible, interpretable, and rigorously validated frameworks for dimensionality reduction and machine learning in neuroscience, advancing both understanding and translational utility in the analysis of complex neural systems. By integrating biologically grounded decomposition methods, open-source analytical software, and data-driven neural network architectures, this work bridges methodological innovation with practical implementation. The tools and techniques developed herein not only facilitate the discovery of latent structure in large-scale neuronal recordings but also enable the automated identification of clinically relevant biomarkers. In doing so, this dissertation empowers researchers to more effectively characterize dynamic brain activity across species, brain states, and experimental paradigms, adding to a computational foundation for reproducible, scalable, and hypothesis-driven neuroscience.

TABLE OF CONTENTS

DEDICATION ... v

ACKNOWLEDGMENTS ... vi

ABSTRACT ... ix

TABLE OF CONTENTS ... xi

LIST OF TABLES ... xv

LIST OF FIGURES ... xvi

LIST OF ABBREVIATIONS ... xx

CHAPTER ONE: Introduction ... 1

- Introduction ... 1
- Challenges with Traditional Analytical Approaches ... 3
- Machine Learning in Neuroscience ... 4
- Motivation, Objectives, Outline, and Implications of This Thesis ... 9

CHAPTER TWO: Nonnegative Matrix Factorization for Analyzing State Dependent Neuronal Network Dynamics in Calcium Recordings ... 13

- Abstract ... 13
- Introduction ... 14
- Results ... 17
  - Perfectly Intraconnected, Independent Nodes ... 19
  - Random Process Propagation ... 24
  - Awake-State Data ... 28
- Discussion ... 31

Methods 35
Dimensionality Reduction (DR) Methods 35
Network Simulation Methods 44
Component Assignment Success 48
CHAPTER THREE: CaNetiCs - An Open-Source Toolbox for Standardized Dimensionality Reductive Analysis of Neuronal Calcium Activity 50
Abstract 50
Introduction 52
Results 55
Neuronal network dynamics of C. elegans under increasing concentrations of isoflurane: 56
Neuronal network dynamics of murine S1 L2/3 under increasing concentrations of isoflurane: 63
Discussion 71
Methods 75
Dimensionality reduction (DR) methods: 75
Other feature implementations in our software wrapper: 80
Validation analysis of previously collected experimental data: 82
Recording analysis pipeline: 83
CHAPTER FOUR: A data augmentation procedure to improve detection of spike ripples in brain voltage recordings 85
Abstract 85

Introduction 86
Methods 89
Human dataset 89
Augmented dataset 90
LSTM Frequency-Domain Method 94
LSTM Time-Domain Method 95
Training and Testing 96
Calculation of Probability Thresholds and Performance Statistics 100
Logistic regression 101
Results 101
Data augmentation produces synthetic spike ripples consistent with in vivo observations 101
Data augmentation moderately improves performance of a frequency-domain spike ripple detector 107
Data augmentation moderately improves performance of a time-domain spike ripple detector 109
Misclassified samples indicate strategies to improve synthetic training data 113
Interpreting signal features related to misclassifications 116
Discussion 117
CHAPTER FIVE: Discussion 121
Dimensionality Reduction 123
Deep and Supervised Learning 125

Constraints Imposed by Sample Size 126

Conclusion 127

BIBLIOGRAPHY 129

CURRICULUM VITAE 143

LIST OF TABLES

Table 1. Model Attributes 44

Table 2. Significance values for statistical comparison of optimal number of components found in Fig. 7b). Wilcoxon Signed-Rank Test, corrected for multiple comparisons, n = 10. 59

Table 3. Significance values for statistical comparison of volume expanse in space (Fig. 8f). Wilcoxon Rank-Sum Test, corrected for multiple comparisons, n = 10. 62

Table 4. Significance values for statistical comparison of optimal number of components found in Fig. 9b). Wilcoxon Signed-Rank Test, corrected for multiple comparisons, n = 36. 63

Table 5. Significance values for statistical comparison of volume expanse in space, found in Fig. 10f). Wilcoxon Rank-Sum Test, corrected for multiple comparisons, n = 36. 68

Table 6. Summary statistics for each spike ripple feature. In vivo statistics are based on 677 observations classified as spike ripples from the six subjects of interest included in the testing sets. 106

Table 7. ROC statistics for the frequency-domain spike ripple detectors. K indicates the number of synthetic spike ripples. 108

Table 8. Performance metrics for the frequency-domain spike ripple detector. Results for (top) in vivo data alone, (middle) synthetic data alone, and (bottom) in vivo data augmented with synthetic data. In each case, we indicate the mean value (standard deviation) across the n=6 subjects used for testing, with two different probability thresholds. PPV: positive predictive value, NPV: negative predictive value. 109

Table 9. ROC statistics for the time-domain spike ripple detectors. K indicates the number of synthetic spike ripples. 110

Table 10. Performance metrics for the time-domain spike ripple detector. Results for (top) no tuning, (second) in vivo data alone, (third) synthetic data alone, and (bottom) in vivo data augmented with synthetic data. In each case, we indicate the mean value (standard deviation) across the n=6 subjects used for testing, with two different probability thresholds. PPV: positive predictive value, NPV: negative predictive value 112

Table 11. Misclassifications resulting from training on synthetic data. 115

LIST OF FIGURES

Figure 1. Representative Model Pipeline. a) Recording Field of View (FOV) with Regions of Interest (ROIs). b) Extracted fluorescent activity traces from ROIs. c) Same recording FOV shown in a) with ROIs colored according to largest component contribution. d) Decomposed activities of the components shown in c). e) Sample truncated heatmap showing neuronal contribution or weight for each component .. 17

Figure 2. Nodal Network Architecture. a) Diagram showing perfect intraconnection between neurons and independence between nodes. b) Randomly selected calcium traces generated by the network. c) Same randomly selected calcium traces shown in b), but sorted according to node 18

Figure 3. Artificial Random Propagating Process Network Architecture. a) Sample connection diagram showing all random processes projecting to all neurons. b) Single histogram of generated connections between processes and neurons from sampled distribution 19

Figure 4. Artificial Nodal Network Analysis. a) Variance Explained for NMF and PCA, left axis (n = 256 networks) per component added to the model, and Akaike Information Criterion (AIC) for NMF, right axis (n = 256 networks). b) Prediction visualization between nodes and components for single network. c) Neuronal weights (y axis) for each model (panel) for each node (color) for each component (x axis). d) Average between neurons inside and outside predicted nodes for each component (n = 256 networks). e) Model assignment accuracy (n = 256 networks)21

Figure 5. Artificial Random Propagating Process Network Analysis. a) Variance Explained for NMF and PCA, left axis (n = 256 networks) per component added to the model, and Akaike Information Criterion (AIC) for NMF, right axis (n = 256 networks). b) Assignment visualization between nodes and components for single network. c) Neuronal weights for NMF model (y axis) for each component (x axis) compared to assigned connection probabilities (color scale) for the predicted random process. d) Correlation between neuronal weights for NMF model (y axis) and the respective assigned connection probability (x axis). e) Average Correlation between neuronal weights for model and the respective assigned connection probabilities (n = 256 networks). f) Model assignment accuracy (n = 256 networks). g) Correlation between component activity from NMF model and assigned random process activity. h) Correlations between component activity from model and assigned random process activity (n = 256 networks) 27

Figure 6. Baseline Awake State Network Analysis a) Variance Explained for NMF and PCA, left axis (n = 36 mice) per component added to the model, and Akaike Information Criterion (AIC) for NMF, right axis (n = 36 mice). b) Correlation between number of neurons (x axis) and number of optimized components (y axis) (n = 36 mice). c) Neuronal weights for NMF model (y axis) for each component (x

axis) for a single mouse. d) Neuronal weight histogram from NMF model for the same mouse as c). e) Component activity for NMF model of the same mouse as in c) and d). f) Pairwise correlations between component activities shown in e) 30

Figure 7. Optimal Dimensionality Models for C. elegans Isoflurane Dataset. a) Activity traces for a subset (n=25) of neurons, for a single animal, across conditions. b) Distribution of optimal components for all animals, across states (n = 10). c) Mean variance explained curve with standard deviation intervals (n=10). d) Mean Akaike Information Criterion (AIC) curve with standard deviation intervals (n=10). e) Visualization of neuronal weights at optimal dimensionality for animal shown in a), standardized across conditions for comparison. f) Aggregated neuronal weights for each animal at optimal dimensionality (n = [3600, 5640, 4080], respectively)....... 57

Figure 8. Three-Dimensional Low-Dimensional Component Space Models for C. elegans Isoflurane Dataset. a) Component activity traces for first three components for the same animal shown in Fig. 7a). b) Decomposed frequency content for components shown in Fig. 8a). c) Visualization of neuronal weights for single three-dimensional models for the same animal in Figs. 8a,b). d) Aggregated neuronal weights for all three-dimensional models, across animals and states (n = 3600). e) Plotted low-dimensional component space of components against each other, with encapsulated volume, for animal shown Fig. 8a,b,c). f) Distribution of volumes, across animals and states (n=10). g) Distribution of directions, across animals and states (n=5380). h) Distributions of speeds across animals and states (n = 5380) 61

Figure 9. Optimal Dimensionality Models for Murine Isoflurane Dataset. a) Activity traces for a subset (n=25) of neurons, for a single animal, across conditions. b) Distribution of optimal components for all animals, across states (n = 36). c) Mean variance explained curve with standard deviation intervals (n=36). d) Mean Akaike Information Criterion (AIC) curve with standard deviation intervals (n=36). e) Visualization of neuronal weights at optimal dimensionality for animal shown in a), standardized across conditions for comparison. f) Aggregated neuronal weights for each animal at optimal dimensionality (n = [27711, 27589, 22354, 26388], respectively). 64

Figure 10. Three-Dimensional Low-Dimensional Component Space Models for Murine Isoflurane Dataset. a) Component activity traces for first three components for the same animal shown in Fig. 9a). b) Decomposed frequency content for components shown in Fig. 10a). c) Visualization of neuronal weights for single three-dimensional models for the same animal in Figs. 10a,b). d) Aggregated neuronal weights for all three-dimensional models, across animals and states (n = 10743). e) Plotted low-dimensional component space of components against each other, with encapsulated volume, for animal shown Fig. 10a,b,c). f) Distribution of volumes, across animals and states (n=36). g) Distribution of directions, across animals and

states (n=188928). h) Distributions of speeds across animals and states (n = 188928) 68

Figure 11. Constrained Nonnegative Matrix Optimization Proof of Constrained Optimization. a) Initialized 25x25 random matrix before optimization. b) Same 25x25 matrix as shown in a), after optimization. c) An initialized binary matrix to constrain the matrix shown in a). d) Difference between weights, before and after optimization, split according to constraint status... 80

Figure 12. Spline construction of the synthetic spike-wave. (A) The waveform is divided into four segments: spike rise (w1), spike descent (w2), wave rise (w3), and wave descent (w4). The waveform amplitude at the transition between segments is characterized by parameters vpeak, vtr, vwave ; parameter tpeak indicates the timing of the spike peak. (B) Three examples of synthetic spike-waves (pink) generated using parameters based on in vivo observations (black); black scale bars below the spike peak indicate 1 standardized unit; intervals are 1 second in duration. 91

Figure 13. Synthetic ripple generation procedure. (A) Schematic representation of the algorithm for generating synthetic ripples (see Methods). (B) Three examples of synthetic ripples (blue) generated based on parameters estimated from in vivo observations (black). Black traces show the in vivo recording bandpass filtered to 80–500 Hz. Panels show CWT scalograms of the in vivo and synthetic ripples; darker colors indicate larger magnitudes; the horizontal (vertical) axis indicates time (frequency); frequencies shown are between 80 and 500 Hz. Orange contour lines in scalograms indicate contours of fitted magnitude surface. Each trace is 200 ms in duration. 93

Figure 14. Event counts for the subject population. Number of spike ripple, spike, and background events for each LO subject (see legend). We test classifier performance on 6 subjects (labels S1, S2, … S6) and trained classifiers on all other subjects. The bar labeled S0 includes event counts for all other subjects (n=28)........................... 98

Figure 15. Synthetic spike ripples mimic in vivo spike ripples. Examples of (A) in vivo and (B) synthetic spike ripples. Scale bars indicate 100 ms (horizontal) and 5 z (vertical)... 103

Figure 16. Data augmentation improves performance of a frequency-domain spike ripple detector. Receiver operator characteristic (ROC) curves with training data consisting of (A) in vivo data alone, (B) synthetic data alone, or (C) in vivo data augmented with synthetic data. Colored curves indicate individual subjects; see legend in (A). Black curves indicate the results for classifications across all subjects................. 107

Figure 17. Data augmentation marginally improves performance of a time-domain method. Receiver operator characteristic (ROC) curves with training data consisting of (A) no tuning, (B) in vivo data alone, (C) synthetic data alone, or (D) in vivo data

after augmentation with synthetic data. Colored curves indicate individual subjects; see legend in (A). Black curves indicate the results for classifications across all subjects. 113

Figure 18. Examples of misclassified observations. Example events correctly classified with training on in vivo data alone (training protocol 1 or TP1) and incorrectly classified with training on synthetic data alone (training protocol 2 or TP2). Blue indicates raw signals; orange indicates band-pass filtered signals (80–500 Hz), scaled by a factor of 5 for visibility. Scale bars (top left) indicate 100 ms (horizontal axis) and 5 z (vertical axis). Examples are given for (A) the frequency-domain detector, (B) the time-domain detector, and (C) when both detectors agree. 114

LIST OF ABBREVIATIONS

AIC .... Akaike Information Criterion

AUC .... Area Under the Curve

CNMF .... Constrained Nonnegative Matrix Factorization

CNN .... Convolutional Neural Network

CWT .... Continuous Wavelet Transform

DR .... Dimensionality Reduction

EEG .... Electroencephalogram

EZ .... Epileptogenic Zone

FIR .... Finite Impulse Response

fMRI .... functional Magnetic Resonance Imaging

FN .... False Negative

FOV .... Field of View

FP .... False Positive

GECI .... Genetically Encoded Calcium Indicator

ICA .... Independent Component Analysis

LOO .... Leave-One-Out

LSTM .... Long Short-Term Memory

ML .... Machine Learning

NMF .... Nonnegative Matrix Factorization

NPV .... Negative Predictive Value

PC .... Principal Component

PCA .... Principal Component Analysis

PCHIP .... Piecewise Cubic Hermite Interpolating Polynomial

PPV .... Positive Predictive Value

RNN .... Recurrent Neural Network

ROC .... Receiver Operating Characteristic

ROIs .... Regions of Interest

SMOTE .... Synthetic Minority Over-sampling Technique

SNR .... Signal-to-Noise Ratio

SOZ .... Seizure Onset Zone

SPW-R .... Sharp Wave Ripple

SVD .... Singular Value Decomposition

TN .... True Negative

TP .... True Positive

UMAP .... Uniform Manifold Approximation and Projection

CHAPTER ONE: Introduction

Introduction

Recent advancements in neuronal recording technologies have led to an unprecedented expansion in both the quantity and complexity of neuroscientific data[1]. Previously developed, and continuously improving, methods such as calcium imaging, electrophysiological arrays, multi-photon microscopy, electroencephalography (EEG), and functional magnetic-resonance imaging (fMRI) now enable researchers to simultaneously capture neural activity across scales ranging from single cells to large-scale brain networks, with spatial resolutions varying from cellular precision to whole-brain coverage, and temporal resolutions extending from milliseconds to extended experimental observational periods[2]. These diverse methodologies have fundamentally transformed neuroscientific research, facilitating more detailed investigations into brain dynamics previously unattainable.

Among these methodologies, calcium imaging has emerged as particularly influential in cellular-level neuroscience[3]. Utilizing genetically encoded calcium indicators (GECIs), such as the GCaMP family, calcium imaging links neuronal activity to fluorescent signals measurable via optical microscopy[4,5], allowing for detailed observation of neuronal populations in single brain regions[6]. Similarly, multi-electrode electrophysiological arrays have enabled precise measurement of neuronal firing patterns, supporting high-resolution analysis of neuronal activity[7]. EEG, on the other hand, offers non-invasive, whole-brain measurement of neural oscillations and biomarkers, providing critical insights into cognitive and pathological brain states over extended periods and across diverse clinical populations[8].

In addition to the increase in neuronal and temporal data points, contemporary neuroscientific research frequently integrates multiple experimental contexts into single studies, further compounding dataset complexity. Researchers routinely explore neural responses across varied conditions including behavioral states[9], sensory stimulations[10,11], pharmacological interventions[12], genetic manipulations[13], and clinical diagnoses[14,15]. Each experimental condition can significantly alter activity patterns, exponentially increasing dataset dimensionality and complexity.

Moreover, the scope of brain regions investigated has vastly broadened. Initially confined to a few accessible areas, current research can span intricate structures, distributed networks, and multiple interconnected brain regions, simultaneously. Technologies such as multi-photon microscopy allow visualization of previously inaccessible deep structures, including hippocampal layers, basal ganglia nuclei, and midbrain circuits[16]. Concurrently, EEG methodologies capture synchronized activity patterns across extensive cortical areas[17]. This broad anatomical coverage substantially increases dataset dimensionality, requiring advanced analytical tools capable of dissecting both localized neuronal dynamics and large-scale network interactions.

Sample sizes within neuroscience studies have likewise increased dramatically, driven by the need for greater statistical robustness and generalizability. Earlier studies often involved limited subjects, whereas contemporary research typically includes larger cohorts spanning multiple experimental groups, genetic backgrounds, or developmental stages. This escalation in sample sizes not only linearly increases data volume but also introduces substantial inter-subject variability, necessitating sophisticated analytical methods capable of identifying robust and biologically meaningful neuronal patterns across individuals and conditions.

Together, these technological and methodological advancements contribute to an exponential growth in neuroscience datasets[1]. Each additional layer of complexity, whether derived from expanded experimental conditions, broader anatomical coverage, or larger subject pools, amplifies the analytical challenges confronting neuroscientists today. Traditional analytical approaches, such as basic statistical summaries or simple averaging techniques, are inadequate for extracting meaningful insights from these extensive, multi-dimensional datasets[18]. Consequently, there is a critical demand for advanced computational frameworks that systematically reduce dimensionality, identify meaningful neuronal patters of activity, and elucidate the dynamics underlying activity across diverse experimental contexts.

Challenges with Traditional Analytical Approaches

Traditional analytical methodologies, characterized primarily by descriptive statistical techniques (mean, median, variance calculations, and spike averaging), face inherent limitations in scaling effectively to handle increasing variant information conferred by expansive neuroscience datasets. Although computationally straightforward, these methods generally aggregate across neurons, time points, or both, obscuring essential fine-scale temporal dynamics and nuanced interactions among neuronal populations. For example, averaging neuronal responses over trials or cell populations inherently discards crucial variance information essential for interpreting dynamic interactions within network activity[19]. These approaches typically simplify complex datasets into singular or aggregate metrics, which significantly reduces the potential insights gained from high-dimensional data. Consequently, traditional analytical strategies struggle to identify subtle yet critical biological phenomena, such as transient neuronal synchronization events, oscillatory dynamics, or state-dependent modulation of neural circuits. Moreover, they lack the flexibility to adapt to diverse experimental designs and large-scale multi-condition datasets, often resulting in overly generalized conclusions or overlooked biological nuances[19].

Machine Learning in Neuroscience

The integration of advanced machine learning (ML) methodologies has fundamentally transformed neuroscience research, profoundly enhancing our capacity to analyze, interpret, and predict intricate neuronal phenomena from large-scale datasets. ML encompasses a wide array of computational frameworks, including deep neural networks, convolutional neural networks (CNNs), recurrent neural networks (RNNs), and more classical supervised and unsupervised learning algorithms[20], each suited to specific analytical tasks within neuroscience. Deep learning methods, such as CNNs, have proven exceptionally adept at identifying complex spatial features in imaging data, providing nuanced classifications of pathological states based solely on imaging signatures[21]. Similarly, RNNs, particularly Long Short-Term Memory (LSTM) models[22], have been effectively utilized for time-series data analysis[23,24], demonstrating superior performance in detecting subtle and transient neuronal events such as spike ripples[25,26] in epilepsy directly from raw electrophysiological signals.

The predictive capabilities inherent in these advanced ML frameworks allow researchers to systematically uncover latent network dynamics, anticipate neuronal activity under varying experimental conditions, and facilitate real-time decision-making in experimental and clinical neuroscientific settings. Additionally, ML techniques enable the extraction of previously hidden neuronal features, elucidating underlying biological processes that traditional analytical methods often overlook. These methodologies leverage computational power and sophisticated algorithmic structures, thereby significantly advancing neuroscientific research beyond the limitations of conventional analytical paradigms.

Amidst the proliferation of large-scale, high-dimensional neuronal datasets, dimensionality reduction (DR) methods have emerged as indispensable analytical tools, offering essential strategies for systematically extracting meaningful, interpretable, low-dimensional representations[10,18,27–30]. These methods significantly mitigate the challenges posed by extensive neuronal datasets by simplifying the complexity of recorded data into components that retain maximal variance, thus preserving crucial biological signals while discarding extraneous noise[31].

Principal Component Analysis (PCA)[32], Independent Component Analysis (ICA)[32], Uniform Manifold Approximation and Projection (UMAP)[33], and notably Nonnegative Matrix Factorization (NMF)[34] represent some of the most prominently employed DR techniques in neuroscience research[18]. Each approach offers distinct advantages derived from their inherent mathematical constraints and operational characteristics. PCA identifies orthogonal directions in high-dimensional data that maximize variance, thereby simplifying complex datasets into uncorrelated components. Despite PCA's wide adoption and computational efficiency, the interpretability of its components can be challenging due to potential negative loadings[32], which do not align with inherently positive neuronal signals. ICA seeks statistically independent components[35], advantageous for disentangling mixed signals within neuronal populations. This technique excels in scenarios where neuronal sources are assumed independent but often yields components that may include biologically ambiguous negative values, complicating interpretability. UMAP, a more recent nonlinear manifold learning technique, effectively captures both global and local data structures, providing rich visualizations that reveal underlying neuronal state transitions[33]. However, its reliance on nonlinear embeddings often complicates quantitative biological interpretations, forgoing linear combinations in lieu of a more sophisticated lower-dimensional nonlinear embedding[18].

Among these methodologies, Nonnegative Matrix Factorization (NMF) stands out due to its inherent mathematical constraints, which mandate nonnegative decompositions, aligning with the intrinsically positive nature of neuronal calcium fluorescence[36], as opposed to the negativity inherent to neural voltage signals. The "parts-based" representation characteristic of NMF[37] facilitates direct biological interpretation by decomposing data into additive, biologically plausible neuronal assemblies or subnetworks. Unlike PCA or ICA, which utilize orthogonality or statistical independence constraints, NMF explicitly models neuronal activity as positive combinations of latent neuronal assemblies, enhancing the biological interpretability and practical applicability of derived components.

Specifically, NMF preserves biological realism through strictly nonnegative component weights, preventing biologically implausible negative neuronal contributions. This enhances interpretability by presenting additive combinations of neuronal sub-networks, providing straightforward mappings to possible biological mechanisms[18]. Moreover, NMF robustly models latent neuronal dynamics across varying experimental conditions, maintaining high explanatory power and interpretability across diverse neuronal datasets. Recent studies have demonstrated the superior performance of NMF in capturing essential organizational structures in simulated and empirical neuronal networks[18,29,30], accurately reconstructing known network architectures and dynamics with high fidelity. Such methodological rigor and robustness reinforce the potential of NMF as a pivotal analytical tool for addressing the intricate challenges inherent in contemporary neuroscience research.

We provide a detailed review of state of the art analyses, for the respective recording modality, at the outset of each chapter, using it as a foundation upon which to build our subsequent analyses. Each use case is contextualized within existing methodologies to highlight both advancements and limitations addressed in this work.

Technical Barriers and the Need for Standardization in Neuroscientific Analysis

Despite their significant analytical power, sophisticated techniques such as DR and ML methods face substantial technical and practical barriers that may hinder their widespread adoption within the neuroscience community[38]. These methods frequently involve complex computational and statistical challenges, necessitating extensive specialized knowledge and technical expertise that may not be readily accessible to many experimental neuroscientists.

Additional significant challenges include ensuring reproducibility, achieving standardization of analytical pipelines, and maintaining computational efficiency. Variability in analytical pipelines, inconsistent methodologies, and differences in software implementations often lead to limited generalizability and considerable difficulties in comparing results across studies, hampering the accumulation of knowledge and consensus within the field[39].

To effectively address these barriers, there is an urgent need for the development of accessible, reproducible, and standardized computational frameworks. Existing analytical toolkits commonly lack comprehensive integration, modularity, or user-friendly implementations, limiting their practical utility. Consequently, creating standardized and intuitive analytical toolboxes is essential to democratizing advanced computational methodologies. Such frameworks would significantly lower entry barriers, thereby facilitating widespread adoption and enabling researchers from diverse computational backgrounds to leverage sophisticated analytical methods effectively in their research.

Motivation, Objectives, Outline, and Implications of This Thesis

The overarching motivation driving this dissertation stems from a critical gap observed in contemporary neuroscience research: the discrepancy between the theoretical capabilities of advanced computational methodologies and their practical implementation and usability within the broader neuroscience community. Despite significant advancements in computational neuroscience, many powerful analytical techniques, particularly dimensionality reduction (DR) methods such as Nonnegative Matrix Factorization (NMF), Principal Component Analysis (PCA), Independent Component Analysis (ICA), and Uniform Manifold Approximation and Projection (UMAP), remain underutilized due to substantial barriers in technical complexity, limited validation regarding applicability, and limited availability of user-friendly analytical platforms. The core of this thesis specifically addresses these barriers by introducing and rigorously validating these methods as tools for analyzing neuronal datasets, and building them into accessible, modular, and comprehensive open-source software toolbox named Calcium Network dynamiCs (CaNetiCs).

CaNetiCs is carefully engineered to bridge the practical divide, facilitating a widespread implementation of sophisticated analytical methodologies within diverse neuroscientific research contexts. At its core, CaNetiCs emphasizes dimensionality reduction techniques, prominently featuring NMF due to its superior biological interpretability for calcium imaging, nonnegativity constraints, and ability to produce meaningful neuronal subnetwork decompositions[18,29,34,36]. Complementing NMF, the toolbox integrates PCA, ICA, and UMAP, providing a robust suite of analytical options suitable for varied research questions and data characteristics. Furthermore, CaNetiCs includes dedicated modules for geometric state-space analyses and neuronal network simulation, enabling researchers to not only analyze empirical data but also systematically benchmark and validate computational methods through controlled simulations. This integrated approach substantially enhances methodological rigor, interpretability, and replicability within neuroscience research.

A central objective of this thesis is not merely methodological advancement but also the empowerment of hypothesis-driven neuroscience. While many computational tools offer exploratory capabilities, CaNetiCs also supports systematic, hypothesis-oriented experimental designs. By allowing neuroscientists to test hypotheses concerning state-dependent neuronal, the toolbox significantly enriches neuroscientific inquiry. For example, researchers investigating neuronal synchronization during experimental modulation under unique contexts can utilize CaNetiCs to precisely dissect relevant neuronal subnetworks and identify differential changes. This systematic integration of computational rigor with experimental hypothesis testing amplifies the depth, precision, and interpretability of neuroscientific findings, thereby advancing the broader understanding of brain dynamics.

The dissertation is structured to build upon the foundational computational framework articulated here. Chapter Two assesses the efficacy of NMF in capturing neuronal network dynamics, explicitly comparing its performance with alternative DR techniques (PCA, ICA, UMAP) using both simulated datasets and empirical calcium imaging data. Detailed analyses establish the superior interpretability and robustness of NMF, highlighting its strengths in identifying biologically relevant neuronal subnetworks and reconstructing known neuronal dynamics with exceptional fidelity.

Chapter Three elaborates on the conceptualization, technical development, and extensive validation of the CaNetiCs toolbox. Through application to diverse experimental datasets, including whole-ganglion imaging in C. elegans and two-photon imaging of murine somatosensory cortex across anesthetic states, the toolbox demonstrates remarkable versatility, accuracy, and ease of use. This comprehensive validation underscores the generalizability of CaNetiCs across distinct experimental paradigms, species, and neuroscientific questions, reinforcing its potential as a standardized analytical platform.

Chapter Four details an orthogonal computational analytical capability by presenting innovative data augmentation techniques aimed at enhancing the detection of subtle neuronal events using machine learning approaches, specifically Long Short-Term Memory (LSTM) neural networks. Employing sophisticated synthetic data generation techniques informed by empirical distributions and expert-driven parameter selections, this chapter significantly advances event detection accuracy and reliability in neuroscientific datasets characterized by limited or sparse ground truth data.

Collectively, the methodologies and tools developed throughout this dissertation embody a substantial advancement in computational neuroscience. By systematically addressing critical technical barriers, promoting standardization, and enhancing interpretability, we aim to significantly democratize sophisticated analytical methods. This democratization not only facilitates deeper exploration and understanding of complex neuronal dynamics but also fosters interdisciplinary collaboration, thereby driving innovative neuroscientific discoveries. Ultimately, this thesis contributes foundational computational infrastructure and methodological rigor to empower the neuroscience community, enhancing collective capacity to elucidate neuronal function, dysfunction, and plasticity across a multitude of experimental conditions and biological scales.

CHAPTER TWO: Nonnegative Matrix Factorization for Analyzing State Dependent Neuronal Network Dynamics in Calcium Recordings

Adapted from Carbonero D, Noueihed J, Kramer MA, White JA. Nonnegative Matrix Factorization for Analyzing State Dependent Neuronal Network Dynamics in Calcium Recordings. Scientific Reports (2024); 14(1), 27899. https://doi.org/10.1038/s41598-024-78448-6

Minor changes in Figures 5 and 6 for clarity

Abstract

Calcium imaging allows recording from hundreds of neurons in vivo with the ability to resolve single cell activity. Evaluating and analyzing neuronal responses, while also considering all dimensions of the data set to make specific conclusions, is extremely difficult. Often, descriptive statistics are used to analyze these forms of data. These analyses, however, remove variance by averaging the responses of single neurons across recording sessions, or across combinations of neurons, to create single quantitative metrics, losing the temporal dynamics of neuronal activity, and their responses relative to each other. Dimensionally Reduction (DR) methods serve as a good foundation for these analyses because they reduce the dimensions of the data into components, while still maintaining the variance. Nonnegative Matrix Factorization (NMF) is an especially promising DR analysis method for analyzing activity recorded in calcium imaging because of its mathematical constraints, which include positivity and linearity. We adapt NMF for our analyses and compare its performance to alternative dimensionality reduction methods on both artificial and in vivo data. We find that NMF is well-suited for

analyzing calcium imaging recordings, accurately capturing the underlying dynamics of the data, and outperforming alternative methods in common use.

Introduction

Recent advances in neural recording methods have made it possible to collect large, complex data sets that can be used to study context-dependent neural dynamics in great detail [1,40]. For example, calcium imaging allows simultaneous recording from hundreds of neurons in a single brain region *in vivo*, while maintaining single-cell resolution [5,6,16,41,42]. Other tools allow cell "tagging" based on factors, including activity during a specified time window and molecular genotype [43–46]. The resultant data set consists of several hundred time series thousands of points long, multiplicatively increasing per experimental condition added, in addition to any tagging data. The process of characterizing such single-cells responses, relating them to those of other cells across time, and accounting for the tagging information, all while considering changes in animal behavior, is a difficult but essential task.

Basic descriptive statistics (mean, median, variance) are a common approach to analyze large neural datasets. However, these approaches are limited because they average over the temporal dynamics of neuronal activity and can obscure interactions among neurons at fine time scales [31,47]. Dimensionality reduction (DR) is a more sophisticated approach that can help to reveal low-dimensional dynamics by transforming high-dimensional data into interpretable, low-dimensional components [20]. Compressing the data in this way can preserve temporal patterns and correlations and focus subsequent analysis on the dynamical patterns of interest.

DR approaches have many applications in the analysis of neural data, including analysis of electrophysiological recordings [31,48,49], automation of trace extraction of neuronal activity in fluorescent recordings [36,40,50–57], and as a pre-processing step to more complex behavioral decoding [58]. To analyze network dynamics in calcium recordings, nonlinear DR methods have emerged to map dynamics to low dimensional manifolds [28,59–61]. While excellent for very low dimensional visualization [62,63], the dynamics are difficult to interpret beyond the manifold being modeled, with manifolds being difficult to interpret beyond three dimensions. Further, dynamical nonlinear methods have been applied to recover neural trajectories as a function of latent variables[64]. Linear DR methods have likewise been employed to either statically cluster cells [29,30,36,65], or to abstractly analyze activity in low-dimensional subspaces [66–69].

While each methods have their own strengths as a function of different low dimensional representations, we aim to decompose our data into a lower dimensional series of neuronal sub-networks and their activities, to detect how neurons form flexible, correlated cell assemblies under different experimental contexts. We argue more generalized linear DR methods will provide a more interpretable representation of the data, at the cost of a few more dimensions than nonlinear methods, dynamical or non-dynamical, as a function of having fewer tunable parameters and simpler mathematical decompositions. We further argue Nonnegative Matrix Factorization (NMF) is an especially promising linear decomposition method for analysis of activity recorded in calcium imaging in this context. NMF operates under the constraint that every element in the data, and the decomposition, be nonnegative [34], as is naturally true for neuronal calcium traces [36]. Further, the positivity constraint, coupled with the linearity of the reduced dimensions, ensures that each of the dimensions sum to form a low dimensional representation of the data, simplifying the interpretation of the data[37]. In neuroscience, specifically in the context of calcium imaging, NMF has emerged as the state of the art method to estimate single neuron activity from fluorescent recordings [50–52,54,56,57,70], demonstrating efficacy signal de-mixing and functional imaging activity extraction[55]. However for analyzing the dynamics of recorded neurons, it has been applied sparingly, either for quantifying single cell pairwise relationships [71] or statically clustering cells under a series of strict assumptions [29]. Here, we develop an application of NMF to analyze holistic network activity (Fig. 1), comparing it to the most widely employed methods that provide a similar decomposition, capturing how precise sub-networks of neuronal dynamics evolve in different experimental contexts. To assess performance of the NMF approach, we generate artificial calcium imaging data from simulated ground-truth networks and assess how well our approach captures the underlying network structures responsible for the emergent dynamics of the network. We then apply the same NMF framework to characterize the dynamics of a representative *in vivo* calcium imaging data set. We conclude that, compared to alternative methods, NMF best captures the ground truth networks that underlie the observed calcium dynamics and provides a useful low-dimensional description of complex physiological data.

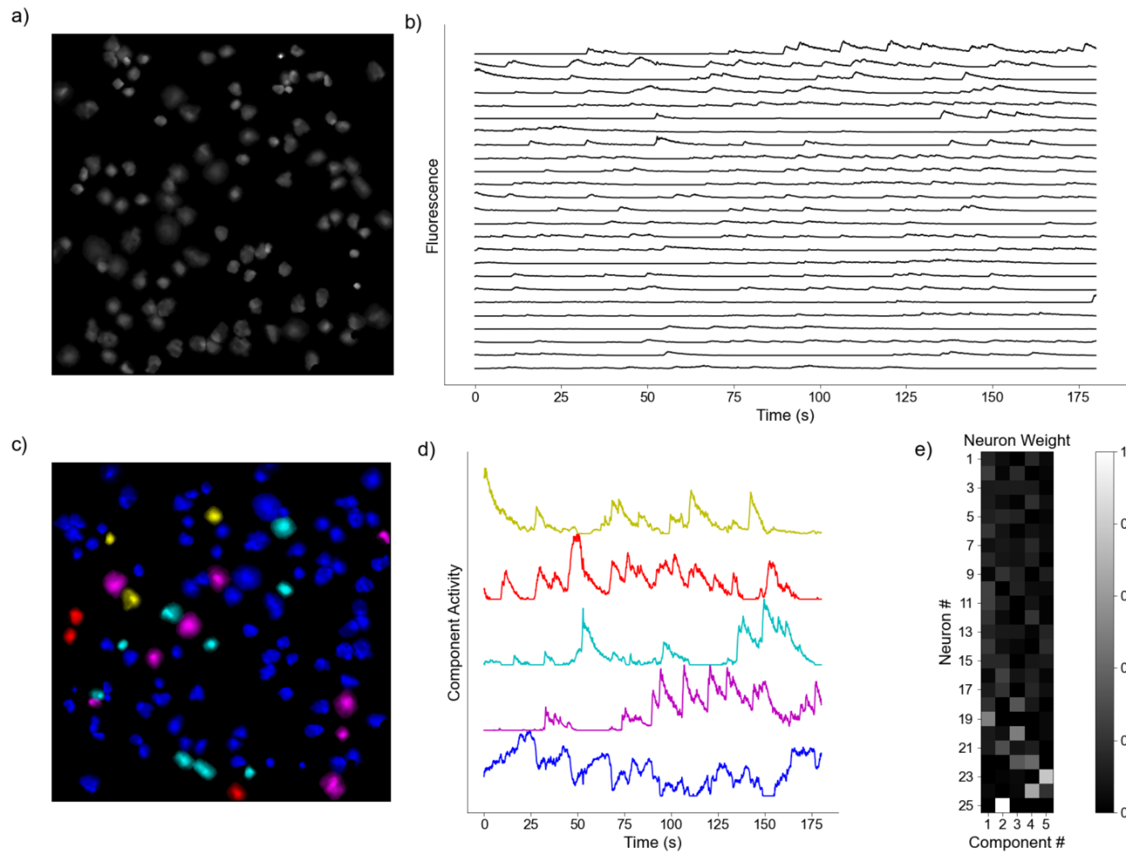

Figure 1. Representative Model Pipeline. a) Recording Field of View (FOV) with Regions of Interest (ROIs). b) Extracted fluorescent activity traces from ROIs. c) Same recording FOV shown in a) with ROIs colored according to largest component contribution. d) Decomposed activities of the components shown in c). e) Sample truncated heatmap showing neuronal contribution or weight for each component

Results

We test NMF against three existing DR approaches in common use: UMAP, PCA, and ICA. We do so using two different simulated neuronal network architectures to assess the relative performance of each DR approach, when the network structure is known. We first build a simple network with five groups of independent neurons, each with 20 neurons, in which the neurons in each independent node are perfectly connected (Fig. 2). In this architecture, a spike in one neuron in a node will drive all other neurons in that node to spike. We use such an architecture to test if, and how well, models can identify the five nodes given traces of simulated activity.

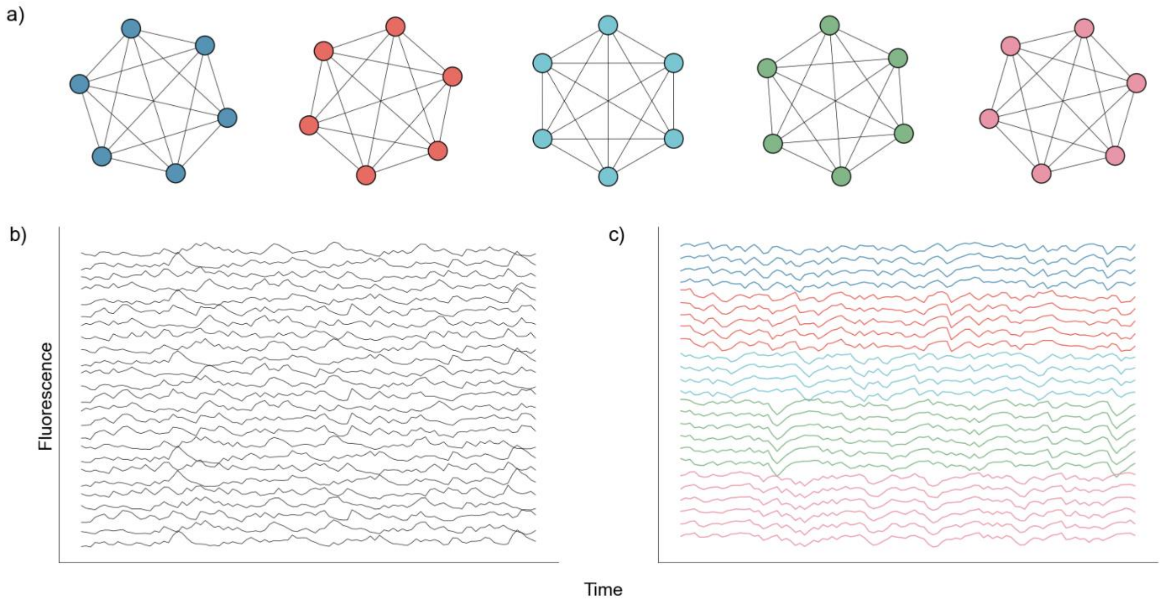

Figure 2. Nodal Network Architecture. a) Diagram showing perfect intraconnection between neurons and independence between nodes. b) Randomly selected calcium traces generated by the network. c) Same randomly selected calcium traces shown in b), but sorted according to node

To further evaluate the performance between DR methods, we examine a more intricate network architecture aligned with our spontaneous calcium recordings. In doing so, we create five distinct random processes driving 150 neurons at varying degrees of strength (Fig. 3). This approach better reflects the dynamics we aim to capture, encompassing underlying activity patterns in diverse contexts driven by weighted sub-networks of neurons. Utilizing this surrogate architecture, our objective is to assess how effectively the models can extract both the underlying activity patterns, and the corresponding contributions by individual neurons, given a clear end target for performance assessment. We present two different analyses for both simulated network architectures. We both

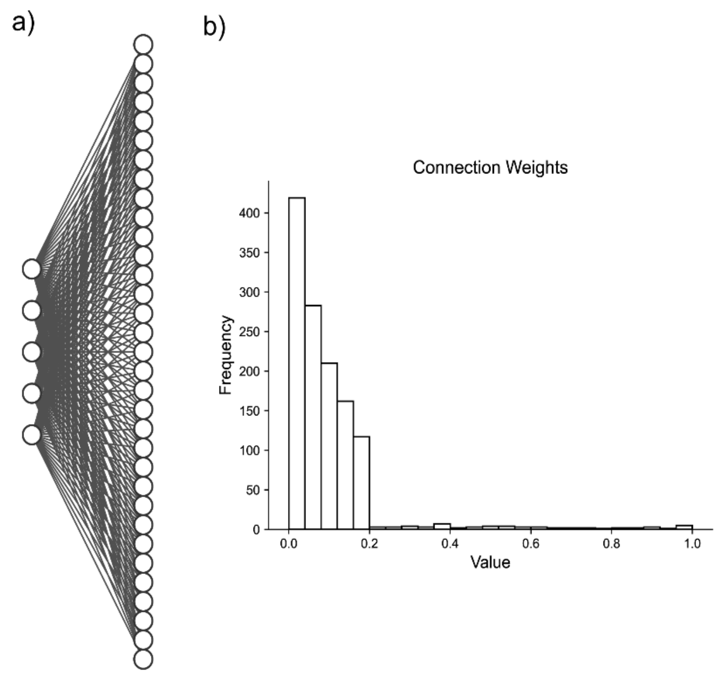

Figure 3. Artificial Random Propagating Process Network Architecture. a) Sample connection diagram showing all random processes projecting to all neurons. b) Single histogram of generated connections between processes and neurons from sampled distribution

perform a qualitative analysis for a single generated network, and a more scaled quantitative approach for a series of 256 random instantiations and simulations of the networks. We find that NMF best captures the dynamics of the simulated networks, with components summing to provide an accurate macroscopic low dimensional representation of the data. We argue the parts-based representation provided inherently by NMF supports the improved performance for these data and the goal of interpretably extracting low dimensional neuronal network dynamics.

Perfectly Intraconnected, Independent Nodes

We generate the nodal networks and fit each of the DR methods to the simulated calcium traces. Initially, we evaluate the average variance explained and the AIC across 256 networks to gauge how well the models capture the original data at different lower dimensionality levels. Both NMF and PCA exhibit similar variance explanation patterns across all components, with the average variance explained stabilizing at five components

(Fig. 4a, left axis), as required to capture the dynamics of simulated activity in five nodes. This stabilization is influenced by the introduction of unique noise increments in each trace (Eq. 13 & Eq. 14). Additionally, the AIC for the NMF method consistently minimizes at five components across all 256 model instantiations, effectively optimizing for the correct number of nodes (Fig. 4a, right axis). We note that the shape of the information criterion curve takes an inverse of the NMF variance explained curve until the inflection point, reflecting the inherent relationship between these two measures, arising from the AIC's balance between model fit and complexity, contrasting with the NMF's focus on maximizing explained variance.

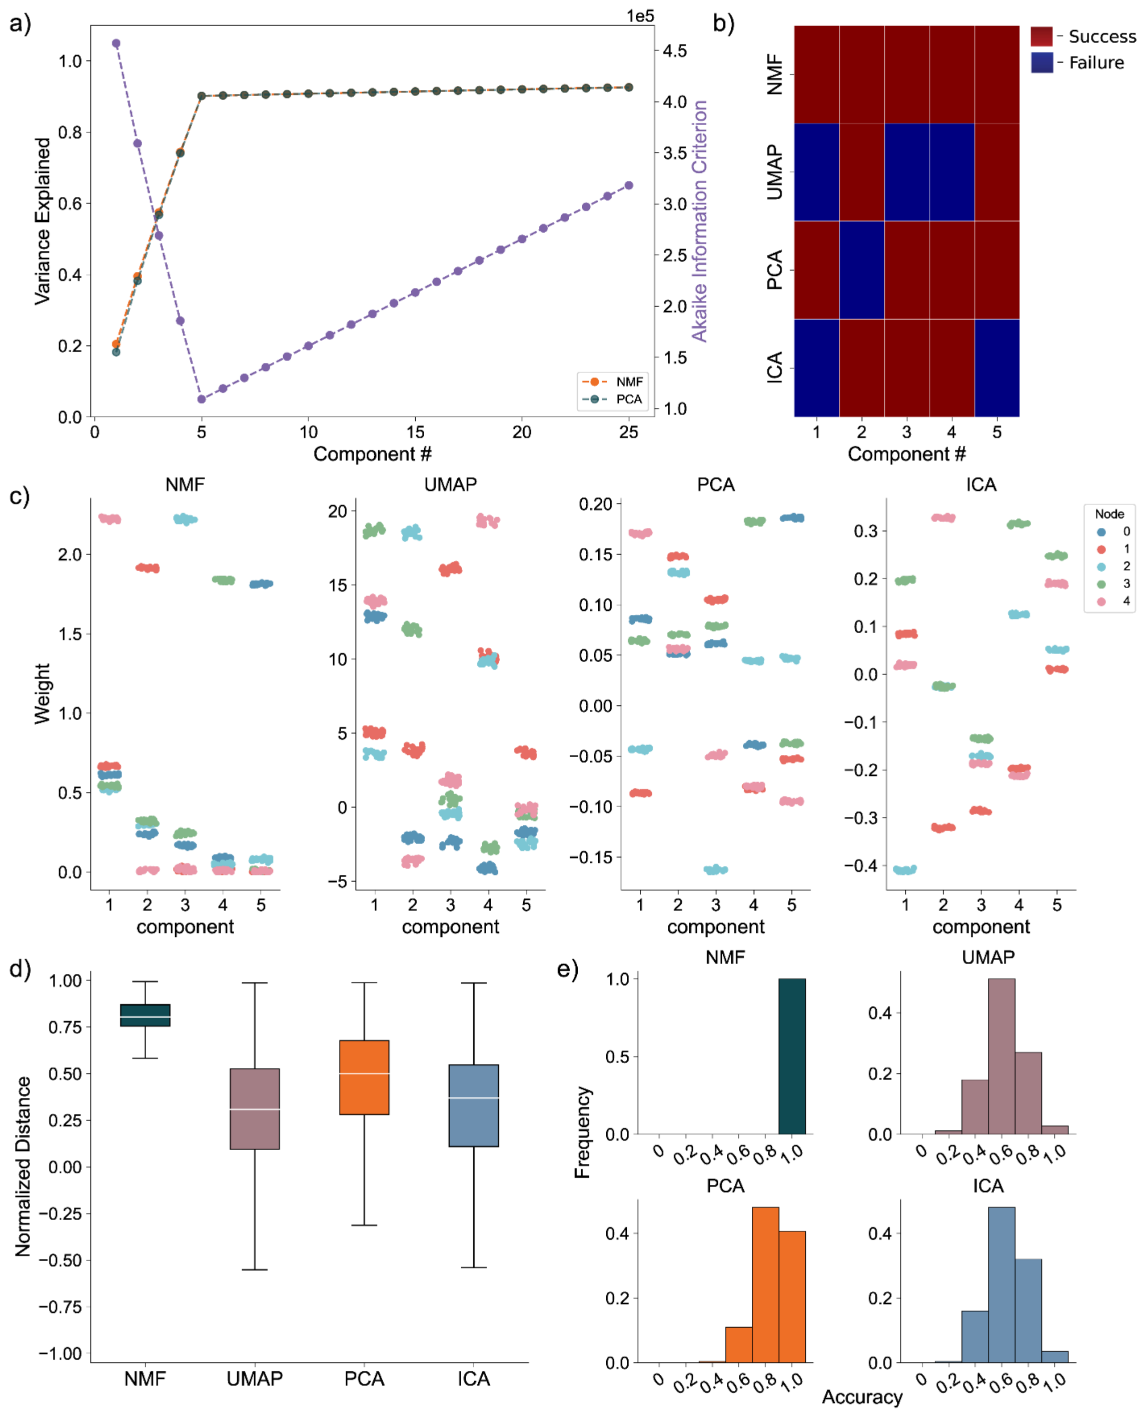

Figure 4. Artificial Nodal Network Analysis. a) Variance Explained for NMF and PCA, left axis (n = 256 networks) per component added to the model, and Akaike Information Criterion (AIC) for NMF, right axis (n = 256 networks). b) Prediction visualization between nodes and components for single network. c) Neuronal weights (y axis) for each model (panel) for each node (color) for each component (x axis). d) Average between neurons inside and outside predicted nodes for each component (n = 256 networks). e) Model assignment accuracy (n = 256 networks)

To evaluate how the different approaches capture the independent nodes of activity as components, we devised a discrete metric to determine success or failure of individual node assignment to components, detailed in the methods, for each architecture. We show the assignment success for a single model for each of the detailed DR methods (Fig. 4b). We find the best performance with the NMF method, which successfully maps each node to a component, followed by PCA, ICA, and UMAP. We repeat this process for 256 replicate models (Fig. 4e) and find that NMF always maps each node to a separate component, followed by PCA which performs with a mean accuracy of 0.8578, and then ICA and UMAP which perform with mean accuracies 0.6445 and 0.6242, respectively, and typically fail to successfully map the entire model.

We then visualize how each of the DR methods map the nodal patterns of activity to each of the components. We do so by showing the weight for each neuron, for each component, for each method, delimited by node (Fig. 4c), for a single instantiation of the simulated neural activity. In this example, we find a very clear separation of weights from each node for the NMF method; each component assigns weights to the neurons within a node. We argue this separation is a result of the parts-based representation that NMF inherently embodies [37]. In a parts-based representation, each component contributes an additive combination of components to faithfully represent the original data. This means that NMF decomposes the overall activity into identifiable parts, enabling a clear separation of weights for each node. Conversely, the PCA and ICA methods produce distributions of node weights with less clear separations within a component. While the first 5 components of the PCA method each identify neurons in each node, the constraints of the method result in counterbalancing variables that obfuscate the underlying input resulting in the data, hindering performance [32]. UMAP poorly separates nodes into five components because a strength of UMAP, and manifold learning in general (e.g., TSNE), is visualization in very low dimensions (restricted to two or three dimensions) [62,63].

To provide a quantitative assessment of model performance, we calculate the average displacement between neuronal weights assigned to a component and all other weights within the same component, for all components of each of the 256 instantiations of the model (Fig. 4d). We do so to provide a quantitative and continuous measurement of model performance (separation between weights in node and not in node, at assigned component) beyond binary success/failure, with the magnitude indicating the degree of separation between the neurons in the assigned node and all other neurons within a given component. Normalized per model, a positive value describes an assigned node whose average weight was higher than the weights of the other node in that component. A negative value would then describe the opposite, with the magnitude of the displacement describing the degree of difference between the weights. Our findings reveal NMF best separates nodes into components, followed by PCA, UMAP, and ICA. While this supports our finding that NMF specifically assigns patterns of nodal activity to components, the other methods do tend to separate a majority of the nodes of neurons individually in each component (the normalized distance tends to exceed 0), with the differing representations being attributable to the different mathematical constraining during fitting (UMAP — very low dimensional manifold, PCA — orthogonality, ICA — independence). However, as we show in our next case, this structure breaks down on the introduction of more complicated data.

Random Process Propagation

We simulate a network with a more complicated architecture to further evaluate model performance. This architecture involves generating five underlying random processes to drive the network activity. Each underlying random process provides input to each of 150 neurons, with a weight randomly sampled from a generated exponential distribution (Fig. 3). By analyzing these artificial networks, we aim to probe how well NMF maps the random process by observations from the received neuronal inputs to components. We similarly aim to compare the performance and low dimensional representation of NMF and the other methods we implement.

Analyzing the average variance explained for all 256 uniquely seeded networks of this architecture, we find that NMF explains significantly more at all components (Fig. 5a, left axis). We argue this is a function of the difference between optimizations. Rather than having a strict orthogonality constraint between dimensions as in PCA, NMF aims to reconstruct the data with as little error as possible into the specified number of dimensions, allowing for explaining more variance with fewer components. The AIC shows a minimum at five components, the number of random processes used to generate the network (Fig. 5a, right axis), seeing a similar inverted relationship between AIC and variance explained.

We quantify assignment success of this model differently from the nodal architecture, due to the more continuous nature of the data, detailed in the methods. For or an example model instantiation (Fig. 5b), only the NMF method successfully assigns each process to one component. Assessing the performance over 256 realizations of the network (Fig 5f), we find that NMF successfully assigns each random process to a component. Consistent with the results for the nodal architecture, the other methods (PCA, ICA, and UMAP) have reduced performance. We conclude that NMF reconstructs the latent network structure from the observed data, capturing the five underlying random processes in a parts-based representation. Beyond discrete prediction, we consider the correlation between the assigned connection probabilities from the latent processes to the observed neurons, and the neuronal weights of the predicted component. To start, we visualize the results for one instance of the NMF model, relating the neuronal weights of each component to the assigned outgoing connection probability from the assigned underlying random process to each neuron (Fig. 5c). We find that large neuronal weights in each component tend to occur for large connection probabilities for the predicted random process with a high correlation of 0.85 (Fig. 5d). Repeating this analysis for 256 realizations of the model, and regardless of prediction success, we find that NMF consistently infers components with high correlations between the connection probabilities and weights (Fig. 5e), and is always able to successfully map all underlying processes to separate components (Fig. 5f). The other methods have reduced performance, consistent with the results for the nodal architecture with PCA, ICA, and UMAP following in performance, respectively.

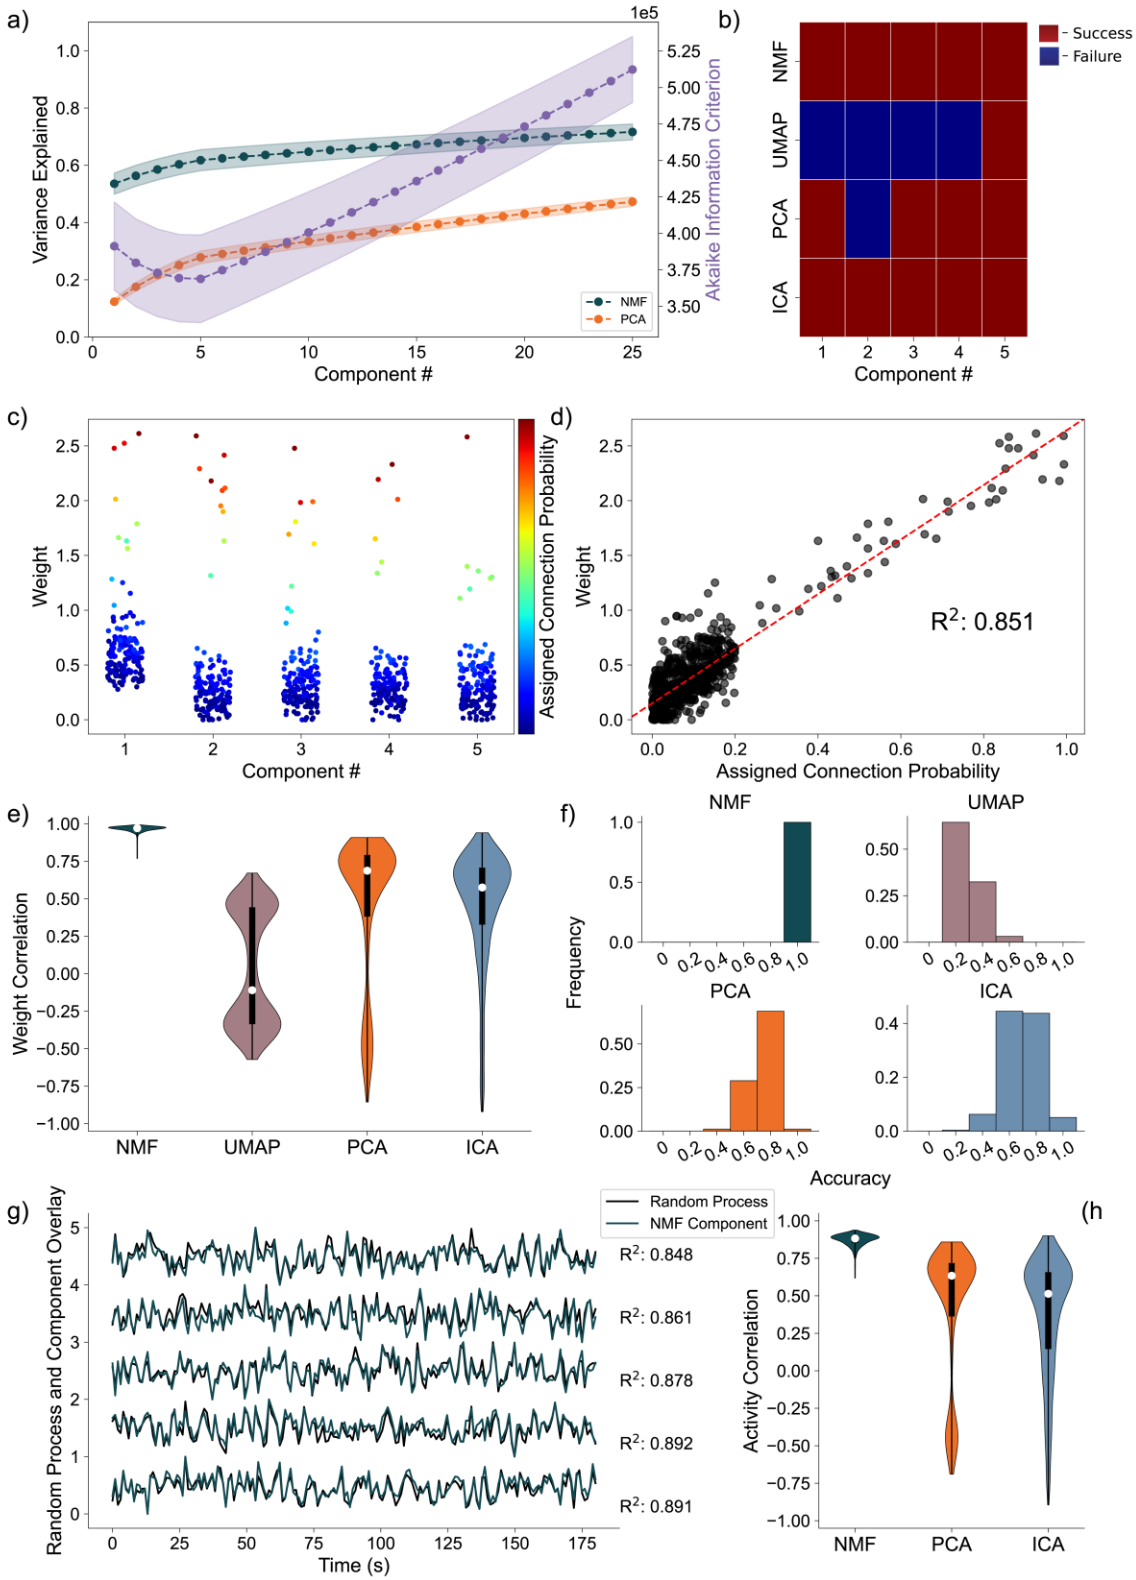
a)
Variance Explained
Component #
Akaike Information Criterion
1e5
NMF
PCA
b)
NMF
UMAP
PCA
ICA
Component #
- Success
- Failure
c)
Weight
Component #
Assigned Connection Probability
d)
Weight
Assigned Connection Probability
R^2: 0.851
e)
Weight Correlation
NMF
UMAP
PCA
ICA
f)
NMF
UMAP
PCA
ICA
Frequency
Accuracy
g)
Random Process and Component Overlay
Time (s)
Random Process
NMF Component
R^2: 0.848
R^2: 0.861
R^2: 0.878
R^2: 0.892
R^2: 0.891
(h
Activity Correlation
NMF
PCA
ICA

Figure 5. Artificial Random Propagating Process Network Analysis. a) Variance Explained for NMF and PCA, left axis (n = 256 networks) per component added to the model, and Akaike Information Criterion (AIC) for NMF, right axis (n = 256 networks). b) Assignment visualization between nodes and components for single network. c) Neuronal weights for NMF model (y axis) for each component (x axis) compared to assigned connection probabilities (color scale) for the predicted random process. d) Correlation between neuronal weights for NMF model (y axis) and the respective assigned connection probability (x axis). e) Average Correlation between neuronal weights for model and the respective assigned connection probabilities (n = 256 networks). f) Model assignment accuracy (n = 256 networks). g) Correlation between component activity from NMF model and assigned random process activity. h) Correlations between component activity from model and assigned random process activity (n = 256 networks)

We finally analyze how the activity of each component correlates with the activity of the underlying random process. Qualitatively, for a single realization of the model, the component activity captures the underlying processes very well (Fig. 5g) with correlations between the activities exceeding 0.85. Repeating this analysis across 256 realizations of the network, we once again find that NMF best captures the activity of the latent input processes, followed by PCA and ICA (Fig. 5h); we note that UMAP does not estimate the times series of the latent processes. To assess computational performance, we compare the median run time for the 256 models of this architecture. Because PCA is a well solved linear algebra problem (singular value decomposition of a matrix)[32], it is the shortest median run time of 0.0245 s. Of the three optimization methods, NMF and ICA have comparable median run times (0.7349 s and 0.7342 s, respectively. UMAP has the longest median run time of 1.4906 s, even after we exclude UMAP's rather high overhead for initialization. This discrepancy be attributed to model complexity (linearity of NMF and ICA, and non-linearity for UMAP).

We propose that the positivity and parts-based representation of the NMF method enable accurate reconstruction of the high-dimensional activity through low dimensional inputs. We propose that the constraints implemented by the PCA and ICA methods may obfuscate inputs to the observed neural network in its low dimensional representations. Finally, we interpret the poor performance of the UMAP method as indicating this method is better adapted for low dimensional visualization of less variant data. We conclude that NMF accurately reconstructs the latent inputs to a biophysically-motivated neuronal network that simulates calcium fluorescence recordings, despite multiple barriers to accurate identification (e.g., the underlying process is unobserved, with random connectivity to the observed neurons whose calcium signal is obfuscated with two types of noise). These results demonstrate NMF outperforms existing methods in common use to extract the underlying dynamics present in a series of neuronal activities recorded in calcium imaging.

Awake-State Data

As a final proof of concept, we analyze application of the DR methods to awake state calcium imaging recordings from 36 mice. In doing so we apply the same analytical framework as for the simulated data. Analysis of the variance explained for each model (Fig. 6a) shows similar performance between the NMF and PCA methods, with a small but highly significant trend towards increased values for the NMF method ($p < 10^{-6}$, Wilcoxon Signed-Rank test). This small advantage in increased variance explained may be due to the NMF method acting to capture the original data as quickly as possible into a set number of components (Eq. 2), rather than using a complex geometrical constraint to analyze the data[35]. We further find that NMF, using AIC, identifies an optimal number of components to describe the data in each of the recordings (Fig. 6b), with number of neurons helping inform the number of components at a correlation of 0.662. We qualitatively characterize an NMF model fit to a single recording, using 18 components as determined by AIC. When characterizing the neuronal weights estimated by the NMF method, we find behavior consistent with the simulated random process network (Fig. 6c); the NMF method results in in the neuronal weights with an approximate exponential distribution (Fig. 6d), with most neurons contributing little to the overall network activity, and few neurons very significantly contributing. We further briefly characterize the correlations between the activity for each decomposed component (Fig. 6e), finding that pairwise correlations are consistently low between them (Fig. 6f). Given that the correlations are low, and we capture a significant portion of the variance, we interpret these results to indicate that the NMF, and the parts-based nature of the model, successfully extracts unique predominant patterns of activity in the data. We conclude that, applied to these *in vivo* calcium imaging recordings, that the NMF method identifies components with distinct (uncoupled) dynamics.

We note the comparison between NMF and PCA for variance explained (Fig. 6a) is the only objective comparison possible here, because the in vivo dataset lacks ground truth, making further comparison between methods confusing and counterproductive.

While we demonstrate here that we successfully decompose statistically similar data from our simulations into statistically similar models, future work will relate inferred components and dynamics to experimentally elucidating biological mechanisms and behavior. We conclude that the NMF method is a promising tool to analyze neuronal network dynamics and identify meaningly sub-network activity via a relatively simple and interpretable approach.

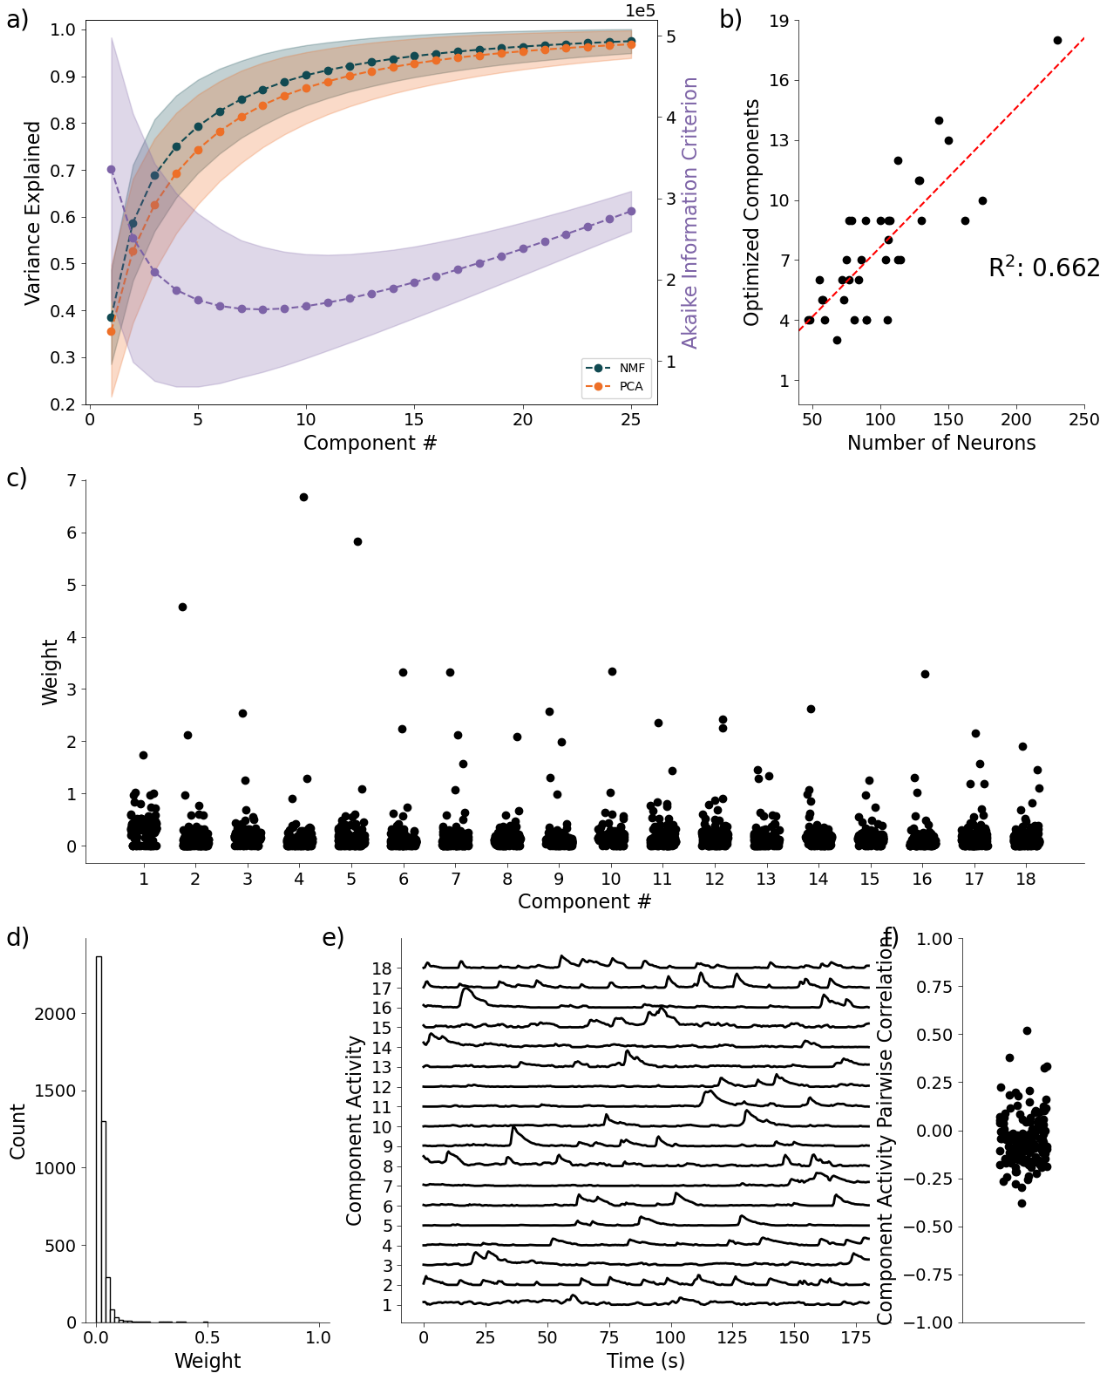

Figure 6. Baseline Awake State Network Analysis a) Variance Explained for NMF and PCA, left axis (n = 36 mice) per component added to the model, and Akaike Information Criterion (AIC) for NMF, right axis (n = 36 mice). b) Correlation between number of neurons (x axis) and number of optimized components (y axis) (n = 36 mice). c) Neuronal weights for NMF model (y axis) for each component (x axis) for a single mouse. d) Neuronal weight histogram from NMF model for the same mouse as c). e) Component activity for NMF model of the same mouse as in c) and d). f) Pairwise correlations between component activities shown in e)

Discussion

We have developed an analytical pipeline to more thoroughly model neuronal network dynamics with NMF, considering all information provided by the model in its low dimensional representation. We generate a series of calcium traces recorded from ground-truth artificial neuronal networks at two degrees of complexity to assess how well our analyses extract the ground-truth responsible for the simulated dynamics, and compare the performance of NMF to UMAP, PCA, and ICA. We find the NMF significantly performs the best, followed by PCA, ICA, and UMAP, respectively. We then apply our NMF pipeline to a series of in vivo baseline recordings and find that NMF confers a similar, and easily interpretable representation to the underlying random process network.

We argue this performance is a result of the "parts-based representation" conferred by NMF [37]. Because every single element of the model is necessarily positive, and the decomposition is aiming to recapture the original data, every component sums to give a low dimensional macroscopic representation of the original data. Interpreting the decomposition as the sum of its parts, rather than a complex cancelling of variables to achieve an algebraic constraint, as in PCA [32], results in NMF being able to provide a better representation of the dynamics. Further, the only fundamental assumption made in NMF is that the data can be represented exclusively as positive values. Our other methods require more stringent assumptions. UMAP makes the assumption that the data can be embedded on a low-dimensional manifold[33]. While very advantageous in less variant data (large scale trial averaged electrophysiological recordings[27]), we argue our data is far too variant with limited scalability (single trial, single recording) for this method to be as effective. ICA assumes that components are both non-normal, and independent [35]. While the assumption of non-normality of traces, variables, is a positive feature given that calcium traces are clearly non-normal data, we argue it is unfair to assume sub-network dynamics are completely independent of each other. Finally, PCA assumes the data is best described by the eigenvectors of the covariance matrix [32]. We note that this algebraic constraint does provide PCA the unique attribute of the model being the same for the first i components, regardless of the number of components fit. NMF, UMAP, AND ICA will differ in their representation of the data, dependent on the number of components fit too. As a result, PCA is guaranteed to perfectly lossless in an inverse transformation of the data, given all components. However, PCA is clearly less efficient than NMF in generating compact descriptions of data like ours.

As we mention in the introduction, a version of NMF, specifically Constrained Nonnegative Matrix Factorization (CNMF) has become nearly ubiquitous for automated trace extraction from calcium recordings[40,50,51,53,57], including a adaptions for microscopic endoscope data[54]. CNMF allows for the incorporation of constraints when there is *a priori* knowledge to build them. With our data, we have no a priori knowledge regarding connectivity, and we show that the unconstrained NMF method performs well. However, it would be feasible to modify our analytical pipeline to account for known constraints in other cases.

While clearly productive for analyzing local calcium recordings, NMF's nonnegativity assumption makes it well-suited for potentially analyzing large-scale anatomical connectome data or functional imaging experiments (given the data were nonnegative). It can identify patterns of co-activation or synchronization among brain regions, revealing the modular organization of brain activity and interactions between assemblies. In the context of anatomical connectome mapping, NMF could uncover patterns of structural connectivity that give rise to modules in the brain. However, the structure in functional imaging and anatomical connectome data is complex and multifaceted. In specific analysis cases such as these constraining towards a particular solution (CNMF) could be advantageous.

Hierarchical clustering has also emerged as a prominent method to analyze these large scale data[72]. While both hierarchical clustering and NMF aim to reveal underlying network structures, there are key differences between these approaches. Hierarchical clustering discretely assigns each neuron to single clusters based on activity similarity, while NMF reduce the data into components with each neuron having a weighted contribution to each component. This allows NMF to capture overlapping and graded membership of neurons in multiple sub-networks. Moreover, NMF's parts-based nature provides a more interpretable representation compared to that of static clustering, which can be more difficult to interpret

NMF's performance is dependent on the statistical characteristics we outline above, and different DR methods may emphasize different aspects of the data structure dependent on their assumptions and the analysis objectives, and. For example, PCA identifies orthogonal components capturing maximum variance, while ICA identifies statistically independent sources. These methods may capture different facets of the brain's functional and structural organization compared to NMF.

Similarly, inhibitory responses, where the fluorescence drops below the neuron's baseline fluorescence, have been observed in calcium imaging[73]. Here, where we focus on extracting the dominant patterns and underlying structure of the neuronal activity, the impact of these minor negative deviations in the fluorescence are statistically very minor. As a result, such minor effects can be lost amongst the other variability present in the activity, which can be captured by alternative statistical methods[73].

This pipeline was developed with the intent of analyzing state dependent neuronal network dynamics, and future work will analyze differential network dynamics that arise from unique experimental contexts. However, NMF, and other DR methods, require an underlying structure in the data and can be especially sensitive to potential extremes in the number of components and rates of activity in the network. An extreme sparsity or extreme saturation of events will lead to a much noisier representation in the model. Further, a fundamental assumption of DR is that the data can be represented with significantly fewer components than there are neurons [31]. Data that would carry hundreds of underlying drivers of activity (or any case where the number of patterns approach the number of neurons) would be obfuscated in a model such as this. Deep learning, on the other hand, requires immense "ground truth" data sets for training. While deep learning models have recently emerged as an excellent tool for supervised learning, generally in neuroscience[24,74–76], and specifically for decoding behavioral events recorded with calcium imaging[58,77], in this case we are seeking to describe data with no such training set or "ground truth" available. Instead, we compare unsupervised methods to solve the specific problem of detecting reducing hundreds of neurons into components, or sub-networks, or cell assemblies." Therefore, this approach will enable more sophisticated and nuanced unsupervised analysis of different states of neuronal networks, and how they shift as a function of context.

In sum, NMF provides a superior method for the holistic analysis of network dynamics recorded in calcium imaging. This analysis is able to go beyond summary statistics, providing a low dimensional representation of dynamics while still considering the activities of each neuron. The mathematical constraints required by NMF, linearity and positivity, complement the nature of fluorescent recording and neuronal activities well. Further, the parts-based nature of NMF provides a simple and interpretable representation of sub-networks of activity summing to drive macroscopic dynamics. As a result, we have developed an NMF pipeline to be an exceptionally valuable tool for elegantly demystifying shifting neuronal network dynamics.

Methods

Dimensionality Reduction (DR) Methods

<u>Nonnegative Matrix Factorization (NMF):</u>

NMF is a linear, matrix-decomposition method requiring all elements be nonnegative [34]. Each reduced dimension, or component, identified in NMF can be interpreted as a specific combination of input features representing distinct sources of variance in the data [37]. The components then sum to represent the original data, rather than canceling variables (using negativity) to achieve a geometrically strict representation [32].

Mathematically, NMF decomposes an original matrix X into two constituent, lower rank, matrices W and H, given by:

$$X = WH$$

(Eq. 1)

with the aim of iteratively updating W and H to minimize an objective function so that the product of the two deconstructed matrices optimally reconstructs the original [34]. Here, X is a column-wise representation of the data, where each column is the time-series fluorescent activity of a neuron (Fig. 1b) extracted from the recording (Fig. 1a), that is decomposed into two lower rank matrices, W and H. For NMF, W can be interpreted as a feature matrix, each vector representing a distinct pattern in the input data, and H interpreted as a coefficient matrix, representing the contribution of each feature vector to each point in the original matrix. When using DR for analyses, especially in the context of neural recordings, the underlying assumption is that the original matrix can be decomposed into a lower rank series of matrices. The rank, k, should be significantly smaller than n (k << n), the number of neurons recorded, while still recapturing a majority of the variance present in the data [78]. Model performance can be assessed by the variance explained as represented by the coefficient of determination (R^2) [79], given by:

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}},$$

(Eq.2)

where,

$$SS_{res} = \sum_{i,j} (X_{i,j} - (WH)_{i,j})^2 \,,$$

$$SS_{tot} = \sum_{i,j} (X_{i,j} - \overline{X})^2 \,.$$

(Eq. 3,4)

Thus, the model captures the original data perfectly in the limit as R^2 approaches one. R^2 increases as components are added, because the rank of the two constituent matrices are increasing. Akaike Information Criterion (AIC) has been adapted for NMF to inform the rank to which the dimensionality should be reduced [37,80], with the optimal model minimizing it [81]; for our implementation of NMF:

$$AIC = 2(\frac{R^2}{2\sigma^2} + k(n + t))$$

(Eq. 5)

Where R^2 is the reconstruction error defined in Eq. 2, σ^2 is the estimated variance of the data set, k is the number of components the model is decomposing into, n is the number of neurons in the recording, and t is the number of time points in each of the time series being decomposed. For our analysis, we interpret each row of H as a sub-network, with each entry of that row as a neuron's contribution of to that sub-network (Fig 1e). Each column of W is then interpreted as the time-series for that sub-network (Fig 1d). Further, the neurons in the original recording FOV can be artificially reconstructed to provide visualization of sub-network contribution (Fig 1c). We provide the pseudocode from an implementation of the NMF analysis pipeline:

```
Input: X [time points x neurons]
Output: NMF_model
normalize data between (0,1)*
initialize list of AICs
for i in range(1, num_neurons)
    W, H = init_WH(n_components = i)
    NMF_model = optimize(X, W, H)
    AIC_list.append(NMF_model.aic)
    if AIC[i-1] < AIC[i]:
        n_components = i-1
        break
W,H = init_WH(n_components = n_components)
NMF_model = optimize(X,W,H)
```

* we normalize using the global maximum and minimum of the trace set to preserve the relative difference between them

The above pipeline describes the sequence of using AIC to find an optimal number of components, and then fit a model to that number of components. Our implementation of NMF is a modification of the implementation found in Python's Scikit-Learn package [82]. We manually initialize W and H, as opposed to the automated method found in the package, due to the wide variety of initialization methods possible for NMF [83,84], and to maintain consistent, precise, and easily reproduceable control over the initial conditions of our models. For the analyses found in this paper, we apply nonnegative dual singular value decomposition (nndsvd), a consistent and efficient initialization method [85]. Further, we implement variance explained and AIC for automatic model selection, as they are not part of the base implementation in the package.

Principal Component Analysis (PCA):

We implement PCA because of its widespread use in neural data analysis [10,36,48,66–68,86,87]. PCA, efficiently calculated by a Singular Value Decomposition (SVD)[32]:

$$X = U\Sigma V^T$$

(Eq. 6)

where X is a column-wise representation of our data (i.e., each column is the time-series fluorescent activity of a neuron, assumed to have zero mean) decomposed into U, Σ, and V^T. V^T is the set of basis or singular vectors, also referred to as the principal components (PCs) of the original matrix, where each row of this matrix is a component of the data. Further, each of the PCs are necessarily orthogonal to each other by construction. U is the projection of the original matrix X onto the principal components in V^T, and Σ is a matrix containing the singular values of X describing the magnitude of the transformation, ordered by decreasing singular values [32]. The variance explained by a PC is the proportion of its singular value compared to the sum of all singular values. For our analysis, we interpret each PC as a sub-network or prevalent pattern of activity, with each entry of the PC being a neuron's weight or contribution to that PC. We then interpret the projection of the data onto each PC as the activity or time series of that sub-network, attributing the dominance of that sub-network to the variance explained. Our implementation of PCA is a modification of the implementation found in Python's Scikit-

Learn package [82]. We provide the pseudocode from an implementation of the PCA pipeline below:

```
Input: X [time points x neurons]
Output: PCA_model
1:  normalize data between (0,1)*
2:  subtract mean of each time series
3:  PCA_model = PCA(X) #data must be in column-format at this step.
* we normalize using the global maximum and minimum of the trace set to preserve the relative
difference between them
```

This is a modification of the implementation of PCA found in sklearn [82]. PCA calls for every variable, or trace in our case, to be centered about zero[32]. However, we remove the uniform normalization from sklearn's implementation to allow for finer tuning of data normalization before decomposition with the model.

<u>Independent Component Analysis (ICA):</u>

We implement ICA because of its ability to extract independent sources from complex signals, which could help separate overlapping patterns of neural activity in a single calcium recording. While ICA has been used for automated trace extraction in calcium recordings [88], this approach has seen limited application in the further analyses of neural signals estimated calcium recordings. ICA is a statistical technique that aims to decompose a multivariate signal into statistically independent components [35], rather than a set of orthogonal components, as found in PCA. ICA operates under the assumption that k independent sources produce the data, and that the data are a linear mixture of these underlying sources. ICA assumes the data to be a linear mixture of sources:

$$x = As$$

(Eq. 7)

where x is the original data, s are the underlying independent sources, and A is a mixing matrix that mixes the components of the sources [35].

In practice, the independent sources found in s, are decomposed using an unmixing matrix, W, such that the linear transformation of the data by W, estimates the underlying independent sources [35]:

$$\hat{s} = Wx$$

(Eq. 8)

Given that W transforms x to estimate the independent sources, we interpret the components of W as the contribution of each neuron to the sub-network activities, which are the independent sources in s. Here, we interpret the s independent sources as the underlying patterns of activity of each sub-network that drive the observed dynamics.

Like PCA, ICA assumes that inputs are zero mean. While PCA is designed to maximize the variance of the data along the principal components [32], ICA attempts to separate the data into k independent sources. We use the fastICA implementation found in sklearn [82], and provide the pseudocode from an implementation of the ICA analysis pipeline below:

Input: X [time points x neurons]

Output: ICA_model (contains A, s, amongst other information)

1: normalize data between (0,1)*

2: subtract mean of each time series

3: ICA_model = ICA(X)

* we normalize using the global maximum and minimum of the trace set to preserve the relative difference between them

While ICA calls for unit variance for each of the features, which in this case would be the calcium time series, we maintain the relative difference in variance between each neuronal time series to maintain the relative difference in activity recorded as a function of the fluorescence.

Uniform Manifold-Approximation Projection (UMAP):

We implement UMAP because of its increasing use to analyze life science data [27,89] and its ability to capture non-linear relationships. UMAP has been shown to capture complex patterns to visualize and cluster data in very low dimensions [33], though the technique remains untested on calcium recordings. Briefly, UMAP aims to optimize between preserving local and global structures, where structure is local between neighboring points and global between data points extending beyond the neighborhood. In neuronal calcium imaging, "neighborhood" refers to local relationships among neurons with similar activity patterns at a time point, while global relationships extend beyond those neighbors at the same point. UMAP subsequently directly maps this structure to a lower dimension [33], resulting in the output being a single matrix (the equivalent of H, V^T, and S), and not a product of matrices. Therefore, UMAP does not provide temporal information about the activity of the lower dimensional patterns or sub-networks (Table

1). Further, there is no well-established way to quantify variance explained in UMAP as in the other linear methods. We use the python package UMAP-learn to implement UMAP, and provide the pseudocode from an implementation of the UMAP analysis pipeline below:

Input: X [neurons x time points] (NOTE TRANSPOSITION FOR THIS CASE, COMPARATIVELY)

Output: UMAP_model (contains low_d_projection, amongst other information)

1: normalize data between (0,1)*

2: UMAP_model = UMAP(X)

* we normalize using the global maximum and minimum of the trace set to preserve the relative difference between them

In addition to the number of components, k, being a manually set hyperparameter, the number of points in the neighborhood is also manually set, in addition to the other 37 tunable parameters in the package [33]. Because our main objective is to identify how neurons group to form sub-networks and shift as a function of experimental contexts, we apply UMAP to decompose the data X into a *k x n* matrix. k is the number of components or subnetworks we determine to decompose into, while n is the number of neurons, aiming to give a similar representation to H, V^T, and S in our other methods.

Model	Variance Explained	Information Criterion	Neuronal Weights	Decomposed Activity
NMF	✓	✓	✓	✓
UMAP	X	X	✓	X
PCA	✓	X	✓	✓
ICA	X	X	✓	✓

Table 1. Model Attributes

Network Simulation Methods

To assess performance of the DR methods, we implement two simulation paradigms. In each paradigm, we simulate the spiking activity in a network of interconnected neurons and estimate calcium imaging traces for each neuron. We then apply each DR approach to the simulated calcium fluorescence data. We describe each simulation paradigm below.

Perfectly Intraconnected, Independent Nodes:

We first consider a network of 100 neurons organized into 5 independent nodes, each consisting of 20 neurons. Neurons within the same node are perfectly connected, such that a spike by any neuron in one node produces a spike in all other neurons in that node. Neurons in different nodes are disconnected, such that activity in two neurons of different nodes are independent (Fig. 2a). The activity of a single neuron is governed by a basic Poisson spike generator [90], given by:

$$spike[t] = x_{rand} < r * dt$$

(Eq. 9)

where a neuron spikes at time t if x_{rand} (uniform [0, 1]) is less than the product of the set firing rate, r, and the time step, dt. We update this generator to include additional network effects, and a refractory period. We represent the connectivity between neurons as an adjacency matrix A. Within any node of 20 neurons, A is 1, while between nodes A is 0; we exclude self-connections by setting the diagonal of A to 0. We also include a refractory period of duration q time steps, such that after a spike a neuron is temporarily unable to generate a subsequent spike.

We simulate the spiking activity of a single neuron in this network n = {1,2, … N} as:

$$spike[t]_n = x_{rand} < \left(\frac{rk}{N} dt + A_{n\,x\,N} * spikes[t-1]_{N\,x\,1}{}^{T} \right) * RS[t]_n$$

(Eq. 10)

where r = 3 is the set nominal rate, N = 100 is the total number of neurons, k = 5 is the number of nodes, and dt is the time step. A is the adjacency matrix of dimensions N x N, where each row represents the incoming connections to a neuron, and each column represents the outgoing connections of a neuron, and $RS[t]_n = \{0,1\}$ is the refractory state of neuron n at time t. A neuron spikes at time t if the randomly generated number x_{rand} (uniform [0,1]) is less than the product in Eq. 10. We note that any multiplication of an adjacency of 1 to neuron n, in conjunction with a spike at time t-1 will result in a spike at time t in neuron n, given the neuron n is not in a refractory state. After simulating the network we convolve each spike train with a calcium kernel [29,91] as described below. The result is a series of 100 fluorescent neuronal activities (Fig. 2b), of which there are five groups of twenty that are perfectly correlated (Fig. 2c).

Random Process Propagation:

We generate a second simulation paradigm to more closely represent a physiologically relevant dynamic in calcium fluorescence data: a small number of independent random spiking patterns simultaneously driving calcium events in all neurons. To do so, we first simulate k = 5 independent spiking patterns as,

$$process\ spike[t]_n = x_{rand} < r\ dt\ .$$

(Eq. 11)

We then consider a population of N = 150 neurons that receive inputs from each of the underlying spiking patters (Fig. 3a) weighted as follow: 20% of the neurons have weights uniformly distributed between 0.2 and 1.0, and 80% of the neurons have weights exponentially distributed (rate parameter $\lambda = 8.0472$) with maximum value of 0.2 (Fig. 3b). In this way, most neurons are connected with weights following an exponential distribution [92], and the few neurons with strong weights drive a majority of the activity. The resulting adjacency matrix A between the k spiking patterns and N neurons has dimensions k x N, where each element describes the weight of the kth underlying process input to neuron n = {1, 2, … N}. We simulate each neuron's spike train as:

$$neuron\ spike[t]_n = x_{rand} < (\ A_{n\ x\ k} * process\ spike[t-1]_{k\ x\ 1}) * RS_n$$

(Eq. 12)

where n indicates the nth neuron, k indicates the kth underlying spike processes, and we include a refractory period for each neuron. Doing so results in N = 150 neuronal spike trains. We convolve the spike train of each neuron with the calcium kernel [29,91], resulting in simulated calcium data similar to our in vivo recordings, mirroring: activity rate, activity distribution, underlying macroscopic drivers, and calcium neuronal time series characteristics.

Calcium Kernel:

We convolve the simulated spike trains with the following calcium kernel to represent the fluorescence recorded from the genetically encoded calcium (GECI) kernel jGCaMP7f:

$$[Ca^{+2}]_t - [Ca^{+2}]_{t-1} = -\frac{\mathrm{dt}}{\tau}([Ca^{+2}]_{t-1} - [Ca^{+2}]_b) + A * spike[t] + \sigma_c\sqrt{dt}\,\epsilon_t$$

(Eq. 13)

where $Ca^{+2}_{t,t-1,b}$ are the calcium concentrations at t, t-1, and the baseline, respectively. Here, dt is the time step, τ is the time constant of the GECI, σ_c is the variance of the calcium noise, and ϵ_t is a random variable following the standard gaussian distribution. We then calculate the fluorescence as:

$$F_t = \alpha([Ca^{+2}]_t) + \beta + \sigma_F\,\epsilon_t ,$$

(Eq. 14)

where F_t is the fluorescence at time t, α is the intensity, β is the bias, σ_F is the variance of the calcium noise, and ϵ_t is a random variable following the standard gaussian distribution. We set τ to 0.265 and σ_c to 0.5 to match the kinetics of jGCaMP7f [4]. We set the calcium baseline to 0.1, A and α to 5, β to 10, and σ_F to 1 [29].

Component Assignment Success

<u>Perfectly Connected, Intraconnected Nodes:</u>

To assess whether nodes were successfully assigned to components, we develop a framework for each of our generated network architectures. For the nodal model, we begin by summing the weights for each node, in each of the decomposed components, described by the pseudocode below:

```
Input: H (matrix of neuronal weights), nodes (array detailing which neurons belong to each node)
Output: aggregated_weight_matrix
aggregated_weight_matrix = zeros(num_nodes, num_components)
for i in range(num_nodes)
    for j in range(num_compnents)
        aggregated_weight_matrix [ i, j ] = sum(H [nodes[ i ] , j )
```

The result is a five by five matrix, where the element [i , j] is the sum of the weights for neurons in the ith node for the jth component. To then determine success of prediction, we establish two conditions that must be met. First, within a particular component, the sum of weights for the node with the highest value should exceed the sums of weights for all other nodes within the same component. Second, this highest sum should also be the greatest for the node, across all components. By applying these conditions, we can assess the assignment for each node individually. In essence, once we obtain the aggregated weight matrix, an ideal model would exhibit the same element as the maximum value for each row and column.

Random Process Propagation:

We calculate a 5 x 5 correlation matrix between the weights in each of the 5 components and the connection probabilities from each of the 5 generated random processes sampled from the defined exponential distribution. To assess assignment, we again consider two conditions. First, for a chosen component, the highest correlation with the connection probabilities from one random process should surpass the correlation to the other random processes in the same component. Second, this highest correlation within each component should identify a different random process; i.e., each component should identify a unique random process. We consider a component successfully assigned to one of the five random processes if its highest correlation is more correlated with that process than the other four processes.

Awake-State Data:

To investigate the application of each DR technique to *in vivo* recordings, we also consider calcium imaging data recorded from 36 awake mice. Briefly, recordings in layer 2/3 of murine S1 [11] were collected during the awake, resting state [46]. Traces representing calcium transients were extracted from the recordings and processed using standard techniques [11].

CHAPTER THREE: CaNetiCs – An Open-Source Toolbox for Standardized Dimensionality Reductive Analysis of Neuronal Calcium Activity

Adapted from Carbonero D, Noueihed J, Gabel CV, Kramer MA, White JA. CaNetiCs – An Open-Source Toolbox for Standardized Dimensionality Reductive Analysis of Neuronal Calcium Activity. (2025) https://doi.org/10.1101/2025.08.01.668030

Abstract

The widespread use of calcium imaging has produced large-scale datasets capturing neuronal population activity across diverse experimental contexts, posing challenges for analyzing complex, high-dimensional data. Dimensionality reduction (DR) methods have been pivotal in addressing these challenges by simplifying data into interpretable, low-dimensional structures, while capturing essential network dynamics. Among DR methods, Nonnegative Matrix Factorization (NMF) can produce biologically meaningful representations through its nonnegativity constraint and parts-based decomposition, making it especially suited for analyzing neuronal calcium signals. To enhance accessibility and standardization in the analysis of state-dependent neuronal dynamics, we introduce **Ca**lcium **Net**work dynami**Cs** (CaNetiCs), an open-source toolbox centered on NMF, integrating standardized DR methods (PCA, ICA, UMAP), geometric low-dimensional component space analyses, and neuronal network simulation modules.
We validate our toolbox by applying it to two diverse experimental datasets that describe responses to graded anesthesia: whole-ganglion cellular calcium imaging of C. elegans and two-photon imaging of murine somatosensory cortex. Our analyses recapitulate previously observed trends, such as network suppression and decorrelation with

anesthesia, while uncovering novel insights into neuronal activity under differing contexts. CaNetiCs provides an accessible, modular, and interpretable framework, facilitating broader adoption of standardized dimensionality reduction methodologies for deeper exploration of neuronal network dynamics across experimental paradigms. The open-source code, along with documentation, is available at https://github.com/dannycarbonero/CaNetiCs.

Introduction

The ability to record activity from large populations of neurons with single-cell resolution has revolutionized neuroscience, enabling new insights into how distributed networks dynamically encode behavioral and physiological states[1]. Calcium imaging, in particular, allows monitoring of hundreds to thousands of neurons simultaneously across different brain regions and experimental conditions[5,6,16,41,42]. However, the complexity and dimensionality of these recordings pose significant challenges for analysis. Standard approaches, such as computing descriptive statistics or averaging across cells and time points, often obscure the rich temporal and population-level structures inherent in the data[31,47].

Dimensionality reduction (DR) techniques provide a powerful solution by projecting high-dimensional neural recordings onto lower-dimensional representations, retaining major sources of variance while facilitating interpretation of collective dynamics[20]. Linear DR methods, such as Principal Component Analysis (PCA)[10,36,61,66,67,86,87] and Independent Component Analysis (ICA)[18,78], have been widely used to identify patterns of coordinated neuronal activity. More recently, nonlinear methods such as Uniform Manifold Approximation and Projection (UMAP) have been employed to capture complex structures in neural population data[27]. Yet, each approach imposes distinct mathematical assumptions, such as orthogonality in PCA, independence in ICA, or manifold continuity in UMAP[32,33,35], that can influence the nature and interpretability of the extracted components.

Nonnegative Matrix Factorization (NMF) represents a particularly promising alternative for calcium imaging analysis[18,36]. By constraining both the data and its decomposition to be nonnegative[34], NMF preserves the additive nature of fluorescence signals and produces parts-based representations, wherein each component corresponds to an interpretable sub-network of activity[18,36]. Unlike PCA or ICA, which impose strict mathematical constraints to construct a less interpretable low-dimensional representation, NMF decomposes neuronal activity into coherent modules that can be directly associated with underlying biological processes[18]. Prior work has demonstrated that NMF robustly captures dominant organizational features of neuronal activity in simulated networks, and outperforms alternative DR methods in reconstructing known ground-truth structures[18].

Although dimensionality reduction methods offer powerful means for interpreting neuronal population dynamics, their application often requires specialized technical knowledge and complex, customized computational workflows. Sophisticated analyses, including variance optimization, information-theoretic model selection, and geometric modeling of state spaces, can be difficult to implement without significant computational expertise. Recent work has led to an increase of open-source software toolboxes emerging to lower the barrier to entry for complex data analyses across fields of neuroscience[36,50,52]. However, available toolboxes remain limited with respect to providing standardized, tunable, and feature-rich user-friendly frameworks specifically tailored for analyzing state-dependent neuronal dynamics from calcium imaging recordings. To address this gap, we developed **Ca**lcium **Net**work dynami**Cs** (CaNetiCs), an integrated toolbox that consolidates validated analytical methods into a flexible and accessible codebase.

Building on our prior work[18], our software offers a standardized implementation of NMF-based dimensionality reduction optimized for calcium recordings, alongside complementary DR approaches (PCA, ICA, UMAP), geometric low-dimensional component space analysis, simulation tools for method benchmarking, and basic neuronal network simulation analysis, providing modularity for hypothesis-driven and exploratory analyses alike.

A key advantage of our NMF decomposition implemented here is the creation of easily interpretable structures: lower-rank components that correspond to distinct patterns of activity, or sub-networks of neuronal activity, and the associated neuronal weights that reflect the extent to which each neuron contributes to each sub-network[18,29,36]. Each neuron's activity is thus approximated as a positive, additive combination of a small number of dominant modes. This parts-based representation[37] provides an intuitive map of network organization, identifies leading drivers of population dynamics, and facilitates quantitative comparisons of neuronal sub-network structures across experimental conditions. Moreover, the nonnegativity constraint ensures that all extracted features remain biologically plausible, avoiding artifacts introduced by cancellation effects inherent to other DR methods[18,32].

We demonstrate the utility of CaNetiCs by building a standard analytical pipeline that applies our DR analysis to two diverse experimental datasets: light-sheet calcium imaging of nearly the entire head ganglia of Caenorhabditis elegans under graded isoflurane anesthesia[10], and two-photon imaging of murine somatosensory cortex (S1 L2/3) across awake, anesthetized, and recovery states[11]. Across both datasets, we find that CaNetiCs not only recapitulates previously reported macroscopic trends in network suppression and decorrelation with anesthesia but also uncovers novel geometric and dynamic features of state transitions at the population level.

By providing an accessible, flexible, and interpretable framework for dimensionality reduction analysis, we seek to enable a broader adoption of systematic and sophisticated network-level approaches for a diversity of calcium recordings under various contexts, and to facilitate finding deeper insights into how large-scale neuronal dynamics are shaped by experimental perturbations, biological state, and species-specific differences.

Results

We first apply our developed pipeline to analyze calcium recordings under increasing concentrations of isoflurane anesthesia in *C. elegans*[10] and in mice[11]. We fit models to the optimal dimensionality informed by an information criterion[80,81,93] and to three dimensions, consistent with previous analyses[10,66–69,94,95], to analyze the neuronal network dynamics under differing concentrations of anesthesia. We show that our analyses naturally capture the different dynamics between the species, support the conclusions of previous more rudimentary analyses, and provide further novel insights to differential dynamics as a function of anesthetic concentration.

Neuronal network dynamics of C. elegans under increasing concentrations of isoflurane

We analyze the neuronal network dynamics of anesthesia recorded in *C. elegans* under no (0% Isoflurane), intermediate (4% Isoflurane), and deep (8% Isoflurane) anesthesia. We re-use these data previously reported in Awal et al[10]. For each animal (n = 10) we apply the CaNetiCs pipeline to the calcium traces (examples in Fig. 7a) as follows.

We first iteratively fit models for 25 components, for each animal, for each condition. Then averaging across animals, we aim to build a profile for variance explained and information criterion across conditions. We begin by analyzing the variance explained across conditions. We find that variance is explained most quickly by the awake state, followed by the deeply anesthetized state, and then the intermediately anesthetized state (Fig. 7c). There are several consistencies in the variance explained profile from our analyses and the previously conducted analyses of the same dataset[10]. Here, we also capture a grand majority of the variance explained within three components. Analysis of individual trace[10] activity determined that neurons are significantly more correlated during the awake state, followed by intermediate anesthesia, and deep anesthesia, with the only statistically significant difference being between the awake and anesthetized states. Because AIC follows from variance explained, optimizing between model complexity and variance captured, AIC follows the same trend as variance explained, but inverted. AIC is lowest during the awake state, followed by the deeply anesthetized state, and finally the intermediary anesthetized state (Fig. 7d). We then fit models to the optimum number of components determined by the Akaike Information Criterion (AIC), for each animal, for each state, where the optimal dimensionality is inferable from the

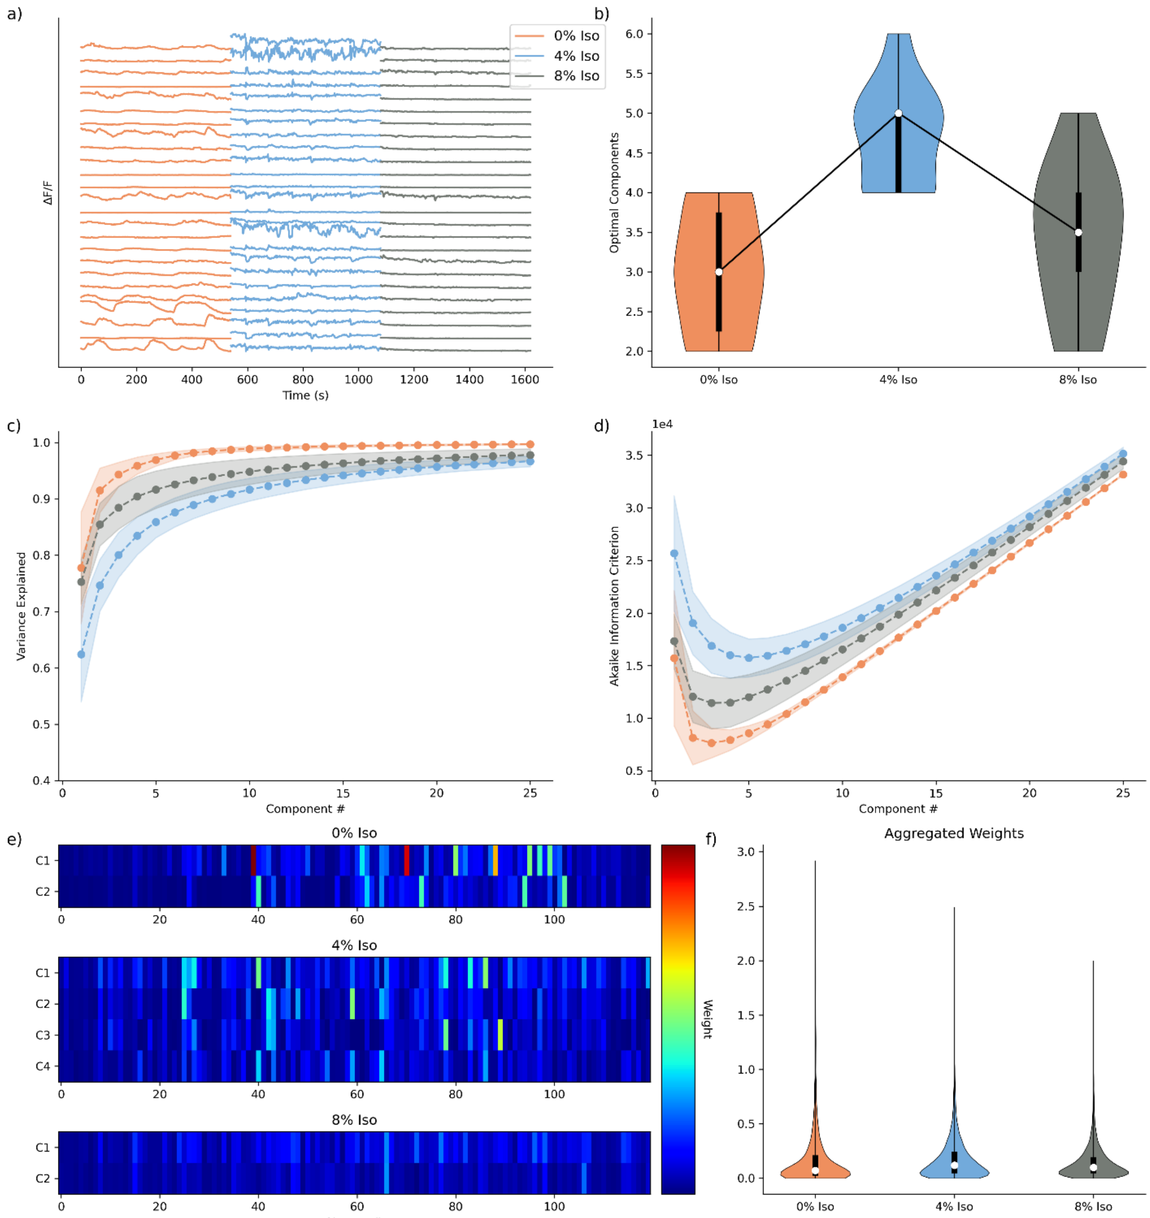

Figure 7. Optimal Dimensionality Models for C. elegans Isoflurane Dataset. a) Activity traces for a subset (n=25) of neurons, for a single animal, across conditions. b) Distribution of optimal components for all animals, across states (n = 10). c) Mean variance explained curve with standard deviation intervals (n=10). d) Mean Akaike Information Criterion (AIC) curve with standard deviation intervals (n=10). e) Visualization of neuronal weights at optimal dimensionality for animal shown in a), standardized across conditions for comparison. f) Aggregated neuronal weights for each animal at optimal dimensionality (n = [3600, 5640, 4080], respectively).

Akaike Information Criterion (AIC), at the minimum of the curve[80,81,93] (Fig. 7d) denotes the optimum number of components to fit the data to. We find that the inferred number of optimal components (Fig. 7b) across states follows the same correlation structure found when analyzing singular neuronal correlations, with the awake state requiring the lowest dimensionality, followed by the deep and intermediate anesthetized states, respectively. We find the difference between the optimal components to be statistically significant (Wilcoxon signed-rank test, corrected for multiple comparisons, Table 1) between the awake state and the intermediate state, with the difference trending towards significance between 4% and 8% anesthesia. This suggests a dimensionality difference between 4% Isoflurane and the other states, while the awake state and deeply anesthetized state can be explained by similar dimensionality, but with different activity profiles. The activity profile differed from that of the correlations in that total activity was highest in the awake state, followed by intermediate anesthesia then deep anesthesia[10].

We next visualize the neuronal weights decomposed under each condition qualitatively for a single model (Fig. 7e) and aggregated across animals for all components (Fig. 7f). In addition to requiring fewer dimensions to optimally explain the awake and deeply anesthetized states, compared to the intermediate state, a few neurons are found to drive most of the activity (Fig 1e). Further, normalizing the scale across the decompositions for the different states, we qualitatively find the highest weights in the model for the awake state, especially for the first component (Fig. 7e). Examining the distributions of aggregated weights for all components across states (Fig. 7f), we find the weights distribute similarly to the total macroscopic trends of activity in the data[10], with the deep anesthetization weights being the most suppressed, the intermediate anesthetization having the widest distribution, and the awake state having the longest tail. These results demonstrate that our analytical pipeline is capturing both the dynamics and the dimensionality of the data in an unsupervised dimensionality reduction model.

Comparison	Components p-value
0% Iso - 4% Iso	0.023
0% Iso - 8% Iso	1.000
4% Iso - 8% Iso	0.094

Table 2. Significance values for statistical comparison of optimal number of components found in Fig. 7b). Wilcoxon Signed-Rank Test, corrected for multiple comparisons, n = 10.

Having chosen the number of dimensions, we proceed by fitting models to three dimensions for all experimental states to build a low-dimensional component space model (Fig. 8), building upon previous analytical methods, and furthering previous analysis[10,66–69,94,95].

We first qualitatively examine neuronal dynamics for a representative animal (the same animal utilized in Fig. 7). By visualizing the activity profiles of the first three components across different anesthetic conditions (0% Isoflurane, 4% Isoflurane, and 8% Isoflurane), we observe clear distinctions that align closely with previous analyses of this dataset[10]. Specifically, the awake state (0% Isoflurane) exhibits a highly structured activity pattern with a distinctly ordered, nearly oscillatory behavior across all three components. In contrast, the intermediate anesthetic state (4% Isoflurane) displays a disrupted and scrambled activity pattern, reflecting a transitional and dynamically unstable neuronal state. The deeply anesthetized state (8% Isoflurane) reveals markedly suppressed neuronal activity, with greatly diminished amplitude across components (Fig. 8a). These qualitative observations corroborate prior findings, offering visual validation of differential neuronal network dynamics across anesthetic states.

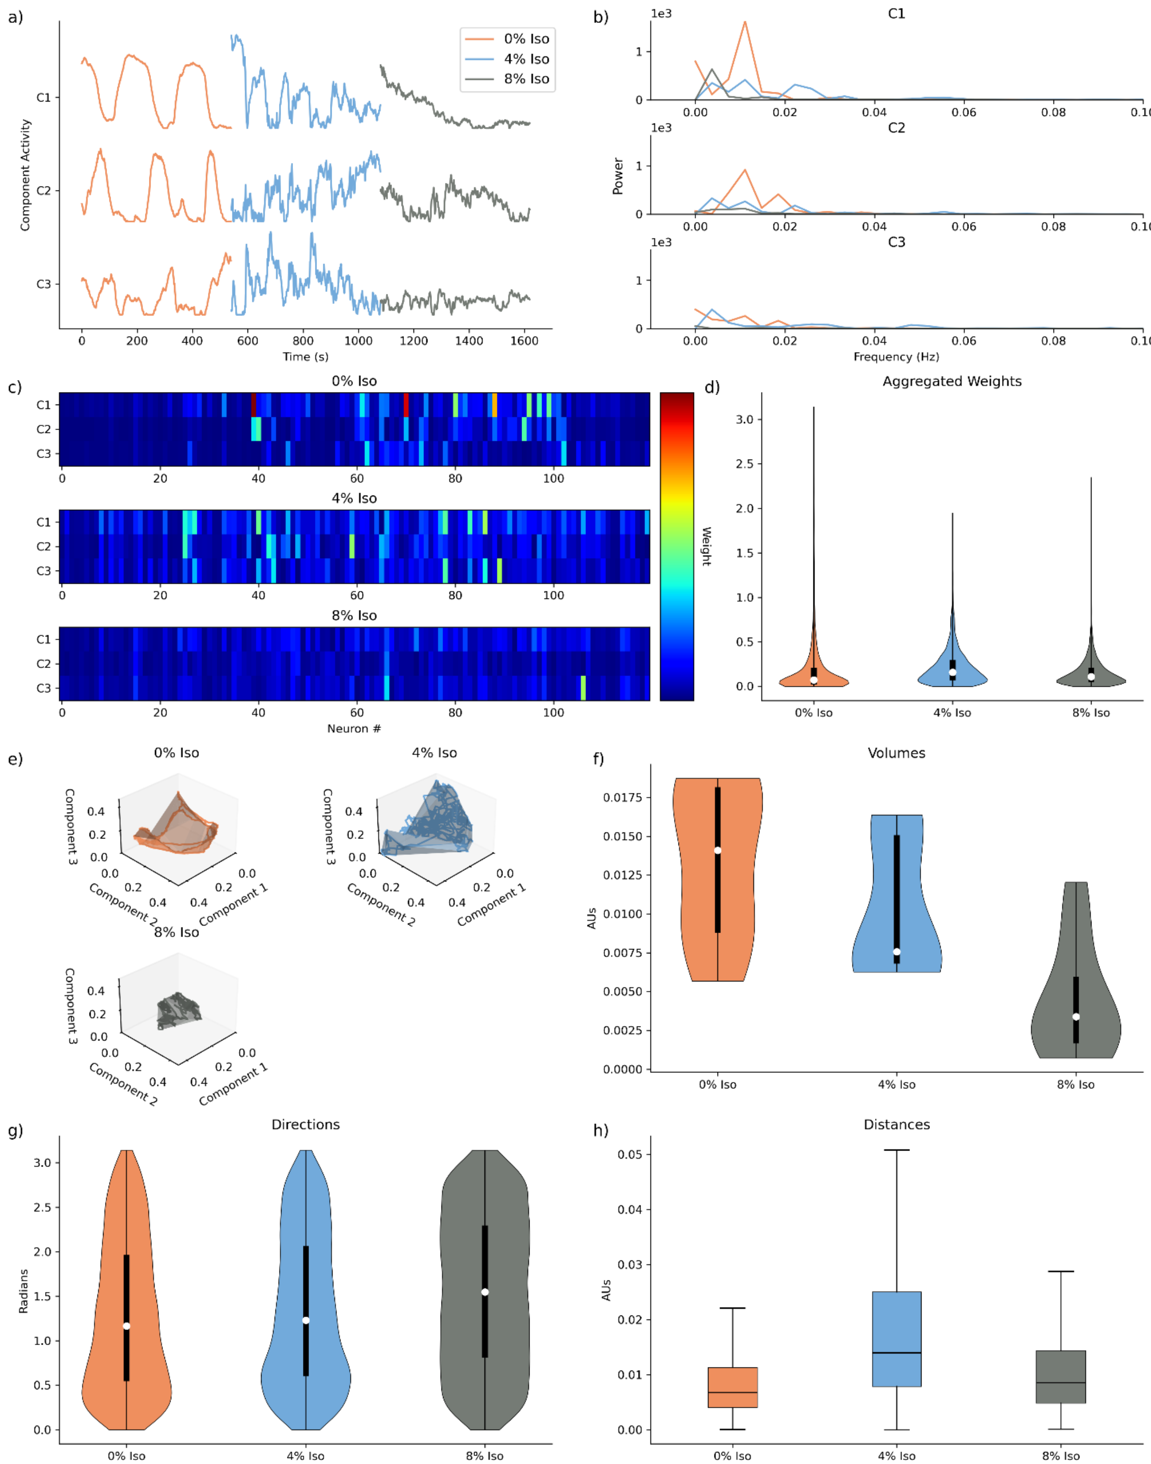
a)
0% Iso
4% Iso
8% Iso
Component Activity
C1
C2
C3
Time (s)
b)
C1
C2
C3
Power
Frequency (Hz)
c)
0% Iso
4% Iso
8% Iso
Neuron #
Weight
d)
Aggregated Weights
e)
0% Iso
4% Iso
8% Iso
Component 1
Component 2
Component 3
f)
Volumes
AUs
g)
Directions
Radians
h)
Distances
AUs

Figure 8. Three-Dimensional Low-Dimensional Component Space Models for C. elegans Isoflurane Dataset. a) Component activity traces for first three components for the same animal shown in Fig. 7a). b) Decomposed frequency content for components shown in Fig. 8a). c) Visualization of neuronal weights for single three-dimensional models for the same animal in Figs. 8a,b). d) Aggregated neuronal weights for all three-dimensional models, across animals and states (n = 3600). e) Plotted low-dimensional component space of components against each other, with encapsulated volume, for animal shown Fig. 8a,b,c). f) Distribution of volumes, across animals and states (n=10). g) Distribution of directions, across animals and states (n=5380). h) Distributions of speeds across animals and states (n = 5380)

To characterize further these oscillatory characteristics, we perform spectral analyses to assess the frequency content within each component across anesthetic conditions. Consistent with our qualitative observations, spectral analysis reveals prominent low-frequency peaks specifically within the awake state, consistent with oscillatory dynamics. These peaks are notably attenuated or absent in both intermediate and deeply anesthetized states, underscoring the disruption and suppression of structured neuronal oscillations as anesthetic depth increases (Fig. 8b).

We next visualize the neuronal weights derived from a three-dimensional decomposition, again for the same representative animal. Qualitatively, we observe a trend consistent with our earlier findings from optimal dimensionality analysis (Fig. 7); neuronal activity in the awake state is driven by larger neuronal weights compared to anesthetized conditions (Fig. 8c). Quantitative analysis of neuronal weights across the full dataset further supports these observations. We detect a similar overall relationship to that observed with optimal component decompositions, though with a notable difference: the intermediate anesthetic state's distribution tail does not exceed that of the deeply anesthetized state. This discrepancy arises from truncating the weight distribution due to the fixed three-dimensional decomposition rather than utilizing state-specific optimal dimensionalities (Fig. 8d).

To further illustrate the dynamics, we plot the dynamics projected to these first three components in three-dimensional space. We find a distinctly ordered movement about this space in the awake state, reflective of the pronounced oscillatory structure. In contrast, the intermediate state displays less well-ordered dynamics, and the deeply anesthetized state shows marked suppression and minimal spatial exploration (Fig. 8e). To characterize the extent of these different dynamics, we compute the volume fit to the curve across all animals and observe the largest volumes in the awake state, followed by intermediate, and the smallest in the deep anesthetic state; the deep state significantly differs from both awake and intermediate states (Wilcoxon Rank sum test, n = 10, corrected for multiple comparisons). These results align with the observed magnitude of neuronal activity differences across conditions (Fig. 8f).

Analyzing the directional trajectories of the movement about space, we find highly consistent, smooth directional changes in the awake state, a marginal decrease in smoothness in the intermediate state, and slightly more irregular directional changes in the deep anesthetic state (Fig. 8g).

Comparison	Volumes p-value
0% Iso - 4% Iso	0.289
0% Iso - 8% Iso	0.004
4% Iso - 8% Iso	0.020

Table 3. Significance values for statistical comparison of volume expanse in space (Fig. 8f). Wilcoxon Rank-Sum Test, corrected for multiple comparisons, n = 10.

Finally, examining the distances traveled through the low-dimensional space, we find that the awake state traverses the shortest distances per unit time, followed by the deep anesthetic state, and then the intermediate anesthetic state (Fig. 8h). These patterns of a smoother, slower, movement in state space are consistent with the previously identified higher neuronal correlation profiles reported in Awal et al.[10].

Comparison	Components p-value
0% Iso - 0.7% Iso	1.000
0% Iso - 1.4% Iso	0.012
0% Iso - Recovery	1.000
0.7% Iso - 1.4% Iso	0.042
0.7% Iso - Recovery	1.000
1.4% Iso - Recovery	0.009

Table 4. Significance values for statistical comparison of optimal number of components found in Fig. 9b). Wilcoxon Signed-Rank Test, corrected for multiple comparisons, n = 36.

Overall, we conclude the awake state's movement is more expansive (i.e., covers more of the lower dimensional space) and more directed and smooth compared to the anesthetized conditions.

Here, by leveraging both qualitative and quantitative analyses across geometric and temporal domains to create a DR framework for calcium recordings, our approach corroborates earlier findings[10] and introduces a unified framework for interpreting complex state-dependent network dynamics while adding novel insight.

Neuronal network dynamics of murine S1 L2/3 under increasing concentrations of isoflurane

As a second illustration of the CaNetiCs pipeline, we analyze calcium recordings from layer 2/3 of murine S1 (n = 36 rodents). The experimental context begins with imaging of an awake or baseline state, increases to an intermediate anesthetic concentration (0.7% Isoflurane), proceeds to a deeper anesthetic concentration (1.4% Isoflurane), and ends with imaging of a recovery state with no anesthetic (examples in Fig. 9a). In Noueihed et al [11], the authors used these data to show that single-cell neuronal activity becomes more sparse and more uniform, as a function of entering deeper levels of anesthetic sedation.

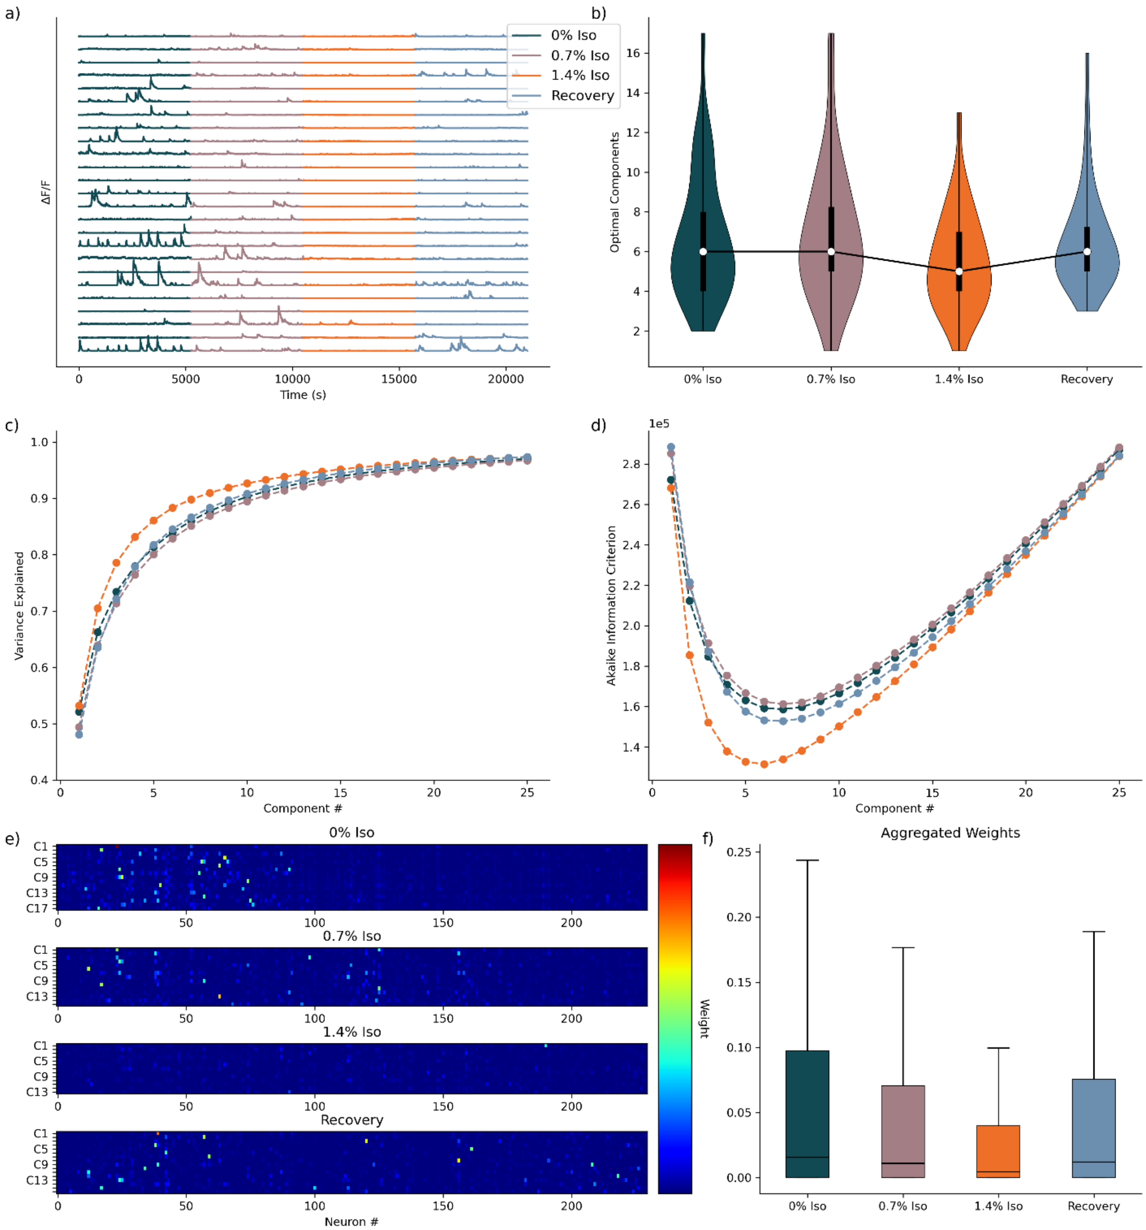

Figure 9. Optimal Dimensionality Models for Murine Isoflurane Dataset. a) Activity traces for a subset (n=25) of neurons, for a single animal, across conditions. b) Distribution of optimal components for all animals, across states (n = 36). c) Mean variance explained curve with standard deviation intervals (n=36). d) Mean Akaike Information Criterion (AIC) curve with standard deviation intervals (n=36). e) Visualization of neuronal weights at optimal dimensionality for animal shown in a), standardized across conditions for comparison. f) Aggregated neuronal weights for each animal at optimal dimensionality (n = [27711, 27589, 22354, 26388], respectively).

Here, we re-use these data to illustrate application of the CaNetiCs pipeline.

We begin by analyzing the variance explained versus the number of components across the three anesthetic conditions. We find that variance was most rapidly captured by the deeply anesthetized state (1.4% Isoflurane), with a distinct gap between this condition and the others (Fig. 9c). This finding aligns closely with prior analyses of single neuronal activities, which demonstrated a marked decrease and increased uniformity or correlation of neuronal activity under deep anesthesia[11]. The variance explained analysis corroborates these findings by rapidly capturing the simplified, more correlated neuronal activity patterns characteristic of deep anesthesia. Computing the Akaike Information Criterion (AIC), we observe a similar trend (Fig. 9d). The deep anesthetic state consistently showed the lowest AIC values, indicating optimal model complexity at fewer components due to the relatively simpler activity profile. Evaluating the optimal number of components inferred from the minimum of the AIC curves, we find that the deeply anesthetized state (1.4% Isoflurane) requires significantly fewer components to represent neuronal activity compared to all other anesthetic conditions (Table 3, Wilcoxon Signed-Rank test, corrected for multiple comparisons, n=36). No significant differences emerged among the other anesthetic conditions regarding optimal dimensionality (Fig. 9b). Examining the neuronal weights for a single representative model, we find results consistent with previous analyses. Specifically, visual inspection (Fig. 9e) suggests neuronal weights are consistently highest during the awake state, indicating active neuronal contributions. Intermediate and recovery states exhibited more moderate neuronal contributing to the model, albeit with increased intermittency, while the deeply anesthetized state demonstrated comparatively minimal neuronal weights. To characterize these observations, we aggregate neuronal weights across all animals and experimental conditions and find neuronal weights are highest in the awake condition, followed sequentially by the recovery and intermediate states, with the deeply anesthetized state consistently demonstrating the lowest neuronal weights (Fig. 9f). Taken together, these results highlight the capacity of our DR approach to robustly capture the underlying structure of neuronal activity in reduced-dimensional representations, adjusting dimensionality as the anesthetic conditions informed the information criterion. Importantly, our analyses consistently suggest the presence of an intermediate state, distinct from both the deeply anesthetized and awake baseline states. This intermediate state is characterized by moderate neuronal activity complexity, driven by lower component weights and reduced interactions between neuronal networks, along with decreased overall neuronal activity. However, it maintains a dimensionality similar to that observed in the awake state.. This intermediate state, partially uncovered here and suggested by previous studies[11], emphasizes the nuanced neuronal dynamics underlying anesthetic-induced changes in cortical activity.

Finally, we fit three-dimensional models to each animal across all experimental states to further analyze the neuronal dynamics. Visual inspection reveals activity patterns consistent with previous analyses of the neuronal activity[11].

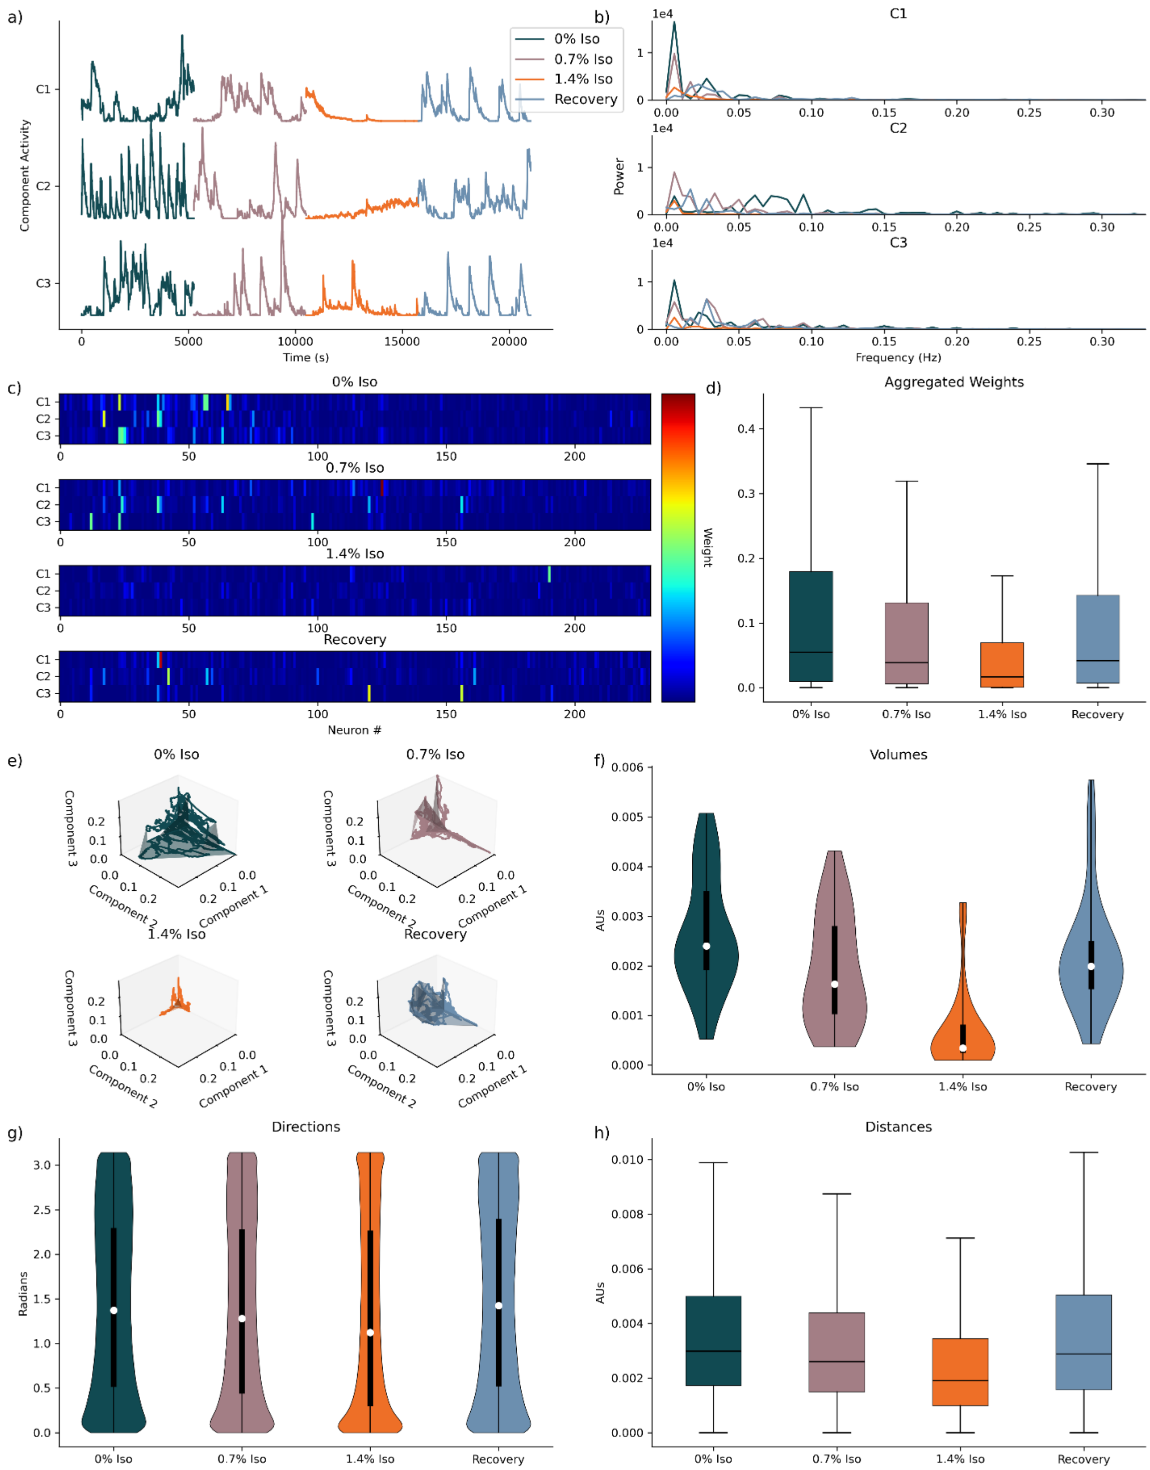

a)
0% Iso
0.7% Iso
1.4% Iso
Recovery
Component Activity
C1
C2
C3
Time (s)
b)
C1
C2
C3
Power
Frequency (Hz)
c)
0% Iso
0.7% Iso
1.4% Iso
Recovery
Neuron #
Weight
d)
Aggregated Weights
e)
Component 1
Component 2
Component 3
f)
Volumes
AUs
g)
Directions
Radians
h)
Distances
AUs

Figure 10. Three-Dimensional Low-Dimensional Component Space Models for Murine Isoflurane Dataset. a) Component activity traces for first three components for the same animal shown in Fig. 9a). b) Decomposed frequency content for components shown in Fig. 10a). c) Visualization of neuronal weights for single three-dimensional models for the same animal in Figs. 10a,b). d) Aggregated neuronal weights for all three-dimensional models, across animals and states (n = 10743). e) Plotted low-dimensional component space of components against each other, with encapsulated volume, for animal shown Fig. 10a,b,c). f) Distribution of volumes, across animals and states (n=36). g) Distribution of directions, across animals and states (n=188928). h) Distributions of speeds across animals and states (n = 188928)

The awake state displays the most extensive activity, followed by intermediate activity levels during intermediate anesthesia (0.7% Isoflurane) and the recovery state, with significantly reduced activity evident in the deeply anesthetized state (1.4% Isoflurane) (Fig. 10a). Subsequent spectral analyses revealed little to no structured differences independent of activity levels across anesthetic conditions, indicating frequency content changes primarily reflect overall activity amplitude (Fig. 10b).

Comparison	Components p-value
0% Iso - 0.7% Iso	0.052
0% Iso - 1.4% Iso	4.05e-09
0% Iso - Recovery	0.360
0.7% Iso - 1.4% Iso	1.41e-06
0.7% Iso - Recovery	1.000
1.4% Iso - Recovery	4.86e-08

Table 5. Significance values for statistical comparison of volume expanse in space, found in Fig. 10f). Wilcoxon Rank-Sum Test, corrected for multiple comparisons, n = 36.

Qualitative visualizations of neuronal weights decomposed into three components for a single representative model indicate trends consistent with the more comprehensive optimal models, exhibiting increased sparsity, with few neurons making meaningful contributions. This sparsity likely reflects the inherent limitation of capturing the full complexity of neuronal dynamics within just three dimensions, underscoring the need for higher-dimensional representations to fully encapsulate the data (Fig. 10c). Expanding this analysis quantitatively by aggregating neuronal weights across all animals confirmed the observed qualitative trends. The awake state again showed the highest neuronal weights, with intermediate and recovery states presenting moderate weights, and the deeply anesthetized state consistently exhibiting minimal contributions. This suggests that although reduced dimensionality captures broad dynamics, it misses details present in higher-dimensional optimal models (Fig. 10d).

Further, plotting neuronal activity in three-dimensional state space revealed that the awake state demonstrated the most expansive and dynamic exploration of this space, consistent with its greater and more diverse neuronal activity (Fig. 10e). Quantitative assessments of volumes encapsulating these trajectories showed statistically significant smaller volumes in the deeply anesthetized state compared to all other states (Table 4, Wilcoxon Rank-Sum test, corrected for multiple comparisons, n=36). Moreover, an emerging, though not statistically significant, trend indicated intermediate spatial volumes for the 0.7% Isoflurane and recovery states, highlighting potential intermediary characteristics between awake and deeply anesthetized conditions (Fig. 10f).

Analyzing directional movement in the three-dimensional space revealed a shift from predominantly stochastic trajectories with a preference for small directional changes in the awake state, towards more distinct bimodal directionality under deep sedation, further positioning intermediate anesthesia and recovery states as transitional (Fig. 10g). Finally, evaluating distances traveled through the three-dimensional space revealed relatively similar trajectories across all states, excluding the deeply anesthetized state, which consistently showed reduced movement.

Collectively, these analyses reinforce the utility of the three-dimensional low-dimensional component space framework in capturing key aspects of neuronal dynamics and revealing nuanced insights, such as a potential intermediate state as a function of entering and exiting anesthesia, that might not immediately be evident from single-neuron analyses alone. The relationships observed among states, dimensionalities, activities, and transitional states underscore complex neuronal dynamics modulated by anesthetic depth and recovery.

Overall, our analytical pipeline extracts low-dimensional representations of neuronal calcium recordings that not only recapitulate findings from previous analyses but also enable novel insights into neuronal network dynamics across experimental conditions. The combination of information-theoretic optimization with dimensionality reduction yields interpretable structures that align with known biological phenomena, such as anesthetic-induced suppression and decorrelation of neural activity. Moreover, we demonstrate that our unsupervised approach can robustly uncover markedly different underlying structures across species, highlighting the method’s generalizability and broad applicability. By providing both quantitative and qualitative validation of network dynamics, we demonstrate that our toolbox offers a unified, flexible framework for analyzing diverse neuronal datasets. These results illustrate that dimensionality-reductive modeling, when standardized and optimized as in our toolbox, not only confirms previous findings but also opens new avenues for hypothesis generation and deeper biological interpretation.

Discussion

Here, we present an integrated analytical pipeline centered around NMF for the dimensionality reductive analysis of neuronal calcium recordings and wrap this framework into an open-source toolbox. Our goal was to provide a standardized, tunable, and interpretable approach for reducing complex neuronal population activity into tractable low-dimensional representation with easily implementable software.

We validated our approach by applying it to two distinct datasets, calcium recordings from C. elegans and L2/3 murine S1 cortex under varying depths of isoflurane anesthesia. Across both datasets, our method robustly captured biologically meaningful differences in neuronal dynamics between experimental states[10,11]. Notably, the low-dimensional structures extracted by NMF closely concurred with previous analyses of these datasets, while providing new avenues for quantitative and geometric analysis, such as characterizing the expanse of network activity in state space and the smoothness of neuronal trajectories moving in a dimensionally reduced space. Importantly, we observed that the extracted low-dimensional representations differed substantially between species, following true underlying biological distinctions, rather than analytical artifacts. This ability to adaptively capture species-specific network structures in an unsupervised manner underscores the generalizability and broad utility of the approach. Further, when comparing three-dimensional decompositions to the optimal dimensionality fits, we observed an important difference between species. In C. elegans, where the underlying neuronal dynamics are comparatively simpler, fitting to three components recapitulated much of the structure captured at the optimally inferred dimensionality. This suggests that the dominant modes of neuronal activity in C. elegans are highly compressible and well-approximated by this lower-dimensionality state space. In contrast, the murine S1 cortex recordings, characterized by more complex and heterogeneous dynamics, revealed that the three-dimensional models began to capture broad trends across anesthetic states but failed to fully encapsulate the richness of the underlying activity. Here, the reduced dimensionality truncated finer structure, particularly in intermediate and recovery states. This highlights a critical consideration. While low-dimensional representations are highly informative, the appropriate dimensionality must be chosen with regard to the biological complexity of the system under study.

Given the hierarchical structure of our data, we focused our statistical analyses on measures aggregated at the level of the animal, such as optimal dimensionality and volumes spanned in state space, rather than on per-point measures like individual neuron weights, speeds, or directions, which exist at vastly higher scales (10^3–10^5 points per condition, detailed in the Fig. captions for each distribution). While the modest number of animals naturally limits the statistical resolution attainable for detecting subtle effects, our aim was not to derive absolute metrics but rather to robustly evaluate relative differences across experimental states. These relative comparisons, such as shifts in dimensionality or space coverage with anesthetic depth, are highly interpretable and biologically meaningful even with moderate sample sizes. Thus, our approach emphasizes the practical utility of low-dimensional analysis for discerning clear network-level changes, without overinterpreting absolute effect sizes beyond the inherent sampling resolution.

Several potential future directions naturally follow from our work. One intriguing avenue would be to leverage the NMF-derived models fit to specific states as "basis sets" for transforming and projecting activity from different states, enabling quantitative comparisons across conditions in a common low-dimensional subspace. This could allow a more nuanced understanding of how network structure shifts with perturbations such as anesthesia, aging, or disease. Moreover, future work could integrate constraints or priors more deeply into the decomposition process to reflect known network motifs or anatomical features, moving toward semi-supervised or guided dimensionality reduction. Additionally, the use of dynamic or adaptive dimensionality, adjusting component number over time in single states, could further improve resolution of state transitions and network remodeling.

While the main analyses in this study focused on spontaneous NMF, fully unsupervised decomposition and dimensionality, we also implemented a constrained version (CNMF) to allow the imposition of specific analytical designs or hypotheses. CNMF has been widely applied to trace extraction and motion correction tasks[40,50,51,53,57]. Here, we implement and reframe as a flexible analytical tool for controlled, hypothesis-driven analysis of low-dimensional structure, opening new possibilities for targeted network investigations.

Beyond NMF and CNMF, we also implemented the standard and widely employed methods of PCA, ICA, and UMAP in our toolbox. While NMF is an exceptional method for optimizing for these low-dimensional representations of macroscopic activity, each of these dimensionality reduction techniques offers distinct advantages depending on analytical goals: PCA identifies orthogonal axes of maximal variance; ICA seeks statistically independent sources; UMAP preserves both local and global relationships through nonlinear manifold embeddings. By integrating these approaches, we provide end-users flexibility to tailor dimensionality reduction to the specific needs of their biological questions.

In addition to building tools for activity analysis, we, further, incorporated simulation modules that allow users to generate synthetic neuronal network recordings with tunable properties. These modules, implementing both interconnected and independent process models[18], provide a powerful exploratory tool for validating dimensionality reduction strategies and understanding model behavior under controlled conditions.

Nonetheless, important limitations must be acknowledged. A major constraint in applying dimensionality reduction methods is the necessity of balancing the sample size of animals against the inherent complexity of the data. While dimensionality reduction is predicated on the assumption that $k \ll n$, where k is the number of components and n is the number of neurons, statistical power ultimately derives from the number of independent biological replicates. In practice, acquiring sufficient sample sizes of high-resolution calcium recordings can be labor-intensive and technically demanding. Thus, interpretations drawn from low-dimensional structures must always be made with careful consideration of the biological sampling resolution.

In conclusion, here we introduce an accessible, flexible, and extensible framework, wrapped into a software toolbox, for applying dimensionality reduction to neuronal calcium imaging data, validated across multiple species and conditions. Our results demonstrate that standardized dimensionality reductive modeling not only reproduces prior observations but extends them, offering new biological insights and frameworks for interpreting network dynamics. We anticipate that this toolbox will facilitate a broader adoption of systematic dimensionality reduction in neuroscience, enabling researchers to more easily reveal structure within complex neuronal populations and recordings.

Methods

Dimensionality reduction (DR) methods

Nonnegative Matrix Factorization (NMF):

We focus our analysis on implementing Nonnegative Matrix Factorization (NMF) because of its superior ability to analyze calcium signals[18]. NMF is a linear decomposition DR technique that aims to factor a given matrix into two constitutive lower rank matrices, where all elements of all matrices are constrained to be nonnegative. Given by a matrix $X \in \mathbb{R}^{m\times n}$, NMF aims to decompose X into two lower rank matrices $W \in \mathbb{R}^{m\times k}$ and $H \in \mathbb{R}^{k\times n}$, such:

$$X = WH$$

(Eq. 15)

where we can capture a significant portion of the variance at $k << \min(m,n)$[78], and the multiplication of W and H optimally approximates X for k[34]. The accuracy of the approximation is described by the variance explained as represented by the coefficient of determination (R^2) [79], and is given by:

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}},$$

(Eq.16)

where,

$$SS_{res} = \sum_{i,j} (X_{i,j} - (WH)_{i,j})^2 ,$$

$$SS_{tot} = \sum_{i,j} (X_{i,j} - \overline{X})^2 .$$

(Eq. 17,18)

with the model better decomposing the original data in X as R^2 increases, perfectly capturing it as the limit of the variance explained approaches one. A major advantage to NMF is the adaptation of Akaike information Criterion (AIC) for more robustly selecting the number of dimensions for decomposition [80,93], where the optimal number of components minimizes[81]:

$$AIC = 2(\frac{SS_{res}}{2\sigma^2} + \psi)$$

(Eq. 19)

where SS_{res} is the sum of the squared residuals (Eq. 3), σ^2 is the estimated variance of the dataset, and ψ are the free parameters in the model (the number of dimensions being decomposed into k, multiplied by the sum of the dimensionality of X, m + n). AIC provides a more objective metric for optimization, as opposed to the "elbow" method used in other DR contexts[27,63], directly informing the optimal dimensionality of the reduction.

Linearity lends itself to more interpretable model representations at the cost of going beyond very low-dimensionality as is common in nonlinear manifold learning[62,63]. We can interpret each entry of each row of H as a neuron's contribution to that component or sub-network of activity, and each column of W is then interpreted as the time-series for that component. Nonnegativity furthers the interpretability conferred by linearity because the data are modeled as the sum of the components, where each component can be interpreted as a specific combination of input features, representing distinct sources of variation in the data[37].

Modeling NMF low-dimensional component spaces in three-dimensions:

Previous research has used linear components and nonlinear components modeled in three-dimensional space, found in W, to further analyze the neuronal dynamics[10,66–69,94,95]. Here, we build an expanded framework for low-dimensional component space analysis.

Fitting geometries to component activity in three-dimensions:

To quantify the extent of neuronal activity represented in three-dimensional space, we fit geometric closed surfaces to the derived curves. The enclosed volume serves as a robust metric to describe the spatial expanse of neuronal activity. Specifically, we apply standard convex hulls[96] or alpha shapes[97], using established computational libraries[98,99] to define precise geometric boundaries around activity patterns. A convex hull is the smallest convex set enclosing all points, often resulting in an overestimation of volume due to its convexity constraint, while alpha shapes provide a more flexible fit, capturing concave features by adjusting a tunable parameter (alpha) that determines the granularity of the boundary. The resulting volume enclosed by these surfaces provides a quantitative measure to characterize the spatial dynamics and expanse of neuronal activity.

Analyzing the movement of component activity in three-dimensions:

Following previous research, we quantify the smoothness of activity trajectories by calculating the angles between consecutive segments, where smaller angles represent smoother transitions[94] indicative of consistent neuronal dynamics. Conversely, larger angles can identify abrupt transitions or instability in neuronal activity. Additionally, we evaluate the dynamics through computed speeds, defined as the Euclidean distance between consecutive points along the curve. Higher speed values indicate rapid shifts in neuronal network states or activity bursts, whereas lower speeds correspond to periods of relative stability or sustained network states.

Constrained Nonnegative Matrix Factorization (CNMF):

Another major advantage of NMF is that it can be constrained after initialization to force optimization towards a more biased or specified model[100]. CNMF has become state of the art for automated trace extraction of calcium recordings[40,50,51,53,57]. However, constraining has yet to be implemented for more targeted analysis of dynamics. We do so here.

Similarly to NMF, CNMF satisfies:

$$X = W(A \odot H)$$

(Eq.20)

where A is a constraint matrix that enforces constraints on selected entries of H. Specifically, A has the same dimensions as H and contains values indicating which elements of H remain free (1) and which are constrained (0). The elementwise multiplication A $\odot$ H then applies these constraints directly, effectively enabling or disabling individual elements of H. We note that for this implementation to remain constrained, multiplicative solvers must be used[100].

We validate our implementation of CNMF using artificially generated datasets. Specifically, we generate random Gaussian data for X, W, and H, and incorporate constraints through a randomly generated binary matrix A. This ensures that each element of H is properly constrained according to A during the optimization process (Fig. 11). We provide the pseudocode for an implementation of CNMF below:

```
Input: X [n_samples x n_features]
1:      W, H = init_WH(n_components = target_component_number)
2:      A = init_A # this is user defined for analysis purposes
3:      model_constrained = optimize_with_constraints(X, W, H, A)
```

*the only differences between an NMF implementation and the above CNMF implementation are in the necessary initializing of the constraint, and the providing of the constraint matrix to the optimization function

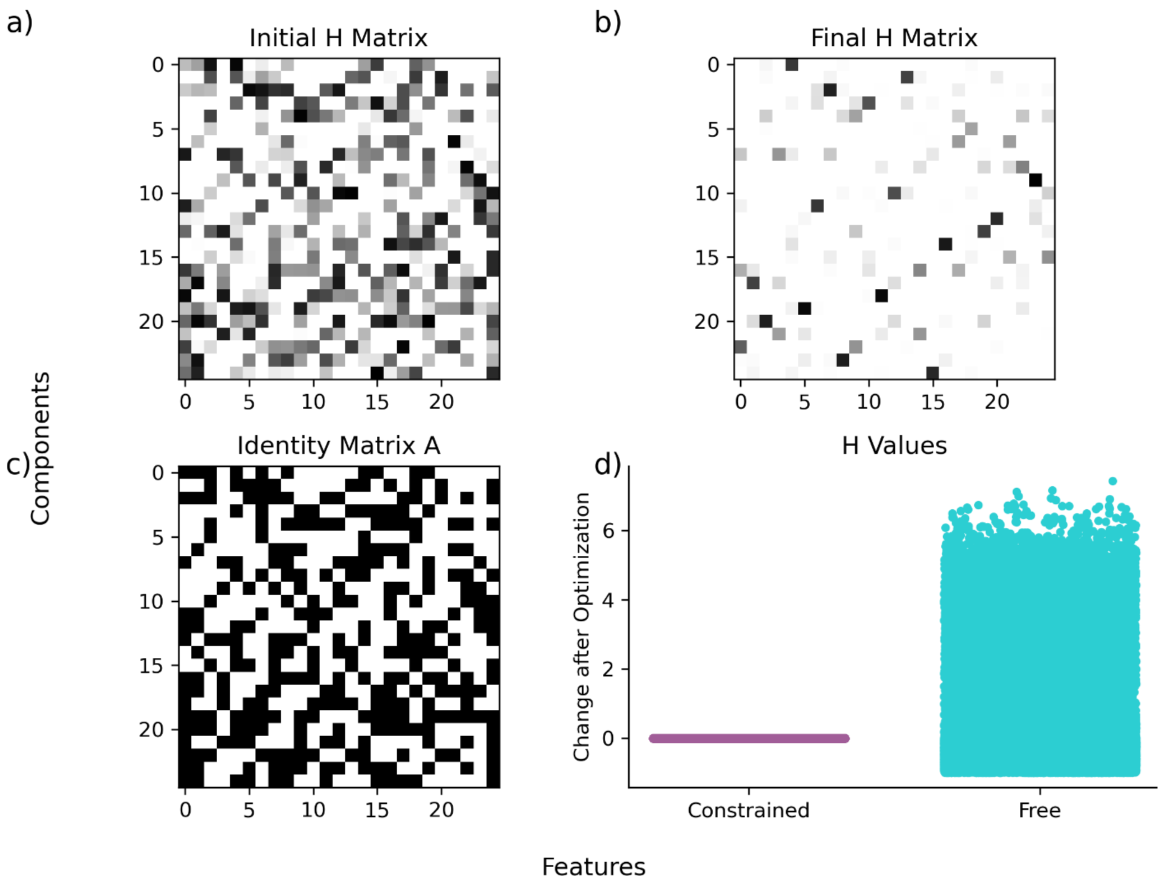

Figure 11. Constrained Nonnegative Matrix Optimization Proof of Constrained Optimization. a) Initialized 25x25 random matrix before optimization. b) Same 25x25 matrix as shown in a), after optimization. c) An initialized binary matrix to constrain the matrix shown in a). d) Difference between weights, before and after optimization, split according to constraint status.

Other feature implementations in our software wrapper

<u>Other NMF initializations:</u>

For our analyses, we initialize using nonnegative singular double value decomposition, a proven method to improve model convergence speed[85]. However, a variety of methods can be used to initialize an NMF with advantages, disadvantages, and different end goals for what components should capture[83,84]. Here we also implement custom, nonnegative PCA, and kmeans clustering initialization methods, in addition to the ones provided by sklearn[82], to allow for more granular user control for analysis, across the general different end goals of initializations[83,84].

Other DR methods:

In addition to NMF and CNMF, our software wrapper integrates a comprehensive set of widespread and previously characterized dimensionality reduction (DR) approaches in the context of this analysis[18]. We employ Principal Component Analysis (PCA), Independent Component Analysis (ICA), and Uniform Manifold Approximation and Projection (UMAP). PCA identifies orthogonal axes of maximal variance, ICA seeks statistically independent components in the data, and UMAP applies a nonlinear approach to preserve both local and global structure in the data. We implement these additional methods to broaden the analytical capabilities of the software and offer a high degree of user customization for more specific or targeted analyses.

Neuronal network simulations:

Previous research has implemented network simulation methods to generate artificial data to validate model performance[18,29]. Here we build a practical and easily tunable way to implement two simulation paradigms, the interconnected networks model in which neurons exhibit user-defined degrees of mutual influence, and the downstream or independent process model, where each neuron (or group of neurons) are primarily driven by separate input streams[18]. While previous research structured a specific set of connection of parameters for validation, our implementation allows full user tunability of parameters across architecture, firing rate, temporal resolution, and connection distributions[18], in addition to full tunability of calcium signal simulation as well, to account for a broad range of sensor dynamics[18,29,91].

Validation analysis of previously collected experimental data

Neuronal network dynamics of C. elegans under increasing isoflurane:

Using NMF, we investigate the neuronal network dynamics of previously collected C. elegans recordings under increasing concentrations of anesthesia (0% Isoflurane, 4% Isoflurane, and 8% Isoflurane), for ten animals. Briefly, researchers encapsulated C. elegans in a permeable hydrogel and used light-sheet microscopy to image nearly all head neurons (n = 150) expressing GCaMP6s (a calcium indicator) at two volumes per second to measure fluorescent activity[10]. Neurons were tracked over time, across conditions, and normalized activity extracted[10,101–104].

Neuronal network dynamics of murine layer 2/3 in murine S1 under increasing isoflurane:

We further investigate neuronal network dynamics of layer 2/3 of murine S1, also under increasing concentrations of isoflurane (0%, 0.7%, 1.4%, and recovery), for 36 animals. Previous work recorded two-photon calcium imaging in head-fixed mice at a frame rate of 29.16 Hz[11], systematically varying the inspired isoflurane concentration. The same neurons were also tracked across conditions[11], and activity also extracted[11,50,53].

Recording analysis pipeline

We develop and implement a standardized pipeline to analyze both recording sets. For each animal in either dataset, we first set the minimum of all traces to 0, to satisfy the constraints of nonnegativity. We then concatenate the recordings across all experimental states to globally normalize activity across each animal, setting the animal minimum to 0 and the maximum to 1. With traces on the same scale across experimental contexts, for each experimental condition, we begin by iteratively fitting models for the first twenty-five components to quantitatively characterize the profile of variance explained and information criterion as a function of adding components. Further, by characterizing the information profile for a particular recording, we can also find the optimal number of components to describe the recording[18], at the minimum of the AIC curve[81].

We continue by fitting the recording to the optimal number of components to analyze neuronal contributions (found in H) [Eq. 15] to patterns of activity in an unsupervised learning fashion.

Finally, we fit a three-component model to apply NMF in three-dimensional space. Recent DR analysis of our data revealed oscillatory dynamics in three-component space[10]. Here, we extend these previous findings by quantifying the oscillatory properties through spectral analysis using a standard Fourier transform[105]. We conclude by finding the space encapsulating the activity in three-dimensional space, volume, and characterizing its movement about space, speed and direction.

We provide the pseudocode for our analytical pipeline below:

```
Input: animal_recordings [time points x neurons x experimental epochs]
make trace minimum for all traces across epochs 0
concatenate traces [time points * experimental epochs x neurons]
normalize traces from (0,1)*
split traces into original dimensions [time points x neurons x experimental epochs]
for epoch in epochs:
    initialize list of variance explained, and AICs
    for i in range(1, num_components) # we use 25
        W, H = init_WH(n_components = i)
        NMF_model = optimize(X, W, H)
        var_ex_list.append(NMF_model.var_ex)
        AIC_list.append(NMF_model.aic)
    optimal_components = min(AIC_list)
    W_optimal, H_optimal = init_WH(n_components = optimal_components)
    model_optimal = optimize(X, W_optimal, H_optimal)
    W_state_space, H_state_space = init_WH(n_components = 3)
    model_state_space = optimize(X, W_state_space, H_state_space)
    volumes, directions, speeds = analyze_state_space(model_state_space.W)
```

* we normalize using the global maximum and minimum of the trace set to preserve the relative difference between them.

CHAPTER FOUR: A data augmentation procedure to improve detection of spike ripples in brain voltage recordings

Adapted from Schlafly ED, Carbonero D*, Chu CJ, Kramer MA. A data augmentation procedure to improve detection of spike ripples in brain voltage recordings. Neuroscience Research (2025); 215, 15–26. https://doi.org/10.1016/j.neures.2024.07.005*

** represents equal contribution from authors, Carbonero written text in maroon*

Abstract

Epilepsy is a major neurological disorder characterized by recurrent, spontaneous seizures. For patients with drug-resistant epilepsy, treatments include neurostimulation or surgical removal of the epileptogenic zone (EZ), the brain region responsible for seizure generation. Precise targeting of the EZ requires reliable biomarkers. Spike ripples — high-frequency oscillations that co-occur with large amplitude epileptic discharges — have gained prominence as a candidate biomarker. However, spike ripple detection remains a challenge. The gold-standard approach requires an expert manually visualize and interpret brain voltage recordings, which limits reproducibility and high-throughput analysis. Addressing these limitations requires more objective, efficient, and automated methods for spike ripple detection, including approaches that utilize deep neural networks. Despite advancements, dataset heterogeneity and scarcity severely limit machine learning performance. Our study explores long-short term memory (LSTM) neural network architectures for spike ripple detection, leveraging data augmentation to improve classifier performance. We highlight the potential of combining training on

augmented and *in vivo* data for enhanced spike ripple detection and ultimately improving diagnostic accuracy in epilepsy treatment.

Introduction

Epilepsy is a significant neurological disorder, with an estimated prevalence rate of 6.38 per 1,000 individuals experiencing active epilepsy at any given time[106]. The diagnosis of epilepsy encompasses a diverse array of neurological syndromes, all characterized by the commonality of recurrent, spontaneous seizures. Seizures are typically unpredictable interruptions of normal brain function, concurrent with abnormally excessive and synchronous neuronal activity[107,108]. In focal epilepsy, seizures begin in a dysfunctional brain region and propagate[109]. For patients with drug-resistant epilepsy (when medications fail to prevent seizure recurrence), therapeutic options include neurostimulation[110–112] or surgical removal of sufficient epileptogenetic cortical tissue, aiming to diminish seizure frequency[109,113].

To identify the epileptogenic zone (EZ) — the cortical tissue responsible for seizure generation — and target subsequent treatment requires a reliable biomarker[114]. Ripples — oscillations characterized by their brief duration, low amplitude, and high frequency (80–250 Hz) — have emerged as a candidate biomarker[115]. Clinical studies demonstrate that ripples occur more often in the seizure onset zone (SOZ)[116–119], suggesting their utility in localizing epileptogenic regions. Moreover, resection of ripple generating tissues has been associated with improved patient prognoses[119–122]. Despite the promise of ripples as a biomarker for the EZ, detection poses significant challenges[75]. The current gold standard to identify ripples relies on manual interpretation of invasive or non-invasive electroencephalogram (EEG) recordings[123] by experts, or manual validation of detection by automatic processes[124]. These manual approaches not only demand substantial time[125], but also suffer from poor inter-rater reliability[126], limiting reproducibility. To address these limitations requires more objective, efficient, and automated detection methods, which could dramatically enhance the precision and reliability of surgical epilepsy treatments.

The task of differentiating between epileptic events and normal brain activity is further complicated by the presence of both physiologically normal and epileptic ripples within cortical tissue. Physiological ripples, associated with human memory processes[127–129], are indistinguishable from epileptic ripples in duration, amplitude, and spectral frequency[130–132]. One approach to distinguish the pathological events associated with epilepsy is to identify ripples that occur with epileptiform spikes — high amplitude (~100 μV), long lasting (~100 ms) deviations in the brain's voltage activity — which serve as a marker of epilepsy with high specificity[14,133,134]. Previous studies suggest that these co-localized spike and ripple events (spike ripples) provide a better marker for brain regions that generate epileptic seizures than either of the events in isolation[15,135,136].

Modern machine learning approaches provide novel tools to automate and improve detection of the interictal (i.e., between seizure) electrophysiological events associated with epilepsy. Existing studies have implemented different deep learning approaches to automate the detection of spikes, ripples, and spike ripples[137]. One common approach employs convolutional neural networks (CNNs), efficient in image classification[21]. Application of CNNs requires converting EEG recordings into images, for example expressing the activity as a spectrogram image. Subsequently, a CNN can be trained to recognize specific features of spike ripples within these spectrogram images[75,138,139]. This approach leverages the ability of CNNs to discern patterns in image data, without requiring specific feature selection.

An alternative deep learning approach implements Long-Short Term Memory models (LSTMs)[22], a variant of Recurrent Neural Networks (RNNs) suitable for time-series data. LSTMs have been applied to classify EEG spikes during the interictal period, using a series of extracted time domain features[140]. Other approaches have focused on Fourier-transformations of these time-series, using frequency-domain features to classify spikes, ripples, and spike ripples in intracranial EEG[23]. Another approach employs an LSTM to provide continuous predictions at each point in a time series, aiming to enhance the temporal resolution and predictive accuracy of event detection[141].

Despite the emergence and growing prominence of these methods to identify the interictal electrophysiological events of epilepsy, their performance and reproducibility are limited by the inherently heterogeneous, variable, and small number of datasets available for training[142,143]. However, data augmentation approaches have emerged to mitigate these limitations. Recent studies have shown that artificially enhancing and expanding datasets improves the performance of epileptiform event detection, helping to address the challenges of data scarcity and variability[144–146].

Here we introduce a new data augmentation strategy to increase the number of spike and spike ripple events available for training a neural network. To generate synthetic traces containing spike ripples and spikes without ripples, we estimate spike and ripple features from *in vivo* data and the existing literature, and then generate new instances of events from these estimates. This approach allows us to create an arbitrarily large training set not constrained to specific features observed within a limited dataset. We then use these data to train two existing LSTM model architectures developed to classify ripple events, one in the frequency domain[23], and another in the time domain[141]. We find that a combination of sequential training on augmented, then *in vivo*, data produces the best performing classifiers. These results demonstrate the potential of data augmentation as a supplement to model training for further improvement of spike ripples detection in human brain voltage recordings.

Methods

Human dataset

In this study, we re-use human data from[15] as described and analyzed in[75]. To summarize, we analyze EEG recordings (sampled at 2035 Hz) from 34 participants: 21 with childhood epilepsy with centrotemporal spikes (CECTS, age 4.9–16.8 years, 17 males) and 13 controls (age 8.7–14.3 years, five males). For each participant, all available data recorded during non-REM sleep were selected for analysis (range 214–2840 s, mean 1030 s). Medications included lacosamide (2 participants), levetiracetam (8 participants), lamotrigine (2 participants), and oxcarbazepine (1 participant); see Table 1 of [15] for additional details. We analyze one EEG channel per hemisphere per participant, prioritizing channels with interictal spikes or selecting C3 and C4 electrodes if no spikes were present. Detailed clinical data are available in [15].

Augmented dataset

We generate synthetic spike ripples by dividing these events into three components: a spike-wave component, a ripple component, and background noise. We generate each component separately, as described below, and then sum these components to generate the composite spike ripple event.

Spike-wave:

We parameterize the spike-wave waveform using Piecewise Cubic Hermite Interpolating Polynomials (PCHIP). PCHIPs are a class of spline comprising cubic polynomials that are monotonic on an interval and C^1 at the knots[147,148].

The monotonicity means that the spline can be constructed to predictably reproduce the prototypical peak-valley-peak morphology of a spike-wave using parameters that are intuitive and interpretable (Fig. 12). Here, we use five knots corresponding to different features of the spike-wave: spike onset, spike peak, post-spike valley, wave peak, and wave termination. The coefficients at the first and last knots are fixed so that the spike-wave evolves on a baseline of 0. As a result, the waveform is characterized by 8 parameters: 5 knot positions ($k_1, \ldots, k_5$) and 3 coefficients to characterize waveform amplitude ($v_{peak}, v_{tr}, v_{wave}$). For interpretability, we transform knot positions to the time of the spike peak (t_{peak}) and a set of four segment widths ($w_1, \ldots, w_4$):

$$t_{peak} = k_2 \quad \rightarrow \quad w_i = k_{i+1} - k_i$$

(Eq. 21)

see Fig. 12.

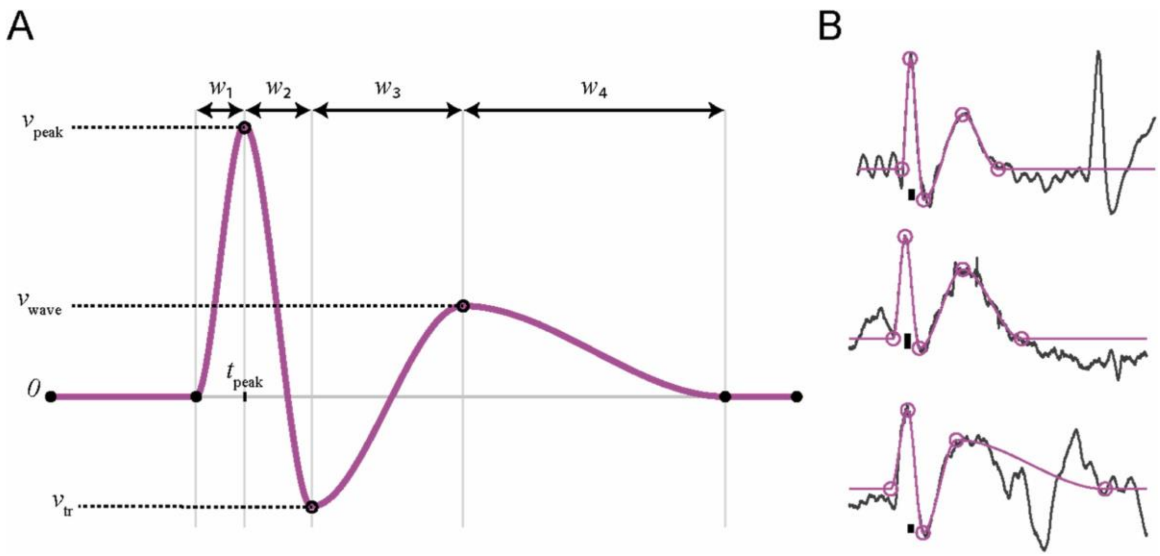

Figure 12. Spline construction of the synthetic spike-wave. (A) The waveform is divided into four segments: spike rise (w1), spike descent (w2), wave rise (w3), and wave descent (w4). The waveform amplitude at the transition between segments is characterized by parameters vpeak, vtr, vwave ; parameter tpeak indicates the timing of the spike peak. (B) Three examples of synthetic spike-waves (pink) generated using parameters based on in vivo observations (black); black scale bars below the spike peak indicate 1 standardized unit; intervals are 1 second in duration.

Ripple:

Ripples have been described as “spectral islands”, or isolated regions of relatively high spectral power in spectrogram images(Chu et al., 2017; Kramer et al., 2019; Nadalin et al., 2021). We model these islands as a function m of distance R from a center point in time-frequency space:

$$R(t, f;\ \mu_t, \mu_f, \alpha, \beta, \gamma) = \alpha(f - \mu_f)^2 + 2\beta(f - \mu_f)(t - \mu_t) + \gamma(t - \mu_t)^2$$

$$m(R;\ A, c) = Ac\exp[-\max(R - 1, 0)]$$

(Eq. 22, 23)

These equations define an ellipsoidal island centered at [μ_t, μ_f] with maximum amplitude Ac (a scalar multiple of c); the island has a plateauing top so that $m(\bullet) = Ac$ for $R(\bullet) \leq 1$.

The function *m* resembles a 2D-Gaussian, except for the plateauing top. Gaussians are often described in terms of variance and correlation, making this a more intuitive description of the spectral island. For this reason, we convert the coefficients α, β, γ to variance and correlation of the corresponding 2D-Gaussian:

$$\begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix} = \begin{bmatrix} {\sigma^2}_f & s_{ft} \\ s_{ft} & {\sigma^2}_t \end{bmatrix}^{-1}$$

$$s_{ft} = \rho \sigma_f \sigma_t$$

(Eq. 24, 25)

To simulate a ripple, we generate a 2s Gaussian white noise trace and then filter (1000 order FIR) the signal to $\mu_f \pm \sigma_f$; this narrowband filtering step minimizes leakage from higher frequencies and produces ripples with sinusoidal shape[149–151]. We then compute the continuous wavelet transform (CWT, Morse wavelet, 10 voices per octave[149]) to obtain the scalogram of the white noise trace and replace the amplitude component of the scalogram with the spectral island estimated from Equation 23 with $c = 0$. Finally, we perform an inverse wavelet transform on the adjusted scalogram to obtain the simulated ripple trace (Fig. 13).

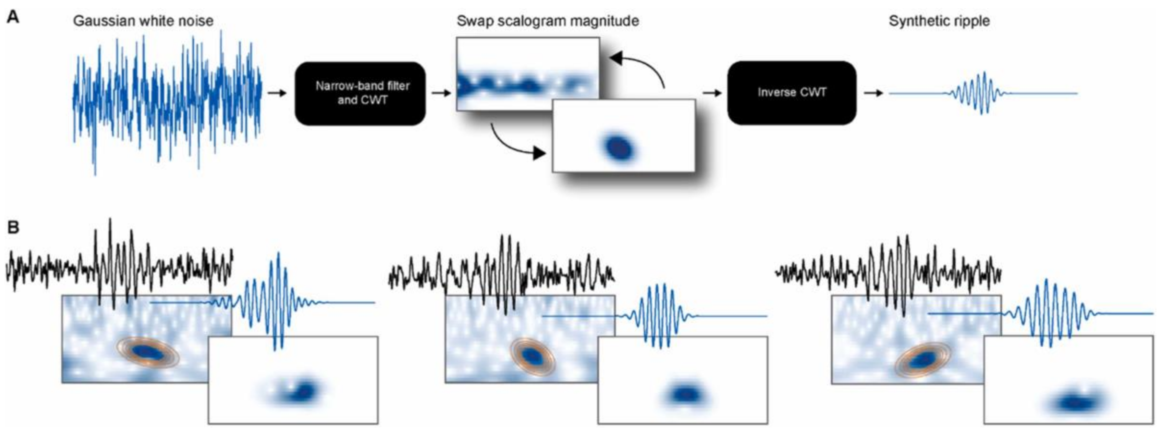

Figure 13. Synthetic ripple generation procedure. (A) Schematic representation of the algorithm for generating synthetic ripples (see Methods). (B) Three examples of synthetic ripples (blue) generated based on parameters estimated from in vivo observations (black). Black traces show the in vivo recording bandpass filtered to 80-500 Hz. Panels show CWT scalograms of the in vivo and synthetic ripples; darker colors indicate larger magnitudes; the horizontal (vertical) axis indicates time (frequency); frequencies shown are between 80 and 500 Hz. Orange contour lines in scalograms indicate contours of fitted magnitude surface. Each trace is 200 ms in duration.

Background signal:

Pink noise, as observed in neural field activity[152], data were simulated as described in[14]. Briefly, pink noise is characterized by the relationship $P(f) \sim f^{-\alpha}$, where P(f) is the sample power spectral density at frequency f. Here, we set α=1 consistent with *in vivo* observations classified as background activity (Median=1.03, IQR=[0.96, 1.08]). The resulting noise trace was scaled so that median power in the high frequency band ([80–500] Hz) matches the ripple background level *c*.

Event selection and signal preprocessing:

Candidate spike ripple events were identified from the *in vivo* data using a semi-automated detector and expert validation, as described in[75]. Background events were undetected by the semi-automated detector and validated by an expert not to contain a spike ripple, as described in[75]. The voltage recordings from each electrode were z-scored relative to the median and median absolute deviation.

LSTM Frequency-Domain Method

In the frequency domain, we implemented the LSTM deep neural network defined in [23] to detect spikes, ripples, and spike ripples. The input features to the network consisted of the power estimated in three consecutive 0.25 s intervals (z-scored data within each interval, Hanning tapered). Within each interval, the power (in decibels) was averaged within eight bands: theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), gamma1 (30–56 Hz), gamma2 (64–116 Hz), rip1 (124–176 Hz), rip2 (184–196 Hz), and rip3 (204–236 Hz). We note that the bands at higher frequencies (>80 Hz) were chosen to avoid electrical noise artifacts at frequencies 120 Hz, 180 Hz, and 200 Hz, as in[23]. For each event in the human and augmented datasets, the size of the input data to the LSTM network was [8 x 3], corresponding to the number of frequency bands (8) by the number of time points (3). The network consisted of 5 layers: (i) a sequence input layer with 3 inputs (the number of time points), (ii) a bidirectional LSTM layer with 200 hidden units, (iii) a fully connected layer with two outputs, (iv) a softmax layer, and (v) a classification layer that computed the cross-entropy loss for classification with two mutually exclusive classes (spike ripple or not). We trained the neural network using the Adam optimization algorithm, with a maximum of 200 epochs, a gradient threshold of 1, and a piecewise learning rate schedule with initial learning rate of 0.001 and learning rate drop period of 200 epochs. Data shuffling was performed every epoch to prevent the model from memorizing the order of samples.

LSTM Time-Domain Method

In the time domain, rather than implementing a new instance of a previously published architecture (as we did for the frequency-domain method), we instead leverage transfer learning[153] to implement a pre-trained LSTM (detailed in [141]) for the detection of spike ripples. This original model was trained to detect hippocampal sharp wave ripples, an electrophysiological phenomenon exhibiting temporal and spectral characteristics analogous to spike ripples. Given this similarity, we hypothesized the existing network would identify spike ripples, with fine-tuning further adapting the network to improve performance. Further, while we implemented previously published architectures for both the frequency- and time-domains, we only reinitialized the frequency-domain network because the time-domain network is more complex and would require re-learning of many more parameters and therefore a much longer training time.

The input to the network consisted of *in vivo* or simulated time series data. The output was also a time series, the same length as the input time series, indicating the probability of a spike ripple at each time. Compared to the frequency-domain LSTM method, the architecture is more complex, consisting of two deep bi-directional LSTM layers, flanked by a series of dropout and batch normalization layers. Before reaching the LSTM layers, the data were first transformed by a Gaussian, and then a series of 1D convolutional layers, also flanked by dropout and batch normalization layers[141]. Here, we simplified the network architecture by replacing the final time distributed layer with a global pooling layer to transform the output from a time series to a single number. We also added a fully connected sigmoid activation layer for binary output. In training the network, we froze all layers except the two newly added layers and the final bi-directional LSTM layer to prevent overfitting.

To train the network, we used the same parameters employed in the original network except for the number of epochs and batch size (here, 128 epochs at a batch size of 32)[141]. We downsampled the EEG data (from 2035 Hz to 1250 Hz) to match the original network[141], using the scipy signal resample[98] method. We applied the network to a 0.5 s interval of data, centered at the time of the event. For each time series, we report a binary label (0 or 1) indicating the absence or presence of a spike ripple, respectively. We shuffled the data to avoid learning of order before training.

Training and Testing

We trained each detector in three ways: (1) using only the *in vivo* data, (2) using only synthetic data, and (3) using a combination of *in vivo* and synthetic data. (1) First, we implemented a leave-one-out (LOO) approach for training and testing each LSTM architecture to evaluate performance across all available data, while minimizing bias towards specific subjects and maximizing performance generalizability. For robust testing of detector performance, we selected six LOO subjects (Fig. 14), each with at least 23 spike ripple events (mean 113, maximum 207). There was a total of 677 spike ripples combining all 6 LOO subjects. The size of the training set depends on the LOO subject. For each excluded LOO subject, the training set consisted of all spike ripples from the remaining subjects (mean 569, minimum 475, maximum 655), and an equivalent number of events without spike ripples. For the latter, events were selected randomly from all available subjects (n=31) with half of the events containing a spike alone (i.e., a spike without a ripple) and half of the events consisting of background activity verified to not contain a spike or spike ripple[75]. In this way, we created a balanced training set. We note that the composition of events from the 28 subjects with fewer than 23 spike ripples consisted of 1246 background events, 393 spike events, and 5 spike ripples (S0 in Fig. 14). We also note that for every LOO subject, we train a newly instantiated model for the frequency-domain detector and tune the best performing instance of the time-domain detector. In this way, we do not iteratively update the neural network model for each LOO subject, but instead tune a new instance of a neural network model for each LOO subject. We note that, for both the frequency-domain and time-domain methods, training resulted in 6 networks, one for each left out subject. We then evaluated the performance of each trained neural network model exclusively on data from the left-out subject.

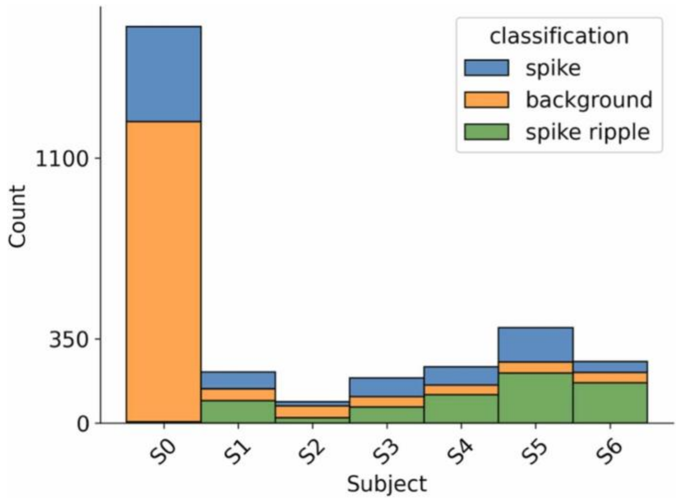

Figure 14. Event counts for the subject population. Number of spike ripple, spike, and background events for each LO subject (see legend). We test classifier performance on 6 subjects (labels S1, S2, … S6) and trained classifiers on all other subjects. The bar labeled S0 includes event counts for all other subjects (n=28).

(2) Second, we trained each LSTM architecture using synthetic data. The training set consisted of K synthetic spike ripples, and K events without a spike ripple, for a total of 2K events, where K = {4000, 6000, 8000, 10000}. Events without a spike ripple included all background events from all subjects not including the six LOO subjects (i.e., all events verified to not contain a spike or spike ripple from subjects with less than 23 spike ripples; total 1246 events) combined with the number of synthetic spike events required to achieve a total of K events. For the frequency-domain method, from the resulting set of 2K events for training, 10% were withheld for validation. For validation of the time-domain method, we first generated another K/10 synthetic spike ripples and selected 10% of background events from all available subjects. We then generated additional synthetic events without a spike ripple so that the sum of these synthetic events and background events totaled K/10; the resulting validation set for the time-domain method therefore consisted of 2K/10 samples or 10% of the training set. We note that, for both the

frequency-domain and time-domain methods, training resulted in one network, and we evaluated performance on each left out subject.

(3) Third, we trained each LSTM architecture using a combination of *in vivo* and synthetic data. To do so, we began with the networks trained using only synthetic data, as described in (2) above. Then, for each LOO subject, we applied transfer learning to update these pretrained networks. For the frequency-domain method, we replaced the fully connected layer and the classification layer. To slow down learning in the transferred layers, we set the initial learning rate to 1e-4, and increased the learning rate to 1e-3 in the new fully connected layer. We also reduced the number of epochs to 50. For the time-domain method, we began with the network trained on the synthetic data and then fine-tuned the network using the *in vivo* recordings with the same parameters and LOO protocol previously described in (1).

Finally, we note that we initially assessed the performance of the time-domain method (RippleNet[141]) applied to the six LOO subjects without additional training. Because the original time-domain method classifies each time point in a signal, we applied the following labeling strategy. For each spike ripple, we labeled the time indices spanning 100 ms before and 50 ms after the event's center as positive for a continuous spike ripple event. To evaluate the Receiver Operating Characteristic (ROC) curve, we considered the highest returned probability value from the original time-domain method around the event center (100 ms before the event center to 50 milliseconds after) for each event. Similar to the approach in[141], we subsequently classified the event as a spike ripple if the returned probability value exceeded a threshold of 0.229 for at least 25 ms. We determined this threshold using all data excluding the LOO subjects, as described below.

Calculation of Probability Thresholds and Performance Statistics

We used two probability thresholds to calculate performance for each network trained on each LOO subject. We chose $p = 0.5$, reported below as $p_{0.5}$, as a naive threshold, consistent with the simple assumption that $p > 0.5$ indicates a positive classification of a spike ripple. We also chose a threshold $p_{validation}$ calculated from the validation set used to train the networks for each LOO subject. We computed this threshold as the optimal operating point of the ROC curve for each validation set (following the algorithm defined in the function *perfcurve* in MATLAB R2023a). We note that this procedure resulted in different probability thresholds for each LOO subject.

We report the mean and standard deviation of 6 performance statistics (specificity, sensitivity, positive predictive value (PPV), negative predictive value (NPV), accuracy, and area under the receiver operator characteristic (ROC) curve) for each left out subject and each probability threshold, according to the following definitions:

$$\text{Sensitivity} = \frac{TP}{TP + FN}$$

$$\text{Specificity} = \frac{TN}{TN + FP}$$

$$\text{Positive Predictive Value} = \frac{TP}{TP + FP}$$

$$\text{Negative Predictive Value} = \frac{TN}{TN + FN}$$

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

(Eq. 26,27,28,29,30)

where TP and TN indicate true positive and true negative, respectively, and FP and FN indicate false positive and false negative, respectively.

Logistic regression

To characterize relationships between signal features and detector classifications, we fit logistic regression models using detector classifications as response variables and estimated signal features as predictors. Models were built stepwise starting with a constant/intercept-only model and allowing linear or interaction terms. Terms were added if doing so significantly reduced the model deviance (p-value of the F-statistic less than 0.05). Fitting was implemented in MATLAB using the *stepwiseglm* function with a binomial distribution and canonical (logit) link function. We report the estimated effect size and p-value of the t-statistic for each coefficient in the complete model; p-values were based on estimated dispersion — the sum of squared Pearson residuals divided by the degrees of freedom for error (DFE).

Results

Data augmentation produces synthetic spike ripples consistent with in vivo observations

Accurate training and assessment of neural networks to detect spike ripples requires sufficient training and testing data. However, creating a database of spike ripples is an extremely time-consuming task; the gold standard for spike ripple detection is manual inspection by a trained expert. Therefore, data for training and testing remain limited, thus limiting detector performance.

One strategy to address this limitation is data augmentation. In general, data augmentation involves generating additional training data, typically by adding noise (as in RippleNet[141]) or through some other manipulation (e.g., image stretch/skew/recolor)[75,154]. Another approach is Synthetic Minority Over-sampling (SMOTE), where new samples are generated by interpolating between existing samples[155].

A limitation of these techniques is the assumption that new samples are near the space of training samples. In practice, such a restriction may not hold. For example, given the relatively small number of subjects available for training, and the variability inherent in brain voltage recordings, the spike ripple features for a novel subject may differ from the training data. Our aim here is to generate an expanded parameter space that includes samples experts would classify as spike ripples, without necessarily having seen a similar example. To this end, we included features of spike ripples derived from the data analyzed here and the literature to guide construction of the synthetic data. Where we did not find descriptions of specific parameters (for example, pertaining to the duration of the rising and falling portions of the wave component), we identified plausible parameter ranges by fitting the spike ripple model to our *in vivo* observations (see *Methods*). We finalized the parameter selection by iteratively adjusting parameters and inspecting the resultant synthetic traces (Fig. 15). The range of parameters is shown in Table 6.

Each simulated spike ripple required 15 parameters (8 spike parameters and 7 ripple parameters). Parameters were selected independently from uniform distributions over the ranges in Table 6. We verified that each parameter set generated relationships that satisfied the constraints given by the 8 computed features in Table 6. Parameter sets that failed to satisfy these constraints were rejected and a new set of parameters was selected. To simulate a signal without a spike-ripple, we reduced the amplitude of the ripple by resampling the parameter A from a uniform distribution in the range [0, 2].

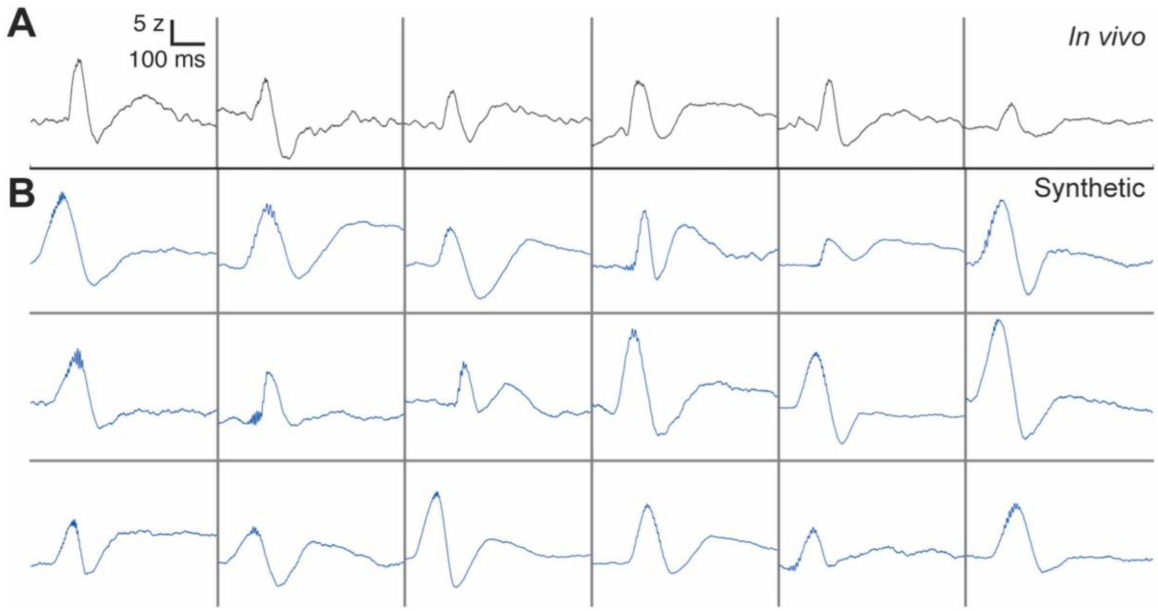

Figure 15. Synthetic spike ripples mimic in vivo spike ripples. Examples of (A) in vivo and (B) synthetic spike ripples. Scale bars indicate 100 ms (horizontal) and 5 z (vertical).

N = 677	Units	Range synthetic	Ref	Mean (StDev; sigma)	Median (IQR)	Range (Q [.05, .95])
SPIKE						
w_1 spike ascent duration	ms	[10, 100]	[15,156–159]	37.4 (3.0; 11.0)	36.2 (12.1)	[23.5, 58.4]
t_{peak} spike peak time	ms	[-50, 50]		9.3 (4.9; 7.6)	7.9 (11.7)	[-4.5, 23.9]
w_2 spike descent duration	ms	[10, 100]	See w_1	51.1 (5.5; 8.0)	52.9 (10.9)	[40.8, 69.7]
w_3 wave ascent duration	ms	[60, 200]		140.1 (19.3; 47.4)	125.7 (43.3)	[83.5, 283.2]
w_4 wave descent duration	ms	[10, 700]		226.0 (33.5; 113.7)	190.4 (94.0)	[71.6, 500.0]
v_{peak} spike amplitude	z	[5, 20]	[156]	8.0 (2.8; 2.2)	7.5 (6.7)	[3.0, 14.1]
v_{tr} trough depth	z	[-8, 1]	[160]	-3.3 (1.1; 1.8)	-3.0 (2.3)	[-6.7, -0.2]
v_{wave} wave amplitude	z	htTrough + [1, 12]	[160]	2.9 (1.3; 2.0)	2.5 (3.3)	[0, 7.6]

COMPUTED						
v_{peak}/w_1 Spike ascent rate	z/ms	<0.5	161	0.24 (0.08; 0.07)	0.22 (0.17)	[0.07, 0.40]
$(v_{peak} - v_{tr})$ / w_2 Spike descent rate	z/ms	<0.5		0.23 (0.08; 0.05)	0.22 (0.15)	[0.10, 0.35]
$(v_{wave} - v_{tr})$ / w_3 Wave ascent rate	z/ms	<0.5		0.05 (0.019; 0.020)	0.05 (0.04)	[0.01, 0.101]
v_{wave} / w_4 Wave descent rate	z/ms	<0.5		0.018 (0.009; 0.02)	0.01 (0.019)	[0, 0.05]
$w_3 + w_4$ Wave width	ms	[150, 785]	162	366.1 (39.1; 118.9)	332.8 (169.4)	[202.4, 635.0]
v_{peak} / (Ac) Spike-to-ripple amplitude ratio	NA	<150		93.3 (43.4; 39.1)	71.5 (51.5)	[30.5, 195.4]
$v_{peak} - v_{tr}$ Peak-trough difference	z	>4		11.3 (3.8; 2.3)	11.4 (8.0)	[5.4, 18.0]
v_{wave} / v_{peak} Wave-spike amplitude ratio	NA	<1		0.36 (0.06; 0.26)	0.35 (0.33)	[0, 0.80]

RIPPLE						
A Ripple amplitude SNR	NA	[5, 25]	[14,149,163]	6.1 (1.1; 1.9)	5.8 (2.9)	[3.3, 10.7]
μ_f mean frequency	Hz	[80, 200]	[14]	135.0 (4.7; 15.2)	135.7 (16.7)	[109.7, 157.0]
μ_t mean time	ms	[-30, 10]	[14]	-9.0 (5.0; 7.6)	-7.01 (10.5)	[-23.8, 3.9]
σ_f width frequency	octaves	[f(fMu), 30]	[118,149]	0.34 (0.038; 0.15)	0.27 (0.12)	[0.20, 0.62]
σ_t width time	ms	0.2	[14,118,149]	16.0 (1.1; 5.9)	14.5 (5.4)	[10.4, 25.9]
ρ time-frequency correlation	NA	0		-0.24 (0.049; 0.32)	-0.26 (0.44)	[-0.68, 0.33]
c background amplitude	z	[0.013, 0.032]		0.017 (0.006; 0.002)	0.017 (0.006)	[0.0078, 0.029]

Table 6. Summary statistics for each spike ripple feature. In vivo statistics are based on 677 observations classified as spike ripples from the six subjects of interest included in the testing sets.

Mean (StDev; sigma): Estimated subject mean and between subject standard deviation (StDev); sigma indicates within subject standard deviation.
Median (IQR): Median and interquartile range (difference between the first and third quartiles, $Q_{0.25} - Q_{0.75}$) across all observations.
Range: Quantiles (Q) 0.05–0.95.
f(fMu) is 1.5/fMu * 1e3 (i.e., we set tSig to accommodate 3 oscillations at frequency fMu within ± 1 tSig; scaling by 1e3 is to convert to units of ms)

Data augmentation moderately improves performance of a frequency-domain spike ripple detector

We first consider a frequency-domain method to detect spike ripples. We implement an LSTM architecture operating on 24 features (the power in 8 frequency bands from three, 250 ms intervals) to assign the probability of a spike ripple for each event (see *Methods: LSTM Frequency Domain Method* and [23]). Without data augmentation, we find good detector performance (Fig. 16A). Using a leave-one-out cross-validation procedure across subjects, we find a mean area under the receiver operator characteristic (ROC) curve (AUC) of 0.80 (standard deviation 0.09). Applying a probability threshold of 0.5, we find good sensitivity (mean 0.67), specificity (mean 0.77), PPV (mean 0.78), NPV (mean 0.67), and accuracy (mean 0.72). We find consistent results for an alternative probability threshold computed in the validation data as the optimal operating point in the ROC curve (Table 7; see Methods).

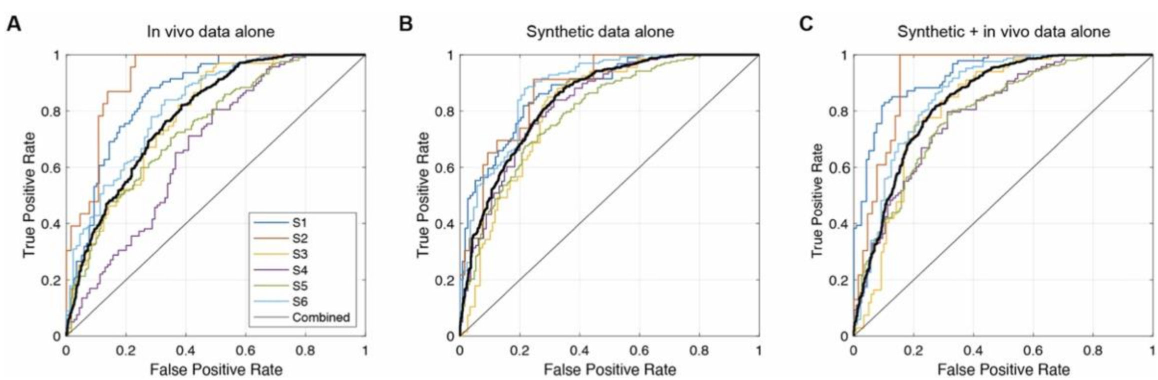

Figure 16. Data augmentation improves performance of a frequency-domain spike ripple detector. Receiver operator characteristic (ROC) curves with training data consisting of (A) in vivo data alone, (B) synthetic data alone, or (C) in vivo data augmented with synthetic data. Colored curves indicate individual subjects; see legend in (A). Black curves indicate the results for classifications across all subjects.

ROC AUC	Mean	Standard Deviation
In vivo data alone	0.80	0.09
Synthetic data alone (K=10,000)	**0.85**	**0.04**
Synthetic data alone (K=8,000)	0.84	0.05
Synthetic data alone (K=6,000)	0.84	0.05
Synthetic data alone (K=4,000)	0.81	0.04
Synthetic + *in vivo* data (K=10,000)	**0.85**	**0.06**
Synthetic + *in vivo* data (K=8,000)	0.86	0.06
Synthetic + *in vivo* data (K=6,000)	0.86	0.06
Synthetic + *in vivo* data (K=4,000)	0.85	0.07

Table 7. **ROC statistics for the frequency-domain spike ripple detectors**. K indicates the number of synthetic spike ripples.

Training with synthetic data produces marginal improvements in detector performance (Fig. 16B). Using synthetic data alone with K=10,000 synthetic spike ripples to train the detector, we find an improved AUC (mean 0.85, standard deviation 0.04), and improved mean values for sensitivity, NPV, and accuracy at both rate thresholds (Table 8). Augmenting the original data with the synthetic data (again with K=10,000 synthetic spike ripples) and performing a leave-one-out cross validation, we find further improvements in performance (Fig. 16C); the AUC (mean 0.85, standard deviation 0.06), and mean sensitivity, specificity, PPV, NPV, and accuracy all increase compared to the *in vivo* data alone (Table 8). We find similar performance when training the detector with 4000, 6000, or 8000 synthetic spike ripples (Table 7), consistent with previous work showing performance plateaus with relatively small sample sizes[26]. We conclude that data augmentation marginally improves performance of the frequency-domain spike ripple detector.

	Sensitivity	**Specificity**	**PPV**	**NPV**	**Accuracy**
		***In vivo* data alone**			
$p_{0.5}$	0.67 (0.12)	0.77 (0.16)	0.78 (0.12)	0.67 (0.08)	0.72 (0.06)
$p_{validation}$	0.67 (0.14)	0.76 (0.18)	0.79 (0.12)	0.67 (0.05)	0.73 (0.07)
		Synthetic data alone			
$p_{0.5}$	0.76 (0.10)	0.74 (0.07)	0.78 (0.07)	0.71 (0.10)	0.76 (0.03)
$p_{validation}$	0.79 (0.08)	0.68 (0.10)	0.74 (0.10)	0.74 (0.07)	0.75 (0.04)
		Synthetic + *in vivo* data			
$p_{0.5}$	0.73 (0.11)	0.81 (0.14)	0.83 (0.10)	0.72 (0.07)	0.78 (0.06)
$p_{validation}$	0.70 (0.16)	0.83 (0.16)	0.85 (0.10)	0.72 (0.09)	0.77 (0.07)

Table 8. Performance metrics for the frequency-domain spike ripple detector. Results for (top) in vivo data alone, (middle) synthetic data alone, and (bottom) in vivo data augmented with synthetic data. In each case, we indicate the mean value (standard deviation) across the n=6 subjects used for testing, with two different probability thresholds. PPV: positive predictive value, NPV: negative predictive value.

Data augmentation moderately improves performance of a time-domain spike ripple detector

We also consider a time-domain method to detect spike ripples. We begin with a network pre-trained to detect sharp wave ripples in rodents[141]. Because spike ripples and sharp wave ripples share similar electrographic features, we expect this existing network will detect spike ripples, and that additional tuning of this network will improve performance.

The input to the network consists of raw time series data (here, an EEG recording), and the output returns the probability that each time point is part of a ripple event. We collapse the output of the network to a binary classification: whether or not a spike ripple event occurs in the time series. In training the network, we freeze most of the network to prevent overfitting and follow a typical tuning protocol (see *Methods: LSTM Time-Domain Method)*. In what follows, we report performance statistics for the original pre-trained network (i.e., with no tuning), after tuning on the *in vivo* data, after tuning on synthetic data alone, and after tuning on both the synthetic data and *in vivo* data.

ROC AUC	Mean	Standard Deviation
No Tuning	0.67	0.12
In vivo data alone	0.81	0.08
Synthetic data alone (K=10,000)	**0.82**	**0.06**
Synthetic data alone (K=8,000)	0.82	0.06
Synthetic data alone (K=6,000)	0.82	0.06
Synthetic data alone (K=4,000)	0.80	0.07
Synthetic + *in vivo* data (K=10,000)	**0.84**	**0.07**
Synthetic + *in vivo* data (K=8,000)	0.82	0.05
Synthetic + *in vivo* data (K=6,000)	0.83	0.07
Synthetic + *in vivo* data (K=4,000)	0.81	0.08

Table 9. ROC statistics for the time-domain spike ripple detectors. K indicates the number of synthetic spike ripples.

As expected, we find that the model without any tuning performs worst, with a mean AUC of 0.67 (standard deviation of 0.12, Fig. 17A) and lowest accuracy across all approaches implemented here, for both probability thresholds (Tables 9 and 10). Models trained on synthetic data alone (with K=10,000 synthetic spike ripples) or *in vivo* data alone both improve the AUC (mean of 0.82 and 0.81, respectively) and accuracy compared to the model without tuning (Tables 9 and 10, Fig. 17B,C). Augmenting the original data with the synthetic data, we find a marginal improvement in AUC (0.84) and accuracy compared to the models trained with the *in vivo* data alone or synthetic data alone (Tables 9 and 10, Fig. 17D). We again find similar performance when training the detector with 4000, 6000, or 8000 synthetic spike ripples (Table 9). We conclude that, consistent with the frequency-domain method, data augmentation marginally improves performance of the time-domain method. However, we note the marginal improvement in AUC and accuracy of the frequency-domain method (AUC 0.85 and accuracy 0.78 for K=10,000; Tables 7 and 8) compared to the time-domain method (AUC 0.84 and accuracy 0.73 for K=10,000; Tables 9 and 10). While a direct comparison between the two methods was not our goal, we propose that the smaller feature space of the frequency-domain method (24 inputs) compared to the time-domain method (625 inputs) may support improved generalizability of the frequency-domain method.

	Sensitivity	**Specificity**	**PPV**	**NPV**	**Accuracy**
No tuning					
$p_{0.5}$	0.71 (0.16)	0.28 (0.14)	0.37 (0.19)	0.66 (0.18)	0.44 (0.18)
$p_{validation}$	0.82 (0.12)	0.18 (0.12)	0.34 (0.18)	0.66 (0.19)	0.39 (0.18)
***In vivo* data alone**					
$p_{0.5}$	0.58 (0.27)	0.81 (0.12)	0.62 (0.28)	0.70 (0.14)	0.71 (0.09)
$p_{validation}$	0.63 (0.30)	0.77 (0.15)	0.60 (0.28)	0.72 (0.13)	0.71 (0.08)
Synthetic data alone					
$p_{0.5}$	0.32 (0.12)	0.94 (0.01)	0.79 (0.11)	0.62 (0.12)	0.66 (0.08)
$p_{validation}$	0.92 (0.04)	0.57 (0.17)	0.65 (0.11)	0.88 (0.06)	0.73 (0.09)
Synthetic + *in vivo* data					
$p_{0.5}$	0.48 (0.25)	0.91 (0.06)	0.86 (0.07)	0.68 (0.12)	0.73 (0.09)
$p_{validation}$	0.67 (0.28)	0.78 (0.15)	0.74 (0.09)	0.76 (0.10)	0.75 (0.09)

Table 10. Performance metrics for the time-domain spike ripple detector. Results for (top) no tuning, (second) in vivo data alone, (third) synthetic data alone, and (bottom) in vivo data augmented with synthetic data. In each case, we indicate the mean value (standard deviation) across the n=6 subjects used for testing, with two different probability thresholds. PPV: positive predictive value, NPV: negative predictive value

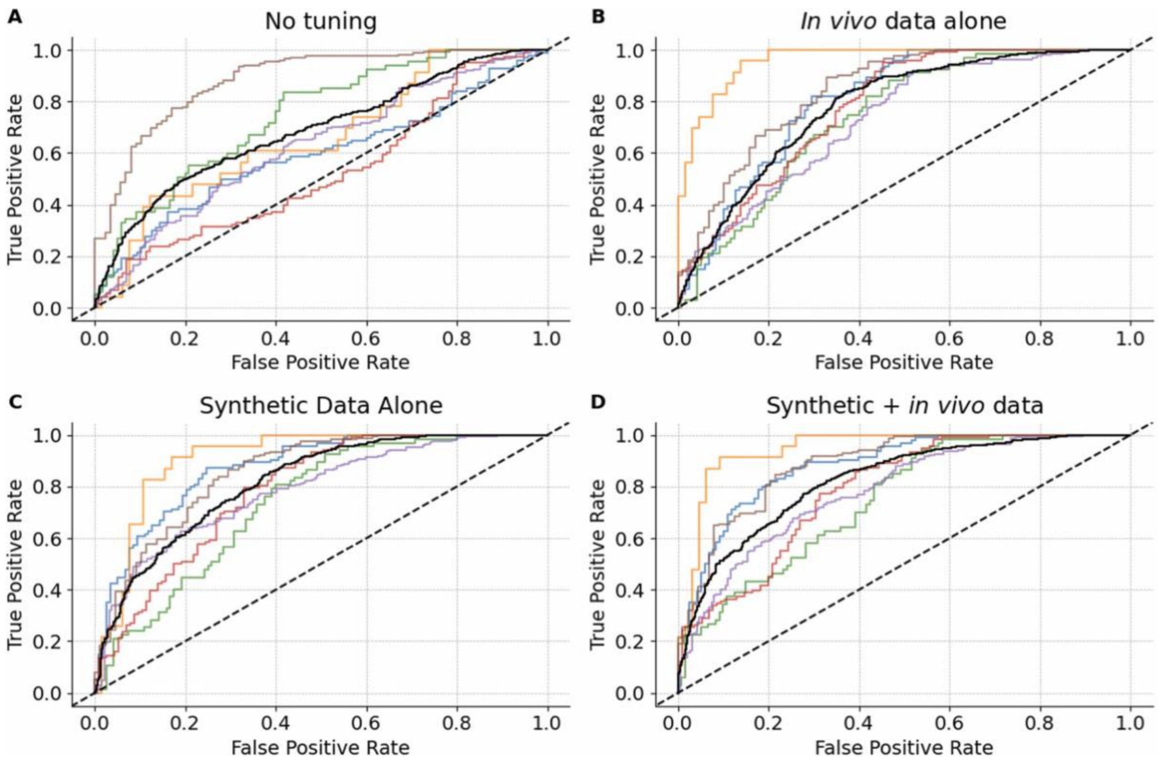

Figure 17. Data augmentation marginally improves performance of a time-domain method. Receiver operator characteristic (ROC) curves with training data consisting of (A) no tuning, (B) in vivo data alone, (C) synthetic data alone, or (D) in vivo data after augmentation with synthetic data. Colored curves indicate individual subjects; see legend in (A). Black curves indicate the results for classifications across all subjects.

Misclassified samples indicate strategies to improve synthetic training data

To understand how the synthetic training data might be improved, we consider observations that were classified correctly using training on *in vivo* data only (training protocol 1 or TP1) and misclassified by training on synthetic data only (training protocol 2 or TP2; TP1 > TP2; Fig. 18). Inspection of these events suggests that the synthetic training set could be improved by including common artifacts and by modeling activity prior to spike onset.

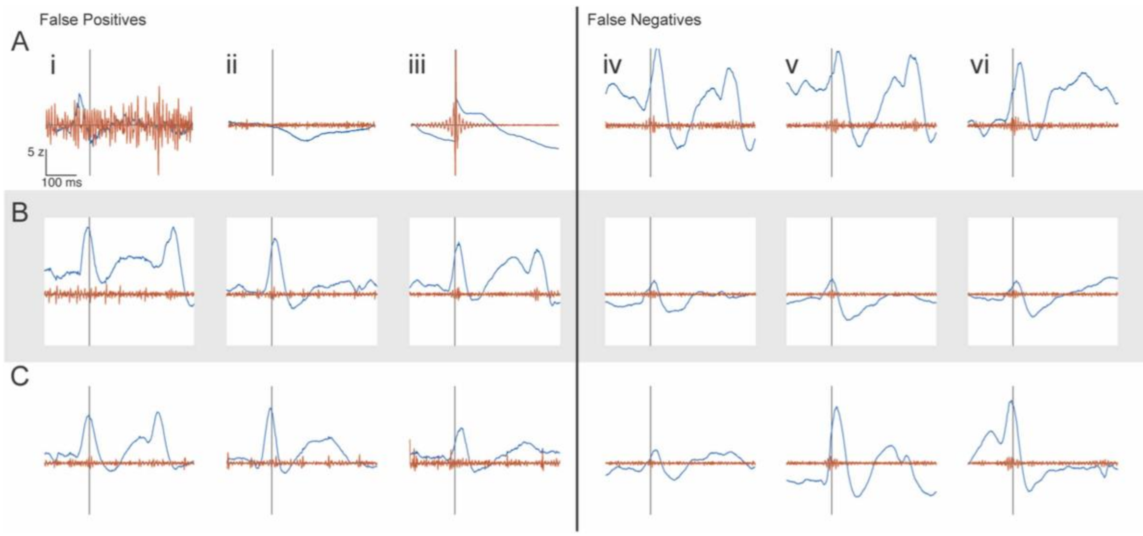

Figure 18. Examples of misclassified observations. Example events correctly classified with training on in vivo data alone (training protocol 1 or TP1) and incorrectly classified with training on synthetic data alone (training protocol 2 or TP2). Blue indicates raw signals; orange indicates band-pass filtered signals (80-500 Hz), scaled by a factor of 5 for visibility. Scale bars (top left) indicate 100 ms (horizontal axis) and 5 z (vertical axis). Examples are given for (A) the frequency-domain detector, (B) the time-domain detector, and (C) when both detectors agree.

In 14% of samples, TP1 > TP2 for the frequency-domain detector (17 false positives, 172 false negatives; Table 11). Inspection of false positives reveals different reasons for misclassification. In one example (Fig. 18A, i), there is a spike wave in the low frequency activity, but the activity in the ripple band is persistent rather than limited to a brief interval during the spike. In another example false positive detection, background activity is classified as a spike-ripple (Fig. 18A, ii). Here, the low frequency activity resembles a low amplitude spike wave. Alternatively, inspection of false negative detections reveals examples of low amplitude spike waves misclassified by the detector as negative for spike ripples (e.g., Fig. 18C, iv). These examples highlight the challenge of interpreting the features identified by deep learning models. In a third example false positive detection, an abrupt shift in the signal causes filter ringing, which the model fails to recognize as artifact (Fig. 18A, iii). Inspection of false negatives reveals multiple examples with high baseline activity before spike onset (Fig. 18A, iv) or large amplitude fluctuations during the wave segment (Fig. 18A, v–vi).

In 17% of samples, TP1 > TP2 for the time-domain detector (143 false positives, 92 false negatives). The most prevalent confound in the false positives is the presence of multiple or repetitive high frequency features, a common EEG artifact[161] (Fig. 18B, i–ii). We also find examples in which the high frequency waveform superimposed on the low-frequency component is too sharp (i.e., more saw-tooth than sinusoid; Fig. 18B, iii). False negatives are typically samples with low amplitudes (Fig. 18B, iv–vi).

In 2% of samples, TP1 > TP2 for both detectors (9 false positives, 15 false negatives; 11). These samples are consistent with the observations noted above. The majority of false positives are samples with multiple high frequency oscillations in the sample window (Fig. 18C, i–iii). False negatives tend to be low amplitude or have large amplitude fluctuations before or after the spike (Fig. 18C, iv–vi).

These observations suggest strategies to improve the synthetic training data, such as: 1) including traces with common artifacts — for example multiple or persistent high frequency oscillations or abrupt shifts in the signal that produce filter ringing[150,151,161,164], 2) modeling of additional low-frequency fluctuations before and after the spike, and 3) decreasing the lower bound on spike peak amplitude (v_{peak}).

N=1372	**LSTM-F**	**LSTM-T**	**Both**
TP1 > TP2 *TP1 is accurate; TP2 misclassifies*	189 (13.8%) 17 FP + 172 FN	235 (17.3%) 143 FP + 92 FN	24 (1.75%) 9 FP + 15 FN

Table 11. Misclassifications resulting from training on synthetic data.
FP: false positive, FN: false negative, TP: training protocol

Interpreting signal features related to misclassifications

An advantage of creating synthetic training data with intuitive parameters is the ability to interpret detector performance. To demonstrate this approach, we estimate a logistic regression model to identify signal features associated with undetected (i.e., false negative) events. Specifically, for all events manually classified by experts as spike ripples ($N = 677$), we define a binary outcome variable consistent with the definition of TP1 > TP2 for both detectors: 1 if the spike ripple was classified correctly using training on *in vivo* data only (TP1) and misclassified by both detectors trained on synthetic data only (TP2), or zero otherwise. In this way, the binary outcome variable represents detections captured by training with *in vivo* data, and missed by training with synthetic data. For this binary outcome variable, we estimate a logistic regression model using the z-scored features estimated for each event as candidate predictors (Table 6). Predictors were added to the model one-at-a-time (stepwise) based on reduction in model deviance (see *Methods*).

We identify three features associated with changes in false negative misclassifications. Decreases in the spike amplitude (v_{peak}) and spike peak time (t_{peak}), and increases in the spike peak to trough difference ($v_{\text{peak}} - v_{\text{tr}}$), increase the odds of an undetected spike ripple (by factors of 5.8 (p = 7e-5), 1.5 ($p = 0.020$), and 2.8 ($p = 0.0044$) for unit changes in the z-scored features, respectively). Consistent with the observations above (Fig. 18C), these results indicate low amplitude spikes and large amplitude fluctuations around the spike — increasing the spike peak to trough difference — increase the odds of an undetected event by detectors trained only on synthetic data. These results also suggest an approach to improve the synthetic training set: including additional spike ripple examples (from *in vivo* data or synthetic) that exhibit these features (e.g., low amplitude spikes) to improve detection of these undetected events.

Discussion

Spike ripples have been proposed as an important biomarker of the epileptogenic zone. However, training neural networks to detect spike ripples remains a challenge due to the difficulty of acquiring gold-standard spike ripple events. To address this challenge, we implemented a new data augmentation approach to generate synthetic spike ripple events. We tested this data augmentation approach using two existing neural networks previously designed to detect spike ripple (or morphologically similar) events. We adapted two Long Short-Term Memory (LSTM) models, differentiated by their architectures and analysis domains — one operating on frequency domain data, and the other on time domain data. We trained these networks systematically across three distinct paradigms: with *in vivo* human EEG recordings alone, with synthetic data alone, and with a sequential approach that initially leverages the synthetic data for training followed by the *in vivo* data. We compared each paradigm's ability to detect spike ripple events. For the frequency-domain and time-domain methods, we found augmentation marginally improved performance, with best performance occurring for networks trained the combined synthetic and *in vivo* data.

Data augmentation is a common step in training a deep learning model. Augmentation commonly involves transforming real observations in a way that does not corrupt the ground truth — for example, by stretching or rotating an image, or by adding a small amount of noise. Another augmentation approach is to interpolate between existing observations (e.g., SMOTE[155]) or generate new data with the same statistics as the training set (e.g., auxiliary classifier generative adversarial networks[26]). Interpolation approaches assume that the training data covers the full space of possible observations, although complete coverage may not be possible for high-dimensional data. Here we attempt to mitigate this issue by generating synthetic data based upon expert knowledge and previous descriptions of spike ripples. Directly comparing the impact of data augmentation methods on spike ripple detection remains an important topic for future work.

Synthetic spike ripples were composed of three components: a low frequency spike component, a high frequency ripple component, and background noise. The spike and ripple components are characterized by 8 and 7 interpretable parameters, respectively. We utilized this number of parameters to capture the most salient, interpretable features of these components (e.g., the spike amplitudes and ripple frequency). While these parameters allow flexibility to capture different types of spike and ripple waveforms, the large number of parameters does make searching the parameter space more difficult. To support a more efficient exploration of the parameter space, future work could reduce the degrees of freedom by eliminating parameters or fixing the ratio between parameters. This interpretability allowed us to generate the synthetic data by referring to previous descriptions, for example to establish reasonable spike widths and ripple amplitudes. This interpretability also allowed us to analyze the detector classifications in the context of these features and identify specific signal features associated with misclassifications.

Understanding these misclassifications suggests strategies to improve detector performance. For instance, we found that increases in spike amplitude reduce the false negative rate, or alternatively that both detectors miss spike ripples with low amplitude spikes. A strategy to improve detector performance, therefore, is to incorporate more examples of observations with lower amplitude spikes in the training data. Finally, we also noted several examples where the detectors mistakenly classified common artifacts as spike ripples (Fig. 18). Therefore, a strategy to improve detector performance is to generate synthetic examples of artifacts and incorporate these into the training data. We examined the impact of data augmentation on two different network architectures. These networks differed in multiple ways, including the type of input feature (time domain versus frequency domain), the degree of complexity, and in the tuning parameters. We found that the frequency-domain method consistently outperformed the time-domain method in classification accuracy. We propose that this difference results from the differing complexity of the approaches, both in the network architectures and in the feature dimensionality. The frequency-domain method compresses the time series data to 8 frequency bands in three 250 ms intervals. The resulting input — with 24 features — creates a more general representation of an event, compared to the specific value at each point of the original time series data. We propose that this more general representation of the input data supports improved generalization of the resulting trained network. In addition, the differing network complexities may contribute to generalizability. Compared to the frequency-domain method, the time-domain method includes an extra bidirectional LSTM layer, and 12 additional normalization, dropout, and convolutional layers. The many additional parameters in the time-domain method may result in overfitting to the training data to minimize loss, degrading generalization. Future work on the time-domain method could involve further optimization, alternative architectures (e.g., removing layers to reduce complexity of the architecture), or alternative training approaches (e.g., training from random instead of pre-trained weights).

Here, we generated synthetic data with the goal of more thoroughly sampling the event space. Generating more synthetic data (e.g., 1 million samples instead of 10,000) would better cover the event space, allowing for the inclusion of rarer event signatures and subtle variations that are less frequently observed. Further, a more diverse training dataset could improve generalizability of the model, reducing the likelihood of model bias towards overrepresented event characteristics. However, reliance on synthetic data assumes that these events accurately represent *in vivo* data, and a much larger training set will increase computational costs.

Here, we demonstrate a new data augmentation procedure to improve detection of spike ripple events. With continued model development and training, coupled with a wider breadth of augmentations, more robust, accurate, and clinically relevant models for spike ripple detection can continue to be developed.

CHAPTER FIVE: Discussion

This dissertation proposes a cohesive set of analytical frameworks designed to significantly enhance neuroscientific data analysis through methodological rigor, interpretability, and computational innovation. Central to these frameworks is the application and validation of Dimensionality Reduction DR techniques and machine learning approaches, systematically developed to uncover meaningful neural dynamics from complex, high-dimensional data.

First, we rigorously validate Nonnegative Matrix Factorization (NMF) as a powerful, interpretable DR method specifically suited to neuronal calcium recordings. By comparing NMF results with those from widely-used DR techniques such as Principal Component Analysis (PCA), Independent Component Analysis (ICA), and Uniform Manifold Approximation and Projection (UMAP), we find that NMF works best to capture underlying biologically relevant network dynamics. Detailed analyses using both simulated neuronal network models and calcium imaging datasets confirm that NMF excels in identifying coherent, biologically plausible neuronal assemblies. Its parts-based, nonnegative decomposition inherently aligns with the physiological characteristics of neuronal calcium signals, enabling clear interpretation of underlying neural structures and interactions. Furthermore, the application of statistical optimization methods, particularly the Akaike Information Criterion (AIC), facilitates consistent unsupervised determination of the optimal number of components for DR, reducing model ambiguity.

Building upon the validated efficacy of NMF, we introduce CaNetiCs, an accessible, modular, and standardized open-source toolbox designed explicitly to democratize advanced DR methods across the neuroscience community. CaNetiCs integrates NMF alongside complementary DR methodologies (PCA, ICA, UMAP) within a unified, user-friendly software environment. CaNetiCs encompasses modules beyond fitting DR models, adding geometric state-space analyses and neuronal network simulations, enabling researchers to not only analyze real datasets but also experiment with analytical strategies through controlled, simulated environments. Crucially, we show that applying modules implemented CaNetiCs to diverse datasets, including neuronal activity from different species such as C. elegans and mice under varying anesthetic conditions, demonstrate its ability to robustly capture fundamentally distinct neuronal dynamics across species. These dynamics align with, and extend, previous analyses, highlighting both species-specific neuronal interactions and common physiological trends such as network suppression and decorrelation under anesthesia. This framework aims to significantly reduce technical barriers, offering neuroscientists without specialized computational expertise a robust platform to implement and interpret sophisticated analyses. Through application to diverse datasets, our software successfully recapitulates known physiological patterns while uncovering novel dynamic features. These capabilities position CaNetiCs as a dynamic tool to enhance both hypothesis-driven and exploratory research approaches in neuroscience.

Finally, we extend an analytical framework beyond DR for single-cell calcium recordings, into a supervised machine learning application, specifically targeting the detection of pathophysiological events in epileptic electrophysiological recordings. By developing data augmentation strategies tailored to electrophysiological characteristics of spike ripples in EEG recordings, our work significantly enhances the accuracy and reliability of deep learning classifiers, notably Long Short-Term Memory (LSTM) neural networks. This methodological innovation addresses critical limitations in current supervised learning approaches, namely the scarcity and variability of labeled training data. The introduction of carefully parameterized synthetic spike ripple data, informed by empirical distributions and neuroscientific expertise, effectively expands the training datasets, thereby improving the robustness and generalizability of trained neural network classifiers. This advancement not only demonstrates deep learning's high efficacy and potential in neuroscientific research but also establishes a replicable strategy for enhancing supervised learning across diverse electrophysiological phenomena. Collectively, by systematically enhancing the interpretability, accessibility, and reliability of neuronal data analyses, these frameworks provide a powerful foundation for deeper exploration of neural network dynamics.

Strengths, Advantages, Limitations, and Practical Considerations of Analytical Approaches

Dimensionality Reduction

One central strength of the DR methodologies evaluated, particularly NMF, lies in their inherent interpretability. Unlike PCA, ICA, and UMAP, which apply orthogonality, independence, or manifold embedding constraints, respectively, NMF offers direct biological interpretability, especially for calcium recordings, due to its nonnegativity and additive decomposition. The "parts-based" representation conferred by NMF ensures components reflect interpretable, biologically plausible neuronal subnetworks. This approach inherently aligns with physiological constraints of neuronal calcium imaging, yielding intuitive mappings between extracted components and underlying neuronal activity. The ability of NMF to represent neuronal activity as combinations of positively weighted subnetworks significantly facilitates biological hypothesis generation and interpretability compared to alternative methods constrained by more complex mathematical constraints.

The CaNetiCs toolbox aims to consolidate validated computational techniques into a unified, accessible platform. The software's modular design aims to enable researchers without extensive computational expertise to easily implement sophisticated DR analyses. By integrating standardized analytical pipelines for NMF, PCA, ICA, and UMAP, in addition to several other modules for standardized activity simulation and analysis, we aim to reduce methodological inconsistencies in neuroscience and promote reproducibility and reliability across diverse experimental contexts.

DR techniques fundamentally assume that underlying neuronal data exhibit intrinsic low-dimensional structures. Such methods rely critically on the assumption that a smaller number of components relative to the total number of neurons effectively captures essential neuronal dynamics. Violations of these assumptions, particularly under conditions of extreme sparsity or saturation in neuronal activity, can significantly compromise the fidelity and interpretability of DR results. Sparse data may lack sufficient activity to robustly define meaningful subnetworks, whereas saturated data can obscure distinct subnetwork structures. Thus, preliminary analyses, such as evaluating variance explained and utilizing information criteria like Akaike Information Criterion

(AIC), are imperative for selecting appropriate dimensionalities and ensuring meaningful, robust analyses.

Deep and Supervised Learning

Deep learning methodologies, notably recurrent neural networks (RNNs) such as Long Short-Term Memory (LSTM) networks, offer exceptional strengths in capturing complex temporal dynamics inherent to neuroscientific data. These approaches excel in identifying subtle yet critical neuronal events, such as spike ripples, that traditional statistical or linear approaches may overlook. Their inherent ability to capture nonlinear relationships and temporal dependencies greatly enhances detection ability. Furthermore, deep learning frameworks benefit substantially from modern data augmentation, transfer learning, and fine-tuning techniques, allowing adaptation to novel or sparse data conditions, thereby significantly bolstering their utility in neuroscience research.

Supervised machine learning inherently depends on accurately labeled training data. Consequently, these approaches are constrained by the reliability, consistency, and precision of initial labeling methodologies. In the context of spike ripple detection, even small labeling inaccuracies or inconsistencies can significantly impact model training and subsequent classification performance. Although advanced mathematical strategies such as data augmentation, active learning, and ensemble labeling can substantially mitigate these limitations, they cannot eliminate label-dependent biases.

Constraints Imposed by Sample Size

The generalizability and robustness of the results presented herein remain fundamentally constrained by sample size limitations. While our analyses validate DR methods and supervised learning approaches, larger datasets incorporating more biological replicates would significantly enhance statistical robustness and generalizability. Expanding sample sizes would strengthen conclusions and lead to novel insights that emerge from the seemingly stochastic neural activity at a granular scale.

Broader Implications for Neuroscience Research

With the analytical frameworks and computational tools developed in this dissertation we aim to facilitate more standardized, reproducible, and widely accessible analytical pipelines. These methodologies are compatible with both hypothesis-driven and exploratory neuroscientific inquiry, fostering systematic investigation of neuronal dynamics across diverse experimental paradigms. Moreover, these frameworks bridge experimental neuroscience with computational biology and clinical research, promoting interdisciplinary collaborations and accelerating the translation of fundamental insights into potential clinical implementations, with advancements in supervised learning for spike ripple detection especially illustrating the potential of these methods in clinical contexts.

Potential Future Directions

While the natural procession for future research would involve application of our designed work to new and diverse datasets, future research directions could also substantially extend these analytical pipelines.

Developing dynamic and adaptive dimensionality reduction methods, which adjust component dimensionality throughout an experimental timeline, present an opportunity to better capture transient neuronal states and rapid shifts in network activity patterns, thereby providing deeper insights into neuronal dynamics' temporal evolution. Additionally, semi-supervised and constrained dimensionality reduction approaches, which integrate anatomical, functional, or connectivity-based priors, could significantly enhance interpretability and biological relevance. By refining identified neuronal subnetworks through such constraints, researchers can advance hypothesis-driven exploration more effectively.

Further, the potential to continue building on a deep-learning foundation using neuronal datasets to predict future neuronal activity patterns, identify neuronal subtypes, or forecast brain states, represents an exciting frontier. Leveraging large-scale neuronal datasets could substantially advance predictive neuroscience, facilitating a deeper understanding of brain function, dysfunction, and dynamic responses across diverse conditions.

Conclusion

In summary, this dissertation represents a substantial advancement in computational neuroscience through the rigorous validation and innovative development of analytical methodologies, comprehensive computational toolkits, and strategic methodological enhancements. By explicitly addressing critical methodological limitations such as dependency on labeled data, intrinsic assumptions in DR, and sample size constraints, this work significantly elevates the interpretability, robustness, and reproducibility of neuroscientific analyses.

The integration of NMF's biologically intuitive, parts-based decomposition alongside the accessible and standardized CaNetiCs platform not only aims to democratize advanced computational methods, but also significantly reduce barriers to interdisciplinary collaboration. Simultaneously, the innovative application of deep learning and tailored data augmentation for detecting subtle electrophysiological events like spike ripples exemplifies the considerable potential of machine learning in translating fundamental neuroscience insights into tangible clinical advances.

Despite inherent limitations, particularly sample size constraints, the presented methodologies lay a foundation for future neuroscientific exploration. Encouraging large-scale collaborative initiatives to pool data and standardize methodologies will further enhance statistical robustness, generalizability, and translational potential. Ultimately, continued expansion, refinement, and implementation of these analytical frameworks promise transformative impacts, facilitating deeper interdisciplinary collaborations, accelerating neuroscientific discoveries, and significantly accelerating our understanding of complex neuronal dynamics and neurological conditions.

BIBLIOGRAPHY

1. Stevenson, I. H. & Kording, K. P. How advances in neural recording affect data analysis. *Nature Neuroscience* **14**, 139–142 (2011).

2. Sejnowski, T. J., Churchland, P. S. & Movshon, J. A. Putting big data to good use in neuroscience. *Nature Neuroscience* **17**, 1440–1441 (2014).

3. Grienberger, C. & Konnerth, A. Imaging Calcium in Neurons. *Neuron* **73**, 862–885 (2012).

4. Dana, H. *et al.* High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nature Methods* **16**, 649–657 (2019).

5. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).

6. Zhang, Y. *et al.* Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature* **615**, 884–891 (2023).

7. Ito, S. *et al.* Large-Scale, High-Resolution Multielectrode-Array Recording Depicts Functional Network Differences of Cortical and Hippocampal Cultures. *PLoS ONE* **9**, e105324 (2014).

8. Huidobro, N. *et al.* Electroencephalographic Biomarkers for Neuropsychiatric Diseases: The State of the Art. *Bioengineering* **12**, 295 (2025).

9. Norman, J. F. *et al.* Reactivation of memory-associated neurons induces downstream suppression of competing neuronal populations. *Proceedings of the National Academy of Sciences of the United States of America* **122**, e2410101122 (2025).

10. Awal, M. R., Wirak, G. S., Gabel, C. V. & Connor, C. W. Collapse of Global Neuronal States in Caenorhabditis elegans under Isoflurane Anesthesia. *Anesthesiology* **133**, 133–144 (2020).

11. Noueihed Noueihed, J. Unsupervised Tracking and Automated Analysis of Multi-Population Neural Activity Under Anesthesia. (Doctoral dissertation, Boston University). https://hdl.handle.net/2144/44762

12. Bhattacharya, S. *et al.* Propofol Anesthesia Alters Cortical Traveling Waves. *Journal of Cognitive Neuroscience* **34**, 1274–1286 (2022).

13. Zhi, W., Li, Y., Wang, L. & Hu, X. Advancing Neuroscience and Therapy: Insights into Genetic and Non-Genetic Neuromodulation Approaches. *Cells* **14**, 122 (2025).

14. Chu, C. J. *et al.* A semi-automated method for rapid detection of ripple events on interictal voltage discharges in the scalp electroencephalogram. *Journal of Neuroscience Methods* **277**, 46–55 (2017).

15. Kramer, M. A. *et al.* Scalp recorded spike ripples predict seizure risk in childhood epilepsy better than spikes. *Brain* **142**, 1296–1309 (2019).

16. Ouzounov, D. G. *et al.* In vivo three-photon imaging of activity of GcamP6-labeled neurons deep in intact mouse brain. *Nature Methods* **14**, 388–390 (2017).

17. Daly, I. *et al.* Electroencephalography reflects the activity of sub-cortical brain regions during approach-withdrawal behaviour while listening to music. *Scientific Reports* **9**, 9415 (2019).

18. Carbonero, D., Noueihed, J., Kramer, M. A. & White, J. A. Nonnegative matrix factorization for analyzing state dependent neuronal network dynamics in calcium recordings. *Scientific Reports* **14**, 27899 (2024).

19. Tlaie, A. *et al.* What does the mean mean? A simple test for neuroscience. *PLoS Computational Biology* **20**, e1012000 (2024).

20. Géron, A. *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*. (O'Reilly Media, Inc, Beijing [China] ; Sebastopol, CA, 2019).

21. Zeimarani, B. *et al.* Breast Lesion Classification in Ultrasound Images Using Deep Convolutional Neural Network. *IEEE Access* **8**, 133349–133359 (2020).

22. Hochreiter, S. & Schmidhuber, J. Long Short-Term Memory. *Neural Computation* **9**, 1735–1780 (1997).

23. Medvedev, A. V., Agoureeva, G. I. & Murro, A. M. A Long Short-Term Memory neural network for the detection of epileptiform spikes and high frequency oscillations. *Scientific Reports* **9**, 19374 (2019).

24. Schlafly, E. D., Carbonero, D., Chu, C. J. & Kramer, M. A. A data augmentation procedure to improve detection of spike ripples in brain voltage recordings. *Neuroscience Research* S0168010224000968 (2024) doi:10.1016/j.neures.2024.07.005.

25. Jing, J. *et al.* Development of Expert-Level Automated Detection of Epileptiform Discharges During Electroencephalogram Interpretation. *JAMA Neurology* **77**, 103 (2020).

26. Geng, D. *et al.* Deep learning for robust detection of interictal epileptiform discharges. *Journal of Neural Engineering* **18**, 056015 (2021).

27. Lee, E. K. *et al.* Non-linear dimensionality reduction on extracellular waveforms reveals cell type diversity in premotor cortex. *eLife* **10**, e67490 (2021).

28. Wenzel, M. *et al.* Reduced Repertoire of Cortical Microstates and Neuronal Ensembles in Medically Induced Loss of Consciousness. *Cell Systems* **8**, 467–474.e4 (2019).

29. Nagayama, M. *et al.* Detecting cell assemblies by NMF-based clustering from calcium imaging data. *Neural Networks* **149**, 29–39 (2022).

30. Ghandour, K. *et al.* Orchestrated ensemble activities constitute a hippocampal memory engram. *Nature Communications* **10**, 2637 (2019).

31. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* **17**, 1500–1509 (2014).

32. Shlens, J. A Tutorial on Principal Component Analysis. Preprint at http://arxiv.org/abs/1404.1100 (2014).

33. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. Preprint at http://arxiv.org/abs/1802.03426 (2020).

34. Lee, D. D. & Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature* **401**, 788–791 (1999).

35. Shlens, J. A Tutorial on Independent Component Analysis. Preprint at http://arxiv.org/abs/1404.2986 (2014).

36. Romano, S. A. *et al.* An integrated calcium imaging processing toolbox for the analysis of neuronal population dynamics. *PLoS Computational Biology* **13**, e1005526 (2017).

37. Devarajan, K. Nonnegative Matrix Factorization: An Analytical and Interpretive Tool in Computational Biology. *PLoS Computational Biology* **4**, 12 (2008).

38. Goldman, M. S. & Fee, M. S. Computational training for the next generation of neuroscientists. *Current Opinion in Neurobiology* **46**, 25–30 (2017).

39. Hoffmann, S. *et al.* The multiplicity of analysis strategies jeopardizes replicability: lessons learned across disciplines. *Royal Society. Open Science.* **8**, rsos.201925, 201925 (2021).

40. Pnevmatikakis, E. A. Analysis pipelines for calcium imaging data. *Current Opinion in Neurobiology* **55**, 15–21 (2019).

41. Stringer, C. & Pachitariu, M. Computational processing of neural recordings from calcium imaging data. *Current Opinion in Neurobiology* **55**, 22–31 (2019).

42. Birkner, A., Tischbirek, C. H. & Konnerth, A. Improved deep two-photon calcium imaging in vivo. *Cell Calcium* **64**, 29–35 (2017).

43. Tonegawa, S., Liu, X., Ramirez, S. & Redondo, R. Memory Engram Cells Have Come of Age. *Neuron* **87**, 918–931 (2015).

44. Josselyn, S. A. & Tonegawa, S. Memory engrams: Recalling the past and imagining the future. *Science* **367**, (2020).

45. Kitamura, T. *et al.* Engrams and circuits crucial for Systems Consolidation of a Memory. *Science* **78**, 73–78 (2017).

46. Norman, J. F., Rahsepar, B., Noueihed, J. & White, J. A. Determining the optimal expression method for dual-color imaging. *Journal of Neuroscience Methods* **351**, 109064 (2021).

47. Sanger, T. D. & Kalaska, J. F. Crouching tiger, hidden dimensions. *Nature Neuroscience* **17**, 338–340 (2014).

48. Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S. I. & Battaglia, F. P. Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *Journal of Computational Neuroscience* **29**, 309–325 (2010).

49. Low, R. J., Lewallen, S., Aronov, D., Nevers, R. & Tank, D. W. *Probing Variability in a Cognitive Map Using Manifold Inference from Neural Dynamics*. http://biorxiv.org/lookup/doi/10.1101/418939 (2018) doi:10.1101/418939.

50. Giovannucci, A. *et al.* CaImAn an open source tool for scalable calcium imaging data analysis. *eLife* **8**, e38173 (2019).

51. Pnevmatikakis, E. A. *et al.* Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron* **89**, 285–299 (2016).

52. Pachitariu, M. *et al. Suite2p: Beyond 10,000 Neurons with Standard Two-Photon Microscopy*. http://biorxiv.org/lookup/doi/10.1101/061507 (2016) doi:10.1101/061507.

53. Pnevmatikakis, E. A. & Giovannucci, A. NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *Journal of Neuroscience Methods* **291**, 83–94 (2017).

54. Zhou, P. *et al.* Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *eLife* **7**, e28728 (2018).

55. Saxena, S. *et al.* Localized semi-nonnegative matrix factorization (LocaNMF) of widefield calcium imaging data. *PLoS Computational Biology* **16**, e1007791 (2020).

56. Bao, Y., Redington, E., Agarwal, A. & Gong, Y. Decontaminate Traces From Fluorescence Calcium Imaging Videos Using Targeted Non-negative Matrix Factorization. *Frontiers in Neuroscience* **15**, 797421 (2022).

57. Zhuang, P. & Wu, J. Reinforcing Neuron Extraction from Calcium Imaging Data via Depth-Estimation Constrained Nonnegative Matrix Factorization. in *2022 IEEE International Conference on Image Processing (ICIP)* 216–220 (IEEE, Bordeaux, France, 2022). doi:10.1109/ICIP46576.2022.9897521.

58. Batty, E. *et al.* BehaveNet: nonlinear embedding and Bayesian neural decoding of behavioral videos. *Advances in Neural Information Processing Systems 32 (NeurIPS 2019.* https://papers.nips.cc/paper_files/paper/2019/hash/a10463df69e52e78372b724471434ec9-Abstract.html.

59. Sotskov, V. P., Pospelov, N. A., Plusnin, V. V. & Anokhin, K. V. Calcium Imaging Reveals Fast Tuning Dynamics of Hippocampal Place Cells and CA1 Population Activity during Free Exploration Task in Mice. *International Journal of Molecular Sciences* **23**, 638 (2022). https://doi.org/10.3390/ijms23020638

60. Rubin, A. *et al.* Revealing neural correlates of behavior without behavioral measurements. *Nature Communications* **10**, 4745 (2019).

61. Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N. & Harvey, C. D. Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell* **170**, 986–999.e16 (2017).

62. Anowar, F., Sadaoui, S. & Selim, B. Conceptual and empirical comparison of dimensionality reduction algorithms (PCA, KPCA, LDA, MDS, SVD, LLE, ISOMAP, LE, ICA, t-SNE). *Computer Science Review* **40**, 100378 (2021).

63. Izenman, A. J. Introduction to manifold learning: Introduction to manifold learning. *Wiley Interdisciplinary Reviews. Computational Statistic*s **4**, 439–446 (2012).

64. Koh, T. H. *et al.* Dimensionality reduction of calcium-imaged neuronal population activity. *Nature Computational Science* **3**, 71–85 (2022).

65. Nagayama, M. *et al.* Sleep State Analysis Using Calcium Imaging Data by Non-negative Matrix Factorization. in *Artificial Neural Networks and Machine Learning – ICANN 2019: Theoretical Neural Computation* (eds. Tetko, I. V., Kůrková, V., Karpov, P. & Theis, F.) vol. 11727 102–113 (Springer International Publishing, Cham, 2019).

66. Briggman, K. L., Abarbanel, H. D. I. & Kristan, W. B. Optical Imaging of Neuronal Populations During Decision-Making. *Science* **307**, 896–901 (2005).

67. Ahrens, M. B. *et al.* Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* **485**, 471–477 (2012).

68. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature Neuroscience* **19**, 1672–1681 (2016).

69. Makino, H. *et al.* Transformation of Cortex-wide Emergent Properties during Motor Learning. *Neuron* **94**, 880–890.e8 (2017).

70. Sheintuch, L. *et al.* Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. *Cell Reports* **21**, 1102–1115 (2017).

71. Miyazaki, T. *et al.* Dynamics of Cortical Local Connectivity during Sleep–Wake States and the Homeostatic Process. *Cerebral Cortex* **30**, 3977–3990 (2020).

72. Pacheco, D. A., Thiberge, S. Y., Pnevmatikakis, E. & Murthy, M. Auditory activity is diverse and widespread throughout the central brain of Drosophila. *Nature Neuroscience* **24**, 93–104 (2021).

73. Vanwalleghem, G., Constantin, L. & Scott, E. K. Calcium Imaging and the Curse of Negativity. *Frontiers in Neural Circuits* **14**, 607391 (2021).

74. Kell, A. J. E., Yamins, D. L. K., Shook, E. N., Norman-Haignere, S. V. & McDermott, J. H. A Task-Optimized Neural Network Replicates Human Auditory Behavior, Predicts Brain Responses, and Reveals a Cortical Processing Hierarchy. *Neuron* **98**, 630–644.e16 (2018).

75. Nadalin, J. K. *et al.* Application of a convolutional neural network for fully-automated detection of spike ripples in the scalp electroencephalogram. *Journal of Neuroscience Methods* **360**, 109239 (2021).

76. Richards, B. A. *et al.* A deep learning framework for neuroscience. *Nature Neuroscience* **22**, 1761–1770 (2019).

77. Zhu, F. *et al.* A deep learning framework for inference of single-trial neural population dynamics from calcium imaging with subframe temporal resolution. *Nature Neuroscience* **25**, 1724–1734 (2022).

78. Cunningham, J. P. & Ghahramani, Z. Linear Dimensionality Reduction: Survey, Insights, and Generalizations. *Journal of Machine Learning Research* **16**(1), 2859–2900. https://jmlr.org/papers/volume16/cunningham15a/cunningham15a.pdf

79. Roh, J., Cheung, V. C. K. & Bizzi, E. Modules in the brain stem and spinal cord underlying motor behaviors. *Journal of Neurophysiology* **106**, 1363–1378 (2011).

80. Cheung, V. C. K., Devarajan, K., Severini, G., Turolla, A. & Bonato, P. Decomposing time series data by a non-negative matrix factorization algorithm with temporally constrained coefficients. in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* 3496–3499 (IEEE, Milan, 2015). doi:10.1109/EMBC.2015.7319146.

81. Akaike, H. Information Theory and an Extension of the Maximum Likelihood Principle. in *Selected Papers of Hirotugu Akaike* (eds. Parzen, E., Tanabe, K. & Kitagawa, G.) 199–213 (Springer New York, NY, 1998). doi:10.1007/978-1-4612-1694-0_15.

82. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830. https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf

83. Esposito, F. A Review on Initialization Methods for Nonnegative Matrix Factorization: Towards Omics Data Experiments. *Mathematics* **9**, 1006 (2021).

84. Hafshejani, S. F. & Moaberfard, Z. Initialization for Nonnegative Matrix Factorization: a Comprehensive Review. Preprint at http://arxiv.org/abs/2109.03874 (2021).

85. Boutsidis, C. & Gallopoulos, E. SVD based initialization: A head start for nonnegative matrix factorization. *Pattern Recognition* **41**, 1350–1362 (2008).

86. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).

87. Kato, H. K., Chu, M. W., Isaacson, J. S. & Komiyama, T. Dynamic Sensory Representations in the Olfactory Bulb: Modulation by Wakefulness and Experience. *Neuron* **76**, 962–975 (2012).

88. Mukamel, E. A., Nimmerjahn, A. & Schnitzer, M. J. Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data. *Neuron* **63**, 747–760 (2009).

89. Cid, E. *et al.* Sublayer- and cell-type-specific neurodegenerative transcriptional trajectories in hippocampal sclerosis. *Cell Reports* **35**, 109229 (2021).

90. Dayan, P., Abbott, L. F. & Abbott, L. F. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. (MIT Press, Cambridge, Mass., 2005).

91. Vogelstein, J. T. *et al.* Spike Inference from Calcium Imaging Using Sequential Monte Carlo Methods. *Biophysical Journal* **97**, 636–655 (2009).

92. Song, S., Sjöström, P. J., Reigl, M., Nelson, S. & Chklovskii, D. B. Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. *PLoS Biology* **3**, e68 (2005).

93. Devarajan, K. & Cheung, V. C. K. On Nonnegative Matrix Factorization Algorithms for Signal-Dependent Noise with Application to Electromyography Data. *Neural Computation* **26**, 1128–1168 (2014).

94. Wirak, G. S., Florman, J., Alkema, M. J., Connor, C. W. & Gabel, C. V. Age-associated changes to neuronal dynamics involve a disruption of excitatory/inhibitory balance in C. elegans. *eLife* **11**, e72135 (2022).

95. Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M. & Harris, K. D. High-dimensional geometry of population responses in visual cortex. *Nature* **571**, 361–365 (2019).

96. Barber, C. B., Dobkin, D. P. & Huhdanpaa, H. The quickhull algorithm for convex hulls. *ACM Transactions on. Mathematical Software* **22**, 469–483 (1996).

97. Edelsbrunner, H. & Mücke, E. P. Three-dimensional alpha shapes. *ACM Transactions on Graphics* **13**, 43–72 (1994).

98. Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* **17**, 261–272 (2020).

99. Bellock, K. alphashape. (2021). https://pypi.org/project/alphashape/1.3.0/

100. Liu, H., Wu, Z., Cai, D. & Huang, T. S. Constrained Nonnegative Matrix Factorization for Image Representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**, 1299–1311 (2012).

101. Jonker, R. & Volgenant, A. A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing* **38**, 325–340 (1987).

102. Huertas, A. & Medioni, G. Detection of Intensity Changes with Subpixel Accuracy Using Laplacian-Gaussian Masks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-8**, 651–664 (1986).

103. Russ, J. C. & Neal, F. B. *The Image Processing Handbook*. (CRC Press, Boca Raton, 2016). doi:10.1201/b18983.

104. Nguyen, J. P., Linder, A. N., Plummer, G. S., Shaevitz, J. W. & Leifer, A. M. Automatically tracking neurons in a moving and deforming brain. *PLoS Computational Biology* **13**, e1005517 (2017).

105. Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362 (2020).

106. Fiest, K. M., Sauro, K. M., Wiebe, S. & Patten, S. B. et al. Prevalence and incidence of epilepsy. *Neurology* **88**(3), 296–303 (2017). https://doi.org/10.1212/wnl.0000000000003509

107. Chabolla, D. R. Characteristics of the Epilepsies. *Mayo Clinic Proceedings* **77**, 981–990 (2002).

108. Fisher, R. S. *et al.* Epileptic Seizures and Epilepsy: Definitions Proposed by the International League Against Epilepsy (ILAE) and the International Bureau for Epilepsy (IBE). *Epilepsia* **46**, 470–472 (2005).

109. Khambhati, A. N. *et al.* Dynamic Network Drivers of Seizure Generation, Propagation and Termination in Human Neocortical Epilepsy. *PLoS Computational Biology* **11**, e1004608 (2015).

110. Geller, E. B. *et al.* Brain ‐ responsive neurostimulation in patients with medically intractable mesial temporal lobe epilepsy. *Epilepsia* **58**, 994–1004 (2017).

111. Heck, C. N. *et al.* Two ‐ year seizure reduction in adults with medically intractable partial onset epilepsy treated with responsive neurostimulation: Final results of the RNS System Pivotal trial. *Epilepsia* **55**, 432–441 (2014).

112. Nair, D. R. *et al.* Nine-year prospective efficacy and safety of brain-responsive neurostimulation for focal epilepsy. *Neurology* **95**, (2020).

113. French, J. A. Refractory Epilepsy: Clinical Overview. *Epilepsia* **48**, 3–7 (2007).

114. Engel, J. Biomarkers in epilepsy: introduction. *Biomarkers in Medicine* **5**, 537–544 (2011).

115. Frauscher, B. *et al.* High‐frequency oscillations: The state of clinical research. *Epilepsia* **58**, 1316–1329 (2017).

116. Guragain, H. *et al.* Spatial variation in high-frequency oscillation rates and amplitudes in intracranial EEG. *Neurology* **90**, (2018).

117. Dümpelmann, M., Jacobs, J. & Schulze‐Bonhage, A. Temporal and spatial characteristics of high frequency oscillations as a new biomarker in epilepsy. *Epilepsia* **56**, 197–206 (2015).

118. Cimbalnik, J. *et al.* Physiological and pathological high frequency oscillations in focal epilepsy. *Annals of Clinical and Translational Neurology* **5**, 1062–1076 (2018).

119. Jacobs, J. *et al.* Interictal high‐frequency oscillations (80‐500 Hz) are an indicator of seizure onset areas independent of spikes in the human epileptic brain. *Epilepsia* **49**, 1893–1907 (2008).

120. Jacobs, J. *et al.* High‐frequency electroencephalographic oscillations correlate with outcome of epilepsy surgery. *Annals of Neurology* **67**, 209–220 (2010).

121. Modur, P. N., Zhang, S. & Vitaz, T. W. Ictal high-frequency oscillations in neocortical epilepsy: implications for seizure localization and surgical resection: Ictal High-Frequency Oscillations. *Epilepsia* **52**, 1792–1801 (2011).

122. González Otárula, K. A., Von Ellenrieder, N., Cuello‐Oderiz, C., Dubeau, F. & Gotman, J. High‐Frequency Oscillation Networks and Surgical Outcome in Adult Focal Epilepsy. *Annals of Neurology* **85**, 485–494 (2019).

123. Nariai, H. *et al.* Prospective observational study: Fast ripple localization delineates the epileptogenic zone. *Clinical Neurophysiology* **130**, 2144–2152 (2019).

124. Boran, E., Sarnthein, J., Krayenbühl, N., Ramantani, G. & Fedele, T. High-frequency oscillations in scalp EEG mirror seizure frequency in pediatric focal epilepsy. *Scientific Reports* **9**, 16560 (2019).

125. Jacobs, J. *et al.* High frequency oscillations (80–500 Hz) in the preictal period in patients with focal seizures. *Epilepsia* **50**, 1780–1792 (2009).

126. Spring, A. M. *et al.* Interrater reliability of visually evaluated high frequency oscillations. *Clinical Neurophysiology* **128**, 433–441 (2017).

127. Buzsáki, G. & Silva, F. L. D. High frequency oscillations in the intact brain. *Progress in Neurobiology* **98**, 241–249 (2012).

128. Kucewicz, M. T. *et al.* High frequency oscillations are associated with cognitive processing in human recognition memory. *Brain* **137**, 2231–2244 (2014).

129. Joo, H. R. & Frank, L. M. The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nature Reviews. Neuroscience* **19**, 744–757 (2018).

130. Alkawadri, R. *et al.* The spatial and signal characteristics of physiologic high frequency oscillations. *Epilepsia* **55**, 1986–1995 (2014).

131. Matsumoto, A. *et al.* Pathological and physiological high-frequency oscillations in focal human epilepsy. *Journal of Neurophysiology* **110**, 1958–1964 (2013).

132. Malinowska, U., Bergey, G. K., Harezlak, J. & Jouny, C. C. Identification of seizure onset zone and preictal state based on characteristics of high frequency oscillations. *Clinical Neurophysiology* **126**, 1505–1513 (2015).

133. Ayala, G. F., Dichter, M., Gumnit, R. J., Matsumoto, H. & Spencer, W. A. Genesis of epileptic interictal spikes. New knowledge of cortical feedback systems suggests a neurophysiological explanation of brief paroxysms. *Brain Research* **52**, 1–17 (1973).

134. Worrell, G. & Gotman, J. High-frequency oscillations and other electrophysiological biomarkers of epilepsy: clinical studies. *Biomarkers in Medicine* **5**, 557–566 (2011).

135. Van Klink, N. E. C. *et al.* Ripples on rolandic spikes: A marker of epilepsy severity. *Epilepsia* **57**, 1179–1189 (2016).

136. Jacobs, J. *et al.* Removing high-frequency oscillations: A prospective multicenter study on seizure outcome. *Neurology* **91**, (2018).

137. Shoeibi, A. *et al.* Epileptic Seizures Detection Using Deep Learning Techniques: A Review. *International Journal of Environmental Research and Public Health* ***18***, 5780 (2021).

138. Zhao, B. *et al.* Integrated Automatic Detection, Classification and Imaging of High Frequency Oscillations With Stereoelectroencephalography. *Frontiers in Neuroscience* **14**, 546 (2020).

139. Zuo, R. *et al.* Automated Detection of High-Frequency Oscillations in Epilepsy Based on a Convolutional Neural Network. *Frontiers in Computational Neuroscience* **13**, 6 (2019).

140. Xu, Z. *et al.* BECT Spike Detection Based on Novel EEG Sequence Features and LSTM Algorithms. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **29**, 1734–1743 (2021).

141. Hagen, E. *et al.* RippleNet: a Recurrent Neural Network for Sharp Wave Ripple (SPW-R) Detection. *Neuroinform* **19**, 493–514 (2021).

142. Korshunova, I. *et al.* Towards Improved Design and Evaluation of Epileptic Seizure Predictors. *IEEE Transactions on Biomedicine Engineering* **65**, 502–510 (2018).

143. Li, X. *et al.* SeizureBank: A Repository of Analysis-ready Seizure Signal Data. *AMIA ... Annual Symposium Proceedings. AMIA Symposium* **2019**, 1111–1120 (2019).

144. Rasheed, K., Qadir, J., O'Brien, T. J., Kuhlmann, L. & Razi, A. A Generative Model to Synthesize EEG Data for Epileptic Seizure Prediction. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **29**, 2322–2332 (2021).

145. Dubey, V. K., Sarkar, S., Shukla, R., Singh, G. & Sahani, A. Epileptic Seizure Stage Classification from EEG Signal Using ResNet18 Model and Data Augmentation. in *2022 IEEE Delhi Section Conference (DELCON)* 1–5 (IEEE, New Delhi, India, 2022). doi:10.1109/DELCON54057.2022.9753352.

146. Wei, Z., Zou, J., Zhang, J. & Xu, J. Automatic epileptic EEG detection using convolutional neural network with improvements in time-domain. *Biomedical Signal Processing and Control* **53**, 101551 (2019).

147. Fritsch, F. N. & Carlson, R. E. Monotone Piecewise Cubic Interpolation. *SIAM Journal on Numerical Analysis* **17**, 238–246 (1980).

148. Heath, M. T. *Scientific Computing: An Introductory Survey*. (McGraw-Hill, Boston, 2002).

149. Worrell, G. A. *et al.* Recording and analysis techniques for high-frequency oscillations. *Progress in Neurobiology* **98**, 265–278 (2012).

150. Otsubo, H. *et al.* High-frequency oscillations of ictal muscle activity and epileptogenic discharges on intracranial EEG in a temporal lobe epilepsy patient. *Clinical Neurophysiology* **119**, 862–868 (2008).

151. Zijlmans, M. *et al.* Ictal and interictal high frequency oscillations in patients with focal epilepsy. *Clinical Neurophysiology* **122**, 664–671 (2011).

152. He, B. J., Zempel, J. M., Snyder, A. Z. & Raichle, M. E. The Temporal Structures and Functional Significance of Scale-free Brain Activity. *Neuron* **66**, 353–369 (2010).

153. Bozinovski, S. Reminder of the First Paper on Transfer Learning in Neural Networks, 1976. *Informatica* **44**(3) (2020). DOI:10.31449/inf.v44i3.2828

154. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning*. (2016). MIT Press.

155. Chawla, N. V., Bowyer, K. W., Hall, L. O. & Kegelmeyer, W. P. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research* **16**, 321–357 (2002).

156. Gotman, J. & Gloor, P. Automatic recognition and quantification of interictal epileptic activity in the human scalp EEG. *Electroencephalography and Clinical Neurophysiology* **41**, 513–529 (1976).

157. Wilson, S. B. & Emerson, R. Spike detection: a review and comparison of algorithms. *Clinical Neurophysiology* **113**, 1873–1881 (2002).

158. Kane, N. *et al.* A revised glossary of terms most commonly used by clinical electroencephalographers and updated proposal for the report format of the EEG findings. Revision 2017. *Clinical Neurophysiology Practice* **2**, 170–185 (2017).

159. Tatum, W. O. *Handbook of EEG Interpretation*. (Demos Medical, New York, 2014).

160. Nonclercq, A. *et al.* Cluster-based spike detection algorithm adapts to interpatient and intrapatient variation in spike morphology. *Journal of Neuroscience Methods* **210**, 259–265 (2012).

161. Park, C. J. & Hong, S. B. High Frequency Oscillations in Epilepsy: Detection Methods and Considerations in Clinical Application. *Journal of Epilepsy Research* **9**, 1–13 (2019).

162. Köksal-Ersöz, E. *et al.* Signal processing and computational modeling for interpretation of SEEG-recorded interictal epileptiform discharges in epileptogenic and non-epileptogenic zones. *Journal of Neural Engineering* **19**, 055005 (2022).

163. Urrestarazu, E., Chander, R., Dubeau, F. & Gotman, J. Interictal high-frequency oscillations (100–500 Hz) in the intracerebral EEG of epileptic patients. *Brain* **130**, 2354–2366 (2007).

164. Blanco, J. A. *et al.* Unsupervised Classification of High-Frequency Oscillations in Human Neocortical Epilepsy and Control Patients. *Journal of Neurophysiology* **104**, 2900–2912 (2010).

CURRICULUM VITAE

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]