

Computational Neuroscience Modeling Frameworks

Oscar Moreau¹, Jonas Popescu², Anna Muller³

¹ Postdoctoral Researcher, Department of Machine Learning, Baltic AI Research University, Tallinn, Estonia. Email: oscar.moreau629@yahoo.com | ORCID: 1658-2792-5735-7659

² Assistant Professor, Department of Computer Science, Central European Tech University, Vienna, Austria. Email: jonas.popescu931@yahoo.com | ORCID: 2181-6318-6560-8017

³ Research Scientist, Department of Computer Science, Mediterranean Institute of Technology, Rome, Italy. Email: anna.muller785@yahoo.com | ORCID: 8403-0112-6755-6107

ABSTRACT

Computational neuroscience seeks to understand brain function through mathematical and computational models that capture neural dynamics at multiple scales -- from individual neuron spikes to whole-brain activity patterns observed in neuroimaging. The choice of modelling framework fundamentally shapes what questions can be addressed: detailed biophysical models (Hodgkin-Huxley, multi-compartment) elucidate single-neuron computation but do not scale to brain-wide simulation, while whole-brain neural mass models capture macroscopic dynamics but abstract away cellular mechanisms. Deep learning models of neural activity predict neural responses with high accuracy but provide limited mechanistic insight. This study developed and benchmarked six computational neuroscience modelling frameworks -- Hodgkin-Huxley biophysical (NEURON), leaky integrate-and-fire spiking networks (NEST), neural mass models (The Virtual Brain, TVB), deep recurrent neural networks (RNN), predictive coding networks, and a proposed Multi-Scale Neural Dynamics Integration Framework (MSNDIF) combining biophysically-constrained cortical microcircuit models, mesoscale population dynamics via mean-field reduction, whole-brain connectivity from the Human Connectome Project, and task-adaptive learned priors from deep predictive coding -- across four neuroscience tasks: single-neuron response prediction (Allen Brain Observatory, 1,248 cortical neurons), cortical microcircuit dynamics (mouse V1, two-photon Ca²⁺ imaging, 8,400 neurons), whole-brain fMRI prediction (HCP task-fMRI, 996 subjects, 7 tasks), and disease state classification (ADNI Alzheimer's dataset, 1,842 subjects). MSNDIF achieved the highest composite accuracy across scales: single-neuron spike prediction $r = 0.84$ (vs Hodgkin-Huxley 0.72, LIF 0.68; $p < 0.001$), microcircuit correlation structure match 0.88 (vs LIF 0.72), whole-brain task prediction AUROC 0.914 (vs TVB 0.824, RNN 0.882), and Alzheimer's classification F1 0.902 (vs RNN 0.862). These results establish multi-scale integrative modelling as the most accurate framework for bridging cellular and brain-wide neural dynamics.

Keywords: Computational neuroscience; Neural modelling; Multi-scale simulation; Spiking neural networks; Neural mass models; Deep learning; Human Connectome Project; fMRI prediction; Brain disease modelling

Citation: Moreau et al. [2024]. Computational Neuroscience Modeling Frameworks. DOI:

<https://doi.org/10.5281/zenodo.19549364>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 May 20 Accepted: 2024 Jul 22 Published: 2024 Sep 15

Research Article: Research Article

1. Introduction

The brain operates across vast spatial and temporal scales: individual neurons fire action potentials in milliseconds, cortical microcircuits integrate signals from thousands of synapses over tens of milliseconds, whole-brain networks exhibit oscillatory dynamics on second-scale timescales, and behavioural adaptation spans minutes to lifetimes. Computational neuroscience provides mathematical frameworks to model these phenomena, bridging molecular neurobiology with systems-level brain function and cognition. The Human Brain Project, the BRAIN Initiative, and the Allen Institute for Brain Science have generated unprecedented datasets -- electrophysiological recordings from tens of thousands of neurons, anatomical connectomes from mouse and human brains, whole-brain functional imaging from millions of scans -- that both enable and demand sophisticated computational models (Sejnowski, 2020).

1.1 The Multi-Scale Modelling Challenge

No single modelling framework adequately captures brain function across all relevant scales. Biophysically detailed models (multi-compartment Hodgkin-Huxley neurons) simulate membrane currents and dendritic integration at biologically realistic accuracy but require approximately 10^6 differential equations per neuron, limiting simulations to small circuits ($<10,000$ neurons). Point neuron spiking models (leaky integrate-and-fire, LIF) simplify each neuron to a single variable, enabling simulations of 10^6 - 10^9 neurons on supercomputers but abstracting away dendritic computation and ion channel dynamics. Neural mass models (NMMs) represent populations of thousands of neurons as mean-field oscillators, enabling whole-brain simulation but losing cellular resolution entirely (Deco et al., 2011). Deep learning models (RNNs, transformers) predict neural activity with high accuracy but offer limited mechanistic interpretation (Pandarinath et al., 2018).

1.2 Data-Driven Neuroscience Modelling

Modern computational neuroscience increasingly uses data-driven approaches that fit model parameters to experimental observations. The Allen Brain Observatory provides patch-clamp electrophysiology for $>1,000$ cortical neurons with standardised stimulus protocols, enabling systematic model fitting. Two-photon calcium imaging captures simultaneous activity of thousands of neurons in behaving animals, providing microcircuit-level ground truth. The

Human Connectome Project (HCP) provides task and resting-state fMRI from 1,200 subjects along with structural connectivity from diffusion MRI, enabling whole-brain model fitting to individual humans (Van Essen et al., 2013). These datasets enable the transition from hand-crafted neuroscience models to data-constrained, individualised brain models that can potentially inform personalised treatment of neurological and psychiatric disorders (Glasser et al., 2016).

1.3 Study Objectives

This study aimed to: (i) develop MSNDIF, a multi-scale neural dynamics integration framework spanning single neurons to whole-brain activity; (ii) benchmark against five established modelling frameworks across four neuroscience tasks at different scales; (iii) evaluate accuracy-biological realism trade-offs; (iv) demonstrate disease-state modelling for Alzheimer's; and (v) assess computational feasibility for individualised brain modelling.

2. Literature Review

Computational neuroscience has developed distinct modelling traditions at different scales, each with complementary strengths and limitations.

2.1 Single-Neuron and Microcircuit Models

The Hodgkin-Huxley (1952) model -- describing action potential generation by voltage-gated sodium and potassium channels -- established biophysical modelling as the gold standard for single-neuron dynamics. NEURON (Carnevale and Hines, 2006) extends HH to multi-compartment dendritic models capturing spatial integration across elaborate neuronal morphologies. The Blue Brain Project simulated a 31,000-neuron reconstruction of rat somatosensory cortex at biophysical detail, reproducing in vivo spiking patterns and responses to stimuli (Markram et al., 2015). Spiking neural networks using LIF or Izhikevich neurons (NEST simulator, Gewaltig and Diesmann, 2007) enable simulation of $>10^6$ neurons on HPC systems, capturing network-level phenomena (oscillations, balanced excitation-inhibition, attractor dynamics) while abstracting single-neuron biophysics.

2.2 Whole-Brain Modelling

Neural mass models (Jansen-Rit, Wilson-Cowan) describe population-averaged firing rates of excitatory and inhibitory populations, reducing each brain region to 2-6 state variables. The

Virtual Brain (TVB, Sanz-Leon et al., 2015) connects NMMs across brain regions using diffusion MRI-derived structural connectivity, simulating whole-brain dynamics that reproduce resting-state fMRI functional connectivity patterns. TVB has been applied to stroke modelling, epilepsy simulation, and consciousness research. The Virtual Epileptic Patient framework fits individualised TVB models to patient MRI and EEG data, predicting seizure propagation with clinical utility (Jirsa et al., 2017). Whole-brain spiking network models (SpiNNaker, BrainScaleS) attempt intermediate scale, simulating 10^5 - 10^6 neurons with approximate biophysics.

2.3 Deep Learning and Predictive Coding

Recurrent neural networks (RNNs) trained on behavioural tasks learn internal representations that match neural activity in prefrontal and parietal cortex, suggesting that optimisation under biological constraints recapitulates cortical computation (Yang et al., 2019). LFADS (Pandarinath et al., 2018) uses variational autoencoders with RNN dynamics to infer latent neural trajectories from population spike recordings, providing data-driven models that decode behaviour with high accuracy. Predictive coding frameworks (Rao and Ballard, 1999; Friston, 2010) propose that the brain implements hierarchical Bayesian inference through predictions flowing top-down and prediction errors flowing bottom-up; deep predictive coding networks (DPC) formalise this computationally and predict neural responses in visual cortex with accuracy matching deep CNNs while providing hierarchical generative structure.

Table 1. Computational neuroscience modelling frameworks evaluated: scale, biological detail, and computational scalability

Framewor k	Scale	Biophys ical Detail	Scala bility	Mul ti-S cale	Reference
Hodg kin-H uxley	Singl e neu ron	Ion cha nnels, d endrites	10^3 neuron s	No	Carnevale and Hines, 2006
LIF (N EST)	Spiki ng ne twork	Point neurons	10^6 - 10^9 neuron s	Limi ted	Gewaltig and Diesmann, 2007
TVB (N MM)	Whol e brain	Populat ion-level	80-400 region s	No	Sanz-Leon et al., 2015

Framewor k	Scale	Biophys ical Detail	Scala bility	Mul ti-S cale	Reference
RNN	Task/ behav iour	Abstrac t units	10^3 - 10^5 units	No	Yang et al., 2019
Predi ctive codin g	Corti cal hi erarc hy	Abstrac t + hier archical	Moder ate	Parti al	Friston, 2010
MSN DIF	Singl e to w hole- brain	Scale-a ppropri ate	Adapti ve	Full inte grati on	This study

LIF: leaky integrate-and-fire. NMM: neural mass model. TVB: The Virtual Brain. Scalability: maximum feasible simulation size. Multi-Scale: ability to simultaneously model multiple scales with consistent interfaces. MSNDIF uses scale-appropriate detail at each level with learned cross-scale couplings.

3. Materials and Methods

3.1 MSNDIF Architecture

MSNDIF integrates four scales through learned cross-scale interfaces. Scale 1 (Single neuron): reduced Hodgkin-Huxley models (3-compartment: soma + apical + basal dendrite) with 9 state variables per neuron, calibrated to Allen Brain Observatory electrophysiological phenotypes (4 cell types: regular spiking, fast spiking, intrinsic bursting, low-threshold spiking). Scale 2 (Microcircuit, 10^3 - 10^4 neurons): cortical column model with empirically measured cell type composition (80% excitatory, 20% inhibitory, 4 inhibitory subtypes), connectivity statistics from Allen Mouse Brain Connectivity Atlas, and receptor kinetics (AMPA, NMDA, GABA_A, GABA_B). Scale 3 (Mesoscale): mean-field reduction of microcircuits produces population-level dynamics for each of 180 cortical regions (HCP multimodal parcellation) with 4 excitatory and 4 inhibitory populations per region. Scale 4 (Whole brain): 180-region network with diffusion MRI-derived structural connectivity (individual-specific for HCP subjects) implementing coupled neural mass dynamics with learned task-adaptive modulation (predictive coding influence on regional excitability). Cross-scale integration: learned interface layers map between scales -- spike trains at Scale 1 drive population rates at Scale 2, microcircuit dynamics at Scale 2 inform mean-field parameters at Scale 3, and whole-brain activity at Scale 4 modulates regional excitability back to Scale 3. Interface parameters learned by matching experimental data at each scale (joint optimisation via gradient

descent on differentiable simulators). Total model parameters: 18.4 M (excluding biophysical ODEs).

3.2 Benchmark Tasks and Datasets

Four neuroscience tasks spanning scales. (1) Single-neuron response prediction: Allen Brain Observatory Cell Types database (1,248 mouse visual cortex neurons with patch-clamp recordings to standardised stimuli: 19 current injection protocols per cell). Task: predict spike timing (Victor-Purpura distance, tolerance 4 ms) given stimulus. (2) Microcircuit dynamics: mouse V1 two-photon Ca2+ imaging (Allen Brain Observatory, 8,400 neurons across 30 recording sessions, natural movie and drifting grating stimuli). Task: predict pairwise neuron correlation structure during spontaneous and stimulus-driven activity (Pearson r between predicted and observed correlation matrices). (3) Whole-brain fMRI: HCP task-fMRI (996 subjects, 7 tasks: working memory, motor, language, social, relational, emotion, gambling). Task: predict task-activated regions from baseline resting-state connectivity (classification AUROC for task identification from whole-brain activity). (4) Disease classification: ADNI Alzheimer's dataset (1,842 subjects: 680 cognitively normal, 896 mild cognitive impairment, 266 AD). Task: classify disease stage from whole-brain fMRI + structural connectivity.

3.3 Baseline Implementations

Five baselines. Hodgkin-Huxley (NEURON): multi-compartment reconstructions with HH channel kinetics fit to Allen Brain Observatory; 10-compartment standard morphologies. LIF (NEST): point neuron spiking networks with empirically calibrated connectivity; 10^5 neurons per microcircuit. TVB (NMM): Wilson-Cowan or Jansen-Rit NMMs on HCP structural connectome; parameters fit to resting-state functional connectivity. RNN: 3-layer gated recurrent units (256 units/layer) trained to reproduce neural activity; end-to-end gradient descent. Predictive coding: hierarchical generative model with 4 levels, variational inference. All models evaluated on identical data splits (80/10/10 train/val/test).

3.4 Evaluation Metrics

Scale-appropriate metrics. Single neuron: Victor-Purpura spike distance (lower = better; converted to correlation $r = 1 - \text{distance}/\text{max}$). Microcircuit: Pearson r between predicted and observed neuron-neuron correlation matrices (8,400 x 8,400); Frobenius norm of matrix difference. Whole-brain fMRI: classification

AUROC for task identification; dynamic functional connectivity correlation. Disease classification: macro-F1 across 3 classes; AUROC per class. Computational efficiency: wall-clock simulation time for 1 second of biological time on standardised hardware (8 x A100 GPU, or 32-core CPU for models without GPU support).

Table 2. Benchmark datasets: four neuroscience tasks spanning single-neuron to whole-brain scales

Task	Scale	Dataset	Subjects/Units	Prediction Target	Key Challenge
Spike prediction	Single neuron	Allen Brain Obs.	1,248 neurons	Spike timing (4ms)	Cell type diversity
Microcircuit	10^3-10^4 neurons	Mouse V1 2P-Ca2+	8,400 neurons	Correlation structure	Network dynamics
Whole-brain fMRI	10^5-10^9 neurons	HCP task-fMRI	996 subjects	Task activation	Individual variability
Disease classif.	Whole-brain	ADNI	1,842 subjects	CN/MCI/AD	Gradual progression

2P-Ca2+: two-photon calcium imaging. HCP: Human Connectome Project. ADNI: Alzheimer's Disease Neuroimaging Initiative. CN: cognitively normal. MCI: mild cognitive impairment. AD: Alzheimer's disease. Tasks span 5 orders of magnitude in scale from single neurons to whole-brain dynamics.

4. Results

4.1 Single-Neuron Spike Prediction

Single-neuron prediction results are presented in Table 3 and Figure 1. MSNDIF achieved spike prediction correlation $r = 0.84$ on the Allen Brain Observatory test set, outperforming Hodgkin-Huxley ($r = 0.72$; $p < 0.001$), LIF ($r = 0.68$; $p < 0.001$), RNN ($r = 0.78$; $p = 0.012$), and predictive coding ($r = 0.66$; $p < 0.001$). MSNDIF's advantage was concentrated in neurons with complex intrinsic dynamics (intrinsic bursting, low-threshold spiking cells) where simplified LIF models fail to capture subthreshold integration. The reduced 3-compartment model with learned cell-type-specific parameters achieved biophysical accuracy comparable to full multi-compartment reconstructions (within 0.03 r) at 20x faster simulation. Per-cell-type performance: regular spiking $r = 0.88$, fast spiking $r = 0.86$, intrinsic bursting $r = 0.82$, low-threshold spiking $r = 0.80$ -- highest for standard pyramidal neurons and

slightly lower for specialised interneurons.

4.2 Microcircuit and Whole-Brain Predictions

Microcircuit and whole-brain results are presented in Table 3 and Figure 2. Microcircuit correlation structure prediction: MSNDIF achieved Pearson $r = 0.88$ with observed two-photon calcium imaging correlation matrices, significantly better than LIF ($r = 0.72$; $p < 0.001$), TVB ($r = 0.64$; $p < 0.001$), and RNN ($r = 0.82$; $p = 0.008$). The improvement reflects MSNDIF's use of empirically measured cell-type-specific connectivity statistics rather than random or uniform connectivity assumed by simpler models. Whole-brain fMRI task prediction: MSNDIF AUROC 0.914 for identifying which of 7 HCP tasks a subject was performing from whole-brain activity, outperforming TVB (0.824; $p < 0.001$), RNN (0.882; $p = 0.006$), and predictive coding (0.864; $p = 0.002$). MSNDIF's task-adaptive modulation module (learned predictive coding priors) captured task-specific regional activation patterns that pure NMM models missed.

4.3 Alzheimer's Disease Classification

Disease classification results are presented in Table 3 and Figure 3. MSNDIF achieved macro-F1 0.902 for classifying ADNI subjects into cognitively normal, MCI, or AD categories, outperforming RNN (0.862; $p = 0.024$), TVB (0.792; $p < 0.001$), and predictive coding (0.786; $p < 0.001$). Per-class performance: CN F1 0.932 (easiest class), MCI F1 0.874 (hardest -- heterogeneous stage), AD F1 0.898. The disease classification accuracy was driven by MSNDIF's integration of whole-brain functional dynamics with structural connectivity: Alzheimer's pathology affects both connectivity (white matter degradation visible in DTI) and function (default mode network dysregulation in fMRI), and models that integrate both modalities outperform structure-only or function-only classifiers.

4.4 Computational Efficiency and Multi-Parameter Comparison

Efficiency and comprehensive comparison are presented in Figure 4. Wall-clock time for 1 second of biological time: MSNDIF 4.6 minutes (whole-brain + cortical column), TVB 2.8 minutes (whole-brain only, no cellular detail), LIF 12.4 minutes (microcircuit scale), Hodgkin-Huxley 48.2 minutes (microcircuit scale), RNN 1.2 minutes (fast but not biophysical), predictive coding 3.8 minutes. MSNDIF's adaptive detail allocation -- detailed biophysics only at the scale of interest, mean-field reduction elsewhere -- provides favourable scalability. MSNDIF dominated across

accuracy metrics while maintaining competitive runtime.

Table 3. Computational neuroscience model performance across four scales and tasks

Metric	HH (NEURON)	LIF (NEST)	TVB	RNN	Pred. Coding	MSNDIF
Spike pred. r	0.72	0.68	N/A	0.78	0.66	0.84*
Microcircuit r	N/A	0.72	0.64	0.82	0.68	0.88*
Whole-brain AUROC	N/A	N/A	0.824	0.882	0.864	0.914*
Alzheimer F1	N/A	N/A	0.792	0.862	0.786	0.902*
Sim time (min/s)	48.2	12.4	2.8*	1.2*	3.8	4.6
Biological detail	Highest	Medium	Low	Minimal	Abstract	Multi-level

*Best in row ($p < 0.05$). N/A: framework not applicable at this scale. Spike pred. r: spike timing correlation (single neuron). Microcircuit r: correlation structure match. AUROC: task classification from whole-brain fMRI. Alzheimer F1: 3-class classification. Sim time: wall-clock minutes per biological second on 8x A100 GPU (or 32-core CPU for CPU-only models).




Figure 1. Single-neuron spike prediction correlation (r) across frameworks




Figure 2. Performance across four neuroscience scales




Figure 3. Alzheimer's disease 3-class classification macro-F1




Figure 4. Multi-parameter neuroscience framework comparison (normalised 0-1)

5. Discussion

This study demonstrates that multi-scale integration -- combining biophysical detail at relevant scales with population-level abstraction elsewhere -- provides the most accurate framework for computational neuroscience modelling across cellular and brain-wide scales.

5.1 The Multi-Scale Advantage

MSNDIF's consistent superiority across all four scales (spike prediction +0.06 r over best alternative, microcircuit +0.06, whole-brain AUROC +0.032, disease F1 +0.040) demonstrates that no single-scale framework can match the integrated approach. The learned cross-scale interfaces enable information flow in both directions: single-neuron dynamics inform microcircuit behaviour, and whole-brain activity feeds back to modulate regional excitability. This bidirectional integration is essential for modelling phenomena that span scales -- for example, Alzheimer's disease begins at the cellular level (amyloid/tau pathology) but manifests as

whole-brain network dysfunction; models that capture only one scale miss the causal chain.

5.2 Toward Personalised Brain Models

The 4.6-minute runtime for whole-brain simulation enables individualised brain modelling at practical timescales: fitting subject-specific MSNDIF parameters to individual HCP data requires approximately 2 hours per subject, making population-scale brain modelling (1,000 subjects) feasible on a modest cluster. Individualised brain models have potential clinical applications: predicting drug response in psychiatric disorders, simulating surgical outcomes for epilepsy, and stratifying neurodegenerative disease patients for personalised intervention. Our Alzheimer's classification F1 0.902 suggests that MSNDIF captures disease-relevant dynamics that could inform clinical trial enrichment and biomarker development.

5.3 Limitations

This study has limitations. The reduced 3-compartment single-neuron model sacrifices some biophysical detail compared to full multi-compartment reconstructions; specific applications (dendritic computation, synaptic plasticity) may require the full NEURON framework. The mouse V1 microcircuit calibration may not transfer to other cortical areas or to human cortex without additional cell-type-specific data. Whole-brain modelling uses HCP's 180-region parcellation; finer parcellation (1,000+ regions) may capture additional dynamics but increases computational cost. The learned cross-scale interfaces lack direct biological interpretation -- they are fit to match experimental data rather than derived from mechanistic principles. Only task-fMRI and ADNI were evaluated for clinical applications; psychiatric disease modelling (schizophrenia, depression) and seizure prediction require separate validation. The 8x A100 hardware requirement limits accessibility.

6. Conclusion

This systematic benchmarking of six computational neuroscience modelling frameworks across four scales establishes MSNDIF -- a Multi-Scale Neural Dynamics Integration Framework combining biophysically reduced single-neuron models, cortical microcircuits with empirical cell-type connectivity, mesoscale mean-field population dynamics, and whole-brain network modelling with task-adaptive predictive coding priors -- as the most accurate framework

for modelling brain dynamics from single neurons to whole-brain activity. MSNDIF achieved spike prediction $r = 0.84$ (vs RNN 0.78; $p = 0.012$), microcircuit correlation $r = 0.88$ (vs RNN 0.82; $p = 0.008$), whole-brain task classification AUROC 0.914 (vs RNN 0.882; $p = 0.006$), and Alzheimer's classification F1 0.902 (vs RNN 0.862; $p = 0.024$). The learned cross-scale interfaces enable bidirectional information flow between cellular and brain-wide dynamics, capturing phenomena that single-scale frameworks miss. Whole-brain simulation in 4.6 minutes per biological second enables individualised brain modelling at practical timescales for clinical applications. MSNDIF code (Python with NEURON + NEST + JAX hybrid), pre-trained models, and HCP/ADNI-calibrated parameters are available at <https://github.com/BAIRU-CompNeuro/MSNDIF-2024> under Apache 2.0 licence. Future work should extend to psychiatric disorder modelling, develop real-time neural interface applications, integrate closed-loop neurostimulation simulation, and pursue clinical validation for presurgical epilepsy planning and neurodegenerative disease progression prediction.

References

- Carnevale, N.T. and Hines, M.L. (2006) The NEURON Book. Cambridge: Cambridge University Press.
- Deco, G., Jirsa, V.K. and McIntosh, A.R. (2011) Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12(1), pp. 43-56.
- Friston, K. (2010) The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), pp. 127-138.
- Gewaltig, M.O. and Diesmann, M. (2007) NEST (NEural Simulation Tool). *Scholarpedia*, 2(4), p. 1430.
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J. and Yacoub, E. (2016) A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615), pp. 171-178.
- Hodgkin, A.L. and Huxley, A.F. (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), pp. 500-544.
- Jirsa, V.K., Proix, T., Perdikis, D., Woodman, M.M., Wang, H. and Gonzalez-Martinez, J. (2017) The Virtual Epileptic Patient: individualized whole-brain models of epilepsy spread. *NeuroImage*, 145, pp. 377-388.
- Markram, H., Muller, E., Ramaswamy, S., Reimann, M.W., Abdellah, M. and Sanchez, C.A. (2015) Reconstruction and simulation of neocortical microcircuitry. *Cell*, 163(2), pp. 456-492.
- Pandarinarath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D. and Kao, J.C. (2018) Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10), pp. 805-815.
- Rao, R.P. and Ballard, D.H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), pp. 79-87.
- Sanz-Leon, P., Knock, S.A., Spiegler, A. and Jirsa, V.K. (2015) Mathematical framework for large-scale brain network modeling in The Virtual Brain. *NeuroImage*, 111, pp. 385-430.
- Sejnowski, T.J. (2020) The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, 117(48), pp. 30033-30038.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E. and Ugurbil, K. (2013) The WU-Minn Human Connectome Project: an overview. *NeuroImage*, 80, pp. 62-79.
- Yang, G.R., Joglekar, M.R., Song, H.F., Newsome, W.T. and Wang, X.J. (2019) Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, 22(2), pp. 297-306.
- Izhikevich, E.M. (2003) Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6), pp. 1569-1572.
- Abbott, L.F. (1999) Lapicque's introduction of the integrate-and-fire model neuron (1907). *Brain Research Bulletin*, 50(5-6), pp. 303-304.
- Deco, G., Kringelbach, M.L., Jirsa, V.K. and Ritter, P. (2017) The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Scientific Reports*, 7(1), p. 3095.
- Yamins, D.L.K. and DiCarlo, J.J. (2016) Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3), pp. 356-365.
- Richards, B.A., Lillicrap, T.P., Beaudoin, P., Bengio, Y., Bogacz, R. and Christensen, A. (2019) A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11), pp. 1761-1770.
- Stimberg, M., Brette, R. and Goodman, D.F.M. (2019) Brian 2: an intuitive and efficient neural simulator. *eLife*, 8, p. e47314.
- Bassett, D.S. and Sporns, O. (2017) Network neuroscience. *Nature Neuroscience*, 20(3), pp. 353-364.

Declarations

Funding

This study was supported by the Estonian Research Council grant PRG2748 (O. Moreau), the Austrian Science Fund (FWF) grant P 34862-N (J. Popescu), and the Italian Ministry of University

and Research (MUR) grant PRIN 2022BK8XP2 (A. Muller). Computational resources provided by the Vienna Scientific Cluster (VSC-5) and University of Tartu HPC Centre. The funding bodies had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

O. Moreau is a co-inventor on a pending patent application (EP24192456.6) covering the MSNDIF multi-scale neural dynamics integration architecture. J. Popescu and A. Muller declare no competing interests.

Data Availability Statement

MSNDIF code (hybrid NEURON + NEST + JAX), pre-trained models, HCP-calibrated parameters, and ADNI-trained disease classifiers are available at <https://github.com/BAIRU-CompNeuro/MSNDIF-2024> under Apache 2.0 licence. Allen Brain Observatory data from alleninstitute.org. HCP from humanconnectome.org. ADNI from adni.loni.usc.edu (requires approved access).

Ethical Approval

This study used publicly available, de-identified datasets under established data use agreements (HCP Restricted Data Use Agreement; ADNI Data Use Agreement). No human subjects were recruited. Animal experiment data (Allen Brain Observatory) collected under Allen Institute IACUC protocols. No additional ethics approval was required.

Appendix A

MSNDIF Multi-Scale Architecture and Cross-Scale Interfaces

This appendix provides the MSNDIF multi-scale architecture, reduced Hodgkin-Huxley formulation, mean-field reduction, and cross-scale interface learning.

A.1 Reduced Hodgkin-Huxley Single-Neuron Model

1a. 3 compartments: soma, apical dendrite, basal dendrite; 9 state variables per neuron (V , m , h , n , Ca , per compartment + 4 coupling)

1b. Ion channels: fast Na^+ , delayed rectifier K^+ , Ca^{2+} , Ca -activated K^+ ; cell-type-specific kinetics from Allen Institute data

2a. 4 cell types: RS (regular spiking), FS (fast spiking), IB (intrinsic bursting), LTS (low-threshold spiking); parameters fit by gradient descent

2b. Simulation: Runge-Kutta 4th order, $dt = 0.025$ ms; GPU-parallel across neurons; 20x faster than NEURON full compartmental at comparable accuracy

A.2 Cross-Scale Interface Learning

3a. Scale 1 $\rightarrow$ 2: spike train statistics (firing rate, ISI distribution, CV) from single neurons aggregated to microcircuit population inputs

3b. Scale 2 $\rightarrow$ 3: microcircuit dynamics (population rates, synchrony) reduced to mean-field 4+4 populations via adiabatic elimination

4a. Scale 3 $\rightarrow$ 4: regional population activity coupled via DTI-derived structural connectivity; delay-weighted neural mass dynamics

4b. Scale 4 $\rightarrow$ 3 feedback: whole-brain state modulates regional excitability (learned task-adaptive prior from predictive coding module)