# GENERATIVE MODELS OF BRAIN DYNAMICS

A REVIEW

**Mahta Ramezanian Panahi**[*]
Mila - Quebec AI Institute
Université de Montréal, Montréal, Quebec, Canada

**Germán Abrevaya**
Departamento de Física, FCEyN, UBA
IFIBA, CONICET, Buenos Aires, Argentina

**Jean-Christophe Gagnon-Audet**
Mila - Quebec AI Institute
Université de Montréal, Montréal, Quebec, Canada

**Vikram Voleti**
Mila - Quebec AI Institute
Université de Montréal, Montréal, Quebec, Canada

**Irina Rish**
Mila - Quebec AI Institute
Université de Montréal, Montréal, Quebec, Canada

**Guillaume Dumas**
Mila - Quebec AI Institute
CHU Sainte-Justine Research Center, Department of Psychiatry,
Université de Montréal, Montréal, Quebec, Canada

## ABSTRACT

The principled design and discovery of biologically- and physically-informed models of neuronal dynamics has been advancing since the mid-twentieth century. Recent developments in artificial intelligence (AI) have accelerated this progress. This review article gives a high-level overview of the approaches across different scales of organization and levels of abstraction. The studies covered in this paper include fundamental models in computational neuroscience, nonlinear dynamics, data-driven methods, as well as emergent practices. While not all of these models span the intersection of neuroscience, AI, and system dynamics, all of them do or can work in tandem as generative models, which, as we argue, provide superior properties for the analysis of neuroscientific data. We discuss the limitations and unique dynamical traits of brain data and the complementary need for hypothesis- and data-driven modeling. By way of conclusion, we present several hybrid generative models from recent literature in scientific machine learning, which can be efficiently deployed to yield interpretable models of neural dynamics.

## Introduction

> *"What I cannot create I do not understand."* — Richard Feynman

The explosion of novel data acquisition and computation methods has motivated neuroscientists to tailor these tools for ad hoc problems. While attempts at pattern detection in enormous datasets are commonplace in the literature–representing a logical first step in applying learning algorithms to complex data—such efforts provide little insight into the observed mechanisms and emission properties. As the above quote from R. Feynman suggests, such methods are *understanding* the brain. The importance of developing interpretable algorithms for biological data– beyond the standard "black-box" models of conventional machine learning–is underscored by the pressing need for superior explainability seen in medical and health-related research. To this end, formal modeling (the practice of expressing some dependent variable unequivocally in terms of some other set of independent variables [1]) is the only way for transparent and reproducible theories [2]. In the present review, we propose that a class of architectures known as generative models constitute an emergent set of tools with superior properties for reconstructing segregated and whole-brain dynamics.

---

[*]Correspondings may be sent to mahtaa [at] gmail [dot] com

A generative model may consist of, for example, a set of equations that determine the evolution of the signals from a human patient based on system parameters. In general, generative models have the benefit over black-box models containing inference mechanisms rather than simple predictive capacity.

**Why prefer inference over prediction?**    The short answer is that arguably, the goal of science is to build upon the prior knowledge and to answer questions of "why" questions rather than some mere forecast of future events and answering "what will be"s. Bzdok and Ioannidis [3] discuss why inference should be prioritized over prediction for building a reproducible and expandable body of knowledge. We argue that this priority should specifically be respected for clinical neuroscience.

It is important to note that modeling is, and should be, beyond prediction [4]. Not only does explicit modeling allow for explanation (which is the main point of science), but it also directs the experiment and engenders new questions.

In this paper, we demonstrate why focusing on the multi-scale dynamics of the brain is essential for biologically plausible and explainable results. For this goal, we review a large spectrum of computational models for reconstructing neural dynamics developed by diverse scientific fields, such as biological neuroscience (biological models), physics and applied mathematics (phenomenological models), and statistics and computer science (data-driven models). On this path, considering the uniqueness of the neural dynamics and the pitfalls in data collection is crucial.

**Neural dynamics is different.**    A neural ensemble is distinctive from the general notion of the dynamical system:

- Unlike chemical oscillations and power grids, the nervous system is a product of biological evolution, which makes it special regarding complexity and organization.
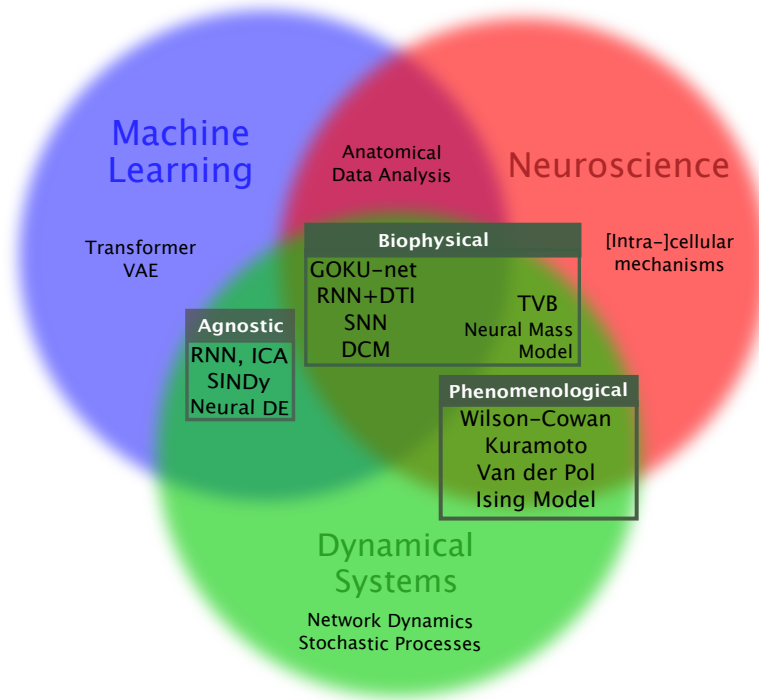


Figure 1: *Venn diagram of the generative models of interest. Based on the abstraction and assumption, methods might belong to one or more of the three worlds of machine learning, neuroscience, and dynamical systems. This review is structured into three main categories that are in fact, intersections of these fields: biophysical (Section 1), phenomenological (Section 2), and agnostic modeling (Section 3). Tools developed independently in each of these fields can be combined to overcome the limitation of data.*

2

- Like most biophysical systems, it is highly dissipative and functions in non-equilibrium regimes (at least while working as a living organ).

- Although the brain exhibits continual neuromodulation, the anatomical structure of the brain is encoded in the genome, hence it is essentially determined [5].

- There are meaningful similarities across species. This is especially good news because, unlike humans, the neural maps of less-complicated species are well-studied [6].

These characteristics help narrow down the search for finding useful models.

**Neural data is different.**   Neural recordings - especially human neural recordings - are often noisy and scarce resources. Difficulties in keeping up with medical certification and regulations and recruiting participants make acquiring these datasets expensive and time-consuming. On the other hand, data is hard to wrangle and contains inconsistent noise across participants and even for a participant at different times (e.g., artifacts, skin condition and time resolution in the case of EEG).

**Overview of generative models**

Current generative models fall into three main categories as shown in Figure 1 with respect to their modeling assumption and objective:

1. **Biophysical models:** Biophysical models are *realistic* models which encapsulate biological assumptions and constraints. Due to the high number of components and the complexity of these systems, these simulations happen on different scales. Blue Brain Project [7] is an instance of this type of modeling.

2. **Phenomenological models:** Analogies and behavioral similarities between neural populations and established physical models open the possibility of using well-developed tools in Statistical Physics and Complex Systems for brain simulations. In such models, some priors of the dynamics are given but not by realistic biological assumptions. A famous example is the model of Kuramoto oscillators [8] in which the goal is to find the parameters that best reconstruct the behavior of the system. These parameters describe the property of the phenomenon (e.g., the strength of the synchrony), although they do not directly express the fabric of the organism.

3. **Agnostic computational:** Data-driven methods that, given a "sufficient" amount of data, can *learn* reconstruct the behavior with little prior knowledge. Examples of such approaches are some self-supervised methods such as latent ODEs [9]. The term "sufficient" carries the main limitation of these approaches: they sometimes need unrealistically large datasets and come with intrinsic biases. The representation that these models give can be analytically far from the Physics of the system or the Phenomenon.

    Figure 2 shows an overview of various generative models and the presence in the literature up to this date.

**Key Contributions:**   The objective is to bridge the gap in the literature of computational neuroscience, dynamical systems, and AI and to review the usability of proposed generative models concerning the limitation of data, the objective of the study and the problem definition, prior knowledge of the system, and sets of assumptions (see Figure 2).

We argue the issues that each set of approaches can resolve and propose hybrid solutions and ideas from other domains (e.g., computer vision and natural language processing) to model brain recording.

# 1   Biophysical Models

Understanding how cognition "emerges" from complex biophysical processes has been one of the main directions of computational neuroscience. Although inferring high-level cognitive tasks from biological processes is far-fetched, different biophysical simulations provide some "explanation" of how neural information relates to behavior. Those attempts are motivated by the need for interpretable realistic models.

While there is no "unifying theory of Neuroscience" yet, biological neuronal models are being developed in different scales and with different degrees of abstraction (see Figure 3). These models are usually grouped into two main categories: the first is the "bottom-up" approach, which emphasizes biophysical details for fine-scale simulation and expects the emergence[2]. An example of this approach is the Blue Brain project [7].

---

[2] Emergence is the manifestation of collective behavior that cannot be deduced from the sum of the behavior of the parts [10, 11].

Conversely, "top-down" schemes focus on explicit high-level functions and designing frameworks based on some targeted behavior. Each of the two approaches works with a different domain of knowledge and has its own pitfalls. The top-down approach can incorporate behavioral insights without worrying about better or hard-to-code biological detail to generate high-level behavior that can be observable. Models of this kind do not provide low-level explanation and are prone to biases related to data collection [12]. The bottom-up perspective, on the other hand, benefits from a customized level of biophysical insight. At the same time, its description is not generalizable to behavior, and it can be difficult to scale (thanks to unknown priors and numerous parallel mechanisms). Also, the reductionist approach to complex systems (here, i.e. brain) is subject to substantial criticism. In particular, while a reductionist approach can help examine causality, it is not enough for understanding how the brain maps onto the behavior [11, 13].

In this section, we review brain models across different scales that are faithful to biological constraints. We focus primarily on the first column from the left in Figure 3, starting from the realistic models with mesoscopic details to more coarse-grained frameworks.

## 1.1 Modeling at the synaptic level

The smallest interacting blocks of the nervous system are proteins [14]. Genetic expression maps and atlases are useful for discovering the functions of these blocks in the neural circuit [15]. However, these maps are not uniformly expressed in the brain [16]. While the expression maps of those proteins continue to unfold [17], combined with connectivity data, they can help quantify dynamics. These maps link the spatial distribution of gene expression patterns, and neural couplings [18] as well as other large-scale dynamics, e.g., dynamic connectivity as a dependent of neurogenetic profile [19].

A notable effort in this regard is the Allen Brain Atlas [20] in which genomic data of mice, humans, and non-human primates [21] have been collected and mapped for understanding structural and functional architecture of the brain [22]. While genomic data by itself is valuable for mapping out connectivity in different cell types, the fifth division, *Allen Institute for Neural Dynamics* was recently announced to start studying the link between neural circuits of laboratory mice and behaviors related to foraging [23].

On a slightly larger scale, a considerable amount of work concerns the relationship between cellular and intracellular events and neural dynamics. Intracellular events and interactions models could generate accurate responses on small [24] and large scales [25]. Some of these models laid the foundation of computational neuroscience and are reviewed in
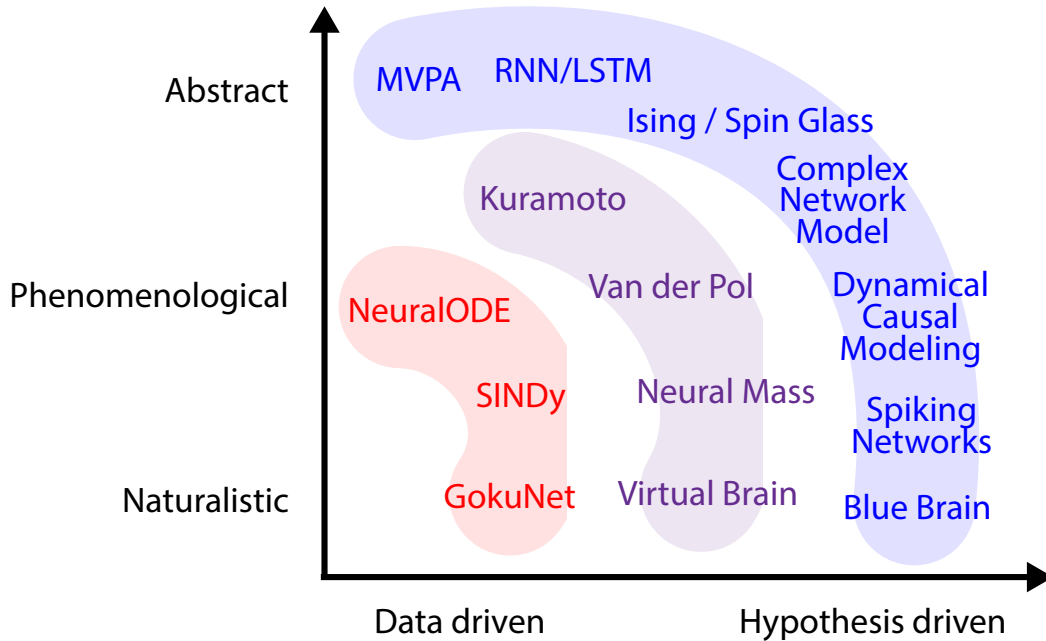


Figure 2: *Overview of generative models: well-developed models (blue), partially-explored approaches (purple), and modern pathways with little or no literature on neural data (red).*

Subsection 1.2. In what follows, we start with neurocomputational models at the mesoscale level (realistic models of small groups of neurons, i.e., the top-left corner of Figure 3), after which we move on towards macro-scale levels with different degrees of abstraction.

## 1.2 Basic biophysics of neurons: A quick history

Zooming a little out of the intra-neuron synaptic level, inter-neuron communications come forth as principal determinants of dynamics. Information transmission is mainly based on the emission of action potentials. The mechanism of this flow of ions was first explained by the influential Hodgkin-Huxley equations and corresponding circuits. The electrical current of the equivalent circuit is described by four differential equations that incorporate membrane capacity and the gating variables of the channels [26]. While the Hodgkin-Huxley model agrees with a wide range of experiments [27, 28] and continues to be a reference for models of ion channels, it needs to be simplified in order to be expandable to the models of the neuronal population. The main difficulty with the Hodgkin-Huxley model is that it requires solving a system of differential equations for each of the gating parameters of each of the single ion channels of a cell while there are more than 300 types of ion channels discovered as of today [29]. Various relaxing assumptions have been proposed, one of which is to dismiss the time dependence of membrane conductance and the dynamics of the action potential by simply assuming the firing happens when the electrical input accumulated at the membrane exceeds a threshold [30]. The latter model is known as integrate-and-fire [31], and it comes in different flavors depending on the form of nonlinearity assumed for the dynamics of leaky or refractory synapses [32].
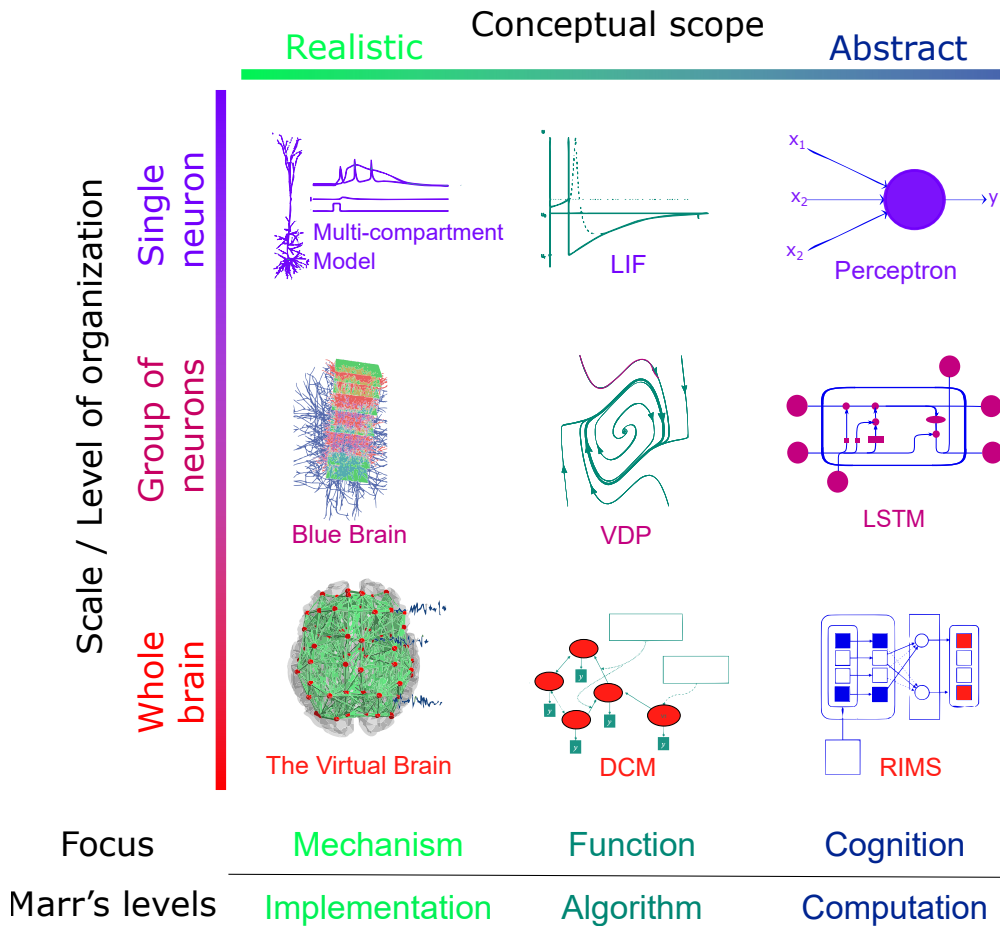


Figure 3: *Instances of modeling across different levels of the organization and problem dimension. The Concept scope is an indicator of biophysical details incorporated in the model. It determines how the focus of the model is directed toward mechanistic reality or the behavioral output. It is also an indicator of where a given model sits on the Marr's level.*

In order to model interesting dynamics of various ion channels, a model of compartments of dendrites, called the multicompartment model, can be employed. An exclusive review article by Herz et al. [33] categorizes compartmental models into five groups based on the level of balance and details involved from Hodgkin-Huxley description to black-box.

While the research on hyper-realistic modeling of many neurons continues, other frameworks focus on simulating the biophysics of the population of neurons. In the Subsection 1.3, we pause on the state of the large-scale synaptic simulations to show how a change of paradigm of computation can overcome their limitation. Models of Neural mass, Wilson-Cowan, and dynamical causation are examples of such alternatives (see Subsections 1.3.1, 1.3.2, and 1.3.3 respectively).

## 1.3   Population-level models

Izhikevich and Edelman [34] describe the first reconstruction of the whole cortex. Their simulation includes a microcircuitry of 22 basic types of neurons with simplified dendrite trees and fewer synapses. The underlying structural data based on the geometry of the white matter is drawn from diffusion tensor imaging (DTI) [35] of the human brain. The microcircuitry of the six-layered neocortex was reconstructed based on cats' visual cortex. The spiking dynamic employed in this model comes from [36] and it is a simplification of the Hodgkin-Huxley model as it outputs the firing rates instead of currents. On a larger scale, some subcortical dynamics (e.g., dopaminergic rewarding from the brainstem) are also implemented.

The significance of this simulation compared to its proceedings is the inclusion of all cortical regions and some of their interplays in the form of cortico-cortical connections. They also considered synaptic plasticity, which is a significant factor in studying developmental changes such as learning. The model also demonstrates several emergent phenomena such as self-sustained spontaneous activity, chaotic dynamics and avalanches, an exhibition of the delta, alpha, and beta waves, and heterogeneous oscillatory activities similar to those in the human brain.

Complexity aside, the model has some shortcomings, including extreme sensitivity to the initial condition. The authors suggest inquiring outcomes on the population level and not for single cells. Despite all the limits, [34] is the first benchmark of whole-cortex modeling and the foundation of future detailed projects such as Blue Brain project [7] and MindScope [37].

Following Izhikevich and Edelman [34], the Blue Brain Project [7] was founded in 2005 as a biological simulation of synapses and neurons of the neocortical microcircuitry. The ambitious goal was to extend this effort to a whole-brain level and build "the brain in a box". The initial simulated subject was only a $2mm$ tall and $210\mu m$ in radius fragment of the somatosensory cortex of a juvenile rat ($\sim 100,000$ neurons). The efforts for further expansions to larger scales, i.e., mouse whole-brain and human-whole brain, are far-fetched by many critics [38].

Far from the initial promise of "understanding" the brain, the Blue Brain Project still has a long way until acquiring a whole map of connections (also known as connectome [39]) in the mouse brain, which is still an order of magnitude smaller than the human brain [40]. That being said, acquiring the connectomic map does not necessarily result in a better understanding of function. Note that while the connectomic structure of the roundworm Caenorhabditis elegans nervous system has been entirely constructed since 1986 [6], research is still unable to explain the behavior of the network, e.g., predicting stimuli based on excitation [41]. Finally, strong concerns regarding the validity of the experiments rise from the fact that the simulation still does not account for the glial cells. Glial cells constitute 90% of the brain cells. They have distinctive mechanisms as they do not output electric impulses [42] but are responsible for inactivating and discharging products of neuronal activities which influence the synaptic properties [43] and consequently learning and cognitive processes [42]. This point of incompleteness sheds extra doubts on the achievability of *brain in silico* from the Human Brain Project.

The above critiques have been calling for a revision of the objectives of the Blue Brain project with more transparency. Hence, new strategies such as the division of Allen Institute, MindScope [44], and the Human Brain Project [45] aim for adaptive granularity, more focused research on human data, and pooling of resources through cloud-based collaboration and open science [46]. Alternatively, smaller teams developed less resource-intensive simulation tools such as Brian [47] and NEST [48].

There are several readily-available simulators of large networks of spiking neurons to reconstruct many-neuron biophysics. Brian [47] is a Python package for defining a customized spiking network. The package can automatically generate the code for simulating a computationally-optimal language (e.g., C++, Python, or Cython). With GPUs available, it can also enable parallelism for faster execution. Brian is more focused on single-compartment models while GENESIS [49] and NEURON [50] center around multicompartment cells.

NEST is another popular package for building ad hoc models of spiking neurons with adjusted parameters. These parameters include the spiking rules (such as IF, Hodgkin-Huxley AdEx), networks (topological or random neural networks), synaptic dynamics (plasticity expressions, neuromodulation) [48].

While working with mid-level packages, Technical limitations and the objective of the study should be considered. These include computational efficiency and the code generation pipeline. Interested researchers are encouraged to refer to the review by Blundell et al. [51] to learn more about the guidelines and proposed solutions.

The steep price of high-resolution computation and the remoteness from high-level cognition can be levitated by replacing detailed dynamics of single neurons with the collective equations of the population. This dimensionality-reduction strategy is the essence of the neural mass models [52], spiking neural network [53], and dynamical causal modeling [54].

### 1.3.1    Neural mass models

Staying faithful to the biophysical truth of the system can happen at scales larger than a few cells. In other words, by reducing the degrees of freedom, one can reduce a massive collection of individual integrate-and-fire equations (mentioned in Subsection 1.2) to a functional DE of the probabilistic evolution of the whole population known as Fokker-Planck DE. However, since Fokker-Planck equations are generally high-dimensional and intractable, a complimentary formalism, known as the mean-field approximation, is proposed for finessing the system [55].

In statistical Physics, the mean-field approximation is a conventional way of lessening dimensions of a many-body problem by averaging over the degrees of freedom. A well-known classic example is the problem of finding collective parameters (such as pressure or temperature) of bulk of gas with known microscopic parameters (such as velocity or mass of the particles) by the Boltzmann distribution. The analogy of the classic gas shows the gist of the neural mass model: the temperature is an emergent phenomenon of the gas *ensemble*. Although higher temperatures correspond to higher *average* velocity of the particles, one needs a computational bridge to map microscopic parameters to the macroscopic one(s). To be clear, remember that each particle has many relevant attributes (e.g., velocity, mass, and the interaction force relative to other particles). Each attribute denotes one dimension in the phase space. One can immediately see how this problem can become computationally impossible even for $1cm^3$ of gas with $\sim 10^{19}$ molecules.

The current state of thermodynamics accurately describes the macroscopic behavior of gas, so why not use this approximation to the many-body problems of neuronal populations? The analogous problem for a neural mass model can be described with the single-neuron activity and membrane potential as the microscopic parameter and the state of the neural ensemble in phase space as the macroscopic parameter. The computational bridge is based on Fokker-Planck equations for separate ensembles.

Neural mass models can be used both for understanding the basic principles of neural dynamics and building generative models [56]. They can also be generalized to neural fields with wave equations of the states in phase space [57] as well as other interesting dynamical patterns [58]. Moreover, these models are applicable across different scales and levels of granularity from subpopulations to the brain as a whole. This generalizability makes them a good candidate for analysis on different levels of granularity, ranging from modeling the average firing rate to decision-making and seizure-related phase transitions. The interested readers are encouraged to refer to the review in [55] to see how neural mass models can provide a unifying framework to link single-neuron activity to emergent properties of the cortex. Neural mass and field models build the foundation for many of the large-scale *in-silico* brain simulations [59] and have been deployed in many of the recent computational environments [60, 61]. Note that the neural mass model can show inconsistency in the limits of synchrony and require complementary adjustments for systems with rich dynamics [62] by mixing with other models of neural dynamics such as Wilson-Cowan [63] as in [59].

### 1.3.2    Wilson-Cowan

Wilson-Cowan is a large-scale model of collective activity of neural population based on mean-field approximation (see Subsection 1.3.1). Seemingly the most influential model in computational neuroscience after Hodgkin-Huxley [26] is Wilson-Cowan [63] with presently over 3000 mentions in the literature.

The significance of this work in comparison to its proceedings (e.g., in [64, 65]) is more than a formal introduction of tools from dynamical systems in neuroscience. This model acknowledges the diversity of synapses by integrating distinct inhibitory and excitatory subpopulations. Consequently, the system is described by two state variables instead of one. Moreover, the model accounts for Dale's principle [66] for a more realistic portrayal. That is to say each neuron is considered purely inhibitory or excitatory. The four theorems proved in the seminal paper [63] conclude the existence

of oscillations as a response to a specific class of stimulus configuration and the exhibition of simple hysteresis for other classes of stimulus.

Wilson-Cowan model lays the foundation for many of the major theoretical advances. Examples of the derivative studies include energy function optimization for formulating associative memory [67], artificial neural networks as a special case with binary spiking neurons [68], pattern formation [69], brain wave propagation [70], movement preparation [71], and Dynamic Causal Modeling [72]. Other studies also demonstrate the possibility of diverse nonlinear behavior of networks of Wilson-Cowan oscillators [73, 74]. More detailed extensions are on the way. For example, second-order approximations [75] and simulation of intrinsic structures such as spiking-frequency adaptation or depressing synapses [76]). For a comprehensive list of continuations, see [77].

### 1.3.3   Dynamical causal models

Deducing the effective connectivity of functionally-segregated brain regions is crucial in developing bio-plausible and explainable models. Dynamical Causal Modeling (DCM), introduced in [54], quantitatively generates the connectivities that fit the observed data by maximizing model evidence, aka marginal likelihood of the model [78].

In a graph where the nodes are functionally-segregated populations, the effective connectivities are found based on three sets of parameters: (1) anatomical and functional couplings, (2) induced effect of stimuli, and (3) the parameters that describe the influence of the input on the intrinsic couplings. Models of the intra-connected regions can be built based on the earlier subsections, e.g., neural mass model, neural fields, or conductance-based models. For a review of such hybrid approaches, see [79].

### 1.3.4   Spiking neural network: Artificial neural networks as a model of natural nervous system

With the introduction of neural networks, the idea of implementing neural circuits and biological constraints into the artificial neural networks (ANN) gained momentum. [80] is an early example that uses ANN with threshold spiking behavior. Despite being oversimplified, their idea formed the basis for a particular type of trainable network known as spiking neural networks (SNN) or biological neural networks (as in [53]). Note that the distinction here with the other forms of spiking networks like Izhinevich's and derivatives (discussed earlier in Subsection 1.3) is that here we are talking about the networks that demonstrate a function approximation as a deep learning algorithm would do.

In contrast to deep neural networks, the activity in this architecture (transmission) is not continuous in time (i.e., during each propagation cycle). Instead, the activities are event-based occurrences with the event being the action potential depolarization[3]. Although ANN architectures that are driven by spiking dynamics have been long used for optimization problems such as pattern recognition [82] and classification [83], they lag behind conventional learning algorithms in many tasks, but that is not the end of the story.

Maass [84] argues that concerning network size, spiking networks are more efficient in computation compared to other types of neural networks such as sigmoidal. Therefore it is worthwhile to implement SNNs in a more agnostic manner as spiking RNNs. Examples of such promising implementations are reservoir computing, liquid, and each state machine. For more on such neuromorphic architectures, see 3.1.2.

### 1.4   Brain Atlases: whole- and population-level modeling

The $21^{st}$ century has been the bursting era of large-scale brain initiatives. The objective of the simulation partly justifies this multitude. As it was previously mentioned, the notion *simulation* is highly versatile in meaning depending on the goal of the project [85], i.e., where it sits on the Figure 3. Some of the projects of this spectrum are listed below.

- BigBrain: a free-access and few-cell-resolution model of human brain in 3D [86].
- Allen Brain Atlas: genome-wide map of gene expression for human adult and mouse brain [20].
- Human Connectome Project: a large-scale structural and functional connectivity map of human brain [coined as connectome in [87]] [88].
- Brain Research through Advancing Innovate Neurotechnologies: BRAIN [89].
- The Virtual Brain (TVB): an open-source neural dynamics simulator using real anatomical connectivity [60].
- Human Brain Project (HBP): aimed to realistically simulate the human brain in supercomputers [90].

The computing power is not the only problem when expanding the regional computations to the whole-brain level, scaling. One new difficulty is the integration of time delays that become significant at the whole-brain level. In local

---

[3]For a more comprehensive overview on types and applications, see Schliebs and Kasabov [81].

connections, the time delays are small enough to be ignored [60] the transmission happens in a variety of finite speeds from 1 to 10 meters per second. As a result of this variation, time delays among different brain parts are no longer negligible. Additional spatial features emerge by the implementation of this heterogeneity [91, 92].

Larger scale approaches could adapt neural mechanisms that rely on intra-region interactions [93] in order to ditch the problems related to the synaptic level studies mentioned earlier. The Virtual Brain (TVB) project is one of these initiatives. TVB captures the network dynamics of the brain by stimulating the neural population structural connectivity, the variant time scales, and noise [94]. TVB allows testing subject-specific hypotheses as the structural connectivity is based on individual DTI. The large-scale activity is an integration of local neural masses connected through large-range dynamics. It has a web platform GUI and can run on a personal computer and has already implemented many types of dynamics for different types of brain signals, namely EEG, MEG, BOLD, fMRI.

With models like TVB, one should note the shift in paradigm from the fine-scale simulations like Blue Brain. Contrary to the Blue Brain, the nodes consist of large groups of neurons (order of a few millimeters), not one or a few neurons. Consequently, the governing equations are the ones for deriving population dynamics and statistical modes. Another essential point is that TVB allows researchers to study the brain's phenomenology parametrically. The following section is dedicated to such studies.

**Neuromorphic compute: Architectures tailored for spiking**

The disparity in energy consumption and computing architecture of biological and silicon neurons are the most important factors that raise eyebrows in assessing *brain-like* algorithms. The brain consumes ∼20 watts of power while this amount for a supercomputer is in order of megawatts [95]. This twist verifies that the processing of information in these simulations is far from the biological truth. Apart from the energy consumption gap, the non-Von Neumann architecture of the brain is another discrepancy that stands in the way of realistic brain simulation in silico. There is no *Von Neumann bottleneck* in the brain as there is no limitation on throughput as a result of separation of memory and computing unit [96]. The brain also has other features that are greatly missed in deep networks. These include synaptic plasticity, high parallelism due to a large number of neurons, high connectivity due to a large number of synapses, resilience to degradation, and low speed and frequency of communication, among other things. Although many of the aspects of biological cognition are complicated to reconstruct (e.g., embodiment and social interaction), the research in neuromorphic computing is addressing the disparities above by targeting hardware design [97].

A potential solution for narrowing this computation gap can be sought at the hardware level. An instance of such a dedicated pipeline is neuromorphic processing units (NPU) that are power efficient and take time and dynamics into the equation from the beginning. An NPU is an array of neurosynaptic cores that contain computing models (neurons) and memory within the same unit. In short, the advantage of using NPUs is that they resemble the brain more realistically than a CPU or GPU because of asynchronous (event-based) communication, extreme parallelism (100-1000000 cores), and low power consumption [98]. Their efficiency and robustness also result from the Physical proximity of the computing unit and memory. Below popular examples of such NPUs are listed. Each of them stemmed from different initiatives.

- **SpiNNaker** or "Spiking Neural Network architecture" is an architecture based on low-power microprocessors and was first introduced in 2005 to help the European Brain Project with computations of large cortical area. The first version could imitate ten thousand spiking neurons and four million synapses with 43 nano Joules of energy per *synaptic event*. [99].

- **TrueNorth** chips are arrays of 4096 neurosynaptic cores amounting to 1 million digital neurons and 256 million synapses. IBM builds TrueNorth primarily as a low-power processor suitable for drones, but it is highly scalable and customizable [100].

- **Loihi** chips have demonstrated significant performance in optimization problems. Intel's fifth NPUs has incorporated biophysical reconstruction of hierarchical connectivity, dendritic compartments, synaptic delays, reward traces. Its circuit is composed of *dandrite units* (for updating state variables), *axon units* (generating feed for the subsequent cores), and *learning unit* (for updating weights based on customized learning rules) [101]

An integrative example of the implementation discussed above is NeuCube. NueCube is a 3D SNN with plasticity that learns the connections among populations from various STBD modulations such as EEG, fMRI, genetic, DTI, MEG, and NIRS. Gene regulatory networks can be incorporated as well if available. Finally, This implementation reproduces trajectories of neural activity. It has more robustness to noise and higher accuracy in classifying STBD than standard machine learning methods such as SVM [102].

Beyond biological alikeness, neuromorphic computing has important technical aspects that are missing in conventional compute units and can revolutionize neural data processing. They demonstrate lower latency, power consumption, and high portability required for real-time interpretation. These attributes make them useful for recent signal collectors like wearable EEG. On the other hand, although they have shown to be highly scalable and adaptable, their high cost per bit is a major pitfall [103, 104].

## 2   Phenomenological Models

In contrast with realistic biological models of *in-vivo* events, phenomenological[4] models offer a way of qualitative simulation of certain observable behavior (or, as it is discussed in Dynamical Systems literature [108], phase trajectories). The key assumption is that although short- and long-range dynamics are dependent on intricate biophysical events, the emerging observables can be encoded in significantly lower dimensions. This dimensionality reduction is thanks to dynamics that are capable of constructing similar statistical features of interest. Since a detailed-enough biophysical

---

[4]Note that here, the notion of "phenomena" here is different than that used by e.g. [105] where phenomenological architecture and properties are regarded as a representation of environment in the first-person mind [106], complementary to "physiological" architecture in the brain–as in [107].

| Connectivity | Measure | Imaging basis | Features |
|---|---|---|---|
| SC | Spatial config. of white matter fibers | Static spatial images (e.g. DTI) | Provides the anatomical architecture |
| FC | Temporal correlations of regional activities | Spatio-temporal images (e.g. fMRI, EEG) | Can be static or dynamic [123] Prone to spurious connections |
| EC | Causal interactions of segregated regions | Spatio-temporal images with a generative model (e.g. Granger Causality [124] or DCM [1.3.3]) | Rules out non-causal correlations |

Table 1: *Complex brain networks are measured through Structural Connectivity (SC), Functional Connectivity (FC), or Effective Connectivity (EC). Computational Connectomics is a common ways of formulating structural and functional networks of the whole brain. Together with theories of dynamical graphs, these representations can provide insights into the collective faith of the system.*

model should eventually exhibit the same collective statistics, one may argue that the phenomenological models offer a detour to system-level reconstruction by ditching lots of cellular and physiological considerations.

Compared to detailed biophysical models, coarse-grained approaches rely on a smaller set of biological constraints and might be considered "too simplistic". However, they are capable of reconstructing many collective phenomena that are still inaccessible to hyper-realistic simulations of neurons [109]. A famous example of emergence at this level is synchronizations in cortex [110]. Moreover, experiments show that the population-level dynamics that are ignorant about the fine-grained detail *better* explain the behavior [111, 112].

The significance of phenomenological models in the reconstruction of brain dynamics is also because of their intuitiveness and reproducibility. They may demonstrate critical properties of the neuronal population. An interesting example is noise-driven dynamics of the brain, which is responsible for multistability and criticality in resting-state [113, 114].

## 2.1 Problem formulation, data, and tools

The idea of using phenomenological models for neural dynamics is mainly motivated by the possibility of using tools from dynamical system theory. The goal is to quantify the evolution of a *state space* built upon the *state variables* of the system. For example, if one can find two population variables $(x, y)$ that determine the state of a neural ensemble, then all the possible pairs of $x$ and $y$s form the basis for the *state space* of the system, let us call this 2-dimensional space $A$. The state of this ensemble at any given time $t$ can always be expressed as a 2-D vector in $A$. In mathematics, $A$ is called a *vector space* defined by the sets of differential equations that describe the evolution of $x$ and $y$ in time. As an intuitive visualization of a vector space, imagine a water swirl: each point of the surface of a water swirl can be represented by a vector with the magnitude and direction of local velocity. One can see how at each point in this space, there is a *flow* that pushes the system in a specific trajectory. Reproducing features of the brain signals or identifying such a sparse state space and the dynamics of a parsimonious set of state variables allows for forecasting the fate of the neural ensemble in future timesteps [115]. The evolution of the state variables is described by differential equations. In what follows in this section, some of the most prominent phenomenological models and their findings are discussed.

Distinct dynamics is observed in a wide set of settings from resting-state activity [109] to task-specific experiments [116]. These insights are also useful in a wide set of modularities including fMRI [117], EEG [118], MEG [119], and Calcium imaging [120]. One common formulation is to build the dynamical graph models of the cortex based on the *anatomical*, *functional*, or *effective connectivity* as described in Table 1. For a more comprehensive review of such networks, refer to [121].

Connectivity matrices introduced in Table 1 are the backbone of the information process pipeline. That being said, this parameter needs to be married to the dynamics of the states in the brain. To date, a large portion of studies have focused on mapping these networks onto the resting-state network, and a lot of structure-function questions remained to be answered by studying the task-related data [122]. In what follows, the models that quantify these dynamics based on the phenomenology of the behavior are discussed.

### 2.1.1 Generative graph models

Recent progress in the science of complex networks and information theory has paved the way for analytical and numerical models of structural and functional connectivity [125]. The network approach to the neural population is a

conventional way to study neural processes as information transmission in time-varying networks. This analogy allows for examining the path and behavior of the system in terms of different dynamical properties.

An insightful interplay of function versus structure is observed along the biologically-plausible line of works of Deco et al. [113]. They reconstructed the emergence of equilibrium states around multistable attractors and characteristic critical behavior like scaling-law distribution of inter-brain pair correlations as a function of global coupling parameters. Furthermore, new studies show that synchrony not only depends on the topology of the graph but also on its hysteresis [126].

Tools from graph theory and network science [127] are used to formulate this relation. Spectral mapping [128] and structure-function topological mapping [129] are proofs of concept in this regard. Generative graph models (traditionally developed by graph theory such as the one for random graph introduced in [130]) are principle tools of inference in this approach and now has been enhanced by machine learning, see for example, deep-network generative models in [131, 132]). Simulations of brain network dynamics and study of controllability [133] has shown how differently regions are optimized for diverse behavior [134].

## 2.2  Inspirations from statistical physics and nonlinear dynamical models

In addition to network science, another axis for interpreting neural data is based on well-established tools initially developed for parametrizing the time evolution of physical systems. Famous examples of these systems include spin-glass [135], different types of coupled oscillators [120, 136], and multistable and chaotic many-body systems [109, 114]. This type of modeling has already offered promising and intuitive results. In the following subsections, we review some of the recent literature with various methodologies.

### 2.2.1  Brain as a complex system

It is not easy to define what a complex system is. [137] defines the *degree of complexity* of a sequence as the minimum length of the program and of the initial data that a Turing machine (aka universal computer) needs to produce that sequence. Despite being a debatable definition, one can conclude that according to it, the spatiotemporal dynamics of the mammalian brain qualifies as a complex system [138, 139]. Therefore, one needs a complex mechanism to reconstruct neural dynamics. In the following few subsections, we review candidate equations for the oscillations in cortical network [140].

**Equilibrium solutions and deterministic chaos**    Whole-brain phenomenological models like the Virtual Brain [94] are conventional generators for reconstructing spontaneous brain activity. There are various considerations to have in mind to choose the right model for the right task. A major trade-off is between the complexity and abstractiveness of the parameters [141]. In other words, to capture the behavior of detailed cytoarchitectural and physiological make-up with a reasonably-parametrized model. Another consideration is the incorporation of noise which is a requirement for multistable behavior [109] i.e., transitions between stable patterns of reverberating activity (aka attractors) in neural population in response to perturbation [142].

**Kuramoto**    Kuramoto model is a mathematical descriptor of coupled oscillators, one that can be written down as simple as a system of ODEs solely based on sinusoidal interactions [143, 144]. Kuramoto model is widely used in physics for studying synchronization phenomena. It is relevant to neurobiological systems as it enables a *phase reduction approach*: Neural populations can be regarded as similar oscillators that are weakly coupled together. These couplings are parameterized in the model. Kuramoto can be extended to incorporate anatomical and effective connectivity and can expand from a low-level model of few-neuron activity to a stochastic population-level model with partial synchrony and rich dynamical properties. One way to do that is to upgrade the classic linear statistics to nonlinear Fokker-Planck equations [145].

There is significant literature on Kuramoto models on neural dynamics on different scales and levels. [146] is a conceptual review of decades of research on the principles of the general form of the Kuramoto model. Numerous studies have found consistency between the results from Kuramoto and other classic models in computational neuroscience like Wilson-Cowan [63] [147]. Kuramoto model is frequently used for quantifying phase synchrony and for controlling unwanted phase transitions in neurological diseases like epileptic seizures and Parkinson's [148]. Still, there are many multistability questions regarding cognitive maladaptation yet to be explored, potentially with the help of Kuramoto models and the maps of effective connectivity. [149] is a review targeting clinical researchers and psychiatrists. It is a good read for learning about the current challenges that could be formulated as a Kuramoto model. Kuramoto is also unique in adaptability to different scales: from membrane resolution with each neuron acting as a delayed oscillator [150] to the social setting where each subject couples with the other one in the dyad by means of interpersonal interactions [151].

**Van der Pol**   Another model relevant to neuroscience is the van der Pol (VDP) oscillator which is probably the simplest relaxation oscillator [152] and a special case of the FitzHugh-Nagumo model, which is, in turn, a simplification of the Hodgkin–Huxley model (see Subsection 1.2) [153]. Through the Wilson-Cowan approximation [154], VDP can also model neural populations. For more information about the Wilson-Cowan model, please see Subsection 1.3.2. Recently, [120] have used coupled VDP oscillators to model a low-dimensional representation of neural activity in different living organisms (larval zebrafish, rats, and humans) measured by different brain imaging modalities, such as calcium imaging (CaI) and fMRI. Besides proposing a method for inferring functional connectivity by using the coupling matrices of the fitted models, it was demonstrated that dynamical systems models could be a valuable resource of data augmentation for spatiotemporal deep learning problems.

Looking at the brain as a complex system of interacting oscillators is a detour for expanding the modeling to larger organization scales. The emergent behavior of the system can be described with "order parameters". Although this is a description with much lower dimensions than the biophysical equations, it still expresses many remarkable phenomena such as phase transitions, instabilities, multiple stable points, metastability, and saddle points [137]. However, parametrizing such models is still an ongoing challenge, and many related studies are limited to the resting-state network. The following section reviews the prospects of recent data-driven methods and how they can leverage the study of system-level behavior.

## 3   Agnostic Computational Models

Jim Garys's framework [155] divides the history of science into four paradigms. Since centuries ago, there have been experimental and theoretical paradigms. Then the phenomena of interest became too complicated to be quantified analytically, so the computational paradigms started with the rise of numerical estimations and simulations. Today, with the bursting advances in recording, storage, and computation capacity of neural signals, neuroscience is now exploring the *fourth paradigm* of Jim Garys's framework [155] i.e., data exploration in which the scientific models are fit to the data by learning algorithms.

In the introduction, we reflected on how scientists should not settle for mere prediction. While the literature on data-driven methods is enormous, this review focuses mainly on the strategies that help gain mechanistic insights rather than those that reproduce data through operations that are difficult to relate to biological knowledge. Instances of these unfavored methods include strict generative adversarial networks with uninterpretable latent spaces or black-box RNN with hard-to-explain parameters. The following section categorizes these methods into established and emerging techniques and discusses some showcases.

### 3.1   Established learning models

Data-driven models have long been used in identifying structure-function relations (similar to the ones mentioned in Table 1) [156, 157]. The shift of studies from single-neuron to neural network, or more precisely, has accelerated in the last decade. This trend is because relying on collective properties of a population of neurons to infer behavior seems more promising than reconstructing the physiological activity of single neurons in hopes of achieving emergence. Yuste [158] argues that the mere representations that relate the state of individual neurons to a higher level of activity have serious shortcomings [32]. However, these shortcomings can be addressed by incorporating temporal dynamics and collective measures into the model. We review the models that satisfy this consideration.

#### 3.1.1   Dimensionality Reduction Techniques

Clustering and unsupervised learning are useful for mapping inputs $(X)$ to features $(Y)$. Later, this set of $(X, Y)$ can be extrapolated to unseen data. There are various methods for identifying this mapping or, in other words, for approximating this function. Principal Component Analysis (PCA) is a primary one. PCA maps data onto a subspace with the maximal variance [159]. It is a common method of dimensionality reduction. However, the orthogonal set of features found by this method are not necessarily statistically-dependent. Therefore, they are not always helpful in finding sources and effective connectivity. Alternatively, Independent Component Analysis, commonly known as ICA, was introduced as a solution to the Blind Source Separation (BSS) problem. Each sample of the data is an ensemble of the state of different sources. However, the characteristics of these sources are the hidden variable [160]. ICA is effective in finding the related source as it maps the data onto the feature space by minimizing the statistical independence for each feature rather than by minimizing the variance. Conventional use of component analysis is with fMRI and EEG recordings. In each time window, each sensor receives a noisy mix of activity in segregated brain regions. One is usually interested in inferring effective connectivity based on such data. Having a large number of electrodes around the scalp enables ICA to identify the independent sources of activities and artifacts. ICA algorithms come in different flavors depending on the dataset and the property of interest. E.g. temporal- [161], spatial- [156], and

13

spatiotemporal-ICA [162, 163] are tailored for different types of sampling. Hybrid approaches, e.g., ICA amalgamated with structural equation modeling (SEM), have shown better performance in given setups with less prior knowledge than SEM alone [164]. The interested reader is encouraged to refer to [165] for a dedicated review of ICA methods.

### 3.1.2   Recurrent neural networks

Recurrent neural networks (RNN) are the Turing-complete [166] algorithms for learning dynamics and are widely used in computational neuroscience. In a nutshell, RNN processes data by updating a "state vector". The state vector holds the memory across steps in the sequence. This state vector contains long-term information of the sequence from the past steps [167].

Current studies validate diverse types of RNNs as promising candidates for generating neural dynamics. Sherstinsky [168] shows how the implicit "additive model", which evolves the state signal, incorporates some of the interesting bio-dynamical behavior such as saturation bounds and the effects of time delays. Several studies modeled the cerebellum as an RNN with granular [169–172] or randomly-connected layers [173]. Moreover, similarities of performance and adaptability to limited computational power (as in biological systems) are observed both in recurrent convolutional neural networks and in the human visual cortex [174].

RNNs vary greatly in architecture. The choice of architecture can be implied by the output of interest (for example text [175] versus natural scenes[176]) or the approaches to overcome the problem with vanishing and exploding gradient (e.g., long short-term memory (LSTM) [177], hierarchical [178], or gated RNNs [179]).

**Hopfield**   Hopfield network [67] is a type of RNN inspired by dynamics of Ising model [180, 181]. In the original Hopfield mechanism, the units are threshold McCulloch and Pitts [80] neurons, connected in a recurrent fashion. The state of the system is described by a vector $V$ which represents the states of all units. In other words, the network is in fact, an undirected graph of artificial neurons. The strength of connection between units $i$ and $j$ is described by $w_{ij}$ which is trained by a given learning rule i.e. commonly Storkey [182] or Hebbian rule (stating that "neurons that fire together, wire together") [183]. After the training, these weights are set, and an energy landscape is defined as a function of $V$. The system evolves to minimize the energy and moves toward the basin of the closest attractor. This landscape can exhibit the stability and function of the network [184].

The Hopfield model can accommodate some biological assumptions and work in tandem with cortical realizations. Similar to the human brain, Hopfield connections are mostly symmetric. Most importantly, since its appearance, it has been widely used for replicating associative memory. However, soon it was revealed that other dynamical phenomena like cortical oscillations and stochastic activity [185] need to be incorporated in order to capture a comprehensive image of the cognition.

**LSTM**   In addition to the problem of vanishing and exploding gradient, other pitfalls also demand careful architecture adjustment. Early in the history of deep learning, RNNs demonstrated poor performance on sequences with long-term dependencies [186]. Long short term memory (LSTM) is specifically designed to resolve this problem. The principle difference of LSTM and vanilla RNN is that instead of a single recurrent layer, it has a "cell" composed of four layers that interact with each other through three gates: input gate, output gate and forget gate. These gates control the flow of old and new information in the "cell state" [177]. On certain scales of computation, LSTM still has considerable performance compared to trendy sequential models like transformers.

**Reservoir Computing**   A reservoir computer (RC) [187] is an RNN with a reservoir of interconnected spiking neurons. Broadly speaking, the distinction of RC among RNNs, in general, is the absence of granular layers between input and output. RCs themselves are dynamical systems that help learn the dynamics of data. Traditionally, the units of a reservoir have nonlinear activation functions that allow them to be universal approximators. Gauthier et al. [188] show that this nonlinearity can be consolidated in an equivalent *nonlinear vector autoregressor*. With the nonlinear activation function out of the picture, the required computation, data, and metaparameter optimization complexity are significantly reduced, interpretability is consequently improved while the performance stays the same.

**Liquid state machine**   LSM can be thought of as an *RNN soup* that maps the input data to a higher dimension that more explicitly represents the features. The word *liquid* come from the analogy of a stone (here an input) dropping into the water (here a spiking network) and propagating waves. LSM maintains intrinsic memory and can be simplified so much that it processes real-time data [189]. Zoubi et al. [190] shows LSM performs notably in building latent space of EEG data (extendable to fMRI). As for the faithfulness to the biological truth, Several studies argue that LSM surpass RNNs with granular layers in matching organization and circuitry of cerebellum [173] and cerebral cortex [187]. Lechner et al. [191] demonstrate the superiority of a biologically-designed LTM on given accuracy benchmarks to other ANNs, including LSTM.

**Physically-informed RNN**   A prominent factor in determining the dynamical profile of the brain is the intrinsic time delays [192]. Integrating these time delays into the artificial networks was initially an inspiration from neuroscience for AI. Later, they came back as a successful tool for integrated sequence modeling for multiple populations. In the last decade, RNN has been used for reconstructing neural dynamics via interpretable latent space in different recording modalities such as fMRI [157] and Calcium imaging [120].

Continuity of time is another extension that can make RNNs more compatible with various forms of sampling and thus neural dynamics from spikes to oscillations. Continous time RNNs (CT-RNNs) are RNNs with activation functions made up of differential equations. They have been proved to be universal function approximators [193] and have surfaced recently in the literature as reservoir computers [188, 194] and liquid time-constant neural networks [195].

Essentially, finding the optimal architecture and hyperparameters for a given problem does not have a straightforward recipe. The loss function in a deep neural network can be arbitrarily complex and usually takes more than a convex optimization. Li et al. [196] shows how parameters of the network can change the loss landscape and trainability. Another more specific issue to the algorithms trained on a temporal sequence is catastrophic forgetting and attention bottleneck. These complications arise from the limitation of memory and attention to the past time steps. New attention models such as transformers and recurrent independent mechanisms (see Subsections 3.1.4 and 3.1.5 respectively) are specifically built to address these issues. As memory-enhanced components, RNN layers appear in other deep and shallow architectures with sequential data as input, including encoder-decoders.

### 3.1.3   Variational autoencoder

Variational autoencoders (VAE for short) is a type of neural network that encodes the ground truth as the input onto a "latent space" and then decodes that space for reconstructing the input[197]. The network is parameterized by minimizing the reconstruction loss, which is, in this case, a metric of information gained by a metric called Kullback–Leibler divergence ([198]. An example of VAE used for regenerating dynamics is by Perl et al. [199] in which the coupling dynamics of the whole brain and the transitions between the states of wake-sleep progression is generated. The goal is to find low (e.g., as low as 2-) dimensional manifolds that can capture the signature structure-function relationship that demonstrates the stage in the wake-sleep cycle [200, 201] and the parameters of generic coupled Stuart-Landau oscillators as in [114]. An idea for regenerating dynamics is to use a deep-network embedded differential equations (as in Subsection 3.2.2) in the latent VAE structure [9].

### 3.1.4   Transformers

The transformer is a relatively new class of ML models that recently has shown state-of-the-art performance on sequence modeling such as natural language processing (NLP) field of research [202]. Beyond NLP, this architecture demonstrates good performance on a wide variety of data, including brain imaging [203–205]. Similar to RNNs, transformers aim to process sequential data such as natural language or temporal signals. It differs from the RNN paradigm because it does not process the data sequentially; instead, it looks at whole sequences with a mechanism called "attention," and by doing so, it alleviates the problem of forgetting long dependencies, which is common in RNN and LSTMs. This mechanism can make both long- and short-term connections between points in the sequence and prioritize them. Transformers are widely used for generating *foundation models* (i.e. models that are pretrained on big data [206]) and they can outperform recurrent networks like LSTM with large models/data [207].

### 3.1.5   Recurrent independent mechanisms

Recurrent independent mechanisms (RIM) is a form of attention model that learns and combineshttps://www.overleaf.com/project/600ee20bdcbe0b4407642e38 independent mechanisms in order to boost generalisability and robustness in executing a *task*. The task in the sense of signal processing can be generating a sequence based on the observed data. The hypothesis is that the dynamics can be learned as a sparse modular structure. In this recurrent architecture, each module independently specializes in a particular mechanism. Then all the RIMs compete through an attention bottleneck so that only the most relevant mechanisms get activated to communicate sparsely with others to perform the task [208].

### 3.2   New frontiers: Scientific ML and Interpretable Models

The independence from prior knowledge sounds interesting as it frees the methodology from inductive biases and makes the models more generalizable by definition. However, this virtue comes at the cost of a need for large training sets. In other words, the trade-off of bias and computation should be considered: Applying lots of prior knowledge and inductive biases result in a lesser need for data and computation. In contrast, little to no inductive bias calls for a great need for big and curated data. It is true that with the advancement of recording techniques, the scarcity of data is less of

a problem than it was before but even with all these advances, having *clean* and *sufficiently large* medical dataset that helps with the problem in hand is not guaranteed.

Total reliance on data is especially questionable when the data has significant complications (as discussed in the introduction). Opting for a methodology guided by *patterns* rather than *prior knowledge* is problematic in particular when the principle patterns of data arise from uninteresting phenomena such as the particular way a given facility may print out the brain images [209].

The thirst for data aside, one of the prominent drawbacks of agnostic modeling and ML, in particular, is that they are famous for providing *opaque blackbox* solutions, meaning that by leaving biological priors out of the picture, the explainability of the outcome is weak. Lack of explainability is a pet peeve for people in science as they are interested in both prediction and the reasoning behind those predictions.

In addition to the implicit assumption of the adequacy of training data, the explicit assumption that these models rely on is that the solution is parsimonious, i.e., there are few descriptive parameters. Despite some possibility of error with this assumption in given problems [210], it is particularly useful in having arbitrarily less complicated descriptions that are generalizable, interpretable, and less prone to over-fitting.

The following sections describe general function approximators that could identify data dynamics without injecting any prior knowledge about the system. They could provide a perfect solution for a well-observed system with unknown dynamics.

### 3.2.1    Sparse identification of nonlinear dynamics

Kaheman et al. [211] proposed a novel approach for quantifying underlying brain dynamics. The key assumption is that the governing multi-dimensional principles can be derived by a system of equations describing the first-order rate of change. In order to use sparse regressions methods such as Sparse Identification of Nonlinear Dynamics (SINDy), one needs to precisely specify the set of parsimonious state variables [212]. That being said, SINDy does not work for small datasets. If it is given less data than possible terms, the system of the governing equations is under-specified. Therefore, the underfitting as a result of insufficient training data is the secondary problem. One approach to address this issue is incorporating the known terms and dismissing the learning for those parts. An example is discussed in Subsection 3.2.3.

### 3.2.2    Differential equations with deep neural networks

A relatively new class of dynamical frameworks combines differential equations with machine learning in a more explicit fashion. This class of frameworks has been used to model the dynamics of time-varying signals. They begin by assuming that the underlying dynamics follow a differential equation. They can then be used to discover the parameters of that differential equation by using standard optimization of deep neural networks. As is evident, such formulations are quite useful in modeling and analyzing brain dynamics, especially using deep networks. Below we describe some of the relevant works in this sub-field.

**Neural Ordinary Differential Equations**    Combining ordinary differential equations (ODEs) with neural networks has recently emerged as a feasible method of incorporating differentiable physics into machine learning. A Neural Ordinary Differential Equation (Neural ODE) [9] provides the use of a parametric model as the differential function in an ODE. This architecture learns the dynamics of a process without explicitly stating the differential function, as has been done previously in different fields. Instead, standard deep learning optimization techniques could be used to train a parameterized differential function that can accurately describe the dynamics of a system. In the recent past, this has been used to infer the dynamics of various time-varying signals with practical applications [9, 213–219].

**Latent ODE**    A dynamic model such as the Neural ODE can be incorporated in an encoder-decoder framework, resembling a Variational Auto-Encoder, as mentioned in Chen et al. [9]. Such models assume that latent variables can capture the dynamics of the observed data. Previous works [9, 213–215] have successfully used this framework to define and train a generative model on time series data.

**Stochastic neural ODEs**    Parametric models can also be incorporated into stochastic differential equations to make Neural Stochastic Differential Equations (Neural SDEs) [216, 217]. Prior works have also introduced discontinuous jumps [218] in the differential equations.

**Neural controlled differential equations**    Further, latent ODE models can add another layer of abstraction. The observed data is assumed to be regularly/irregularly sampled from a continuous stream of data, following dynamics described by a continuously changing hidden state. Both the dynamics of the hidden state and the relationship between

the interpolated observations and the hidden state can be described by neural networks. Such systems are called Neural Control Differential Equations (Neural CDE) [219]. Broadly speaking, they are the continuous equivalent of RNNs.

### 3.2.3 Differential equations enhanced by deep neural networks

The above methods use deep neural networks to define the differential function in ordinary differential equations. In contrast, UDEs and GOKU-nets (described below) take the help of deep neural networks to enhance differential equations. UDEs replace only the unknown parts of a known partial differential equation, while GOKU-nets use explicit differential equations as part of deep neural network pipelines.

**Universal differential equations:** Universal Differential Equations (UDE) offer an alternate way of incorporating neural networks into differential equations while accounting for prior knowledge. In their seminal work, Rackauckas et al. [220] demonstrate how it is possible to aid a partial differential equation model by learning the unknown terms with universal approximators such as neural networks. Furthermore, they show how by combining this approach with a symbolic regression, such as SINDy, these models can accelerate the discovery of dynamics in limited data with significant accuracy.

**Generative ODE Modeling with Known Unknowns:** Another promising approach is the case of the Generative ODE Modeling with Known Unknowns aka GOKU-nets [221]. GOKU-net consists of a variational autoencoder structure with ODEs inside. In contrast with Latent ODEs, here, the ODEs are not parameterized but given explicit forms. Hence, it is possible to make use of some prior knowledge of the dynamics governing the system, such as in SINDy and UDEs, but there is no need to have direct observations of the state variables as in those cases. For example, one could hypothesize that the latent dynamics of a system follow some particular differential model such as Kuramoto or van der Pol. This model then jointly learns the transformation from the data space to a lower-dimensional feature representation and the parameters of the explicit differential equation.

The machine learning techniques are now routinely used for classification and regression of brain states (see Wein et al. [121] for a review). However, they have much more potential than black-box, data-intensive classifiers. This is because new sequential models are sometimes designed to identify the missing pieces of the puzzle of dynamics. They can also act as generative models and provide a broad potential for testing the biophysical and system-level hypotheses. Some of the methods introduced in this section are explained in detail in Kutz [222] textbook. Moreover, extremely helpful tutorials can be found in Brunton [223] YouTube™ channel.

## 4 Conclusion

The key purpose of this review was to dive into samples of already-popular paradigms or the ones authors found most promising for reconstructing neural dynamics with all the special considerations. To achieve this, we sorted the computational models with respect to two indicators: the scale of organization and the level of abstraction (Figure 3).

The scope of our study is broadly generative models of neural dynamics in biophysics, complex systems, and AI with some limitations. This paper is an interdisciplinary study that covers a time span from the mid-twentieth century when the pioneer models like Wilson-Cowan [63], and Hodgkin-Huxley [26] arose, up until the recent decade when gigantic brain atlas initiatives, groundbreaking research in ML, and unprecedented computation power became available. Given the rate of publication in the related fields, a systematic review was impossible. Therefore, this paper is a starting point for gaining an eagle-eye view of the current landscape. It is up to the reader to adjust the model scale and abstraction depending on the problem at hand (see Fig. 3).

We emphasized the distinctiveness of the problems in computational neuroscience and cognitive science. One key factor is the trade-off of complexity and inductive bias with the availability of data and prior knowledge of the system. While there is still no ultimate recipe yet, hybrid methods could simultaneously tackle explainability, interpretability, plausibility, and generalizability.

## 5 Acknowledgments

## 6   Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## References

[1] Andy J Wills and Emmanuel M Pothos. On the adequacy of current empirical evaluations of formal models of categorization. *Psychological bulletin*, 138(1):102, 2012.

[2] Olivia Guest and Andrea E Martin. How computational modeling can force theory building in psychological science. 2020.

[3] Danilo Bzdok and John PA Ioannidis. Exploration, inference, and prediction in neuroscience and biomedicine. *Trends in neurosciences*, 42(4):251–262, 2019.

[4] Joshua M Epstein. Why model? *Journal of artificial societies and social simulation*, 11(4):12, 2008.

[5] Mikhail I Rabinovich, Pablo Varona, Allen I Selverston, and Henry DI Abarbanel. Dynamical principles in neuroscience. *Reviews of modern physics*, 78(4):1213, 2006.

[6] John G White, Eileen Southgate, J Nichol Thomson, and Sydney Brenner. The structure of the nervous system of the nematode caenorhabditis elegans. *Philos Trans R Soc Lond B Biol Sci*, 314(1165):1–340, 1986.

[7] Henry Markram. The blue brain project. *Nature Reviews Neuroscience*, 7(2):153–160, 2006.

[8] Yasaman Bahri, Jonathan Kadmon, Jeffrey Pennington, Sam S. Schoenholz, Jascha Sohl-Dickstein, and Surya Ganguli. Statistical Mechanics of Deep Learning. *Annual Review of Condensed Matter Physics*, 11:501–528, 2020. ISSN 19475462. doi: 10.1146/annurev-conmatphys-031119-050745.

[9] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 2018.

[10] Steven Johnson. *Emergence: The connected lives of ants, brains, cities, and software*. Simon and Schuster, 2002.

[11] John W. Krakauer, Asif A. Ghazanfar, Alex Gomez-Marin, Malcolm A. MacIver, and David Poeppel. Neuroscience needs behavior: Correcting a reductionist bias. *Neuron*, 93(3):480–490, Feb 2017. ISSN 1097-4199. doi: 10.1016/j.neuron.2016.12.041.

[12] Megha Srivastava, Tatsunori Hashimoto, and Percy Liang. Robustness to spurious correlations via human annotations. In *International Conference on Machine Learning*, pages 9109–9119. PMLR, 2020.

[13] Philip W Anderson. More is different. *Science*, 177(4047):393–396, 1972.

[14] Martijn P. van den Heuvel, Lianne H. Scholtens, and René S. Kahn. Multiscale neuroscience of psychiatric disorders. *Biological Psychiatry*, 86(7):512–522, Oct 2019. ISSN 0006-3223. doi: 10.1016/j.biopsych.2019.05. 015.

[15] John C Mazziotta, Arthur Toga, Alan Evans, Peter Fox, Jack Lancaster, and Roger Woods. A probabilistic approach for mapping the human brain: the international consortium for brain mapping (icbm). In *Brain mapping: the systems*, pages 141–156. Elsevier, 2000.

[16] Ed S Lein, Michael J Hawrylycz, Nancy Ao, Mikael Ayres, Amy Bensinger, Amy Bernard, Andrew F Boe, Mark S Boguski, Kevin S Brockway, Emi J Byrnes, et al. Genome-wide atlas of gene expression in the adult mouse brain. *Nature*, 445(7124):168–176, 2007.

[17] Michael J Hawrylycz, Ed S Lein, Angela L Guillozet-Bongaarts, Elaine H Shen, Lydia Ng, Jeremy A Miller, Louie N Van De Lagemaat, Kimberly A Smith, Amanda Ebbert, Zackery L Riley, et al. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature*, 489(7416):391–399, 2012.

[18] Jonas Richiardi, Andre Altmann, Anna-Clare Milazzo, Catie Chang, M Mallar Chakravarty, Tobias Banaschewski, Gareth J Barker, Arun LW Bokde, Uli Bromberg, Christian Büchel, et al. Correlated gene expression supports synchronous activity in brain networks. *Science*, 348(6240):1241–1244, 2015.

[19] Ibai Diez and Jorge Sepulcre. Neurogenetic profiles delineate large-scale connectivity dynamics of the human brain. *Nature communications*, 9(1):1–10, 2018.

[20] Allan R Jones, Caroline C Overly, and Susan M Sunkin. The allen brain atlas: 5 years and beyond. *Nature Reviews Neuroscience*, 10(11):821–828, 2009.

[21] Michael Hawrylycz, Lydia Ng, David Feng, Susan Sunkin, Aaron Szafer, and Chinh Dang. The allen brain atlas. *Springer Handbook of Bio-/Neuroinformatics*, pages 1111–1126, 2014.

[22] Terri L Gilbert. The allen brain atlas as a resource for teaching undergraduate neuroscience. *Journal of Undergraduate Neuroscience Education*, 16(3):A261, 2018.

[23] Announcing the allen institute for neural dynamics, a new neuroscience division of the allen institute, Nov 2021. URL https://alleninstitute.org/what-we-do/brain-science/news-press/press-releases/announcing-allen-institute-neural-dynamics-new-neuroscience-division-allen-institute.

[24] Daniel P Dougherty, Geraldine A Wright, and Alice C Yew. Computational model of the camp-mediated sensory response and calcium-dependent adaptation in vertebrate olfactory receptor neurons. *Proceedings of the National Academy of Sciences*, 102(30):10415–10420, 2005.

[25] Miles A Whittington, Roger D Traub, and John GR Jefferys. Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, 373(6515):612–615, 1995.

[26] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500–544, 1952.

[27] Roger D Traub, Robert K Wong, Richard Miles, and Hillary Michelson. A model of a ca3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *Journal of neurophysiology*, 66(2):635–650, 1991.

[28] JOSEPH B Patlak and MAURICIO Ortiz. Slow currents through single sodium channels of the adult rat heart. *The Journal of general physiology*, 86(1):89–104, 1985.

[29] Irene S Gabashvili, Bernd HA Sokolowski, Cynthia C Morton, and Anne BS Giersch. Ion channel gene expression in the inner ear. *Journal of the Association for Research in Otolaryngology*, 8(3):305–328, 2007.

[30] LF Abbott and Thomas B Kepler. Model neurons: from hodgkin-huxley to hopfield. In *Statistical mechanics of neural networks*, pages 5–18. Springer, 1990.

[31] RB Stein and Alan Lloyd Hodgkin. The frequency of nerve action potentials generated by applied currents. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 167(1006):64–86, 1967.

[32] Jonathan A Michaels, Benjamin Dann, and Hansjörg Scherberger. Neural population dynamics during reaching are better explained by a dynamical system than representational tuning. *PLoS computational biology*, 12(11): e1005175, 2016.

[33] Andreas VM Herz, Tim Gollisch, Christian K Machens, and Dieter Jaeger. Modeling single-neuron dynamics and computations: a balance of detail and abstraction. *science*, 314(5796):80–85, 2006.

[34] Eugene M. Izhikevich and Gerald M. Edelman. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, 105(9):3593–3598, 2008. ISSN 0027-8424. doi: 10.1073/pnas.0712231105.

[35] Christopher J Honey, Olaf Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.

[36] Eugene M Izhikevich. Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6): 1569–1572, 2003.

[37] Christof Koch, Clay Reid, Hongkui Zeng, Stefan Mihalas, Mike Hawrylycz, John Philips, Chinh Dang, and Allan Jones. Project mindscope. In *The future of the brain*, pages 25–39. Princeton University Press, 2014.

[38] Alison Abbott. Documentary follows implosion of billion-euro brain project. *Nature*, 588(7837):215–216, 2020.

[39] Andreas Horn, Dirk Ostwald, Marco Reisert, and Felix Blankenburg. The structural–functional connectome and the default mode network of the human brain. *NeuroImage*, 102:142–151, 2014. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2013.09.069. URL https://www.sciencedirect.com/science/article/pii/S1053811913010057. Multimodal Data Fusion.

[40] Yves Frégnac and Gilles Laurent. Neuroscience: Where is the brain in the human brain project? *Nature News*, 513(7516):27, 2014.

[41] Christof Koch. Neuroscience: The connected self. *Nature*, 482(7383):31–31, 2012.

[42] R Douglas Fields, Alfonso Araque, Heidi Johansen-Berg, Soo-Siang Lim, Gary Lynch, Klaus-Armin Nave, Maiken Nedergaard, Ray Perez, Terrence Sejnowski, and Hiroaki Wake. Glial biology in learning and cognition. *The neuroscientist*, 20(5):426–431, 2014.

[43] Fritz A. Henn and Anders Hamberger. Glial cell function: Uptake of transmitter substances. *Proceedings of the National Academy of Sciences*, 68(11):2686–2690, Nov 1971. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.68.11.2686.

[44] Michael Hawrylycz, Costas Anastassiou, Anton Arkhipov, Jim Berg, Michael Buice, Nicholas Cain, Nathan W. Gouwens, Sergey Gratiy, Ramakrishnan Iyer, Jung Hoon Lee, and et al. Inferring cortical function in the mouse visual system through large-scale systems neuroscience. *Proceedings of the National Academy of Sciences*, 113 (27):7337–7344, Jul 2016. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1512901113.

[45] Katrin Amunts, Christoph Ebell, Jeff Muller, Martin Telefont, Alois Knoll, and Thomas Lippert. The human brain project: Creating a european research infrastructure to decode the human brain. *Neuron*, 92(3):574–581, Nov 2016. ISSN 0896-6273. doi: 10.1016/j.neuron.2016.10.046.

[46] Benedikt Fecher and Sascha Friesike. Open science: one term, five schools of thought. *Opening science*, pages 17–47, 2014.

[47] Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *Elife*, 8:e47314, 2019.

[48] Marc-Oliver Gewaltig and Markus Diesmann. Nest (neural simulation tool). *Scholarpedia*, 2(4):1430, 2007.

[49] James M Bower and David Beeman. *The book of GENESIS: exploring realistic neural models with the GEneral NEural SImulation System*. Springer Science & Business Media, 2012.

[50] Nicholas T Carnevale and Michael L Hines. *The NEURON book*. Cambridge University Press, 2006.

[51] Inga Blundell, Romain Brette, Thomas A Cleland, Thomas G Close, Daniel Coca, Andrew P Davison, Sandra Diaz-Pier, Carlos Fernandez Musoles, Padraig Gleeson, Dan FM Goodman, et al. Code generation in computational neuroscience: a review of tools and techniques. *Frontiers in neuroinformatics*, 12:68, 2018.

[52] Olivier David and Karl J Friston. A neural mass model for meg/eeg:: coupling and neuronal dynamics. *NeuroImage*, 20(3):1743–1755, 2003.

[53] Jilles Vreeken. Spiking neural networks, an introduction, 2003.

[54] K.J. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *NeuroImage*, 19(4):1273 – 1302, 2003. ISSN 1053-8119. doi: https://doi.org/10.1016/S1053-8119(03)00202-7. URL http://www.sciencedirect.com/science/article/pii/S1053811903002027.

[55] Gustavo Deco, Viktor K Jirsa, Peter A Robinson, Michael Breakspear, and Karl Friston. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput Biol*, 4(8):e1000092, 2008.

[56] Karl Friston. Mean-fields and neural masses. *PLoS Comput Biol*, 4(8):e1000081, 2008.

[57] S. Coombes. Waves, bumps, and patterns in neural field theories. *Biological Cybernetics*, 93(2):91–108, Aug 2005. ISSN 1432-0770. doi: 10.1007/s00422-005-0574-y.

[58] Stephen Coombes, Peter beim Graben, Roland Potthast, and James Wright. *Neural fields: theory and applications*. Springer, 2014.

[59] Stephen Coombes and Áine Byrne. Next generation neural mass models. *Nonlinear Dynamics in Computational Neuroscience*, page 1, 2019.

[60] Viktor Jirsa, Olaf Sporns, Michael Breakspear, Gustavo Deco, and Anthony Randal McIntosh. Towards the virtual brain: network modeling of the intact and the damaged brain. *Archives italiennes de biologie*, 148(3): 189–205, 2010.

[61] Junji Ito, Andrey R. Nikolaev, and Cees van Leeuwen. Dynamics of spontaneous transitions between global brain states. *Human Brain Mapping*, 28(9):904–913, Feb 2007. ISSN 1065-9471. doi: 10.1002/hbm.20316.

[62] Nicolás Deschle, Juan Ignacio Gossn, Prejaas Tewarie, Björn Schelter, and Andreas Daffertshofer. On the validity of neural mass models. *Frontiers in Computational Neuroscience*, 14:118, 2021. ISSN 1662-5188. doi: 10.3389/fncom.2020.581040. URL https://www.frontiersin.org/article/10.3389/fncom.2020.581040.

[63] Hugh R Wilson and Jack D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1–24, 1972.

[64] PA Anninos, B Beek, TJ Csermely, EM Harth, and G Pertile. Dynamics of neural structures. *Journal of Theoretical Biology*, 26(1):121–148, 1970.

[65] Raymond L Beurle. Properties of a mass of cells capable of regenerating pulses. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, pages 55–94, 1956.

[66] John C Eccles, P Fatt, and K Koketsu. Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurones. *The Journal of physiology*, 126(3):524–562, 1954.

[67] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, Apr 1982. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.79.8.2554.

[68] Geoffrey E Hinton and Terrence J Sejnowski. Optimal perceptual inference. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, volume 448. Citeseer, 1983.

[69] Shun-ichi Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological cybernetics*, 27(2):77–87, 1977.

[70] James A Roberts, Leonardo L Gollo, Romesh G Abeysuriya, Gloria Roberts, Philip B Mitchell, Mark W Woolrich, and Michael Breakspear. Metastable brain waves. *Nature communications*, 10(1):1–17, 2019.

[71] Wolfram Erlhagen and Gregor Schöner. Dynamic field theory of movement preparation. *Psychological review*, 109(3):545, 2002.

[72] Sadjad Sadeghi, Daniela Mier, Martin F. Gerchen, Stephanie N. L. Schmidt, and Joachim Hass. Dynamic causal modeling for fmri with wilson-cowan-based neuronal equations. *Frontiers in Neuroscience*, 14:1205, 2020. ISSN 1662-453X. doi: 10.3389/fnins.2020.593867. URL https://www.frontiersin.org/article/10.3389/fnins.2020.593867.

[73] Hugh R. Wilson. Hyperchaos in wilson–cowan oscillator circuits. *Journal of Neurophysiology*, 122(6):2449–2457, Dec 2019. ISSN 0022-3077. doi: 10.1152/jn.00323.2019.

[74] James MacLaurin, Jamil Salhi, and Salwa Toumi. Mean field dynamics of a wilson–cowan neuronal network with nonlinear coupling term. *Stochastics and Dynamics*, 18(06):1850046, Dec 2018. ISSN 0219-4937. doi: 10.1142/S0219493718500466.

[75] Sami El Boustani and Alain Destexhe. A master equation formalism for macroscopic modeling of asynchronous irregular activity states. *Neural computation*, 21(1):46–100, 2009.

[76] Bolun Chen and Paul Miller. Attractors in networks of bistable neuronal units with depressing synapses. page V06.001, 2018.

[77] Alain Destexhe and Terrence J Sejnowski. The wilson–cowan model, 36 years later. *Biological cybernetics*, 101 (1):1–2, 2009.

[78] J. Daunizeau, O. David, and K. E. Stephan. Dynamic causal modelling: A critical review of the biophysical and statistical foundations. *NeuroImage*, 58(2):312–322, Sep 2011. ISSN 1053-8119. doi: 10.1016/j.neuroimage. 2009.11.062.

[79] Rosalyn Moran, Dimitris A. Pinotsis, and Karl Friston. Neural masses and fields in dynamic causal modeling. *Frontiers in Computational Neuroscience*, 7:57, May 2013. ISSN 1662-5188. doi: 10.3389/fncom.2013.00057.

[80] Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, Dec 1943. ISSN 1522-9602. doi: 10.1007/BF02478259. Company: Springer Distributor: Springer Institution: Springer Label: Springer number: 4 publisher: Kluwer Academic Publishers.

[81] Stefan Schliebs and Nikola Kasabov. Evolving spiking neural network—a survey. *Evolving Systems*, 4(2):87–98, Jun 2013. ISSN 1868-6486. doi: 10.1007/s12530-013-9074-9.

[82] Nikola K Kasabov. *Evolving connectionist systems: the knowledge engineering approach*. Springer Science & Business Media, 2007.

[83] Snjezana Soltic, Simei Gomes Wysoski, and Nikola K Kasabov. Evolving spiking neural networks for taste recognition. In *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 2091–2097. IEEE, 2008.

[84] Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671, 1997.

[85] Hugo de Garis, Chen Shuo, Ben Goertzel, and Lian Ruiting. A world survey of artificial brain projects, part i: Large-scale brain simulations. *Neurocomputing*, 74(1):3–29, Dec 2010. ISSN 0925-2312. doi: 10.1016/j. neucom.2010.08.004.

[86] Esther Landhuis. Neuroscience: Big brain, big data. *Nature*, 541(76387638):559–561, Jan 2017. ISSN 1476-4687. doi: 10.1038/541559a.

[87] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: A structural description of the human brain. *PLOS Computational Biology*, 1(4):e42, Sep 2005. ISSN 1553-7358. doi: 10.1371/journal.pcbi.0010042.

[88] David C. Van Essen, Stephen M. Smith, Deanna M. Barch, Timothy E. J. Behrens, Essa Yacoub, and Kamil Ugurbil. The wu-minn human connectome project: An overview. *NeuroImage*, 80:62–79, Oct 2013. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2013.05.041.

[89] Anna Devor, Peter A. Bandettini, David A. Boas, James M. Bower, Richard B. Buxton, Lawrence B. Cohen, Anders M. Dale, Gaute T. Einevoll, Peter T. Fox, Maria Angela Franceschini, and et al. The challenge of connecting the dots in the b.r.a.i.n. *Neuron*, 80(2):270–274, Oct 2013. ISSN 0896-6273. doi: 10.1016/j.neuron. 2013.09.008.

[90] Greg Miller. Blue brain founder responds to critics, clarifies his goals. *Science*, 334(6057):748–749, Nov 2011. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.334.6057.748.

[91] Victor K Jirsa and JA Scott Kelso. Spatiotemporal pattern formation in neural systems with heterogeneous connection topologies. *Physical Review E*, 62(6):8462, 2000.

[92] Spase Petkoski and Viktor K. Jirsa. Transmission time delays organize the brain network synchronization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 377 (2153):20180132, Sep 2019. doi: 10.1098/rsta.2018.0132.

[93] Fernando Lopes da Silva. Neural mechanisms underlying brain waves: from neural membranes to networks. *Electroencephalography and clinical neurophysiology*, 79(2):81–93, 1991.

[94] Paula Sanz Leon, Stuart A Knock, M Marmaduke Woodman, Lia Domide, Jochen Mersmann, Anthony R McIntosh, and Viktor Jirsa. The virtual brain: a simulator of primate brain network dynamics. *Frontiers in neuroinformatics*, 7:10, 2013.

[95] Ming ZHANG, Gu Zonghua, and Pan Gang. A survey of neuromorphic computing based on spiking neural networks. *Chinese Journal of Electronics*, 27(4):667–674, 2018.

[96] Wm A Wulf and Sally A McKee. Hitting the memory wall: Implications of the obvious. *ACM SIGARCH computer architecture news*, 23(1):20–24, 1995.

[97] Zongyuan Cai and Xinze Li. Neuromorphic brain-inspired computing with hybrid neural networks. In *2021 IEEE International Conference on Artificial Intelligence and Industrial Design (AIID)*, pages 343–347. IEEE, 2021.

[98] URL https://books.google.com/books/about/How_to_Build_a_Brain.html?id=BK0YRJPmuzgC.

[99] Thomas Sharp, Francesco Galluppi, Alexander Rast, and Steve Furber. Power-efficient simulation of detailed cortical microcircuits on spinnaker. *Journal of neuroscience methods*, 210(1):110–118, 2012.

[100] Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, et al. Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE transactions on computer-aided design of integrated circuits and systems*, 34(10):1537–1557, 2015.

[101] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee Micro*, 38(1):82–99, 2018.

[102] Nikola K Kasabov. Neucube: A spiking neural network architecture for mapping, learning and understanding of spatio-temporal brain data. *Neural Networks*, 52:62–76, 2014.

[103] Mike Davies. Lessons from loihi: Progress in neuromorphic computing. In *2021 Symposium on VLSI Circuits*, pages 1–2. IEEE, 2021.

[104] Mohammadali Sharifshazileh, Karla Burelo, Johannes Sarnthein, and Giacomo Indiveri. An electronic neuromorphic system for real-time detection of high frequency oscillations (hfo) in intracranial eeg. *Nature communications*, 12(1):1–14, 2021.

[105] Antti Revonsuo. *Inner presence: Consciousness as a biological phenomenon.* Mit Press, 2006.

[106] David Woodruff Smith. Phenomenology. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2018 edition, 2018.

[107] Andrew A. Fingelkurts, Alexander A. Fingelkurts, and Carlos F. H. Neves. Phenomenological architecture of a mind and operational architectonics of the brain: the unified metastable continuum. *New Mathematics and Natural Computation*, 05(01):221–244, Mar 2009. ISSN 1793-0057. doi: 10.1142/S1793005709001258.

[108] Steven H Strogatz. *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering.* CRC press, 2018.

[109] Juan Piccinini, Ignacio Perez Ipiña, Helmut Laufs, Morten Kringelbach, Gustavo Deco, Yonatan Sanz Perl, and Enzo Tagliazucchi. Noise-driven multistability vs deterministic chaos in phenomenological semi-empirical models of whole-brain activity. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(2):023127, 2021.

[110] Alex Arenas, Albert Diaz-Guilera, Jurgen Kurths, Yamir Moreno, and Changsong Zhou. Synchronization in complex networks. *Physics Reports*, 469(3):93–153, December 2008. ISSN 0370-1573. doi: 10.1016/j.physrep. 2008.09.002.

[111] Kevin L Briggman, Henry DI Abarbanel, and William B Kristan. Optical imaging of neuronal populations during decision-making. *Science*, 307(5711):896–901, 2005.

[112] Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012.

[113] Gustavo Deco and Viktor K Jirsa. Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *Journal of Neuroscience*, 32(10):3366–3375, 2012.

[114] Gustavo Deco, Morten L Kringelbach, Viktor K Jirsa, and Petra Ritter. The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Scientific reports*, 7(1):1–14, 2017.

[115] Maria Luisa Saggio and Viktor Jirsa. Phenomenological mesoscopic models for seizure activity. *arXiv preprint arXiv:2007.02783*, 2020.

[116] Ajay S. Pillai and Viktor K. Jirsa. *Symmetry Breaking in Space-Time Hierarchies Shapes Brain Dynamics and Behavior*, volume 94. Cell Press, Jun 2017. doi: 10.1016/j.neuron.2017.05.013.

[117] R. Matthew Hutchison, Thilo Womelsdorf, Elena A. Allen, Peter A. Bandettini, Vince D. Calhoun, Maurizio Corbetta, Stefania Della Penna, Jeff H. Duyn, Gary H. Glover, Javier Gonzalez-Castillo, and et al. Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, 80:360–378, 2013. ISSN 10538119. doi: 10.1016/j.neuroimage.2013.05.079.

[118] Selen Atasoy, Gustavo Deco, Morten L. Kringelbach, and Joel Pearson. Harmonic brain modes: A unifying framework for linking space and time in brain dynamics. *The Neuroscientist*, 24(3):277–293, Jun 2018. ISSN 1073-8584. doi: 10.1177/1073858417728032.

[119] Luke Tait, Ayşegül Özkan, Maciej J. Szul, and Jiaxiang Zhang. A systematic evaluation of source reconstruction of resting meg of the human brain with a new high-resolution atlas: Performance, precision, and parcellation. *Human Brain Mapping*, n/a(n/a), 2021. ISSN 1097-0193. doi: 10.1002/hbm.25578. URL https://onlinelibrary. wiley.com/doi/abs/10.1002/hbm.25578.

[120] Germán Abrevaya, Guillaume Dumas, Aleksandr Y. Aravkin, Peng Zheng, Jean-Christophe Gagnon-Audet, James Kozloski, Pablo Polosecki, Guillaume Lajoie, David Cox, Silvina Ponce Dawson, and et al. Learning brain dynamics with coupled low-dimensional nonlinear oscillators and deep recurrent networks. *Neural Computation*, page 1–40, May 2021. ISSN 0899-7667. doi: 10.1162/neco_a_01401.

[121] Simon Wein, Gustavo Deco, Ana Maria Tomé, Markus Goldhacker, Wilhelm M. Malloni, Mark W. Greenlee, and Elmar W. Lang. Brain connectivity studies on structure-function relationships: A short survey with an emphasis on machine learning. *Computational Intelligence and Neuroscience*, 2021:e5573740, May 2021. ISSN 1687-5265. doi: 10/gk87qg.

[122] Joana Cabral, Morten L. Kringelbach, and Gustavo Deco. Functional connectivity dynamically evolves on multiple time-scales over a static structural connectome: Models and mechanisms. *NeuroImage*, 160:84–96, Oct 2017. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2017.03.045.

[123] Sriniwas Govinda Surampudi, Joyneel Misra, Gustavo Deco, Raju Surampudi Bapi, Avinash Sharma, and Dipanjan Roy. Resting state dynamics meets anatomical structure: Temporal multiple kernel learning (tmkl) model. *NeuroImage*, 184:609–620, Jan 2019. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2018.09.054.

[124] Mingzhou Ding, Yonghong Chen, and Steven L Bressler. 17 granger causality: basic theory and application to neuroscience. *Handbook of time series analysis: recent theoretical developments and applications*, 437, 2006.

[125] Daniel J. Lurie, Daniel Kessler, Danielle S. Bassett, Richard F. Betzel, Michael Breakspear, Shella Kheilholz, Aaron Kucyi, Raphaël Liégeois, Martin A. Lindquist, Anthony Randal McIntosh, and et al. Questions and controversies in the study of time-varying functional connectivity in resting fmri. *Network Neuroscience*, 4(1): 30–69, Jan 2020. ISSN 2472-1751. doi: 10.1162/netn_a_00116.

[126] William Qian, Lia Papadopoulos, Zhixin Lu, Keith Wiley, Fabio Pasqualetti, and Danielle S. Bassett. Path-dependent dynamics induced by rewiring networks of inertial oscillators. *arXiv:2009.13466 [nlin]*, Sep 2020. URL http://arxiv.org/abs/2009.13466. arXiv: 2009.13466.

[127] Mark Ed Newman, Albert-László Ed Barabási, and Duncan J Watts. *The structure and dynamics of networks*. Princeton university press, 2006.

[128] Cassiano O Becker, Sérgio Pequito, George J Pappas, Michael B Miller, Scott T Grafton, Danielle S Bassett, and Victor M Preciado. Spectral mapping of brain functional connectivity from diffusion imaging. *Scientific reports*, 8(1):1–15, 2018.

[129] Hualou Liang and Hongbin Wang. Structure-function network mapping and its assessment via persistent homology. *PLoS computational biology*, 13(1):e1005325, 2017.

[130] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

[131] Tamara G Kolda, Ali Pinar, Todd Plantenga, and Comandur Seshadhri. A scalable generative graph model with community structure. *SIAM Journal on Scientific Computing*, 36(5):C424–C452, 2014.

[132] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.

[133] Thomas Kailath. *Linear systems*, volume 156. Prentice-Hall Englewood Cliffs, NJ, 1980.

[134] Evelyn Tang, Chad Giusti, Graham L. Baum, Shi Gu, Eli Pollock, Ari E. Kahn, David R. Roalf, Tyler M. Moore, Kosha Ruparel, Ruben C. Gur, and et al. Developmental increases in white matter network controllability support a growing diversity of brain dynamics. *Nature Communications*, 8(11):1252, Nov 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-01254-4.

[135] Gustavo Deco, Mario Senden, and Viktor Jirsa. How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Frontiers in computational neuroscience*, 6:68, 2012.

[136] Joana Cabral, Henry Luckhoo, Mark Woolrich, Morten Joensson, Hamid Mohseni, Adam Baker, Morten L Kringelbach, and Gustavo Deco. Exploring mechanisms of spontaneous functional connectivity in meg: how delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. *Neuroimage*, 90:423–435, 2014.

[137] Hermann Haken. *Information and Self-Organization: A Macroscopic Approach to Complex Systems*. Springer Series in Synergetics. Springer-Verlag, 3 edition, 2006. ISBN 978-3-540-33021-9. doi: 10.1007/3-540-33023-2. URL https://www.springer.com/gp/book/9783540330219.

[138] Francesco Sforazzini, Adam J Schwarz, Alberto Galbusera, Angelo Bifone, and Alessandro Gozzi. Distributed bold and cbv-weighted resting-state networks in the mouse brain. *Neuroimage*, 87:403–415, 2014.

[139] R Matthew Hutchison, L Stan Leung, Seyed M Mirsattari, Joseph S Gati, Ravi S Menon, and Stefan Everling. Resting-state networks in the macaque at 7 t. *Neuroimage*, 56(3):1546–1555, 2011.

[140] György Buzsáki and Andreas Draguhn. Neuronal oscillations in cortical networks. *science*, 304(5679):1926–1929, 2004.

[141] Michael Breakspear. Dynamic models of large-scale brain activity. *Nature neuroscience*, 20(3):340–352, 2017.

[142] J. A. Scott Kelso. Multistability and metastability: understanding dynamic coordination in the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1591):906–918, Apr 2012. ISSN 0962-8436. doi: 10.1098/rstb.2011.0351.

[143] Yoshiki Kuramoto. Chemical turbulence. pages 111–140, 1984.

[144] Naoko Nakagawa and Yoshiki Kuramoto. From collective oscillations to collective chaos in a globally coupled oscillator system. *Physica D: Nonlinear Phenomena*, 75(1-3):74–80, 1994.

[145] Michael Breakspear, Stewart Heitmann, and Andreas Daffertshofer. Generative models of cortical oscillations: Neurobiological implications of the kuramoto model. *Frontiers in Human Neuroscience*, 4:190, 2010. ISSN 1662-5161. doi: 10.3389/fnhum.2010.00190. URL https://www.frontiersin.org/article/10.3389/fnhum.2010.00190.

[146] Steven H. Strogatz. From kuramoto to crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1):1–20, 2000. ISSN 0167-2789. doi: https://doi.org/10.1016/S0167-2789(00)00094-4. URL https://www.sciencedirect.com/science/article/pii/S0167278900000944.

[147] Frank C. Hoppensteadt and Eugene M. Izhikevich. Thalamo-cortical interactions modeled by weakly connected oscillators: could the brain use fm radio principles? *Biosystems*, 48(1):85–94, 1998. ISSN 0303-2647. doi: https://doi.org/10.1016/S0303-2647(98)00053-7. URL https://www.sciencedirect.com/science/article/pii/S0303264798000537.

[148] B.R.R. Boaretto, R.C. Budzinski, T.L. Prado, J. Kurths, and S.R. Lopes. Protocol for suppression of phase synchronization in hodgkin–huxley-type networks. *Physica A: Statistical Mechanics and its Applications*, 528:121388, 2019. ISSN 0378-4371. doi: https://doi.org/10.1016/j.physa.2019.121388. URL https://www.sciencedirect.com/science/article/pii/S0378437119308088.

[149] Chinedu I Anyaeji, Joana Cabral, and David Silbersweig. On a quantitative approach to clinical neuroscience in psychiatry: Lessons from the kuramoto model. *Harvard Review of Psychiatry*, 29(4):318–326, 2021.

[150] D. Hansel, G. Mato, and C. Meunier. Synchrony in Excitatory Neural Networks. *Neural Computation*, 7(2): 307–337, 03 1995. ISSN 0899-7667. doi: 10.1162/neco.1995.7.2.307. URL https://doi.org/10.1162/neco.1995.7.2.307.

[151] Guillaume Dumas, Mario Chavez, Jacqueline Nadel, and Jacques Martinerie. Anatomical connectivity influences both intra-and inter-brain synchronizations. *PloS one*, 7(5):e36414, 2012.

[152] John Guckenheimer and Philip Holmes. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, volume 42. Springer Science & Business Media, 2013.

[153] Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961. ISSN 0006-3495. doi: https://doi.org/10.1016/S0006-3495(61)86902-6. URL https://www.sciencedirect.com/science/article/pii/S0006349561869026.

[154] T. Kawahara. Coupled van der pol oscillators—a model of excitatory and inhibitory neural interactions. *Biological cybernetics*, 39(1):37–43, 1980.

[155] Anthony JG Hey et al. *The fourth paradigm: data-intensive scientific discovery*, volume 1.

[156] Martin J McKeown, Scott Makeig, Greg G Brown, Tzyy-Ping Jung, Sandra S Kindermann, Anthony J Bell, and Terrence J Sejnowski. Analysis of fmri data by blind separation into independent spatial components. *Human brain mapping*, 6(3):160–188, 1998.

[157] Georgia Koppe, Hazem Toutounji, Peter Kirsch, Stefanie Lis, and Daniel Durstewitz. Identifying nonlinear dynamical systems via generative recurrent neural networks with applications to fMRI. 15(8):e1007263. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1007263. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1007263. Publisher: Public Library of Science.

[158] Rafael Yuste. From the neuron doctrine to neural networks. *Nature reviews neuroscience*, 16(8):487–497, 2015.

[159] Panos P. Markopoulos, George N. Karystinos, and Dimitris A. Pados. Optimal algorithms for 1-subspace signal processing. *IEEE Transactions on Signal Processing*, 62(19):5046–5058, Oct 2014. ISSN 1941-0476. doi: 10.1109/TSP.2014.2338077.

[160] Barak A Pearlmutter and Lucas C Parra. Maximum likelihood blind source separation: A context-sensitive generalization of ica. In *Advances in neural information processing systems*, pages 613–619, 1997.

[161] Vince Daniel Calhoun, T Adali, GD Pearlson, and James J Pekar. Spatial and temporal independent component analysis of functional mri data containing a pair of task-related waveforms. *Human brain mapping*, 13(1):43–53, 2001.

[162] Markus Goldhacker, P Keck, A Igel, Elmar Wolfgang Lang, and Ana Maria Tomé. A multi-variate blind source separation algorithm. *Computer methods and programs in biomedicine*, 151:91–99, 2017.

[163] Zhi Wang, Maogeng Xia, Zhen Jin, Li Yao, and Zhiying Long. Temporally and spatially constrained ica of fmri data analysis. *PLoS One*, 9(4):e94211, 2014.

[164] Jagath C Rajapakse, Choong Leong Tan, Xuebin Zheng, Susanta Mukhopadhyay, and Kanyan Yang. Exploratory analysis of brain connectivity with ica. *IEEE engineering in medicine and biology magazine*, 25(2):102–111, 2006.

[165] Vince D Calhoun and Tülay Adali. Multisubject independent component analysis of fmri: a decade of intrinsic networks, default mode, and neurodiagnostic discovery. *IEEE reviews in biomedical engineering*, 5:60–73, 2012.

[166] Joe Kilian and Hava T. Siegelmann. The dynamic universality of sigmoidal neural networks. *Information and Computation*, 128(1):48–56, Jul 1996. ISSN 0890-5401. doi: 10.1006/inco.1996.0062.

[167] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[168] Alex Sherstinsky. Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306, 2020.

[169] Dean V Buonomano and Michael D Mauk. Neural network model of the cerebellum: temporal discrimination and the timing of motor responses. *Neural computation*, 6(1):38–55, 1994.

[170] Constanze HofstoÈtter, Matti Mintz, and Paul FMJ Verschure. The cerebellum in action: a simulation and robotics study. *European Journal of Neuroscience*, 16(7):1361–1376, 2002.

[171] Javier F Medina, Keith S Garcia, William L Nores, Nichole M Taylor, and Michael D Mauk. Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *Journal of Neuroscience*, 20(14): 5516–5525, 2000.

[172] Tadashi Yamazaki and Shigeru Tanaka. Neural modeling of an internal clock. *Neural computation*, 17(5): 1032–1058, 2005.

[173] Tadashi Yamazaki and Shigeru Tanaka. The cerebellum as a liquid state machine. *Neural Networks*, 20(3): 290–297, Apr 2007. ISSN 0893-6080. doi: 10.1016/j.neunet.2007.04.004.

[174] Courtney J Spoerer, Tim C Kietzmann, Johannes Mehrer, Ian Charest, and Nikolaus Kriegeskorte. Recurrent neural networks can explain flexible trading of speed and accuracy in biological vision. *PLoS computational biology*, 16(10):e1008215, 2020.

[175] Ilya Sutskever, James Martens, and Geoffrey E Hinton. Generating text with recurrent neural networks. In *ICML*, 2011.

[176] Richard Socher, Cliff Chiung-Yu Lin, Andrew Y Ng, and Christopher D Manning. Parsing natural scenes and natural language with recursive neural networks. In *ICML*, 2011.

[177] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[178] Salah Hihi and Yoshua Bengio. Hierarchical recurrent neural networks for long-term dependencies. 8, 1996. URL https://proceedings.neurips.cc/paper/1995/file/c667d53acd899a97a85de0c201ba99be-Paper.pdf.

[179] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

[180] William A Little. The existence of persistent states in the brain. *Mathematical biosciences*, 19(1-2):101–120, 1974.

[181] Stephen G Brush. History of the lenz-ising model. *Reviews of modern physics*, 39(4):883, 1967.

[182] Amos Storkey. Increasing the capacity of a hopfield network without sacrificing functionality. In *International Conference on Artificial Neural Networks*, pages 451–456. Springer, 1997.

[183] Donald Olding Hebb. *The organisation of behaviour: a neuropsychological theory*. Science Editions New York, 1949.

[184] Han Yan, Lei Zhao, Liang Hu, Xidi Wang, Erkang Wang, and Jin Wang. Nonequilibrium landscape theory of neural networks. *Proceedings of the National Academy of Sciences*, 110(45):E4185–E4194, 2013.

[185] Xiao-Jing Wang. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological reviews*, 90(3):1195–1268, 2010.

[186] Jürgen Schmidhuber. Learning complex, extended sequences using the principle of history compression. *Neural Computation*, 4(2):234–242, 1992.

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

[188] Daniel J Gauthier, Erik Bollt, Aaron Griffith, and Wendson AS Barbosa. Next generation reservoir computing. *arXiv preprint arXiv:2106.07688*, 2021.

[189] Anvesh Polepalli, Nicholas Soures, and Dhireesha Kudithipudi. Digital neuromorphic design of a liquid state machine for real-time processing. In *2016 IEEE International Conference on Rebooting Computing (ICRC)*, pages 1–8, 2016. doi: 10.1109/ICRC.2016.7738687.

[190] Obada Al Zoubi, Mariette Awad, and Nikola K. Kasabov. Anytime multipurpose emotion recognition from eeg data using a liquid state machine based framework. *Artificial Intelligence in Medicine*, 86:1–8, 2018. ISSN 0933-3657. doi: https://doi.org/10.1016/j.artmed.2018.01.001.

[191] Mathias Lechner, Ramin Hasani, Manuel Zimmer, Thomas A Henzinger, and Radu Grosu. Designing worm-inspired neural networks for interpretable robotic control. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 87–94. IEEE, 2019.

[192] Bo Chang, Lili Meng, Eldad Haber, Lars Ruthotto, David Begert, and Elliot Holtham. Reversible architectures for arbitrarily deep residual neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[193] Ken-ichi Funahashi and Yuichi Nakamura. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural networks*, 6(6):801–806, 1993.

[194] David Verstraeten, Benjamin Schrauwen, Michiel D'Haene, Dirk Stroobandt, et al. A unifying comparison of reservoir computing methods. *Neural Networks*, 20:391–403, 2007.

[195] Ramin Hasani, Mathias Lechner, Alexander Amini, Daniela Rus, and Radu Grosu. Liquid time-constant networks. *arXiv preprint arXiv:2006.04439*, 2020.

[196] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL `https://proceedings.neurips.cc/paper/2018/file/a41b3bb3e6b050b6c9067c67f663b915-Paper.pdf`.

[197] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[198] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.

[199] Yonatan Sanz Perl, Hernán Boccacio, Ignacio Pérez-Ipiña, Federico Zamberlán, Helmut Laufs, Morten Kringelbach, Gustavo Deco, and Enzo Tagliazucchi. Generative embeddings of brain collective dynamics using variational autoencoders. *arXiv preprint arXiv:2007.01378*, 2020.

[200] Pablo Barttfeld, Lynn Uhrig, Jacobo D Sitt, Mariano Sigman, Béchir Jarraya, and Stanislas Dehaene. Signature of consciousness in the dynamics of resting-state brain activity. *Proceedings of the National Academy of Sciences*, 112(3):887–892, 2015.

[201] Justin L Vincent, Gaurav H Patel, Michael D Fox, Abraham Z Snyder, Justin T Baker, David C Van Essen, John M Zempel, Lawrence H Snyder, Maurizio Corbetta, and Marcus E Raichle. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, 447(7140):83–86, 2007.

[202] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[203] Jiayao Sun, Jin Xie, and Huihui Zhou. Eeg classification with transformer-based models. In *2021 IEEE 3rd Global Conference on Life Sciences and Technologies (LifeTech)*, pages 92–93, 2021. doi: 10.1109/LifeTech52111.2021.9391844.

[204] Yonghao Song, Xueyu Jia, Lie Yang, and Longhan Xie. Transformer-based spatial-temporal feature learning for eeg decoding, 2021.

[205] Demetres Kostas, Stéphane Aroca-Ouellette, and Frank Rudzicz. Bendr: Using transformers and a contrastive self-supervised learning task to learn from massive amounts of eeg data. *Frontiers in Human Neuroscience*, 15:253, 2021. ISSN 1662-5161. doi: 10.3389/fnhum.2021.653659. URL `https://www.frontiersin.org/article/10.3389/fnhum.2021.653659`.

[206] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.

[207] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

[208] Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. Recurrent independent mechanisms. In *International Conference on Learning Representations*, 2020.

[209] Andrew Ng. Andrew ng x-rays the ai hype. ai pioneer says machine learning may work on test sets, but that's a long way from real world use, May 2021. URL `https://spectrum.ieee.org/andrew-ng-xrays-the-ai-hype`.

[210] Weijie Su, Małgorzata Bogdan, and Emmanuel Candes. False discoveries occur early on the lasso path. *The Annals of statistics*, pages 2133–2150, 2017.

[211] Kadierdan Kaheman, J Nathan Kutz, and Steven L Brunton. Sindy-pi: A robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *arXiv preprint arXiv:2004.02322*, 2020.

[212] M Quade, M Abel, J Kutz Nathan, and SL Brunton. Sparse identification of nonlinear dynamics for rapid model recovery. *Chaos (Woodbury, NY)*, 28(6):063116–063116, 2018.

[213] Yulia Rubanova, Ricky T. Q. Chen, and David Duvenaud. Latent odes for irregularly-sampled time series. *Advances in Neural Information Processing Systems*, 32, 7 2019. ISSN 10495258. URL `http://arxiv.org/abs/1907.03907`.

[214] David Kanaa, Vikram Voleti, Samira Kahou, Christopher Pal, and Mila CIFAR. Simple video generation using neural odes. In *Workshop on Learning with Rich Experience, Advances in Neural Information Processing Systems*, volume 32, 2019.

[215] Cagatay Yildiz, Markus Heinonen, and Harri Lähdesmäki. Ode2vae: Deep generative second order odes with bayesian neural networks. In *Advances in Neural Information Processing Systems*, 2019.

[216] Xuanqing Liu, Tesi Xiao, Si Si, Qin Cao, Sanjiv Kumar, and Cho-Jui Hsieh. How does noise help robustness? explanation and exploration under the neural sde framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[217] Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3870–3882. PMLR, 26–28 Aug 2020. URL `http://proceedings.mlr.press/v108/li20i.html`.

[218] Junteng Jia and Austin R Benson. Neural jump stochastic differential equations. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL `https://proceedings.neurips.cc/paper/2019/file/59b1deff341edb0b76ace57820cef237-Paper.pdf`.

[219] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *arXiv preprint arXiv:2005.08926*, 2020.

[220] Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal differential equations for scientific machine learning. Jan 2020. URL `https://arxiv.org/abs/2001.04385v3`.

[221] Ori Linial, Neta Ravid, Danny Eytan, and Uri Shalit. Generative ode modeling with known unknowns. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 79–94, 2021.

[222] J Nathan Kutz. *Data-driven modeling & scientific computation: methods for complex systems & big data*. Oxford University Press, 2013.

[223] Steve Brunton. Steve brunton's youtube channel. URL `https://www.youtube.com/c/Eigensteve/featured`.