



Subject Areas:

applied mathematics, mathematical biology, biophysics, complex systems

Keywords:

holistic mathematical modelling, biological complexity, multi-scale modelling, data-driven dynamics

Author for correspondence:

Ramón Nartallo-Kaluarachchi
e-mail: nartallokalu@maths.ox.ac.uk

From reductionism to realism: Holistic mathematical modelling for complex biological systems

Ramón Nartallo-Kaluarachchi^{1,2}, Renaud Lambiotte¹, and Alain Goriely¹

¹Mathematical Institute, University of Oxford, Woodstock Road, Oxford, OX2 6GG, United Kingdom

²Centre for Eudaimonia and Human Flourishing, University of Oxford, 7 Stoke Pl, Oxford, OX3 9BX, United Kingdom

At its core, the physics paradigm adopts a reductionist approach, aiming to understand fundamental phenomena by decomposing them into simpler, elementary processes. While this strategy has been tremendously successful in physics, it has often fallen short in addressing fundamental questions in the biological sciences. This arises from the inherent *complexity* of biological systems, characterised by heterogeneity, polyfunctionality and interactions across spatiotemporal scales. Nevertheless, the traditional framework of complex systems modelling falls short, as its emphasis on broad theoretical principles has often failed to produce predictive, empirically-grounded insights. To advance towards actionable mathematical models in biology, we argue, using neuroscience as a case study, that it is necessary to move beyond reductionist approaches and instead embrace the complexity of biological systems—leveraging the growing availability of high-resolution data and advances in high-performance computing. We advocate for a holistic mathematical modelling paradigm that harnesses rich representational structures such as annotated and multilayer networks, employs agent-based models and simulation-based approaches, and focuses on the inverse problem of inferring system dynamics from observations. We emphasise that this approach is fully compatible with the search for fundamental biophysical principles, and highlight the potential it holds to drive progress in mathematical biology over the next two decades.

1. Introduction

A fundamental aspect of human nature and success as a species is the desire to make sense of the world around us. Over the course of thousands of years and through the scientific revolution, mathematics, specifically *mathematical modelling* (MM), has become the dominant and ultimate paradigm used to abstract physical systems and describe their behaviour over time and space. This has resulted in the unparalleled success of MM in domains like physics and engineering, and has profoundly changed human society. The relationship has become naturally symbiotic with new mathematics constantly being discovered through the development of models of real-world phenomena.




Figure 1. The spectrum of interaction structures in mathematical models. Models of physical system exist on a spectrum from the regular to the random. For regular structures, continuum assumptions are often valid leading to a low number of degrees of freedom. In the random case, behaviour is typically described by a statistical distribution rather than individual elements. Between these extrema are complex systems, where interactions are heterogeneous leading to high-dimensional dynamics.

Many such examples have become canon in the field of MM such as Poincaré's attempts to model celestial mechanics [1], which laid the ground-work for the study of dynamical systems, divergent series, and resonance phenomena. In an effort to understand the bridges of Königsberg, Euler developed concepts that would later culminate in graph and network theory [2]. More recently, whilst deriving simple equations to describe the dynamics of the climate, Lorenz discovered the first example of a chaotic attractor [3]. Whilst we are still far from a complete theory, it can be argued that 20th century was the century of physics when mathematical abstractions led to a significant maturity in the field, with many fundamental problems being solved and theories confirmed [4].

However, biological systems, spanning from entire ecosystems to sub-cellular processes, have proved to be more impregnable to mathematical analysis. This is why it was first developed as a mostly descriptive science before moving to theoretical concerns. Yet, many fundamental theoretical questions in biology remain unanswered. How does life emerge from a precise arrangement of individually inanimate atoms? How does evolution shape species? How are cognition and consciousness orchestrated by the noisy, electrical signals of the human brain? How can we leverage the mechanisms of biology to eradicate disease? These problems demand solutions and, as of yet, mathematics has failed to provide a clear and consensus framework promising to yield them.

Nevertheless, it would be inaccurate to characterise mathematical biology as an immature discipline, as numerous models have demonstrated great success and provided valuable insights. However, these breakthroughs often exist in isolation and are typically limited to specific biological systems or processes. To understand why this limitation arises from the intrinsic nature of biological systems, we must first identify the characteristics of the dominant *reductionist* approach that has resulted in such

rampant success in physics. Throughout this *Perspective*, we will build our arguments with examples primarily from neuroscience and the human brain, which serves as the prototypical example of a complex biological system due to its inherent complexity, intricate organisation, and central importance to all aspects of human life and society.

We start with one of the most important breakthroughs in mathematical biology, the Hodgkin-Huxley (HH) model for electrophysiological excitability [5]. Some species of squid possess a so-called 'giant axon', up to 1.5mm in diameter, that controls water propulsion [6]. By isolating this large axon, Hodgkin and Huxley in the 1950s were able to insert voltage clamp electrodes and vary extracellular ionic concentrations to develop a complete set of *ordinary differential equations* (ODEs) that describe the propagation of electrical signals through a neuronal axon, providing the first complete description of the action potential, the basis of all neuronal communication. This is a canonical example of *reductionism*, a philosophical approach that assumes that complex phenomena can be described in terms of a series of simpler, more fundamental processes. In the case of the HH model, the key assumption was that neuronal activity could be modelled as an electrical circuit, with each component of the cell being represented by an electrical element. This, in turn, abstracts away other subcellular and biochemical processes.

The reductionist approach has proved extremely successful across science, but particularly in physics and chemistry, where simple experiments with only a few independent variables can reveal fundamental principles, and fundamental principles can be assembled together like building blocks to explain more complex situations. However, this approach may also be responsible for the lack of theoretical success that has plagued the biological sciences [7–11]. For example, whilst there are only a few elementary units for physical systems e.g. protons, neutrons and electrons, biological systems are characterised

by an array of unique elements e.g. the zoo of different lymphocytes in the immune system, the collections of proteins, or the extremely diverse populations of neurons and glial cells in the nervous system. Returning to Hodgkin and Huxley, whilst their result represented the crucial breakthrough in understanding the subcellular mechanisms governing neural firing, it brought us only marginally closer to understanding cognition or computational processes in the brain.

It is well known that many physical systems can be organised onto a spectrum spanning from the *regular* to the *random* [12], as illustrated in Fig. 1. It is well appreciated that the reductionist approach appears to be most effective at the extrema of this spectrum, where interactions present certain types of symmetries. At one end of the spectrum, for systems characterised by regular, symmetric organisation, such as atoms in a metal or crystal, individual elements can be abstracted away through continuum assumptions. For example, the flow of heat can be modelled using a *partial differential equation* (PDE) framework, without a need to model individual atoms [13]. At the other end of the spectrum, in order to describe random systems composed of enormous numbers of noisy, interacting components, such as the behaviour of molecules in a gas, physicists developed the framework of *statistical mechanics* [14]. Under this framework, it is typically assumed that interacting elements are identical. Moreover, one does not describe the behaviour of each individual element but instead the dynamics of a probability distribution that describes the ensemble, typically through mean-field approximations. This reduces the number of degrees of freedom and renders tractable mathematical descriptions.

Between these two extremes lie systems that, for lack of a better definition, are best described as *complex*. Such systems are characterised by topologically irregular interactions, heterogeneous elements, poly-functional processes, and evolve over multiple time- and length-scales. These features are particularly ubiquitous in biological processes. Returning to the brain, consider the activity in a large population of interconnected neurons. An accurate model would incorporate the complex connectivity between units, the heterogeneity of different cell-types, as well as the neurochemical processes underpinning their mechanism. Consequently, modelling their behaviour would require an intricate model with an enormous number of degrees of freedom, invariably leading to high-dimensional dynamics. Of course, the idea that biology is characterised by its complexity is not new [15, 16], and much of the emerging field of *complex systems science* is motivated by the study of biological systems [17, 18]. Nevertheless, despite the sustained attention over the past two decades, complex systems approaches to biology have yet to yield convincing and tangible results.

In this *Perspective*, we focus on the problem of biological

complexity and MM. In particular, we present a unified framework of complementary research directions, which together form what we call *holistic mathematical modelling*. We make the case that in modern sciences the classical *reductionist approach* cannot work in isolation, and it is necessary to develop models that consolidate different spatial and temporal scales, use multimodal data, and consider complex representational structures. This holistic approach goes beyond isolated functions to integrate multiple biophysical processes within a single model. Moreover, we argue that the theoretical approaches to complex systems that have been dominant for the past two decades, especially those models that are inspired by biological problems, are limited and often cannot yield empirically useful or actionable insights. Yet, the extent of the data and computational requirements needed for high-fidelity models is often prohibitive. Despite this, we believe that the development of such models is a necessity for progress in mathematical biology to match the accomplishments of experimental biology. Looking to the future, we argue that modern tools in data collection, *machine learning* (ML), *artificial intelligence* (AI), and *high-performance computing* (HPC) will not replace modelling, but instead facilitate it at unprecedented scale. By adopting a modern approach, we present a holistic modelling framework that embraces complexity and empiricism, whilst synergising with existing approaches, to develop a collective arsenal capable of answering the most pressing scientific challenges.

2. Rich and realistic representations for complex systems

In an effort to accurately represent real-world systems, from biology and beyond, the study of *networks* has become one of the most popular methods for abstracting elements and their interactions [19]. In the last 20 years, this has led to a number of important studies such as that of ‘scale-free’ [20], ‘small-world’ [12], or community-structured networks [21]. As a result, the study of complex systems has become, for many, synonymous with *network science*. This abstraction has been particularly useful in the modelling of a range of biological systems such as the brain [22], ecosystems [23], biomedical systems [24], protein interactions [25] and gene regulation [26]. By introducing weighted, directed and signed connections, as well as *temporal networks*, where links vary in time, these abstractions can be extended to model a variety of different biological processes.

Whilst intrinsically motivated by applications, the study of complex systems has, to date, remained largely theoretical, with a substantial focus on general principles that aim to transcend the confines of a specific system or process [27]. Similarly, many studies aimed directly at biological complexity, focus on ‘top-down’ models, which aim at reproducing qualitative *emergent* phenomena [28], as opposed to predictive accuracy. Whilst this is no doubt




Figure 2. Richer representations for complex systems. Networks have become the dominant paradigm for representing interactions in complex systems. However, a range of extensions exist that have the expressiveness to represent more complex interactions as well as their dynamics. Examples include hypergraphs, simplicial complexes, as well as temporal, multilayer and annotated networks.

important, it has broadly failed to represent biological systems with the required authenticity to posit or confirm theories, or often to yield practical results. As a result, complex systems science can often exist in a vacuum, adjacent to the mainstream of biology.

In recent times, in an effort to generalise networks to more expressive, realistic structures, significant attention has turned towards *higher-order* representations of complex systems [29–31] such as *hypergraphs*, *simplicial complexes* and *multilayer networks*, illustrated in Fig. 2. Multilayer networks have proved to be particularly useful representations for multi-scale biological systems [32] as they allow for modelling a number of different, yet coupled processes, evolving simultaneously [33, 34]. For example, neural activity in the human brain network is facilitated by blood-flow through the cerebral vasculature. Moreover, ageing results in structural changes to the networks that facilitate cognition. Whilst these processes have different mechanisms, dynamics and time-scales, they remain intrinsically coupled, reducing the efficacy of any model that attempts to consider them in isolation. Despite this, many models on multilayer networks remain largely theoretical, as they are typically overparameterised, requiring large amount of multimodal data for fitting. Nevertheless, the increasing availability and quality of multimodal data raises the opportunity for this expressive modelling technique to be integrated into more powerful, predictive methods, as shown by the recent *MINT* toolbox for neuroimaging data in disease [35]. We believe that the next 20 years will see an increased focus on this type of data and modelling.

Somewhat surprisingly, one higher-order representation of a complex system whose dynamics are rarely considered is the *annotated network*¹ [37]. Unlike the simple structure found in graph theory of a typical network, $G = (V, E)$, defined by a vertex set, V , and edge set, E , annotated networks can encode nodal *metadata* in the form of an additional vector of features $\mathbf{d}_i \in \mathbb{R}^d$ that is attached to each vertex $i \in V$ (or to each edge) [38, 39]. In particular, this extension allows us to model the heterogeneity of

individual elements, adding a crucial layer of additional expressiveness to our mathematical models. Whilst still nascent, such an approach has proved useful in neuroscience in the form of *biologically annotated connectomes* [40], which are human brain networks obtained from structural imaging where each node (region) is imbued with additional information such as cell-type density, gene transcription, receptor density, and myelination. Open-source software, such as *neuromaps* [41], facilitates the integration of this *multimodal data* into richer structural representations. Such heterogeneity has been used to develop biophysical models of whole-brain dynamics that incorporate regional heterogeneity through variations in gene expression [42] or synaptic strength [43] for added biological complexity and closer agreement with observed neural dynamics.

In an attempt to tackle the complexity of biological systems, the time has come to apply these frameworks more concretely to problems in biology. We are not arguing here that complex systems approaches to biology are entirely novel, but instead that the primary focus on emergent phenomena and the emphasis on universal phenomena have limited the accuracy and predictive power of these models. In turn, the mainstream within these fields may replace reductionist abstractions with richer representations that are capable of incorporating the heterogeneity seen in experiments. To achieve this goal, integration with multimodal data is crucial, but it is worth noting the scaling challenge that has proved limiting to this point. Given a system with N elements, the number of pairwise interactions scales as N^2 . If we introduce higher-order or multilayer interactions, the parameter space scales at an even faster rate $\sim N^b$. As a result, such over-parametrised models are plagued by both a lack of sufficient multimodal data and the inference algorithms needed to fit them. However, as we progress into the age of *big data*, modern experimental techniques are able to record multimodal data on an unprecedented scale, particularly in biological systems [44, 45]. Moreover, novel techniques in computational statistics and ML are able to perform parameter inference for such complex

¹It is worth noting that whilst annotated networks have been rarely studied in terms of dynamics, they are prevalent in some fields, such as *graph-representation learning* [36].

models using observed data [46, 47]. Unlike the purely statistical models typically found in ML, these systems have parameters that remain both physically constrained and meaningful. This ensures the interpretability of complex dynamical models and allays the usual concern about von Neumann's elephant² and his wiggly trunk [48]. With accurate calibration on experimental data and efficient simulation on HPCs, detailed, heterogeneous and elaborate models of complex systems can make the leap from generic and descriptive to mechanistic and predictive.

(a) Integrating many time- and length-scales

Whilst rich structures such as multilayer, annotated, or hyper-networks offer the expressiveness necessary to model a range of interacting biological processes, one remaining challenge is that many of these processes evolve over multiple time- and length-scales. In domains like physics, such scales of interaction can often be separated, hence the success of reductionist single-scale models. For example, to model the motion of planets in a solar system, we do not need to consider the atmospheric dynamics that occur at a smaller spatial scale or the gravitational field of stars that are light-years away. Similarly, a celestial model can assume that stars and planets are of fixed mass, ignoring stellar evolution and possible collision events, which occur on a much slower time-scale, thus are insignificant. These processes are mostly decoupled which allows for a remarkably precise modelling of planet motions over millennia.

The situation is completely different for biological systems that are best characterised by interacting spatiotemporal scales [49]. Focusing on a single scale, naturally limits our ability to explain certain phenomenon. For example, neurodegenerative diseases, like Alzheimer's disease (AD), evolve due to the spread of toxic proteins over the course of many years [50]. However, the symptoms of these conditions manifest themselves in cognitive decline stemming from pathological neural activity, on the order of $10^{-3} - 10^0$ seconds [51]. In order to understand the mechanisms by which AD causes cognitive decline, it is insufficient to only model either neural activity or toxic protein spread in isolation. Fig. 3 highlights the many processes underlying the development of AD along with their respective time- and length-scales. Recent studies have presented an approach that attempts to integrate these time-scales [52, 53]. First, they model network degradation due to toxic-protein spread, freezing the state of the network at distinct snapshots to then form the topology underlying a model of fast neural activity.

Whilst rich representations of complex systems can capture the intricate structure of interactions in biological processes, the full characterisation of their dynamics requires multi-scale models that couple interacting processes. As with

many of these proposed frameworks, this naturally complicates mathematical analysis and elegance in favour of greater biological realism.

3. Towards simulation-based science

(a) Agent-based models

Mathematical models are often appreciated for their elegance and parsimony. In particular, such models can often be treated analytically leading to great insight. As models become more complex, they lose analytical tractability and numerical methods of analysis must be used. One framework for modelling complex systems that hinges on direct simulation, is the use of *agent-based models* (ABMs), which have been embraced in economics [54, 55], computational social science [56], and mathematical biology [57, 58].

The typical structure of ABMs is the time evolution of a micro-scale model that defines the characteristics of individual elements and their interactions. Examples of such systems include models of the spread of disease through populations and ecological dynamics or collective motion with agents being individuals or populations (Panel A, Fig. 4). Unlike traditional models, ABMs call for a different approach to mathematical analysis, one that is often realised through exhaustive simulation and data-driven calibration. For this reason, such models were previously found to be computationally restrictive. However, following advancements in scientific computing, large-scale simulations of ABMs have become both tractable and attractive [59].

The adoption of ABMs requires a shift in attitudes towards mathematical models. In particular, the proponents of such models argue that the complexity of many systems requires a shift towards a purely *simulation-based science*, where additional features and realism are favoured over parsimony and elegance [60]. The current inclination toward a minimal model is often justified via *Occam's razor*, the philosophical argument that the simplest model is often the best. However, the lack of success of such minimal models in biology, compounded with the success of overparameterised ML models, has called into question this long-standing perspective [61, 62]. Whilst ML models are able to discover a multiplicity of latent features that have predictive power, a parsimonious mathematical modeller must instead identify the handful that may be important. Whilst not without critics [63], more complex models relieve the modeller of some of this choice by allowing them to include sufficient complexity at the level of model definition. Moreover, it is important to note the differences between an ABM and a ML model. Whilst both lack analytical tractability and rely on computationally-intensive processes, the latter may be considered to be a 'black-box',

²Whilst often attributed to von Neumann, the quote "...with four parameters I can fit an elephant, and with five I can make him wiggle his trunk." first appears in an article by Freeman Dyson, where he quotes Enrico Fermi, who is in turn quoting von Neumann [48].

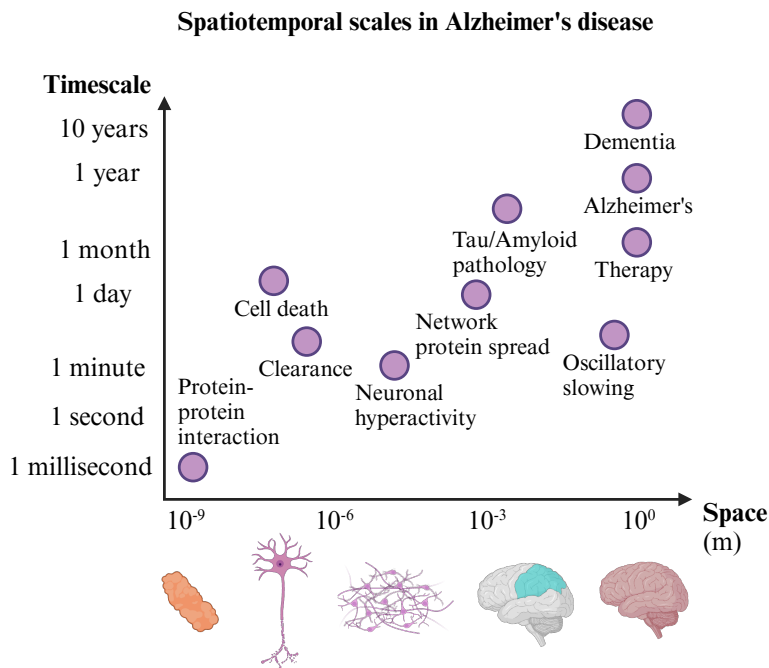


Figure 3. Spatiotemporal scales in Alzheimer's disease. Complex biological processes, such as the development of neurodegenerative diseases such as Alzheimer's disease (AD), evolve over a range of interacting spatiotemporal scales. Focussing on AD, protein misfolding ultimately leads to the spread of toxic tau and amyloid- β proteins through axonal fibres, which is partially cleared through cell regulation. Ultimately, this leads to cell death and pathological neural dynamics. This manifests in tau and amyloid pathology and, finally, cognitive decline which can be combatted with therapy. A MM of this process must take into account the full range of spatiotemporal scales in order to build a mechanistic understanding of AD.

whose mechanisms remain somewhat opaque, whilst the former attempts to model the basic physical mechanisms governing a system. This poses additional advantages, such as the ability to test realistic interventions in a system due to the model's interpretability [64].

Let us consider the example of an epidemic model. A traditional compartmental model, such as the 'SIR' model [65], is clearly overly simplistic and cannot express the heterogeneity of individuals in a population. Yet, it is mathematically tractable due to its ODE formulation and can be analysed to obtain fundamental insight and identify key quantities such as the *basic reproduction number*, R_0 . Moreover, with simple parameter modifications, the model can incorporate the effect of interventions such as social distancing or medication. On the other hand, *deep learning* based models are able to analyse large amounts of data, such as digital, genomic, environmental and behavioural data, in order to make accurate predictions about the development of the epidemic [66]. However, we cannot test the effects of interventions in such a model due to its opacity. Finally, we can also consider an ABM, which can model population heterogeneity, experiment with a variety of interventions, and make realistic predictions. Again, one major drawback is the need for exhaustive simulation and data for calibration. With modern digital data collection methods and improved HPC, though, such restrictions will

loosen in the coming years.

Despite these advantages, the necessity to specify many local rules of interaction can cause ABMs to suffer from *ad-hoc* selection of many *hidden* parameters. Typically at the modeller's discretion, such parameters are chosen inline with expected behaviour, leading to biases towards modeller's hypotheses. The modeller may often tinker with local rules until an expected behaviour is found on the screen, leading to the danger that the modelling process resembles a video game with the sole purpose of mimicking reality rather than understanding emergent behaviours. In search of a transparent framework for ABMs, modellers should aim to minimise hidden parameters and choose priors in a principled, rational manner rather than performing endless hand-tuning. These, in turn, may be updated with a Bayesian mechanism, leveraging available data.

(b) Digital twins

Another simulation-based framework for the modelling of complex systems, is the *digital twin* (DT). A DT is an adaptive model that is designed to virtually emulate the behaviour of a physical system [67]. Typically such a model is implemented on a computer. DTs have been considerably useful in engineering where they can be used to simulate the behaviour of a system under a range of conditions.




Figure 4. Simulation-based modelling. A. Agent-based models are powerful tools for describing the behaviour of heterogeneous populations of agentic units. Examples include the collective dynamics of swarms and flocks, where each agent can be modelled with its own dynamical properties. B. Digital twins are simulation-based models of physical systems implemented in a computer. One application of biomedicine is to use personalised recordings of neural activity or heartbeat data to calibrate individual in-silico models. Ultimately, this will lead to a virtual version of the biological system where interventions can be tested via simulation.

Successful examples include spacecraft during the Apollo missions [68], as well as other mechanical and aerospace engineering systems [69]. In some cases, a DT is composed of an ABM, such as in models of cities. Similarly to ABMs, whilst DTs have an underlying mathematical formulation, they require exhaustive simulation and HPC. Moreover, they are often characterised by more engineered interfaces such as *virtual* or *augmented reality* (VR/AR), for use by stakeholders other than the mathematical modeller.

In biology, DTs have sparked particular interest in the field of *precision medicine* [70]. Precision medicine refers to a treatment paradigm that considers the individual characteristics and data of a patient. This can be realised through the use of a personalised DT that uses biomedical data to calibrate a simulation of a specific patient's system to test interventions *in-silico*³, as illustrated in Panel B, Fig. 4. Two particular systems showing great promise are the heart [71] and the brain [72]. In particular, *The Virtual Brain* (TVB), more specifically *The Virtual Epileptic Patient* (VEP), is a project leveraging detailed biophysical models and HPCs, that predicts personalised estimates of epileptogenic zone networks [73]. The empirical, actionable success of this approach stands in contrast with the dominant paradigm of complex systems neuroscience that focuses on simple oscillatory models of neural dynamics [74] and has yet to realise a concrete medical application [75]. Nevertheless, it would be naive to assume that the pure emulation of a system is either sufficient or necessary to understand it. *The Human Brain Project*⁴, was one such attempt to emulate the whole brain from individual neurons upwards [76], that, despite some notable results at the micro-scale [77],

completely failed to deliver on its original promises [78]. On the other hand, whilst models like the VEP can have strong predictive power, this may come at the cost of interpretability. Important neuroscientific phenomena such as synchronisation or state-switching can be obscured by the more complex model whilst they would be obvious in an oscillator network or hidden Markov model. Moreover, they require large datasets of personalised neural recordings, therefore relying primarily on EEG, as opposed to more expensive, but much richer, modalities such as fMRI or MEG⁵. Finally, the VEP requires large computational resources, as opposed to simplified models, which can be yield some interesting and important insights at the 'laptop-level'.

At their best, DTs incorporate the necessary complexity to answer specific questions, test interventions, and simulate different environments, limited by the available data, rather than a perfectly faithful reconstruction of an entire system which achieves very little in scientific terms.

Both DTs and ABMs represent a new perspective to modelling that relies on computational simulation and empirical data. Blending this approach with modern technology, such as VR and AR, can lead to novel discoveries, out of the reach of simple mathematical models. To illustrate this, we highlight a recent study by Sayin et al [79] that investigated collective behaviour by analysing real locusts interacting in VR with a swarm of simulated counterparts. This experiment highlighted the insufficiency of the elegant, dominant, but reductionist, *Vicsek flocking model* [80] and brought new insight into

³As opposed to experiments *in-vitro* (in glass i.e. in a laboratory) or *in-vivo* (in a living system), *in-silico* refers to a third paradigm where experiments can be run on a computer.

⁴Continuing now as the *Blue Brain Project*.

⁵Electroencephalography (EEG), functional magnetic resonance imaging (fMRI) and magnetoencephalography (MEG) are methods of recording neural activity in humans.

the collective motion of swarming insects. At its core, this approach remains MM. However, the main-streaming of such frameworks requires the uprooting of the long-entrenched dogmas in applied mathematics that favour elegant solutions over realism and complexity.

(c) Automatic modelling

Advances in simulation-based science open the door for a new paradigm that we call *automatic modelling* (AM). By combining techniques in ML, system identification, and model selection, AM is an encompassing term for methods that are able to leverage simulations to perform ‘theory-discovery’ alongside *a posteriori reductionism*. We envisage a framework where modellers initially construct detailed models that incorporate all available complexities. Subsequently, employing optimisation strategies from computational statistics and ML, these models are systematically simplified by fitting them to experimental data while simultaneously penalising model complexity. This approach would allow AM to extract, in a sparse manner, the most relevant features or underlying physical principles essential for accurately explaining the observed system behaviour. Depending on the required predictive power, users of AM could adjust model complexity against model interpretability. It would answer the crucial question: how much of the data can be explained by a reduced set of key mechanisms?

Whilst this framework remains to be formally established, it shows strong alignment with existing model-selection and regularisation techniques, thus suggesting that it will become feasible within the coming decades, especially as LLM prompts for scientific discovery become increasingly conversational. Moreover, there exist some preliminary related approaches such as that of the *Bayesian machine scientist* [81], which explores and selects between candidate models using Monte-Carlo methods and Bayesian inference, penalising model-complexity, in order to perform system identification. *AI-Feynman* [82], *AI-Descartes* [83], and *AI-Hilbert* [84] are three related approaches that attempt to discover physical laws directly from observed data by blending theory and experimental data using ML techniques. AM represents a future ideal where rather than relying on subjective decisions by modellers or opaque nonlinear transformations, the underlying engine could objectively identify models with essential features. This would reveal the dominant mechanism responsible for explaining experimental data through an automated process.

(d) Generative ML and AI

It would be remiss to discuss the future of MM without making reference to the incredible developments in generative forms of ML and AI. In particular, innovations such as the transformer architecture [85], derivatives such as large language models (LLM), or other techniques,

such as diffusion models [86], seem posed to revolutionise society, including science, even after taking into account the unreasonable media hype. In particular, generative ML stands in stark contrast with the discriminative models of the past, as these methods learn a dynamical process that is capable of sampling new candidates, including protein structures [87], molecules [88], and human-level language capabilities [89].

Despite their generative and predictive power, these uninterpretable, black-box models seem at odds with MM, nevertheless there are a number of complementary directions. Whilst transformers lack the biological plausibility of spiking or recurrent networks, their function resembles that of grid-cells in the hippocampus [90], suggesting that they could yet teach us something about the brain [91], or simply prompt new methods in dimensionality reduction and visualisation for biological data [92]. Moreover, transformers have led to the development of so called *foundation models*, which can subsequently be fine-tuned for a number of scientific tasks, as diverse as predicting chaotic [93] and stochastic dynamics [94] or analysing single-cell data [95]. As such algorithms develop, they become a new tool in the arsenal of the modern mathematical modeller, opening up a plethora of synergistic methods [96], including improved model calibration or data-driven approaches to dynamical systems.

4. Modelling as an inverse-problem

The traditional approach to modelling empirical observations of a dynamical system is to attempt to write down differential equations that describe the behaviour of the system. Next, one attempts to find the parameters of this model that best explain the available data. However, in large-scale complex systems, writing down a mechanistic model becomes a difficult challenge. For example, whilst Hodgkin and Huxley were able to isolate and study the dynamics of a single neuron in a model organism, the same cannot be performed for a macroscopic brain region in a human. However, modellers now have access to an unprecedented amount of empirical data i.e. neural recordings. This has inspired the development of a modern *data-driven* approach to dynamical systems, specifically the area of *system identification* [97]. In particular, both parameter inference and system identification are examples of an *inverse problem* [98], which involves calculating a set of causal factors, be it parameter- or model-selection, from a set of observations.

These techniques usually begin with a *multivariate time-series* (MVTs), $\{x_i(t_j) : i = 1, \dots, N, t_j = j\Delta t\}$, which records snapshots of the state, $\mathbf{x} = (x_1, \dots, x_N)$, at discrete time-points, t_j . Typically, we assume that the time-series was

generated from a dynamical system of the form,

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t) + \xi(\mathbf{x}, t), \quad (4.1)$$

where \mathbf{f} represents the mechanistic dynamics of the system, known as the *drift*, and ξ represents the additive noise, known as the *diffusion*. The goal of a data-driven method is to infer \mathbf{f} from the MVTs⁶ [99]. Here methods can vary, with some attempting to infer the exact mathematical terms in \mathbf{f} , typically known as *sparse regression* methods, and others approximating \mathbf{f} with deep learning methods. Popular existing techniques include *SINDy* [100], which performs sparse regression for nonlinear, deterministic dynamics, as well as its extensions to stochastic differential equations (SDEs) [101] and PDE [102] models. For SDEs, a number of other methods exist such as *stochastic force inference* [103, 104], *InferenceMAP* [105] and maximum-likelihood estimation [99]. These methods have particular applications to Brownian dynamics in soft-matter systems [106–109]. More recently, extensions have been considered for high-dimensional networked systems [110, 111].

Beyond sparse force inference methods, the dynamics, \mathbf{f} , can be approximated using *generative machine learning* [112]. Deep learning architectures such as *recurrent neural networks*⁷ (RNNs) and *transformers* have been used to learn and predict dynamics from observations with applications to neural and molecular data [94, 114, 115].

(a) Latent representations for dynamical systems

Many methods for analysing nonlinear dynamics from high-dimensional data rely on the *manifold hypothesis* [116], which argues that high-dimensional data arising from complex systems are typically clustered around a lower-dimensional manifold [117, 118]. As a result, the dynamics evolve in a latent space with fewer dimensions. This hypothesis justifies the use of dimensionality reduction methods [119] such as *Principal Component Analysis* (PCA), *t-distributed Stochastic Neighbour Embedding* (t-SNE), and *Uniform Manifold Approximation and Projection* (UMAP), which are ubiquitous in biological data analysis. Such an assumption is also prevalent in data-driven approaches to dynamical systems, where much focus is placed on learning low-dimensional latent representations of high-dimensional dynamics. Architectures such as *variational auto-encoders* (VAEs) are particularly adept at projecting high-dimensional time-series into low-dimensional spaces by learning nonlinear transformations [120], as illustrated in Panel A, Fig. 5. Examples of this approach include the discovery of whole-brain dynamics from neural recordings, which has been used to develop heterogenous models of brain regions [121].

Whilst latent dimensions no longer correspond to physical

measurements from a system, they can remain interpretable. For example, in the study of motor-control, so called *neural manifolds* are hypothesised to encode representations of physical space in neural activity [122], which can be decoded using deep learning [123]. Latent variables can also be constrained to discrete-state spaces such as *hidden Markov models* (HMMs). HMMs can be fit to time-series data using variational inference [124] or *VAMPNets* [125]. These have been applied extensively in biology to uncover discrete latent states, such as in genomic sequence analysis [126], or in neural recordings [127, 128].

Opposite to dimensionality reduction, ‘*lifting*’, projection of dynamics into higher-dimensional space, is also particularly useful for solving inverse problems. In particular, the *Koopman operator* lifts nonlinear dynamics into an infinite-dimensional system where the dynamics are linear [129–131]. Finite approximations of the Koopman operator can be learned using *dynamic mode decomposition* [132] and deep learning [133], as illustrated in Panel B, Fig. 5.

Data-driven approaches to dynamical systems represent a modern framework for the modelling of complex systems. When dealing with high-dimensional and complex systems, the construction of mechanistic dynamical models from first principles has proved an insurmountable challenge. On the other hand, data-driven methods leverage the enormous amounts of available data, alongside modern ML techniques, to reframe modelling as a data-led, inverse problem. Such techniques are not only predictive, but often informative and interpretable. For example, lifted linear embeddings of nonlinear dynamics are more tractable for mathematical analysis, such as eigendecomposition, than their nonlinear counterparts. Moreover, low-dimensional latent dynamics can give insight into the effective dynamics of complex systems that are obscured by dimensionality. The modelling of complex biological systems can benefit greatly from an inferential approach that focuses on predicting observed dynamics rather than theoretical models, particularly towards the aim of system control through perturbation and intervention.

5. A fundamental physics for biology

We have argued that theoretical biology should turn towards complex, empirical models with predictive capabilities, even if these models occasionally complicate traditional mathematical analyses. Nevertheless, this shift does not negate the importance of identifying fundamental theoretical principles that govern biological systems, but instead, supports it. Unlike physics, biology still lacks clear translational, foundational principles. Thus, the field urgently requires the development of a ‘theoretical physics for living systems’ [134]. Whilst significant progress remains limited, several promising directions have emerged in recent years, where holistic MM will play a crucial role.

⁶Stochastic methods often also attempt to infer the diffusion tensor $\mathbf{D}(\mathbf{x}, t) = \frac{1}{2}\xi(\mathbf{x}, t)\xi^T(\mathbf{x}, t)$ [99].

⁷*Reservoir computers* [113] are also often used for predicting nonlinear dynamics, however they are a form of RNN.

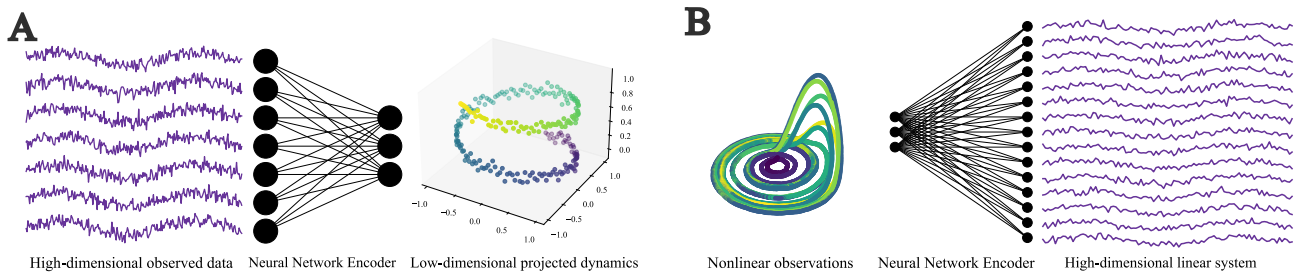


Figure 5. Latent representations for observed dynamical systems. A. High-dimensional data from complex biological systems is often clustered onto a low-dimensional manifold. Using data science techniques such as variational auto-encoders, we can discover the low-dimensional latent dynamics of the process. B. Following the theory of the Koopman operator, nonlinear dynamics can be ‘lifted’ into high-dimensional spaces where the dynamics are approximately linear. This projection can be achieved with deep learning techniques such as neural networks.

Models of collective phenomena often exhibit *phase-transitions* occurring at *critical points* [135]. In 1987, Bak et al. hypothesised that complex systems may *self-organise* to the critical point, where both correlation-length across the system and sensitivity to perturbations are maximised [136, 137]. This theory sparked particular interest in biological sciences, where evidence of *self-organised criticality* has been found in neuronal firing [138] and genetic expression [139, 140], leading to more general theories about how critical dynamics facilitate biological function [141, 142]. Moreover, recent work formalising the concept of *Griffith’s phases*, parameter regions showing critical-like dynamics as opposed to isolated points, promise new avenues for the furtherance of this theory [143], which could prove to be both a governing principle for biological systems, and explain observed biological complexity. The study of criticality in neural systems has motivated a number of data-driven approaches to dynamical systems, in particular to the problem of inferring collective properties from a fraction of total observables [144, 145] - techniques with implications far beyond neuroscience.

Another prominent theory accounting for self-organisation in biological systems is the *free-energy principle* (FEP) [146]. Originally proposed as a theory of computation in the brain, the FEP suggests that biological organisms encode sensory information in the form of a probabilistic world model that is updated according to Bayesian mechanics⁸. Next, it posits that organisms act in order to minimise the *surprise*⁹ i.e. the difference between sensory information and the world model. Whilst the FEP is formulated to account for perception-actions loops in the brain, it represents a more general theory of decision making that has been recently used to model collective behaviour in swarms [147]. More generally, this theory, amongst other theories of agentic decision making, can be used to construct complex interacting models, such as ABMs and DT, that begin from a cognitive, mechanistic basis.

Nonequilibrium thermodynamics has long been regarded as central to an understanding of the energetic constraints of biological processes [15, 148]. However, the theoretical formalism for describing the *stochastic* thermodynamics of fluctuating, microscopic processes is particularly recent [149, 150]. Subsequently, stochastic thermodynamics has become extremely useful for describing the activity of biological processes, inspiring a plethora of techniques for analysing the thermodynamics of observed biological activity [151–154], with recent interest in neural activity [155–157]. In particular, a major focus has been placed on the reconstruction of force fields from stochastic trajectories, to estimate thermodynamic quantities [103, 158]. In addition, recent work has shown that ML algorithms can learn thermodynamic laws, such as the arrow of time, opening new avenues for the analysis of biological data [159]. Despite these advancements, there remain a number of open questions regarding the thermodynamic limits of complexity and information processing [160], and how this pertains to biological processes. Nevertheless, investigating the thermodynamic constraints on biological processes promises to be an avenue ripe for discovery.

Finally, in an effort to describe the mechanisms by which complex structures, such as amino acids, evolve, Sharma et al. recently proposed *assembly theory* [161]. Assembly theory quantifies the complexity of a structure in terms of the amount of selection required to produce it. Whilst theoretical in nature, assembly theory offers an avenue for understanding the constraints on the complexity of molecules, giving insight into evolution, but also yielding a testable indicator for the existence of extraterrestrial life [162]. Moreover, its algorithmic nature makes it well suited for simulation-based discovery and in-silico experimentation, ultimately offering new directions towards artificial or digital life [163].

In the near future, mathematical biology and complex systems should prioritise developing empirically predictive and practical models, while continuing the pursuit of

⁸It is worth noting that the ‘energy’ in the FEP is not usually a physical energy but a larger unifying quantity that plays a similar role.

⁹The negative log probability of an outcome.

underlying theoretical principles. This area of investigation requires bold theorising that, whilst controversial at times, promises to move the field forward. Modellers should err towards theoretical predictions that are easily testable on existing data, and are rooted in established and well-tested physical principles. Progress on this front will only serve to improve the mechanistic understanding and empirical usefulness of the complex models that we advocate for. Again, these goals may be symbiotic. As mathematical and computational models begin to yield useful, actionable insights, the apparent resistance in empirical biological sciences to theoretical hypotheses may shift, ultimately blurring the edges between domains and leading to a truly interdisciplinary approach. Moreover, the holistic MM advanced here will naturally lead to the development of further theories. For example, system identification will reveal novel mechanisms governing biological processes, whilst simulation-based approaches amount to in-silico experimentation where theories can be efficiently tested and remoulded. Together, these techniques yield a novel approach to science that is better fit for both our unsolved problems and the modern world.

6. Back to the future

While it is foolish to predict the future, it is childish to deny its inevitability. As we observe the fast developments of complexity science, perspicuous challenges emerge. Clearly, developing a holistic framework for modelling complex biological systems that explains empirical data and yields useful predictions, remains one of the most daunting yet pressing challenges in interdisciplinary science. In this Perspective, we lay out a roadmap of promising research directions and necessary conceptual shifts required to achieve significant advances in the field over the coming decades. The unique challenges of biological complexity call for innovative tools that move beyond the prevalent reductionist approaches in current mathematical modelling.

We highlight three particularly promising avenues that have the potential to significantly enhance mathematical biology, by adopting a holistic viewpoint integrating multiple biophysical processes and scales within a single model.

First, we advocate adopting rich and sophisticated modelling frameworks, such as multilayer, annotated, temporal, and hyper-networks, which possess sufficient expressive power to represent biological complexity accurately. Second, we recommend the adoption of simulation-based methodologies, including agent-based models and digital twins, to incorporate detailed mechanisms and rigorously test interventions. Third, we encourage the use of data-driven modelling for dynamical processes, emphasising the inverse problem—inferring underlying mechanisms from empirical data—over the traditional forward-problem approach typically used in mathematical modelling.

Nonetheless, we reaffirm that the identification of fundamental theoretical principles should remain the ultimate goal of biology. We highlight recent key advances offering translational, foundational insights into longstanding biological questions, building upon established theories from physics. We believe that holistic mathematical modelling promises to yield new paradigms capable of constructing and testing biological theories efficiently, hence achieving the dream of a truly quantitative theory of life.

One can appreciate the substantial progress achieved over recent decades while recognising that significant further advances are both attainable and essential. Successfully addressing these challenges would dramatically transform interdisciplinary research in mathematics and biology, paving the way for breakthroughs in synthetic biology, artificial life, and numerous other biotechnologies currently confined to science fiction.

Data Accessibility. No new data were created in this study.

Authors' Contributions. R.N.-K, R.L and A.G designed the study. R.N.-K wrote the manuscript. R.N.-K, R.L and A.G reviewed and edited the manuscript.

Competing Interests. The authors declare no competing interests

Funding. R.N.-K. is supported by an Engineering and Physical Sciences Research Council (EPSRC) doctoral scholarship from grants EP/T517811/1 and EP/R513295/1, and The Alan Turing Institute through an Enrichment Community Award which is funded under grant EP/Z532861/1. R.L. acknowledges support from the EPSRC grants EP/V013068/1, EP/V03474X/1 and EP/Y028872/1.

References

- ¹H. Poincaré, *Les Methodes Nouvelles de la Mécanique Celeste* (Gauthier-Villars, 1892).
- ²L. Euler, "Solutio problematis ad geometriam situs pertinentis", *Commentarii Academiae Scientiarum Petropolitanae*, 128–140 (1741).
- ³E. N. Lorenz, "Deterministic Nonperiodic Flow", *Journal of the Atmospheric Sciences* **20**, 130–141 (1963).
- ⁴J. Agar, *Science in the Twentieth Century and Beyond* (Polity Press, 2012).
- ⁵A. L. Hodgkin and A. F. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve", *Journal of Physiology* **117** (1952).
- ⁶J. Z. Young, "The Functioning of the Giant Nerve Fibres of the Squid", *Journal of Experimental Biology* **15**, 170–185 (1938).

- ⁷T. Nagel, *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False* (Oxford University Press, 2012).
- ⁸P. Anderson, "More Is Different", *Science* **177**, 393–396 (1972).
- ⁹P. Ball, "Should biology put complexity first?", *Cell Systems* **16** (2025).
- ¹⁰S. Brenner, "Sequences and consequences", *Proceedings of the Royal Society B* **365**, 207–212 (2010).
- ¹¹L. Gyllingberg, A. Birhane, and D. J. Sumpter, "The lost art of mathematical modelling", *Mathematical Biosciences* **362** (2023).
- ¹²D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks", *Nature* **393**, 440–442 (1998).
- ¹³J. Fourier, *Théorie analytique de la chaleur* (Firmin Didot Père et Fils, 1822).
- ¹⁴J. W. Gibbs, *Elementary Principles in Statistical Mechanics* (Charles Scribner's Sons, 1902).
- ¹⁵E. Schrödinger, *What is Life? The Physical Aspect of the Living Cell and Mind* (Cambridge University Press, Cambridge, United Kingdom, 1944).
- ¹⁶M. Gell-Mann, *The Quark & the Jaguar: Adventures in the Simple & the Complex* (Abacus, 1995).
- ¹⁷A. Ma'ayan, "Complex systems biology", *Journal of the Royal Society Interface* **14** (2017).
- ¹⁸D. C. Krakauer, *The Complex World: An Introduction to the Foundations of Complexity Science* (SFI Press, 2024).
- ¹⁹M. Newman, *Networks* (Oxford University Press, 2018).
- ²⁰A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks", *Science* **286**, 509–512 (1999).
- ²¹M. Girvan and M. E. Newman, "Community structure in social and biological networks", *Proceedings of the National Academy of Sciences* **99**, 7821–7826 (2002).
- ²²D. S. Bassett and O. Sporns, "Network neuroscience", *Nature Neuroscience* **20**, 353–364 (2017).
- ²³J. A. Dunne, R. J. Williams, and N. D. Martinez, "Food-web structure and network theory: The role of connectance and size", *Proceedings of the National Academy of Sciences* **99**, 12917–12922 (2002).
- ²⁴A.-L. Barabási, N. Gulbahce, and J. Loscalzo, "Network Medicine: A Network-based Approach to Human Disease", *Nature Reviews Genetics* **12**, 56–68 (2011).
- ²⁵N. Safari-Alighiarloo, M. Taghizadeh, M. Rezaei-Tavirani, B. Goliaei, and A. A. Peyvandi, "Protein-protein interaction networks (PPI) and complex diseases", *Gastroenterology and Hepatology From Bed to Bench* **7**, 17–31 (2014).
- ²⁶E. Davidson and M. Levin, "Gene regulatory networks", *Proceedings of the National Academy of Sciences* **102**, 4935 (2005).
- ²⁷G. Bianconi et al., "Complex systems in the spotlight: next steps after the 2021 Nobel Prize in Physics", *Journal of Physics: Complexity* **4** (2023).
- ²⁸G. Pezzulo and M. Levin, "Top-down models in biology: explanation and control of complex living systems above the molecular level", *Journal of the Royal Society Interface* **13** (2016).
- ²⁹R. Lambiotte, M. Rosvall, and I. Scholtes, "From networks to optimal higher-order models of complex systems", *Nature Physics* **15**, 313–320 (2019).
- ³⁰L. Torres, A. S. Blevins, D. Bassett, and T. Eliassi-Rad, "The Why, How, and When of Representations for Complex Systems", *SIAM Review* **63**, 435–485 (2021).
- ³¹F. Battiston et al., "The physics of higher-order interactions in complex systems", *Nature Physics* **17**, 1093–1098 (2021).
- ³²M. Gosak, R. Markovič, J. Dolensek, M. S. Rupnik, M. Marhl, A. Stožer, and M. Perc, "Network science of biological systems at different scales: A review", *Physics of Life Reviews* **24**, 118–135 (2018).
- ³³M. Kivelä, A. Arenas, M. Barthélemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks", *Journal of Complex Networks* **2**, 203–271 (2014).
- ³⁴R. F. Betzel and D. S. Bassett, "Multi-scale brain networks", *NeuroImage* **160**, 73–83 (2017).
- ³⁵S. Sarraf, B. Avelar-Pereira, S. M. H. Hosseini, and ADNI, "Multilayer Integration of Networks Toolbox (MINT)", *Communications Biology* **8** (2025).
- ³⁶W. L. Hamilton, *Graph Representation Learning* (Morgan & Claypool, 2020).
- ³⁷M. E. J. Newman and A. Clauset, "Structure and inference in annotated networks", *Nature Communications* **7** (2016).
- ³⁸L. Peel, D. B. Larremore, and A. Clauset, "The ground truth about metadata and community detection in networks", *Science Advances* **5** (2017).
- ³⁹A. Bassolas, A. Holmgren, A. Marot, M. Rosvall, and V. Nicosia, "Mapping nonlocal relationships between metadata and network structure with metadata-dependent encoding of random walks", *Science Advances* **8** (2022).
- ⁴⁰V. Bazinet, J. Y. Hansen, and B. Misić, "Towards a biologically annotated brain connectome", *Nature Reviews Neuroscience* **24**, 747–760 (2023).
- ⁴¹R. D. Markello et al., "neuromaps: structural and functional interpretation of brain maps", *Nature Methods* **19**, 1472–1479 (2022).
- ⁴²G. Deco, M. L. Kringelbach, A. Arnatkeviciute, S. Oldham, K. Sabarodin, N. C. Rogasch, K. M. Aquino, and A. Fornito, "Dynamical consequences of regional heterogeneity in the brain's transcriptional landscape", *Science Advances* **7** (2021).
- ⁴³M. Demirtaş, J. B. Burt, M. Helmer, J. L. Ji, B. D. Adkinson, M. F. Glasser, D. C. V. Essen, S. N. Sotiropoulos, A. Anticevic, and J. D. Murray, "Hierarchical Heterogeneity across Human Cortex Shapes Large-Scale Neural Dynamics", *Neuron* **20**, 1181–1194 (2019).

- ⁴⁴V. Marx, “The big challenges of big data”, *Nature* **498**, 255–260 (2013).
- ⁴⁵T. J. Sejnowski, P. S. Churchland, and J. A. Movshon, “Putting big data to good use in neuroscience”, *Nature Neuroscience* **17** (2013).
- ⁴⁶T. Gaskin, G. A. Pavliotis, and M. Girolami, “Neural parameter calibration for large-scale multiagent models”, *Proceedings of the National Academy of Sciences* **120**, e2216415120 (2023).
- ⁴⁷K. Cranmer, J. Brehmer, and G. Louppe, “The frontier of simulation-based inference”, *Proceedings of the National Academy of Sciences* **117**, 30055–30062 (2020).
- ⁴⁸F. Dyson, “A meeting with Enrico Fermi”, *Nature* **427** (2004).
- ⁴⁹J. O. Dada and P. Mendes, “Multi-scale modelling and simulation in systems biology”, *Integrative Biology* **2** (2011).
- ⁵⁰C. Soto, “Unfolding the role of protein misfolding in neurodegenerative diseases”, *Nature Reviews Neuroscience* **4**, 49–60 (2003).
- ⁵¹W. de Haan, K. Mott, E. C. W. van Straaten, P. Scheltens, and C. J. Stam, “Activity Dependent Degeneration Explains Hub Vulnerability in Alzheimer’s Disease”, *PLOS Computational Biology* **8**, e1002582 (2012).
- ⁵²A. Goriely, E. Kuhl, and C. Bick, “Neuronal Oscillations on Evolving Networks: Dynamics, Damage, Degradation, Decline, Dementia, and Death”, *Physical Review Letters* **125** (2020).
- ⁵³C. G. Alexandersen, W. de Haan, C. Bick, and A. Goriely, “A multi-scale model explains oscillatory slowing and neuronal hyperactivity in Alzheimer’s disease”, *Journal of the Royal Society Interface* **20** (2023).
- ⁵⁴R. Axtell and J. D. Farmer, “Agent-Based Modeling in Economics and Finance: Past, Present, and Future”, *Journal of Economic Literature* **63**, 197–287 (2025).
- ⁵⁵J. D. Farmer and D. Foley, “The economy needs agent-based modelling”, *Nature* **460**, 685–686 (2009).
- ⁵⁶A. Volkening, “A primer on data-driven modeling of complex social systems”, arXiv **2210.08636** (2022).
- ⁵⁷G. An, Q. Mi, J. Dutta-Moscato, and Y. Vodovotz, “Agent-based models in translational systems biology”, *WIREs Systems Biology and Medicine* **1**, 151–182 (2009).
- ⁵⁸A. Goriely, *The Mathematics and Mechanics of Biological Growth* (Springer Verlag, 2017).
- ⁵⁹E. N. Bodine, R. M. Panoff, E. O. Voit, and A. E. Weisstein, “Agent-Based Modeling and Simulation in Mathematics and Biology Education”, *Bulletin of Mathematical Biology* **82** (2020).
- ⁶⁰A. Lavin et al., “Simulation Intelligence: Towards a New Generation of Scientific Methods”, arXiv **2112.03235** (2022).
- ⁶¹P. Ball, “The Tyranny of Simple Explanations”, *The Atlantic* (2016).
- ⁶²M. Dubova et al., “Is Ockham’s razor losing its edge? New perspectives on the principle of model parsimony”, *Proceedings of the National Academy of Sciences* **122**, e2401230121 (2025).
- ⁶³I. Mazin, “Inverse Occam’s razor”, *Nature Physics* **18**, 367–368 (2022).
- ⁶⁴R. E. Baker, J.-M. Peña, J. Jayamohan, and A. Jérusalem, “Mechanistic models versus machine learning, a fight worth fighting for the biological community?”, *Biology Letters* **14** (2018).
- ⁶⁵W. O. Kermack and A. G. McKendrick, “A contribution to the mathematical theory of epidemics”, *Proceedings of the Royal Society A* **115** (1927).
- ⁶⁶A. Rodríguez, H. Kamarthi, P. Agarwal, J. Ho, M. Patel, S. Sapre, and B. A. Prakash, “Machine learning for data-centric epidemic forecasting”, *Nature Machine Intelligence* **6**, 1122–1131 (2024).
- ⁶⁷D. Gelernter, *Mirror Worlds: Or: The Day Software Puts the Universe in a Shoebox...How It Will Happen and What It Will Mean* (Oxford University Press, 1991).
- ⁶⁸B. D. Allen, *Digital Twins and Living Models at NASA*, technical report (NASA, 2021).
- ⁶⁹A. Ferrari and K. Willcox, “Digital twins in mechanical and aerospace engineering”, *Nature Computational Science* **4**, 178–183 (2024).
- ⁷⁰F. S. Collins and H. Varmus, “A New Initiative on Precision Medicine”, *The New England Journal of Medicine* **372**, 793–795 (2015).
- ⁷¹G. Coorey et al., “The health digital twin to tackle cardiovascular disease—a review of an emerging interdisciplinary field”, *npj Digital Medicine* **5** (2022).
- ⁷²V. Jirsa, H. Wang, P. Triebkorn, M. Hashemi, J. Jha, J. Gonzalez-Martinez, M. Guye, J. Makhlova, and F. Bartolomei, “Personalised virtual brain models in epilepsy”, *The Lancet Neurology* **22**, 443–454 (2023).
- ⁷³H. E. Wang et al., “Delineating epileptogenic networks using brain imaging data and personalized modeling in drug-resistant epilepsy”, *Science Translational Medicine* **15** (2023).
- ⁷⁴P. Ashwin, S. Coombes, and R. Nicks, “Mathematical frameworks for oscillatory network dynamics in neuroscience”, *The Journal of Mathematical Neuroscience* **6**, 1–92 (2016).
- ⁷⁵F. E. Turkheimer, F. E. Rosas, O. Dipasquale, D. Martins, E. D. Fagerholm, P. Expert, F. Váša, L.-D. Lord, and R. Leech, “A Complex Systems Perspective on Neuroimaging Studies of Behavior and Its Disorders”, *Neuroscientist* **28**, 382–399 (2022).
- ⁷⁶H. Markram, “The Blue Brain Project”, *Nature Reviews Neuroscience* **7**, 153–160 (2006).
- ⁷⁷H. Markram et al., “Reconstruction and Simulation of Neocortical Microcircuitry”, *Cell* **163**, 456–492 (2015).
- ⁷⁸S. Theil, “Why the Human Brain Project Went Wrong—and How to Fix It”, *Scientific American* (2015).

- ⁷⁹S. Sayin et al., “The behavioral mechanisms governing collective motion in swarming locusts”, *Science* **387**, 995–1000 (2025).
- ⁸⁰T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, “Novel Type of Phase Transition in a System of Self-Driven Particles”, *Physical Review Letters* **75** (1995).
- ⁸¹R. Guimerà, I. Reichardt, A. Aguilar-Mogas, F. A. Massucci, M. Miranda, J. Pallarès, and M. Sales-Pardo, “A Bayesian machine scientist to aid in the solution of challenging scientific problems”, *Science Advances* **6** (2020).
- ⁸²S.-M. Udrescu and M. Tegmark, “AI Feynman: A physics-inspired method for symbolic regression”, *Science Advances* **6** (2020).
- ⁸³C. Cornelio, S. Dash, V. Austel, T. R. Josephson, J. Goncalves, K. L. Clarkson, N. Megiddo, B. E. Khadir, and L. Horesh, “Combining data and theory for derivable scientific discovery with AI-Descartes”, *Nature Communications* **14** (2023).
- ⁸⁴R. Cory-Wright, C. Cornelio, S. Dash, B. E. Khadir, and L. Horesh, “Evolving scientific discovery by unifying data and background knowledge with AI Hilbert”, *Nature Communications* **15** (2024).
- ⁸⁵A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention Is All You Need”, in 31st Conference on Neural Information Processing Systems (NIPS 2017) (2017).
- ⁸⁶J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep Unsupervised Learning using Nonequilibrium Thermodynamics”, in *Proceedings of the 32nd International Conference on Machine Learning* (JMLR, 2015).
- ⁸⁷J. Jumper et al., “Highly accurate protein structure prediction with AlphaFold”, *Nature* **596**, 583–589 (2021).
- ⁸⁸D. D. Martinelli, “Generative machine learning for de novo drug discovery: A systematic review”, *Computers in Biology and Medicine* **145** (2022).
- ⁸⁹*GPT-4 Technical Report*, technical report (Open AI, 2024).
- ⁹⁰J. C. Whittington, J. Warren, and T. E. Behrens, “Relating transformers to models and neural representations of the hippocampal formation”, in *International Conference on Learning Representations 2022* (2022).
- ⁹¹L. Muller, P. S. Churchland, and T. J. Sejnowski, “Transformers and cortical waves: encoders for pulling in context across time”, *Trends in Neuroscience* **47**, 788–802 (2024).
- ⁹²W. Gilpin, “The cell as a token: high-dimensional geometry in language models and cell embeddings”, arXiv **2503.20278** (2025).
- ⁹³J. Lai, A. Bao, and W. Gilpin, “Panda: A pretrained forecast model for universal representation of chaotic dynamics”, arXiv **2505.13755** (2025).
- ⁹⁴C. Casert, I. Tamblin, and S. Whitelam, “Learning stochastic dynamics and predicting emergent behavior using transformers”, *Nature Communications* **15** (2024).
- ⁹⁵A. Chevaliera et al., “TEDDY: A Family Of Foundation Models For Understanding Single Cell Biology”, arXiv **2503.03485** (2025).
- ⁹⁶B. Noordijk, M. L. G. Gomez, K. H. W. J. ten Tusscher, D. de Ridder, A. D. J. van Dijk, and R. W. Smith, “The rise of scientific machine learning: a perspective on combining mechanistic modelling with machine learning for systems biology”, *Frontiers in Systems Biology* **4** (2024).
- ⁹⁷S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control* (Cambridge University Press, 2022).
- ⁹⁸A. M. Stuart, “Inverse problems: A Bayesian perspective”, *Acta Numerica* **19**, 451–599 (2010).
- ⁹⁹R. Friedrich, J. Peinke, M. Sahimi, and M. R. R. Tabar, “Approaching complexity by stochastic methods: From biological systems to turbulence”, *Physics Reports* **506**, 87–162 (2011).
- ¹⁰⁰S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems”, *Proceedings of the National Academy of Sciences* **113**, 3932–3937 (2016).
- ¹⁰¹J. L. Callahan, J.-C. Loiseau, G. Rigas, and S. L. Brunton, “Nonlinear stochastic modelling with Langevin regression”, *Proceedings of the Royal Society A* **477** (2021).
- ¹⁰²S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Data-driven discovery of partial differential equations”, *Science Advances* **3** (2017).
- ¹⁰³A. Frishman and P. Ronceray, “Learning Force Fields from Stochastic Trajectories”, *Physical Review X* **10** (2020).
- ¹⁰⁴D. B. Brückner, P. Ronceray, and C. P. Broedersz, “Inferring the Dynamics of Underdamped Stochastic Systems”, *Physical Review Letters* **125** (2020).
- ¹⁰⁵M. E. Beheiry, M. Dahan, and J.-B. Masson, “InferenceMAP: mapping of single-molecule dynamics with Bayesian inference”, *Nature Methods* **12**, 594–595 (2015).
- ¹⁰⁶L. P. García, J. D. Pérez, G. Volpe, A. V. Arzola, and G. Volpe, “High-performance reconstruction of microscopic force fields from Brownian trajectories”, *Nature Communications* **9** (2018).
- ¹⁰⁷D. B. Brückner and C. P. Broedersz, “Learning dynamical models of single and collective cell migration: a review”, *Reports on Progress in Physics* **87** (2024).
- ¹⁰⁸F. S. Gnesotto, G. Gradziuk, P. Ronceray, and C. P. Broedersz, “Learning the non-equilibrium dynamics of Brownian movies”, *Nature Communications* **11** (2020).
- ¹⁰⁹P. Ronceray, *Learning dynamical models from stochastic trajectories* (Thèse d’Habilitation à Diriger les Recherches, Aix-Marseille Université, 2024).

- ¹¹⁰T.-T. Gao, B. Barzel, and G. Yan, “Learning interpretable dynamics of stochastic complex systems from experimental data”, *Nature Communications* **15** (2024).
- ¹¹¹T.-T. Gao and G. Yan, “Autonomous inference of complex network dynamics from incomplete and noisy data”, *Nature Computational Science* **2**, 160–168 (2022).
- ¹¹²W. Gilpin, “Generative learning for nonlinear dynamics”, *Nature Reviews Physics* **6**, 194–206 (2024).
- ¹¹³D. J. Gauthier, E. Bollt, A. Griffith, and W. A. S. Barbosa, “Next generation reservoir computing”, *Nature Communications* **12** (2021).
- ¹¹⁴D. Durstewitz, G. Koppe, and M. I. Thurm, “Reconstructing computational system dynamics from neural data with recurrent neural networks”, *Nature Reviews Neuroscience* **24**, 693–710 (2023).
- ¹¹⁵G. Koppe, H. Toutounji, P. Kirsch, S. Lis, and D. Durstewitz, “Identifying nonlinear dynamical systems via generative recurrent neural networks with applications to fMRI”, *PLOS Computational Biology* **15**, e1007263 (2019).
- ¹¹⁶C. Fefferman, S. Mitter, and H. Narayanan, “Testing the manifold hypothesis”, *Journal of the American Mathematical Society* **29**, 983–1049 (2016).
- ¹¹⁷V. Thibeault, A. Allard, and P. Desrosiers, “The low-rank hypothesis of complex systems”, *Nature Physics* **20**, 294–302 (2024).
- ¹¹⁸B. Prasse and P. V. Miegheem, “Predicting network dynamics without requiring the knowledge of the interaction graph”, *Proceedings of the National Academy of Sciences* **119** (2022).
- ¹¹⁹L. van der Maaten, E. Postma, and J. van den Herik, “Dimensionality Reduction: A Comparative Review”, *Journal of Machine Learning Research* **10**, 66–71 (2009).
- ¹²⁰L. Girin, S. Leglaive, X. Bie, J. Diard, T. Hueber, and X. Alameda-Pineda, “Dynamical Variational Autoencoders: A Comprehensive Review”, *Foundations and Trends in Machine Learning* **15**, 1–175 (2021).
- ¹²¹V. Sip, M. Hashemi, T. Dickscheid, K. Amunts, S. Petkoski, and V. Jirsa, “Characterization of regional differences in resting-state fMRI with a data-driven network model of brain dynamics”, *Science Advances* **9** (2023).
- ¹²²J. A. Gallego, M. G. Perich, L. E. Miller, and S. A. Solla, “Neural Manifolds for the Control of Movement”, *Neuron* **94**, 978–984 (2017).
- ¹²³A. Gosztolai, R. L. Peach, A. Arnaudon, M. Barahona, and P. Vandergheynst, “MARBLE: interpretable representations of neural population dynamics using geometric deep learning”, *Nature Methods* **22**, 612–620 (2025).
- ¹²⁴M. J. Beal, “Variational algorithms for approximate Bayesian inference”, PhD thesis (University of London, 2003).
- ¹²⁵A. Maradt, L. Pasquali, H. Wu, and F. Noé, “VAMPnets for deep learning of molecular kinetics”, *Nature Communications* **9** (2018).
- ¹²⁶B.-J. Yoon, “Hidden Markov Models and their Applications in Biological Sequence Analysis”, *Current Genomics* **10**, 402–415 (2009).
- ¹²⁷D. Vidaurre, S. M. Smith, and M. W. Woolrich, “Brain network dynamics are hierarchically organized in time”, *Proceedings of the National Academy of Sciences* **114**, 12827–12832 (2017).
- ¹²⁸D. Vidaurre, L. T. Hunt, A. J. Quinn, B. A. E. Hunt, M. J. Brookes, A. C. Nobre, and M. W. Woolrich, “Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks”, *Nature Communications* **9** (2018).
- ¹²⁹S. L. Brunton, M. Budisic, E. Kaiser, and J. N. Kutz, “Modern Koopman Theory for Dynamical Systems”, *SIAM Review* **64** (2022).
- ¹³⁰I. Mezić, “Koopman Operator, Geometry, and Learning of Dynamical Systems”, *Notices of the American Mathematical Society* **67** (2021).
- ¹³¹I. Mezić, “The Operator is the Model”, *SIAM News* **57** (2024).
- ¹³²M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, “A Data-Driven Approximation of the Koopman Operator: Extending Dynamic Mode Decomposition”, *Journal of Nonlinear Science* **25**, 1307–1346 (2015).
- ¹³³B. Lusch, J. N. Kutz, and S. L. Brunton, “Deep learning for universal linear embeddings of nonlinear dynamics”, *Nature Communications* **9** (2018).
- ¹³⁴P. Mehta, “Theory is dead. Long live theory! For a 21st century statistical physics of life”, arXiv **2410.20506** (2024).
- ¹³⁵H. Stanley, *Introduction to phase transitions and critical phenomena* (Oxford University Press, 1987).
- ¹³⁶P. Bak, C. Tang, and K. Wiesenfeld, “Self-organized criticality: An explanation of the $1/f$ noise”, *Physical Review Letters* **59** (1987).
- ¹³⁷P. Bak and M. Paczuski, “Complexity, contingency, and criticality”, *Proceedings of the National Academy of Science* **92**, 6689–6696 (1995).
- ¹³⁸J. M. Beggs and D. Plenz, “Neuronal Avalanches in Neocortical Circuits”, *Journal of Neuroscience* **23**, 11167–11177 (2003).
- ¹³⁹B. Vidiella, A. Guillamon, J. Sardanyés, V. Maull, J. Pla, N. Conde, and R. Solé, “Engineering self-organized criticality in living cells”, *Nature Communications* **12** (2021).
- ¹⁴⁰M. Nykter, N. D. Price, M. Aldana, S. A. Ramsey, S. A. Kauffman, L. E. Hood, O. Yli-Harja, and I. Shmulevich, “Gene expression dynamics in the macrophage exhibit criticality”, *Proceedings of the National Academy of Science* **105**, 1897–1900 (2008).

- ¹⁴¹M. A. Muñoz, “Colloquium: Criticality and dynamical scaling in living systems”, *Reviews of Modern Physics* **90** (2018).
- ¹⁴²T. Mora and W. Bialek, “Are Biological Systems Poised at Criticality?”, *Journal of Statistical Physics* **144**, 268–302 (2011).
- ¹⁴³P. Moretti and M. A. Muñoz, “Griffiths phases and the stretching of criticality in brain networks”, *Nature Communications* **4** (2013).
- ¹⁴⁴A. Levina, V. Priesemann, and J. Zierenberg, “Tackling the subsampling problem to infer collective properties from limited data”, *Nature Reviews Physics* **4**, 770–784 (2022).
- ¹⁴⁵J. Wilting and V. Priesemann, “Inferring collective dynamical states from widely unobserved systems”, *Nature Communications* **9** (2018).
- ¹⁴⁶K. Friston, “The free-energy principle: a unified brain theory?”, *Nature Reviews Neuroscience* **11**, 127–138 (2010).
- ¹⁴⁷C. Heins, B. Millidge, L. D. Costa, R. P. Mann, K. J. Friston, and I. D. Couzin, “Collective behavior from surprise minimization”, *Proceedings of the National Academy of Sciences* **127**, e2320239121 (2024).
- ¹⁴⁸G. Nicolis and I. Prigogine, *Self-Organisation in Nonequilibrium Systems* (Wiley & Sons, 1977).
- ¹⁴⁹U. Seifert, “Stochastic thermodynamics, fluctuation theorems and molecular machines”, *Reports on Progress in Physics* **475** (2012).
- ¹⁵⁰C. P. Broedersz and P. Ronceray, “25 years of nanoscale thermodynamics”, *Nature* **604**, 46–47 (2022).
- ¹⁵¹C. Battle, C. P. Broedersz, N. Fakhri, V. F. Geyer, J. Howard, C. F. Schmidt, and F. C. Mackintosh, “Broken detailed balance at mesoscopic scales in active biological systems”, *Science* **352**, 604–607 (2016).
- ¹⁵²J. Li, J. M. Horowitz, T. R. Gingrich, and N. Fakhri, “Quantifying dissipation using fluctuating currents”, *Nature Communications* **10** (2019).
- ¹⁵³I. D. Terlizzi, M. Gironella, D. Herraes-Aguilar, T. Betz, F. Monroy, M. Baiesi, and F. Ritort, “Variance sum rule for entropy production”, *Science* **383**, 971–976 (2024).
- ¹⁵⁴J. M. Horowitz and T. R. Gingrich, “Thermodynamic uncertainty relations constrain non-equilibrium fluctuations”, *Nature Physics* **16**, 15–20 (2019).
- ¹⁵⁵C. W. Lynn, E. J. Cornblath, L. Papadopoulos, and D. S. Bassett, “Broken detailed balance and entropy production in the human brain”, *Proceedings of the National Academy of Sciences* **118** (2021).
- ¹⁵⁶R. Nartallo-Kaluarachchi, L. Bonetti, G. Fernández-Rubio, P. Vuust, G. Deco, M. L. Kringelbach, R. Lambiotte, and A. Goriely, “Multilevel irreversibility reveals higher-order organisation of non-equilibrium interactions in human brain dynamics”, *Proceedings of the National Academy of Sciences* **122**, e2408791122 (2025).
- ¹⁵⁷R. Nartallo-Kaluarachchi, M. Kringelbach, G. Deco, R. Lambiotte, and A. Goriely, “Nonequilibrium physics of brain dynamics”, *arXiv* **2504.12188** (2025).
- ¹⁵⁸R. Nartallo-Kaluarachchi, P. Expert, D. Beers, A. Strang, M. L. Kringelbach, R. Lambiotte, and A. Goriely, “Decomposing force fields as flows on graphs reconstructed from stochastic trajectories”, *Proceedings of the Third Learning on Graphs Conference (LoG 2024)* **PMLR** **269** (2024).
- ¹⁵⁹A. Seif, M. Hafezi, and C. Jarzynski, “Machine learning the thermodynamic arrow of time”, *Nature Physics* **17**, 105–113 (2021).
- ¹⁶⁰D. H. Wolpert, “The stochastic thermodynamics of computation”, *Journal of Physics A: Mathematical and Theoretical* **52** (2019).
- ¹⁶¹A. Sharma, D. Czégel, M. Lachmann, C. P. Kempes, S. I. Walker, and L. Cronin, “Assembly theory explains and quantifies selection and evolution”, *Nature* **622**, 321–328 (2023).
- ¹⁶²S. M. Marshall et al., “Identifying molecules as biosignatures with assembly theory and mass spectrometry”, *Nature Communications* **12** (2021).
- ¹⁶³S. I. Walker and P. C. W. Davies, “The algorithmic origins of life”, *Journal of the Royal Society Interface* **10** (2013).