Open Problems in Mathematical Biology

Sean T. Vittadello[1,2] & Michael P.H. Stumpf[1,2,3]

[1]Melbourne Integrative Genomics, University of Melbourne, Australia
[2]School of BioSciences, University of Melbourne, Australia
[3]School of Mathematics and Statistics, University of Melbourne, Australia

June 22, 2022

Abstract

Biology is data-rich, and it is equally rich in concepts and hypotheses. Part of trying to understand biological processes and systems is therefore to confront our ideas and hypotheses with data using statistical methods to determine the extent to which our hypotheses agree with reality. But doing so in a systematic way is becoming increasingly challenging as our hypotheses become more detailed, and our data becomes more complex. Mathematical methods are therefore gaining in importance across the life- and biomedical sciences. Mathematical models allow us to test our understanding, make testable predictions about future behaviour, and gain insights into how we can control the behaviour of biological systems. It has been argued that mathematical methods can be of great benefit to biologists to make sense of data. But mathematics and mathematicians are set to benefit equally from considering the often bewildering complexity inherent to living systems. Here we present a small selection of open problems and challenges in mathematical biology. We have chosen these open problems because they are of both biological and mathematical interest.

1 Introduction

In 1202 Leonardo Pisano (aka Fibonacci) studied the population dynamics of rabbits in his book *Liber Abaci*. By some, admittedly glib, measure mathematical biology could be said to predate theoretical physics — which also started in Pisa with the work of Galileo Galilei — by some 400 years. Arguably this head start has been lost and for the past century physics and mathematics have become deeply entwined. Biology, by contrast, was often dominated by observational or reductionist approaches and mathematical methods have for a long time only played a negligible role. Biologists have, however, been quick to adopt quantitative methods and tools as soon as there was a clear use for them. Biological problems and analyses have, for example, been important drivers of statistical methodology

1

development, and statistics has become a main feature — often not well liked — of biology undergraduate and graduate education. Mathematics, less so.

The reasons why mathematics is a useful tool for many (not all) biological questions have been superbly laid out before [1, 2, 3]. The principal reason, arguably, is that through the lens of mathematics we can make accessible the structures and mechanisms underlying heterogeneous and often disparate data and observations. If, for example, we can connect two observations or measurements in one model (say the level of growth factor in a Petri dish or well containing cells and the down-stream change in phosphorylation of nuclear ERK [4, 5]) we have a non-trivial mechanistic hypothesis from which we initiate further investigations. The analogy between mathematics and a microscope [2] in making the unobserved or unobservable visible holds surprisingly often.

Our focus here is therefore different: we are interested in what are the fundamental challenges for mathematical biology right now. Any such discussion will be subjective and incomplete. In our case it is clearly also influenced by our respective backgrounds (pure mathematics and theoretical/statistical physics), our past and current research interests (which are pretty broad, but mostly concerned with cellular processes and biophysics), and the conversations that we have had with researchers over the years. One key person was Edmund Crampin, especially for MPHS. The selection of material here cannot claim completeness, but we do think that all the problems touched upon here would have been of interest to Edmund. We hope that this will also be of wider interest.

Our discussion below is broken into four parts: (i) A brief discussion as to the different demands on modelling and the modeller in biology and physics; (ii) A selection of open problems in mathematical biology. These questions were partly crowd-sourced, and we attempt to put them into the wider context of mathematical modelling in biology and the biomedical sciences. This part contains open problems that are concerned with making mathematical modelling more useful and relevant to biological problems, and open problems where new mathematics will be essential to find answers to long-standing open questions at the core of biological research. We then (iii) provide an, arguably opinionated, overview of the opportunities that exist for ideas and concepts primarily from pure mathematics, but also theoretical physics to enrich the toolbox of mathematical biologists and study aspects that are particular to living systems.

2   Why is Modelling Biology so Difficult?

The question really refers to a comparison with physics (and parts of chemistry): why is modelling biology apparently so much more difficult than modelling physics? There are three aspects to our attempt at an answer: (i) physicists can often rely on first principles to model systems of interest; (ii) in biology we often have to start from heuristic models; and finally, (iii) differences in education between physics and biology.

Our attempts at finding an answer to this question should start by pointing out that

not all of physics is easy to model. For many physical problems, however, it is possible to write down the structure of a model, often starting from first principles [6, 7]. While it took major intellectual feats to write down, for example, Kepler's laws, Newton's laws, Maxwell's equations, Einstein's theories of relativity, and the laws of quantum mechanics and quantum field theory, the structure of a model is typically derivable from first principles. Noether's Theorem [8] states that every conservation law has an associated symmetry: for example, conservation of energy is related to invariance with respect to time. Invariance with respect to position, or translational invariance, implies the conservation of momentum. In both a real and an abstract sense, Noether's theorem underlies much of what is often perceived as the beauty of physics; and being able to rely on it (the Ward–Takahashi identities generalise the classical picture into the domain of quantum field theory) makes the life of the working physicist much easier.

For dissipative systems, that is for systems where energy is consumed, Noether's theorem does not apply in its simple form. Biology is replete with systems that consume and turn over energy, a fundamental aspect of life. This fact alone suffices to account for the profound difficulties that arise in finding mathematical descriptions of biological systems and processes.

In practical terms the main difference between modelling physics and modelling biology is that the mathematical form of the model is known in most physical cases. It may be difficult to solve the many particle Dirac Equation, or the Navier–Stokes equation in 3D for realistic viscosity, but the form or type of equations is rarely uncertain. If it is, however, for example systems with broken symmetry [9], then it is typically of utmost physical interest to identify and impose the correct form of the model.

In biology, by contrast, we have to start from a different level of abstraction, and the best model structure is rarely if ever known with any degree of certainty. We can make pragmatic assumptions to develop some simplified models, and the canonical model of this approach is clearly the Lotka–Volterra model

$$\frac{dx}{dt} = \alpha x - \beta xy,$$

$$\frac{dy}{dt} = \gamma xy - \delta y$$

where x and y denote a prey and a predator species respectively[1]. The rate parameters, α, β, γ, and δ represent, in turn, the prey growth rate, the decline in the prey population due to predation, the growth of the predator population, and the reduction in the predator population due to predator death. This model has been successful in explaining many aspects of population biology and ecology. But it is highly abstracted and idealised, and fundamental features of reproduction, predation, and so forth are purposefully ignored in this simple model. But this simple Lotka–Volterra model can form the basis for more complicated descriptions, including spatially structured populations, overlapping niches, and multiple trophic levels. We can use these simple models to obtain qualitative insights

into fundamental aspects of biological processes and phenomena, including bet-hedging and the evolution of virulence. The Lotka–Volterra model is sufficiently easy to allow mathematical analysis, and simple enough to be accessible to people with little direct modelling experience. Progress in physics has often relied on a small canon of models that are equally relevant and accessible to theoreticians and experimentalists: the single harmonic oscillator, Ising Model, particle in a box, H_2 molecule, etc. Biology needs more of these [10].

Finally, the training tradition in physics is different: the typical physics curriculum is heavy on mathematics and mathematical methods (with considerable differences between institutions and even nations) that prepare students for theoretical physics. Over the past generation the mathematics curricula of physics degrees appear to have changed comparatively little in substance, but all physicists are exposed to the mathematical techniques that are necessary to make progress in research. The situation in biology is very different: life-scientists often receive training in qualitative and statistical methods required to make sense of observational and experimental data, whereas there is generally no consideration of modelling the behaviour of biological systems. The mathematical tools required by mathematical biology, meanwhile, are still proliferating; the diversity of problems considered and the complexity of data being generated pose considerable technical and conceptual challenges which require a broad mathematical tool-set. We pick up this point in Section 4.

3   Open Problems in Mathematical Biology: A Small Selection.

3.1   A Note on Notation

We are considering models, $f(x; \theta)$, describing the state $x \in \Omega \subseteq \mathbb{R}^N$ of a biological system; the states x are confined to the state space, Ω, of the system. Here $\theta \in \Theta \subseteq \mathbb{R}^d$ is the d-dimensional vector of model parameters (e.g. reaction rate parameters) in the parameter space Θ. We denote by x_i and θ_i the ith components of the state and parameter vectors, respectively.

When $f(x; \theta)$ represents a dynamical system, for example an ordinary differential equation $dx/dt = f(x; \theta)$, we use $x_0 \in \Omega$ to denote the initial conditions. Alternative models will be represented by subscripts, j, therefore $f_j(x; \theta)$.

We represent any experimental observations or data by $\mathcal{D} = \{d_1, d_2, \ldots, d_m\}$; here d_i is a vector with number of coordinates possibly less than N to account for observations that do not capture the whole state space, therefore where $\dim(d_i) < \dim(x)$.

3.2   Inverse Problems and Parameter Sensitivity

In many situations we have a candidate model, $f(x; \theta)$, or at least its structure. What we then require are the parameters, θ, and on occasion initial conditions, $x_0 = x(0)$. Choice

of parameters is crucial if we want to compare model output with observed behaviour or data $\mathcal{D}$. Often parameters are chosen from the literature, but this approach is fraught with potential problems [11]. Literature information can be imprecise about the conditions under which parameters were determined; if, for example, ambient pH is not comparable then the reaction rate may differ from the conditions under consideration. Combining reaction rate parameters from different sources is problematic if the corresponding experimental conditions were not identical or at least comparable. So in many instances we will have to infer parameters from data [12, 13, 14, 15, 16].

We typically have three options to delimit the range for the parameter θ. We can identify the value of θ that produces model output that is most similar to the observed data. The likelihood,

$$L(\theta) = \Pr(\mathcal{D}|\theta), \tag{1}$$

is the probability of the data for a given parameter θ. The maximum-likelihood estimate (MLE), $\hat{\theta} = \mathrm{argmax}_{\theta}(L(\theta))$, corresponds to the parameter that has the highest probability of producing the observed data, and it can be obtained through optimisation. We usually use optimisation to estimate $\hat{\theta}$. In addition to this point-estimate we can determine confidence intervals for the value of θ.

In a Bayesian framework we determine the posterior probability of a parameter,

$$\Pr(\theta|\mathcal{D}) = \frac{\Pr(\mathcal{D}|\theta)\pi(\theta)}{\Pr(\mathcal{D})}, \tag{2}$$

where the prior, $\pi(\theta)$, represents our knowledge (or assumptions) about the parameter value prior to looking at any data. The posterior is thus determined from both prior information and the available data, $\mathcal{D}$, via the likelihood, $\Pr(\mathcal{D}|\theta)$. The probability distribution over parameters rather than a single estimate tends to be the focus of Bayesian inference. Both frequentist and Bayesian inference are vast areas of research in their own right.

What really drives inference, irrespective of the framework, is the information content of the data (assuming that the model structure, that is the functional form of $f(x;\theta)$, is sufficiently close to the truth). One measure of this is how much varying θ changes the probability of observing the data, therefore the likelihood or the posterior probability. Information geometry arguably provides the most rigorous way of assessing certainty in an estimate, especially the so called Fisher information [17]. Under mild regularity conditions the Fisher information matrix is given by

$$[I(\theta)]_{i,j} = -E\left[\frac{\partial^2}{\partial\theta_i\partial\theta_j}\log L(\theta)\,|\theta\right]. \tag{3}$$

The curvature of the the log-likelihood, or posterior function (or the curvature of a different cost function), is an appropriate measure of certainty for uni-modal likelihoods. For multi-modal likelihoods/posteriors *ad-hoc* measures have been proposed, but no easy or universally satisfactory solution exists [18, 19].

What we do find in many practical applications is that for any given dataset and model we can only infer some of the parameters (more generally some combinations of parameters); the rest are simply not inferable with any precision, for example their marginal posteriors do not differ substantially from the posteriors. Identifying which parameter combinations can be learned [20] from prior knowledge plus data will become a crucial part of parameterising large models, including whole cell models [21]. In the short term sensitivity and robustness analysis — therefore varying parameters and assessing the resulting differences in model output — can be used to assess which parameters can be inferred from data [22].

3.3 Model Selection

Parameter estimation is already a formidable problem. But generally, as we have shown above and will argue further below, the structure and nature of the correct model are unknown. Choosing which model M_i from a set of models, $\mathcal{M} = \{M_1, \ldots, M_N\}$, does the best job of explaining data has therefore become a central focus in some areas of mathematical biology. A host of methods are available to enforce something akin to Ockham's razor: generating a model that is as complex as required but not more so. These methods have been reviewed extensively elsewhere [23]. When parameters can be estimated from the data using likelihood or Bayesian methods, information criteria such as the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) are popular because they explicitly penalise against models with more parameters:

$$\text{AIC}_i = -2\log L(\theta_i; M_i) + 2k_i, \tag{4}$$

$$\text{BIC}_i = -2\log L(\theta_i; M_i) + k_i \log n, \tag{5}$$

where k_i is the number of parameters of model i and n is the size of the dataset. They are both, however, approximations (and in the case of the AIC biased towards overly complex models as the amount of data increases) to the marginal likelihood of a model,

$$\Pr(\mathcal{D}|M_i) = \int_\Omega \Pr(\theta|M_i)\Pr(\mathcal{D}|\theta, M_i)d\theta, \tag{6}$$

and the model posterior probability,

$$\Pr(M_i|\mathcal{D}) = \frac{\Pr(\mathcal{D}|M_i)\pi(M_i)}{\Pr(\mathcal{D})}. \tag{7}$$

Here the number of model parameters is *implicitly* penalised.

Model selection is challenging because of the high-dimensional integration problem in Equation (6). Any advances in constraining and guiding parameter inference — such as by identifying subspaces of the full parameter space Θ containing parameters that do not affect the model fit to data — will potentially help, but these subspaces will differ between

6

alternative models [24]. Specifying the prior probability of models is notoriously difficult: in principle there is no clear *a priori* bound on the number of models, which could be (and arguably ought to be) infinite. This is a fundamental problem of, and a recurring theme in, Bayesian model selection.

As a general rule, problems encountered in parameter estimation tend to be exacerbated in the context of model selection. Computational demands obviously increase with the number of models. Perhaps not so obvious is that many computational schemes that apply straightforwardly in the context of parameter estimation — such as Markov Chain Monte Carlo sampling [25], or approximate Bayesian computation (ABC) [26, 27] — require significant modification and adaptation for model selection.

As in the case of parameter estimation, appropriate robustness analysis can help. This often involves extending analysis over many related models [28, 29]. In topological sensitivity analysis [30], for example, the structure of the model is perturbed by adding, deleting, or otherwise changing one or more terms in the equations. That is, we go from candidate model M to another model M'. We then fit model M', or in reality many alternative models with perturbed structures, to the data to see if alternative models can explain the data. The results of such an analysis can often be sobering, with many alternative model structures being virtually indistinguishable. This level of topological robustness could imply two very different things: (i) the observed behaviour is robust to variation in the model structure; or (ii) the large number of models may simply reflect that the data are not sufficient to discriminate between models.

3.4 Design Principles

Even for mathematical modellers in biology the model should be of subsidiary interest compared with the biological phenomena under investigation. Model comparison methods, including model selection, multi-model averaging, and topological sensitivity analysis can be used to group models that provide comparable descriptions of a certain type of biological behaviour, henceforth denoted by Q. For biological systems, instances of a behaviour Q could be oscillations, switch-like characteristics [31], signal filtering [32], robust perfect adaptation [33, 34], or Turing pattern behaviour [35, 36].

For some types of behaviour, Q, it is possible to derive precise mathematical statements that express either the necessary or sufficient conditions (ideally both) for a system to exhibit Q. For example, we know that: (i) multi-stability depends on the presence of positive feedback [37]; (ii) for robust perfect adaptation (RPA), algebraic conditions have been derived that establish the existence of RPA [34]; (iii) a Turing instability is required for spontaneous pattern formation in reaction diffusion systems [35]. But generally we still have to validate whether a given model M exhibits behaviour Q or not. Further, for many behaviours Q we may not even have easy criteria to consider.

Given a model universe, $\mathcal{M}$, we can use simulations to identify a subset, $\widetilde{\mathcal{M}} \subseteq \mathcal{M}$, of models that exhibit a particular behaviour Q. From the subset $\widetilde{\mathcal{M}}$ we may be able to distill

7

the design principles using a comparative analysis of the models $M \in \widetilde{\mathcal{M}}$. This process, however, is not without difficulties: structural properties, of for example the reaction network, do not always suffice to guarantee the behaviour Q; boundary conditions and parameter regimes shape the model output, dynamics, and behaviour, so must also be considered within the comparison methodology. Visual inspection is prohibitive, unreliable, and unrigorous.

The necessity for more flexible, universal, and automatable methodologies to compare models is evident. With this objective in mind, employing simplicial complexes (which are a generalisation of combinatorial graphs) has allowed us to develop flexible representations of both quantitative and qualitative models that are described with any formalism [38, 39]. Comparison of the simplicial complexes, by our notions of distance or equivalence, provides a meaningful conceptual comparison for the corresponding models. Our model comparison methodology allows for clustering of models according to flexible conceptual considerations, which can lead to the discovery of novel design principles for biological systems.

3.5   Model Development

When we cannot rely on first principles for model development we have to use heuristics or intuition to construct models. As we have shown above, this can work but is in no way guaranteed. Considering large model universes, $\mathcal{M}$, and model selection or topological sensitivity analysis offers a potential solution. But every model $M \in \mathcal{M}$ still has to be written down and translated into code to be executed by the computer. When this has to happen for many models, or for very large models, the potential for errors is considerable. If basic properties of models (number and types of interactions, structural properties, kinetics, and so forth) can be specified it is often possible to construct these models automatically, for example by exhaustive enumeration. Metaprogramming provides particularly elegant ways of automatically and on the fly developing and analysing mathematical model structures [40]. In metaprogramming, software is considered as part of the data and can therefore be manipulated during run-time and on the fly.

If we can guide the way in which software 'writes and rewrites itself', for example by combining it with model selection or robustness analysis, we could take steps towards developing larger and larger, including whole cell, models. In this framework we would start from a model seed, $M^{(0)}$, and propose a new model according to

$$M^{(i)} \xrightarrow{\ G(M^{(i)})\ } M^{(i+1)}, \tag{8}$$

where $G(M^{(i)})$ is used to propose an update to the model that results in a new model $M^{(i+1)}$; this could, for example, be a stochastic context-free grammar.

In addition to being, we would argue, less prone to errors than manual curation, this would also facilitate reproducible model development. We already have the ability to share models between groups using file exchange formats [41, 42]. But SBML and CellML

model specifications do not allow us to deduce how and on what basis certain modelling choices were made; even reading the original publications rarely suffices. If models are created algorithmically then the choices that go into a model-development pipeline can be summarised and communicated efficiently. The development of the model thus becomes reproducible if the updating mechanism and the initial model, $M^{(0)}$, are known.

The procedure in Equation (8) can be interpreted as constituting a type of random walk through model space. The process needs to be guided by model selection after every iteration of Equation (8), or every few iterations. The statistical challenges are considerable: (i) conventional parameter estimation at each step would quickly become prohibitive; and (ii) the dimension of the model space can change at each step which makes model selection — the statistical comparison of successive models, $M^{(i)}$ and $M^{(i+1)}$ — challenging. Techniques such as reversible jump Markov chain Monte Carlo are required to construct appropriate samplers that can explore varying dimensional spaces.

3.6 Multiscale, Hierarchical, and Spatio-Temporal Models

Most biological problems intrinsically span scales: from molecules to cells; from cells to tissues and whole organisms; from species to ecosystems; from genetic variants to populations. Dynamics at one scale have repercussions at other scales, both higher and lower. Sometimes we can only observe data at a higher level and have to draw inferences at the lower level: molecular dynamics drive cell-fate decisions; cellular processes lead to tissue formation; within-host dynamics drive population-level epidemics. The literature in this field is enormous and rapidly growing; we therefore limit the discussion in this section to two aspects: (i) mathematical and statistical challenges related to development and analysis of multiscale models; and (ii) a brief outline of bond graphs as a flexible framework for thermodynamically correct models.

The joint consideration of biological processes across scales opens up new possibilities to capture complex processes. Often we can take a pragmatic approach and use coarse descriptions of the lower-scale dynamics to explore the resulting behaviour at a higher scale. For example, we can use simplified cellular dynamics to model the formation of patterns in tissues. In this manner, agent-based models, for example, allow us to test our hypotheses about intra-cellular dynamics by comparisons with the behaviour observed at the tissue level. The critical evaluation of agreement between multiscale models and biological reality requires careful consideration: just like for models at the same scale there may be many alternative model realisations that give rise to the same behaviour. Testing the sensitivity to model details, a multiscale modelling equivalent of topological sensitivity analysis, does not seem to be widely used. Agent-based models can be computationally prohibitively expensive — this is already a problem for parameter estimation for agent-based models. Here there is room for the development of emulation methods that replace expensive simulation by statistical models.

One of the challenges of multi-scale and hierarchical models is to describe the interfaces

between levels or modules correctly; here thermodynamics imposes tight constrains [43] that are generally not incorporated in mathematical biology. Bond graphs have been conceived with that in mind [44]; they are also an area to which Edmund Crampin made major contributions [45, 46, 47, 48, 49, 50, 51].

In the bond graph formalism, energy and energy flow are tracked explicitly to ensure that reactions are commensurate with the laws of thermodynamics. This has two advantages: (i) resulting models incorporate stoichiometry, conservation of energy, and molecule numbers (mass) from the outset; (ii) thermodynamic constraints lead to models that are automatically modular, like real biological systems. Therefore bond graphs offer a rigorous and well defined framework to develop biophysically and thermodynamically consistent models of molecular processes. Here we will build on this work and extend it to develop whole cell models.

The critical point in model development is to link up sub-models of, for example, gene regulation, metabolism, and the cell cycle. This modularity turns out to be a natural consequence when thermodynamic conservation laws are applied to the component models. Furthermore, the bond graph formalism provides a straightforward pathway for modelling processes at different levels of complexity, according to available knowledge and data. It can be combined with both multi-physics and multiscale approaches to aid model construction. Here multi-physics refers to models that can capture different physical processes — such as heat transfer, dynamics of biochemical reactions over space and time, and transport phenomena inside the cells as well as between the cell and the extracellular environment — in a consistent manner. Bond graphs allow for this, and therefore provide a natural framework to consider situations where cell mechanics, electrophysiology, and metabolic processes need to be considered jointly.

3.7 Stochastic Nonlinear Dynamics

The theories of stochastic processes and dynamical systems have developed largely independently and in parallel. Many investigations in mathematical biology have also considered deterministic dynamics with scant reference to stochasticity. Classical work by Waddington, for example, does not consider noise and random dynamics, with the exception of fluctuating environments. Many important processes in ecology, population biology, epidemiology, cell and developmental biology, and physiology, have been studied with great success using purely deterministic descriptions. Stability analysis and bifurcation theory can give us qualitative, even global qualitative, insights into the complex dynamics underlying ecological dynamics or cell cycle regulation.

Often the average of a stochastic process can be described adequately using deterministic models. Diffusion and reaction-diffusion systems are good examples for which the underlying stochastic dynamics are sufficiently simple and deterministic partial differential equations can model the system sufficiently: solutions to the diffusion equation have been used, for example, to describe ecological migration, tumour growth, and biological pattern

formation.

In population and evolutionary genetics, by contrast, stochasticity — random sampling of alleles/individuals — has generally provided the accepted null model. The neutral model of evolution, for example, assumes that most genetic variation has no effect on fitness, the expected number of offspring. Historically, the development of statistics, probability theory, and population genetics are tightly linked. But in ecology and cell biology, for example, stochastic analyses have sometimes been conducted as *ad hoc* additions to deterministic analyses.

In many biological problems chance plays an important role, as do nonlinearity and feedback. The analysis of these systems is fraught with challenges. For illustration let us consider stochastic differential equations. Because random processes are not differentiable, we write this in the form,

$$dX = f(X;\theta)dt + g(X;\eta)dW_t, \tag{9}$$

where the first term, $f(X;\theta)$, captures the deterministic dynamics and the second term, $g(X;\eta)$, the stochastic contributions to the dynamics; dW_t denotes the Wiener process increment; and θ and η are the parameters of deterministic and noisy dynamics, respectively. For most biological systems, especially biochemical reaction systems, $f(X;\theta)$ and $g(X;\eta)$ are not independent but related. This relationship can be derived from the microscopic dynamics; frequently, however, $g(X;\eta)$ is chosen for convenience alone, and then typically in the form of white noise, therefore $g(X;\eta) = \sigma^2 = $ constant; this type of noise, which is independent of the state variable X is often referred to as "additive noise", in contrast to noise that depends explicitly on the state, which we refer to as "multiplicative noise".

The extent to which qualitative features of the deterministic dynamics affect the behaviour of a model once noise enters the equations requires careful consideration and depends a lot on the type and form of the noise function, $g(X;\eta)$. For example, the fact that the deterministic system is bi-stable in no way guarantees that the corresponding stochastic system has a bi-modal distribution over state space. This lack of conservation of qualitative features of dynamical systems — including number and nature of stationary states, existence of bifurcations, etc, — as noise is introduced, or more precisely our inability to predict *a priori* if qualitative features are conserved upon the introduction of noise, is becoming a limiting factor in many analyses. In the 1970, for example, René Thom and colleagues introduced catastrophe theory[52], which describes the types of sudden change (catastrophes) that can be encountered in dynamical systems. This found applications in population and developmental biology, but fell quickly out of favour [53]. The past two years have experienced renewed interest in catastrophe theory[54, 55, 56] and stochastic dynamics, if accounted for correctly, can change the qualitative behaviour [57].

In short, we do not understand anywhere near enough how to model, analyse, and understand complex dynamical systems that are subject to stochastic effects. Conceptually there are many problems that need to be considered: for example, we would really like to understand how the effects of noise are mitigated or attenuated in order to give rise to highly

stable and predictive phenotypes. Computationally, there are also a host of important and worthwhile problems to consider. Simulating stochastic dynamics is challenging and costly. Powerful approximations to exact simulation methods exist, and continue to proliferate [58], but ultimately we lack reliable approximations that continue to work when dynamics become nonlinear.

3.8   Natural Selection

Evolution provides the organising principle for the life sciences. Dobzhansky's statement "Nothing in biology makes sense except in the light of evolution" has become a rallying cry for biologists. The more recent version by Michael Lynch [59], "Nothing in evolution makes sense except in light of population genetics", makes this more precise as genetic change and selection on new variants occurs at the population level. Where, when and how selection operates has been at the core of evolutionary analyses and discussions [60]. Here, as in the previous section, there is a tension between deterministic and stochastic drivers, which need to be considered jointly to understand. The deterministic dynamics reflect the effects of natural selection, and the stochastic dynamics are due to the inherent randomness in the demographic and genealogical dynamics [61].

Probabilistic arguments have been dominant in the population genetics literature [61, 62]: random sampling of gametes appears to be remarkably effective at describing many real population genetic diversity patterns. Even the neutral model, which assumes that most genetic variants are not subject to natural selection (or, in its modified form, that most genetic variation is neutral or mildly deleterious), has underpinned the bulk of population genetic analysis. Neutrality makes the problem of studying evolutionary processes and data much easier. Most importantly, neutrality is a powerful null model against which we can test our hypotheses on experimental data.

In population dynamics, deterministic approaches have been used to study, for example, pathogen evolution and early pre-biotic evolutionary processes. One argument that has been put forward is that for sufficiently large population sizes the random evolutionary (demographic sampling) effects become less important. The effects of finite population size have also been largely ignored in game theoretic approaches to evolution.

The discussion of the previous section is, of course, relevant here as well. The tension between, and combined effects of, deterministic and stochastic dynamics are clearly important in understanding evolutionary dynamics. Simple models are useful, but not always enough [62]; as evolutionary arguments continue to increase in importance in very applied and data-rich areas, such as oncology and host-pathogen evolution, there is a need for both nuance and increased precision. In cancer biology the evolutionary characteristics and the disease aetiology have been shown to be tightly linked [63, 64, 65, 66, 67]. These studies have shown how complicated evolutionary processes are, and there are some salient lessons for evolutionary analysis more generally. Other evidence for the diverse molecular factors affecting evolution come from long term evolution experiments [68, 69, 70, 71, 72].

12

The upshot of the wealth of data, and the complicated interplay between deterministic and stochastic effects in driving evolution, is that we need more sophisticated maps from genotype to phenotype, especially if the view of omnigenic inheritance persists [73, 74, 75]. Having better understanding of how genes, their expression, and their interactions affect phenotypes may help in resolving long standing questions in both fundamental biology, but also in understanding the causes and mechanisms of complex diseases. Developing mathematical models for whole cells, and linking these in multi-scale models of organisms, is likely to form an important step in that direction.

3.9 Hybrid and Data Driven Modelling

Uncertainty about the correct model structure has been a common theme to many of the instances touched upon above. We often lack the knowledge to construct meaningful mathematical models; model selection can help to rule out some models. Automated model development and multi-scale models are opening up new opportunities to increase our universe of mathematical models for biological systems. Hybrid models [76, 77], see Figure 1, offer another route towards building better models. These models contain mech-
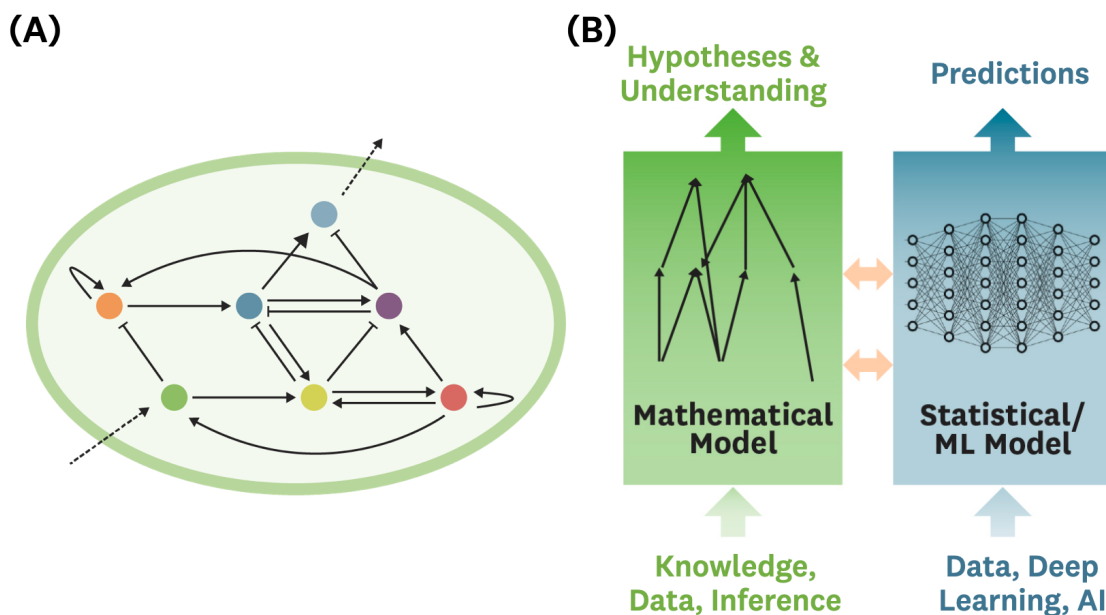

Figure 1: (A) We are used to modelling biological systems as, for example, systems of coupled ordinary differential equations. (B) Hybrid models combine statistical models with mechanistic models, and allow investigation of complex systems where only part of the mechanistic relationship is sufficiently well developed to be modelled mechanistcially.

anistic elements that interface with purely data driven, statistical, or black-box models, to combine the benefits of both approaches. Purely data driven statistical models can provide a baseline against which we compare — using for example information-theoretic measures such as the Kullback–Leibler divergence [78] — mechanistic and hybrid models. Further, we can develop data driven (or deductive) modelling approaches as potential emulators to feed back into, for example, the statistical inference approaches where computationally costly simulation is replaced by statistical ("black box") models [79].

Flexibility in choosing between mechanistic and data-driven modelling approaches, and the ability to combine them can open up powerful new ways to study, explain, and predict the behaviour of biological systems. Statistical "black-box" and machine learning models, by themselves, provide little direct insight into the underlying causal mechanisms that link genetic components to organismic behaviour. Analysing the resulting models with "white box" approaches including saliency maps or input-masked gradients provide important mechanistic, testable, and predictive insights into the critical regulatory regions and their interactions. Knowing the statistical dependencies should ultimately serve as a guide towards generating causal hypotheses for mechanistic/mathematical/functional relationships. We can use hybrid models to test the robustness of learned design principles: the mechanistic model embeds the inferred design principle, and the statistical model the physiological or environmental context (akin to, for example, extrinsic noise in gene expression).

The concept of neural differential equations [80, 81], where the right-hand side of the differential equation, $dx/dt = f(x; \theta)$ is replaced by a neural network,

$$\frac{dx}{dt} = \text{NN}, \tag{10}$$

holds the potential to develop new forms of hybrid models, all from within the same dynamical systems perspective.

Both mechanistic modelling and data-driven/machine learning modelling are well developed in their own right. Hybrid modelling, by contrast, is still in its infancy and the particular combination of hypothesis-driven and data-driven modelling required to develop useful models is only beginning to be studied. There are, however, exciting opportunities to bring methods from, for example, control engineering [82, 83, 84] to bear on hybrid modelling attempts. The use of a hybrid-Ansatz in modelling may also resolve some of the outstanding issues in parameter and model identifiability [13, 85], where the non-identifiable parts could, or should, perhaps be modelled using a machine learning approach.

4 A New Mathematical Biology

In this section we change pace and consider some broad mathematical and theoretical ideas for new approaches to mathematical biology. While this discussion is largely from

the perspective of pure mathematics, it provides a reflective analysis of the possibilities for the role of mathematics in biology.

4.1   Overview

Mathematical biology has predominantly avoided the explicit consideration of life. This paradoxical situation reflects a fundamental deficiency not only of mathematical biology but also, though perhaps to a lesser extent, of theoretical and experimental biology. Of primary concern here is the absence of a general theoretical definition of life and consequent realisation within an appropriate mathematical framework. While progress has been made, including the *relational biology* by theoretical physicist Nicolas Rashevsky [86] and theoretical biologist Robert Rosen [87], the *hypercycle* by biophysical chemist Manfred Eigen [88], and the *chemoton* by theoretical biologist Tibor Gánti [89], such a key concern is inexplicably the province of only a select few.

Definitions serve to identify and provide the foundation for explanation and knowledge. That a general definition of life is wanting is attendant to the complexity of the emergent phenomenon of life, though arguably more so to the methodological reductionism that has dominated scientific practice, particularly within biology and related fields, in recent decades. Reductionism has been preeminent in experimental biology, and consequently biologists collect enormous quantities of experimental data often without an integrated perspective of the living organism, thus without progression towards knowledge of the phenomenon of life [90].

While the ultimate ambition of biology is a general theory of life, this remains an open problem: even general concepts and principles are absent [91]. Mathematical biology, along with theoretical biology, can assume a fundamental role in developing a theory of life. Much of current mathematical biology, however, is subject to the same shortcomings as experimental biology. To make progress mathematical biology may need to consider three main aspects, namely emergent behaviour, mathematical maturity, and unification, which together form a basis for a more functional and effective approach to the study of living organisms with mathematics.

4.2   More is different

In 1972, Philip W. Anderson[1] published his article "More is different" [9], which became a highly influential exposition on the limitations of methodological reductionism and the importance of emergence in understanding complex phenomena. Anderson relates the hierarchical progression from simplicity to complexity with the progression from reduction to emergence [92, Chapter II, page 110]. New levels of system complexity are then associated with the emergence of novel phenomena, concepts, principles, and functions that

---

[1]In 1977 P.W. Anderson shared the Nobel Prize in Physics with Edmund's maternal grand father, Sir Nevill F. Mott, and J.H. Van Vleck.

are consistent with, though not readily explained by, the lower levels of the system. While emergent properties are often associated with the hierarchical, or multiscale, structure of a system [92, Chapter III, page 139], there are alternative perspectives that attempt to avoid any potential circularity arising from this association [93].

Reductionism, where the intention is to explain phenomena via upward causation in relation to underlying processes, has been very effective in science. The limitations of reductionism, however, are readily apparent for both simple and complex systems. For example, copper metal has the properties of being red, shiny, malleable, ductile, and conductive, while individual copper atoms exhibit none of these properties that emerge when many copper atoms form the metal structure [92, Chapter III, page 143]. Superconductivity provides an example of emergence in a more complex system: detailed computation of the motions of all electrons and ions in a superconducting material could describe the known phenomenology of superconductivity, however the computation cannot reveal the cause of superconductivity, namely the phenomenon of symmetry breaking. This is simply not captured by the framework at this lower physical level [92, Chapter III, page 136]. At much greater complexity, the behaviour of a living cell is an emergent property of the enormous number of molecules of which it consists [94]. Cells are regarded as the smallest functional unit of life, a property not associated with, and not readily predictable from, their constituent molecules.

Different types of emergence have been defined, including pragmatic, nominal, weak, and strong emergence. Strong emergence describes the appearance of an emergent phenomenon that cannot be described with a reductive account, and may be associated with downward causation and emergent forces or physical laws. Weaker forms of emergence are not necessarily associated with novel forces or physical laws. In particular, weak emergence describes a phenomenon that arises in an unexpected or very complicated manner at a higher level of complexity, and would not likely be evident without direct observation, even though a reductive understanding of the physics of the lower levels is in principle sufficient to describe the phenomenon [95]. A proposition for the existence of an emergent phenomenon must be accompanied by a rigorous explanation, via a detailed physical mechanism at the level of emergence, for the qualitative difference between the emergent phenomenon and the underlying levels. In Figure 2 we illustrate the complementary perspectives of methodological reductionism and weak emergence.

A fundamentally important reason for identifying weak emergence is that it may allow for the explanation of a phenomenon without the need to consider the reductive understanding of the lower levels, thereby allowing for a much simplified mathematical description or simulation. Identifying emergence thereby provides for potential flexibility and simplification in describing a complex system. For example, we can consider the motility of a living cell without knowing the movements of all constituent molecules.

While the concept of emergence has been incorporated into many areas of physics, the same cannot be said for biology and related fields, notably mathematical biology. Biologists generally believe that, given a sufficient understanding of chemistry, molecular biology, and
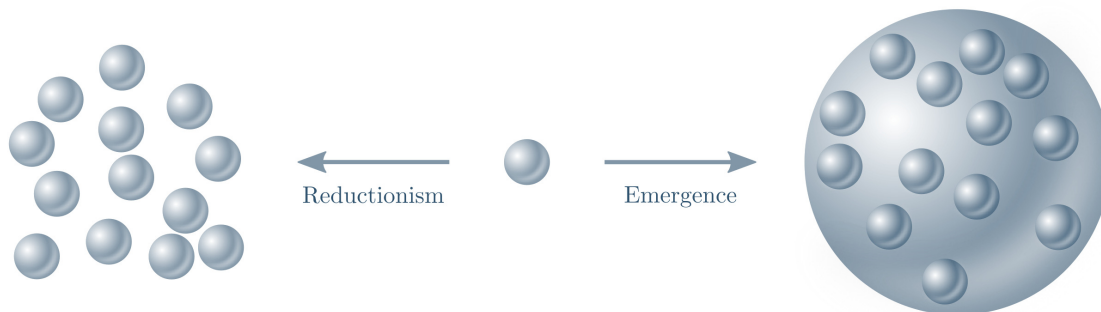
Figure 2: Illustration of methodological reductionism and weak emergence. Reductionism assumes that all complex behaviour can be described from knowledge of the underlying constituents. While weak emergence is consistent with the properties of the underlying constituents the emergent behaviour, which may correspond to novel principles of structure, organisation, or function, is not readily described from knowledge of the individual constituents.

cell biology, a complete explanation of living organisms will be achieved through reductive investigation. This is particularly evident in genetics and related areas where the tendency is to believe that understanding gene or protein sequences will lead to a complete explanation of life. To the contrary, the biologist Lynn J. Rothschild illustrates the numerous examples of form and function in biological systems for which the concept of emergence may, if not should, be employed [96, Chapter 6, page 151]. Moreover, the evolutionary biologist Ernst Mayr recognised the ubiquity of emergent behaviour in biological systems, and therefore included emergence as one of six reasons underlying the difficulties in predicting the behaviour of biological systems [97, Chapter 2, pages 58–59].

While reductionism has been very successful in yielding knowledge of many biochemical and cellular processes underlying biological systems, the limitations of reductionism are now apparent from the seemingly underestimated complexity of these systems [98]. By identifying emergent phenomena we can explain the properties of a biological system without the limitations of accounting for the complexity of the underlying constituents. Whether or not living organisms can be explained in principle from biochemistry, and thereby the laws of physics, is not of concern from the pragmatic perspective of utilising emergence to simplify and expedite the process of understanding biological systems.

Embracing emergence in biology does not dispel the need for much of current mathematical biology, but requires the additional consideration of identifying emergent phenomena at higher levels of system complexity. The latter is a very difficult program that requires new theoretical and mathematical ideas, some of which we have outlined above. The result, however, will be a much more sophisticated understanding and description of biological systems that yields mathematical models which avert prohibitively complicated computa-

tions. For this, rigorous mathematical definitions of emergence and emergent properties are required, and while there are some initial considerations [99, 93], further work is required.

To summarise with the words of Anderson on the "triumph of emergence over reductionism": "... large objects such as ourselves are, in myriad ways, the product of principles of organization and of collective behavior which in no meaningful sense can be reduced to the behavior of our elementary constituents. Large objects are often more constrained by those principles than by what the principles act upon." [92, Chapter 4, Pages 195–196].

4.3   Mathematical Maturity

The list of articles in any issue of the Journal of Mathematical Physics illustrates the mathematical maturity enjoyed by theoretical physics, notably the tendency for mathematical rigour including the use of the definition-theorem-proof style of pure mathematics [100], and the use of sophisticated mathematical theory. Currently, only a relatively small proportion of the research literature in mathematical biology exhibits such maturity.

Nicolas Rashevsky, one of the pioneers of mathematical biology, envisioned that the relationship between mathematical biology and experimental biology would correspond to that between mathematical physics and experimental physics [101]. That the vision of Rashevsky has been realised in such limited ways is probably due to several reasons, which we have already touched on above: biological systems are generally more complex than inanimate systems; avoidance by mathematicians who regard biology as unsophisticated and lacking rigorous definitions; strictly reductionist approaches, possibly influenced by experimental biology; insufficiency of deep theorising about living organisms, and in particular the formation of rigorous definitions; imprudent importation of mathematical models from other fields, such as physics and engineering; complacency with modelling simple biological systems with simple models and simple mathematics; indifference, or inability, toward applying more sophisticated mathematics or developing new mathematics.

The complexity of living organisms relative to physical systems suggests the need for the application of mathematical theory that is even more sophisticated, as well as the need for new mathematical ideas leading to new areas of mathematics [2, 102, 103]. For this, rigorous definitions of biological concepts are a necessity, and without such rigour the progression of mathematical biology, and how it supports experimental biology, is clearly limited. Mathematical biology requires new mathematical theory [103], and the passive importation of mathematical models from other fields, particularly physics and engineering, and their application in superficial ways to biological systems runs counter to the constructive progression of mathematical and theoretical biology. Capturing the intricacies of living systems, especially their emergent properties, will require but also inspire such new mathematics.

There is, we feel, increasing need of employing the definition-theorem-proof style of mathematics in mathematical biology. The general lack of rigorous definitions in biology is of fundamental concern for mathematical biology since all mathematical assertions origi-

nate from definitions, and the requirement to state clear definitions in mathematical biology would stimulate much needed theoretical thought regarding biological concepts. Further, a theorem asserts a general principle, and formulating a mathematical result as a theorem encourages the discovery of more general ideas rather than individual cases, demonstrates rigour, allows for immediate connection with other mathematical theory, simplifies the exposition, and may occasion a deeper understanding of the mathematics and the corresponding system. Rigorous mathematics, consisting of the definition-theorem-proof style is indubitably of importance for the success of mathematical biology [104].

4.4   Unification

Given the lack of general theoretical frameworks in biology [105] it is of no surprise that biological knowledge is fragmented, even to the extent that different subfields routinely use distinct terms for the same concept [106, 107]. In principle, mathematical and theoretical biology should be placed to remedy this situation, however they are arguably in a similar position which may only serve to further this fragmentation. The fundamental need for unification, or integration, within mathematical biology has long been realised, however little progress has been made. Here we discuss three particular areas associated with mathematical biology for which unification is imperative: biological data, mathematical models, and system perspective.

In his 2002 Nobel lecture, molecular biologist Sydney Brenner said that "We are all conscious today that we are drowning in a sea of data and starving for knowledge" [108]. Rather than mere descriptions and routine analysis of biological data, deep theorising and integration of the often heterogeneous data is required to yield knowledge [109]. Mathematical modelling can be of much assistance here, particularly regarding the assessment of consistency between a hypothesis and the experimental data, or for the identification of design principles. Mathematical models are fundamental for developing knowledge of an observed biological phenomenon, such as the identification of design principles. Given a causative hypothesis for a phenomenon we can construct a mathematical model that represents the hypothesis explicitly, and then compare model predictions with experimental data to determine whether the hypothesis is consistent, or not, with the evidence. In opposition to much of current practice in mathematical biology, however, integration of heterogeneous data requires the consideration of multiple insightful hypotheses, and therefore multiple models, or the consideration of a single hypothesis for consistency with a broad range of experimental data. Consideration of a single hypothesis/model within the context of a single data set, or multiple data sets that are closely related, achieves little toward data integration and novel biological theory. The objective is the development of new theoretical knowledge: being able to simulate a phenomenon does not in itself correspond to an explanation of the phenomenon.

The unification of mathematical models is at the frontier of mathematical biology, and general methodologies for comparing the underlying conceptual structure of models,

irrespective of mathematical formalism, provide essential facilitation for the process [38, 39]. Model comparison can partition a set of models into categories where similar models, regarded as sharing an essential qualitative characteristic, belong to the same category. Similar models may simply be incremental variations of each other, or may originally appear distinct and unrelated. Discovering similarities between seemingly distinct models can lead to novel theoretical insights into biological phenomena [38], and ultimately to a unified mathematical description of a biological system. A perspective of unification can encourage the avoidance of excessive incrementalism in model development and, along with a rigorous definition-theorem-proof style, yield general mathematical frameworks for biological systems, rather than the essentially algorithmic approach of developing simple models of simple systems.

It is widely recognised that to understand a biological phenomenon it must be perceived as an integrated whole [86]. For example, a biological function has meaning for an organism at the level at which the function is integrated, and below this level the function does not exist [110, 111, 112]. Ecology has, of course, long recognised this. Comprehending the integration is particularly difficult since at a given level of organisation in a biological system the biological constituents are generally heterogeneous [2], and naturally span several scales. Most of current mathematical biology is either concerned with "life-less" systems such as biochemical networks, or considers living organisms (such as a cell) as inanimate objects. The latter is a pragmatic and justified choice when the aim is to develop models that replicate, rather than explain, the behaviour of living systems, for example cell migration in response to gradients. But it falls way short of describing essential hallmarks of life. Here there appears to be genuine scope for a new mathematical biology that captures such hallmarks, including life as an emergent phenomenon.

5   Conclusion

Open problems in research can be interpreted as gaping holes in our understanding, or as exciting opportunities for future research. Most of the problems discussed here have aspects of both. There are considerable technical and computational challenges, but also some conceptual challenges of a mathematical nature, related to model development and model calibration. There is a definite appeal of using and combining methods from dynamical systems theory, computer science, statistical inference and machine learning to develop better, more reproducible methods of model development. Multiscale modelling, hybrid models, model selection, and automated generation of models individually, but especially when taken together, offer enormous potential to expand the scope of mathematical biology. We believe that the ability to explain, predict and control living systems requires a broadening of the toolbox of mathematical biologists.

In our view the appeal of combining hypothesis-driven and data-driven modelling frameworks is considerable. The technical challenges in doing so are also considerable. Data-

driven modelling has seen a rapid rise over the past decade and deep neural networks have become powerful tools in biology [113, 114, 115]; and their predictive performance is perhaps more immediately persuasive than are the mechanistic insights that modelling can provide. But both are needed [109].

We believe that a new mathematical biology will be required to understand life at the fundamental level we discussed in the previous section. Since this requires the substantial development of new biological and mathematical theory, we have little concrete idea of what this will look like. Anderson's work [9], maybe coupled to thermodynamic considerations [43], provides a starting perspective: emergent phenomena, including life itself, are not readily explained from principles that are detected by reductionist approaches.

Acknowledgements

References

[1] May. Uses and abuses of mathematics in biology. *Science*, 303:790, 2004.

[2] Cohen. Mathematics is biology's next microscope, only better; biology is mathematics' next physics, only better. *PLoS Biol*, 2(12):e439, 2004.

[3] Mariano Bizzarri, Douglas E Brash, James Briscoe, Verônica A Grieneisen, Claudio D Stern, and Michael Levin. A call for a better understanding of causation in cell biology. *Nat Rev Mol Cell Biol*, 20(5):261–262, 2019.

[4] Joseph Cursons, Jerry Gao, Daniel G Hurley, Cristin G Print, P Rod Dunbar, Marc D Jacobs, and Edmund J Crampin. Regulation of erk-mapk signaling in human epidermis. *BMC Syst Biol*, 9:41, Jul 2015.

[5] S Filippi, Chris P Barnes, P D W Kirk, Takamasa Kudo, Katsuyuki Kunida, Siobhan S McMahon, Takaho Tsuchiya, Takumi Wada, Shinya Kuroda, and Michael P H Stumpf. Robustness of MEK-ERK Dynamics and Origins of Cell-to-Cell Variability in MAPK Signaling. *Cell reports*, 15(11):2524–2535, 2016.

[6] Kip S. Thorne and Roger D. Blandford. *Modern classical physics: optics, fluids, plasmas, elasticity, relativity, and statistical physics*. Princeton University Press, 2017.

[7] Hermann Weyl. *Symmetry*. Princeton University Press, Princeton, 1952.

[8] Herbert Goldstein, Charles P Poole, and John L Safko. *Classical mechanics*. Addison Wesley, San Francisco, 3rd ed edition, 2002.

[9] P. W. Anderson. More is different. *Science*, 177:393–396, 1972.

[10] Ben D MacArthur. The geometry of cell fate. *Cell Syst*, 13(1):1–3, 01 2022.

[11] Paul D W Kirk, Ann C Babtie, and Michael P H Stumpf. Systems biology (un)certainties. *Science*, 350(6259):386, 2015.

[12] Alejandro F Villaverde, Fabian Fröhlich, Daniel Weindl, Jan Hasenauer, and Julio R Banga. Benchmarking optimization methods for parameter estimation in large kinetic models. *Bioinformatics*, 35(5):830–838, 03 2019.

[13] Alejandro F Villaverde, Antonio Barreiro, and Antonis Papachristodoulou. Structural identifiability of dynamic systems biology models. *PLoS Comput Biol*, 12(10):e1005153, Oct 2016.

[14] Attila Gábor and Julio R Banga. Robust and efficient parameter estimation in dynamic models of biological systems. *BMC Syst Biol*, 9:74, Oct 2015.

[15] Juliane Liepe, P D W Kirk, S Filippi, T Toni, Chris P Barnes, and Michael P H Stumpf. A framework for parameter estimation and model selection from experimental data in systems biology using approximate Bayesian computation. *Nature Prot*, 9(2):439–456, 2014.

[16] Toni, Welch, Strelkowa, Ipsen, and Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J Roy Soc Interface*, 6:187, 2009.

[17] M Komorowski, Maria J Costa, David A Rand, and Michael P H Stumpf. Sensitivity, robustness, and identifiability in stochastic chemical kinetics models. *Proc Natl Acad Sci*, 108(21):8645–8650, 2011.

[18] Maria Secrier, Tina Toni, and Michael P H Stumpf. The abc of reverse engineering biological signalling systems. *Mol Biosyst*, 5(12):1925–35, Dec 2009.

[19] Bryan C Daniels, Yan-Jiun Chen, James P Sethna, Ryan N Gutenkunst, and Christopher R Myers. Sloppiness, robustness, and evolvability in systems biology. *Curr Opin Biotechnol*, 19(4):389–95, Aug 2008.

[20] Juliane Liepe, Sarah Filippi, Michał Komorowski, and Michael P H Stumpf. Maximizing the information content of experiments in systems biology. *PLoS Comput Biol*, 9(1):e1002888, 2013.

[21] Ann C Babtie and Michael P.H. Stumpf. How to deal with parameters for whole-cell modelling. *J Roy Soc Interface*, 14(133):20170237, 2017.

[22] Kee-Myoung Nam, Benjamin M Gyori, Silviana V Amethyst, Daniel J Bates, and Jeremy Gunawardena. Robustness and parameter geography in post-translational modification systems. *PLoS Comput Biol*, 16(5):e1007573, 05 2020.

[23] Paul D W Kirk, Thomas W Thorne, and Michael P H Stumpf. Model selection in systems and synthetic biology. *Current opinion in biotechnology*, 24(4):767, 2013.

[24] Chris P Barnes, Daniel Silk, X Sheng, Michael P H Stumpf, and Michael P H Stumpf. Bayesian design of synthetic biological systems. *Proc Natl Acad Sci USA*, 108(37):15190–15195, 2011.

[25] Ivo Siekmann, James Sneyd, and Edmund J Crampin. Mcmc can detect nonidentifiable models. *Biophys J*, 103(11):2275–86, Dec 2012.

[26] Tina Toni and Michael P H Stumpf. Simulation-based model selection for dynamical systems in systems and population biology. *Bioinformatics*, 26:104, 2010.

[27] Chris P Barnes, Sarah Filippi, Michael P H Stumpf, and Thomas W Thorne. Considerate approaches to constructing summary statistics for abc model selection. *Stat and Comp*, 22:1181–1197, 2012.

[28] Dayanne M Castro, Nicholas R de Veaux, Emily R Miraldi, and Richard Bonneau. Multi-study inference of regulatory networks for more accurate models of gene regulation. *PLoS Comput Biol*, 15(1):e1006591, 01 2019.

[29] Michael P H Stumpf. Multi-model and network inference based on ensemble estimates: avoiding the madness of crowds. *J R Soc Interface*, 17(171):20200419, 10 2020.

[30] Ann C Babtie, Paul D W Kirk, and Michael P H Stumpf. Topological sensitivity analysis for systems biology. *Proc Natl Acad Sci*, 111(52):18507, 2014.

[31] Miriam Leon, Mae L Woods, Alex J H Fedorec, and Chris P Barnes. A computational method for the investigation of multistable systems and its application to genetic switches. *BMC Syst Biol*, 10(1):130, 12 2016.

[32] Lingxia Qiao, Wei Zhao, Chao Tang, Qing Nie, and Lei Zhang. Network topologies that can achieve dual function of adaptation and noise attenuation. *Cell Syst*, 9(3):271–285.e7, 09 2019.

[33] Wenzhe Ma, Ala Trusina, Hana El-Samad, Wendell A Lim, and Chao Tang. Defining network topologies that can achieve biochemical adaptation. *Cell*, 138(4):760–73, Aug 2009.

[34] Robyn P Araujo and Lance A Liotta. The topological requirements for robust perfect adaptation in networks of any size. *Nat Commun*, 9(1):1757, 2018.

[35] Natalie S Scholes, David Schnoerr, Mark Isalan, and Michael P H Stumpf. A comprehensive network atlas reveals that turing patterns are common but not robust. *Cell Syst*, 9(3):243–257.e4, 09 2019.

[36] Thomas Leyshon, Elisa Tonello, David Schnoerr, Heike Siebert, and Michael P H Stumpf. The design principles of discrete turing patterning systems. *J Theor Biol*, 531:110901, 12 2021.

[37] Onn Brandman, James E Ferrell, Jr, Rong Li, and Tobias Meyer. Interlinked fast and slow positive feedback loops drive reliable cell decisions. *Science*, 310(5747):496–8, Oct 2005.

[38] Sean T. Vittadello and Michael P. H. Stumpf. Model comparison via simplicial complexes and persistent homology. *Royal Society Open Science*, 8:211361, 2021.

[39] Sean T. Vittadello and Michael P. H. Stumpf. A group theoretic approach to model comparison with simplicial representations. *arXiv:2111.02170v1*, 2021.

[40] Michael P H Stumpf. Statistical and computational challenges for whole cell modelling. *Current Opinion in Systems Biology*, 26:58–63, 2021.

[41] Nicolas Le Novère, Andrew Finney, Michael Hucka, Upinder S Bhalla, Fabien Campagne, Julio Collado-Vides, Edmund J Crampin, Matt Halstead, Edda Klipp, Pedro Mendes, Poul Nielsen, Herbert Sauro, Bruce Shapiro, Jacky L Snoep, Hugh D Spence, and Barry L Wanner. Minimum information requested in the annotation of biochemical models (miriam). *Nat Biotechnol*, 23(12):1509–15, Dec 2005.

[42] Dagmar Waltemath, Richard Adams, Daniel A Beard, Frank T Bergmann, Upinder S Bhalla, Randall Britten, Vijayalakshmi Chelliah, Michael T Cooling, Jonathan Cooper, Edmund J Crampin, Alan Garny, Stefan Hoops, Michael Hucka, Peter Hunter, Edda Klipp, Camille Laibe, Andrew K Miller, Ion Moraru, David Nickerson, Poul Nielsen, Macha Nikolski, Sven Sahle, Herbert M Sauro, Henning Schmidt, Jacky L Snoep, Dominic Tolle, Olaf Wolkenhauer, and Nicolas Le Novère. Minimum information about a simulation experiment (miase). *PLoS Comput Biol*, 7(4):e1001122, Apr 2011.

[43] G F Oster, A S Perelson, and A Katchalsky. Network thermodynamics: dynamic modelling of biophysical systems. *Q Rev Biophys*, 6(1):1–134, Feb 1973.

[44] Peter J Gawthrop and Edmund J Crampin. Modular bond-graph modelling and analysis of biomolecular systems. *IET Syst Biol*, 10(5):187–201, Oct 2016.

[45] Peter J Gawthrop, Michael Pan, and Edmund J Crampin. Modular dynamic biomolecular modelling with bond graphs: the unification of stoichiometry, thermodynamics, kinetics and data. *J R Soc Interface*, 18(181):20210478, 08 2021.

[46] Peter J Gawthrop, Peter Cudmore, and Edmund J Crampin. Physically-plausible modelling of biomolecular systems: A simplified, energy-based model of the mitochondrial electron transport chain. *J Theor Biol*, 493:110223, 05 2020.

[47] Peter Cudmore, Michael Pan, Peter J Gawthrop, and Edmund J Crampin. Analysing and simulating energy-based models in biology using bondgraphtools. *Eur Phys J E Soft Matter*, 44(12):148, Dec 2021.

[48] Michael Pan, Peter J Gawthrop, Joseph Cursons, and Edmund J Crampin. Modular assembly of dynamic models in systems biology. *PLoS Comput Biol*, 17(10):e1009513, 10 2021.

[49] Niloofar Shahidi, Michael Pan, Soroush Safaei, Kenneth Tran, Edmund J Crampin, and David P Nickerson. Hierarchical semantic composition of biosimulation models using bond graphs. *PLoS Comput Biol*, 17(5):e1008859, 05 2021.

[50] Peter J Gawthrop and Edmund J Crampin. Energy-based analysis of biomolecular pathways. *Proc Math Phys Eng Sci*, 473(2202):20160825, Jun 2017.

[51] Michael Pan, Peter J Gawthrop, Kenneth Tran, Joseph Cursons, and Edmund J Crampin. Bond graph modelling of the cardiac action potential: implications for drift and non-unique steady states. *Proc Math Phys Eng Sci*, 474(2214):20180106, Jun 2018.

[52] Rene Thom. *Structural stability and morphogenesis.* Addison-Wesley, 1989.

[53] V. I Arnold. *Catastrophe theory.* Springer-Verlag, Berlin, 3rd rev. and expanded ed edition, 1992.

[54] Elena Camacho-Aguilar, Aryeh Warmflash, and David A Rand. Quantifying cell transitions in c. elegans with data-fitted landscape models. *PLoS Comput Biol*, 17(6):e1009034, 06 2021.

[55] Rand, Raju, Sáez, Corson, and Siggia. Geometry of gene regulatory dynamics. *Proc Natl Acad Sci*, 118(38):e2109729118, 2021.

[56] Meritxell Sáez, Robert Blassberg, Elena Camacho-Aguilar, Eric D Siggia, David A Rand, and James Briscoe. Statistically derived geometrical landscapes capture principles of decision-making dynamics during cell fate transitions. *Cell Syst*, 13(1):12–28.e3, 01 2022.

[57] Megan A Coomer, Lucy Ham, and Michael P H Stumpf. Noise distorts the epigenetic landscape and shapes cell-fate decisions. *Cell Syst*, 13(1):83–102.e6, 01 2022.

[58] David Schnoerr, G Sanguinetti, and Ramon Grima. Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. *Journal of Physics A*, 50(9):093001, 2017.

[59] Michael Lynch. *The origins of genome architecture.* Sinauer Associates, Sunderland, Mass., 2007.

[60] George C Williams. *Natural selection: domains, levels, and challenges.* Oxford University Press, New York, 1992.

[61] W. J Ewens. *Mathematical population genetics*, volume v. 27. Springer, New York, 2nd ed edition, 2004.

[62] J Wakeley. Recent trends in population genetics: more data! more math! simple models? *J Hered*, 95(5):397–405, 2004.

[63] Steven A. Frank. *Dynamics of cancer: incidence, inheritance, and evolution*. Princeton University Press, Princeton, N.J., 2007.

[64] Arne Traulsen, Tom Lenaerts, Jorge M Pacheco, and David Dingli. On the dynamics of neutral mutations in a mathematical model for a homogeneous stem cell population. *J R Soc Interface*, 10(79):20120810, Feb 2013.

[65] Adam L MacLean, Sarah Filippi, and Michael P H Stumpf. The ecology in the hematopoietic stem cell niche determines the clinical outcome in chronic myeloid leukemia. *Proc Natl Acad Sci U S A*, 111(10):3883–8, Mar 2014.

[66] Marc J Williams, Benjamin Werner, Chris P Barnes, Trevor A Graham, and Andrea Sottoriva. Identification of neutral tumor evolution across cancer types. *Nat Genet*, 48(3):238–244, Mar 2016.

[67] Eszter Lakatos, Marc J Williams, Ryan O Schenck, William C H Cross, Jacob Househam, Luis Zapata, Benjamin Werner, Chandler Gatenbee, Mark Robertson-Tessi, Chris P Barnes, Alexander R A Anderson, Andrea Sottoriva, and Trevor A Graham. Evolutionary dynamics of neoantigens in growing tumors. *Nat Genet*, 52(10):1057–1066, 10 2020.

[68] Zachary D Blount, Richard E Lenski, and Jonathan B Losos. Contingency and determinism in evolution: Replaying life's tape. *Science*, 362(6415), 11 2018.

[69] Kyle J Card, Thomas LaBar, Jasper B Gomez, and Richard E Lenski. Historical contingency in the evolution of antibiotic resistance after decades of relaxed selection. *PLoS Biol*, 17(10):e3000397, 10 2019.

[70] Michael Doebeli, Yaroslav Ispolatov, and Burt Simon. Towards a mechanistic foundation of evolutionary theory. *Elife*, 6, 02 2017.

[71] Richard E Lenski. What is adaptation by natural selection? perspectives of an experimental microbiologist. *PLoS Genet*, 13(4):e1006668, 04 2017.

[72] Dustin J Marshall, Martino Malerba, Thomas Lines, Aysha L Sezmis, Chowdhury M Hasan, Richard E Lenski, and Michael J McDonald. Long-term experimental evolution decouples size and production costs in escherichia coli. *Proc Natl Acad Sci U S A*, 119(21):e2200713119, 05 2022.

[73] Evan A Boyle, Yang I Li, and Jonathan K Pritchard. An expanded view of complex traits: From polygenic to omnigenic. *Cell*, 169(7):1177–1186, Jun 2017.

[74] Xuanyao Liu, Yang I Li, and Jonathan K Pritchard. Trans effects on gene expression can drive omnigenic inheritance. *Cell*, 177(4):1022–1034.e6, 05 2019.

[75] Iain Mathieson. The omnigenic model and polygenic prediction of complex traits. *Am J Hum Genet*, 108(9):1558–1563, 09 2021.

[76] Ruth E Baker, Jose-Maria Peña, Jayaratnam Jayamohan, and Antoine Jérusalem. Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biol Lett*, 14(5), 05 2018.

[77] Bo Yuan, Ciyue Shen, Augustin Luna, Anil Korkut, Debora S Marks, John Ingraham, and Chris Sander. Cellbox: Interpretable machine learning for perturbation biology with application to the design of cancer combination therapy. *Cell Syst*, 12(2):128–140.e4, 02 2021.

[78] Siobhan S Mc Mahon, Aaron Sim, Sarah Filippi, Robert Johnson, Juliane Liepe, Dominic Smith, and Michael P H Stumpf. Information theory and signal transduction systems: from molecular information processing to network inference. *Semin Cell Dev Biol*, 35:98–108, Nov 2014.

[79] Yongxiang Feng, Zhen Cheng, Huichao Chai, Weihua He, Liang Huang, and Wenhui Wang. Neural network-enhanced real-time impedance flow cytometry for single-cell intrinsic characterization. *Lab Chip*, 22(2):240–249, 01 2022.

[80] Ward Quaghebeur, Elena Torfs, Bernard De Baets, and Ingmar Nopens. Hybrid differential equations: Integrating mechanistic and data-driven techniques for modelling of water systems. *Water Res*, 213:118166, Feb 2022.

[81] Elisabeth Roesch, Christopher Rackauckas, and Michael P H Stumpf. Collocation based training of neural ordinary differential equations. *Stat Appl Genet Mol Biol*, 20(2):37–49, 07 2021.

[82] Kirubhakaran Krishnanathan, Sean R Anderson, Stephen A Billings, and Visakan Kadirkamanathan. A data-driven framework for identifying nonlinear dynamic models of genetic parts. *ACS Synth Biol*, 1(8):375–84, Aug 2012.

[83] Gerardo Aquino, Luke Tweedy, Doris Heinrich, and Robert G Endres. Memory improves precision of cell sensing in fluctuating environments. *Sci Rep*, 4:5688, Jul 2014.

[84] Eszter Lakatos and Michael P H Stumpf. Control mechanisms for stochastic biochemical systems via computation of reachable sets. *Roy Soc Open Sci*, page 079723, October 2016.

[85] Heather A Harrington, Kenneth L Ho, Thomas W Thorne, Michael P H Stumpf, and Michael P H Stumpf. Parameter-free model discrimination criterion based on steady-state coplanarity. *Proc Natl Acad Sci USA*, 109(39):15746–15751, 2012.

[86] N. Rashevsky. Topology and life: In search of general mathematical principles in biology and sociology. *The Bulletin of Mathematical Biophysics*, 16:317–348, 1954.

[87] Robert Rosen. A relational theory of biological systems. *The Bulletin of Mathematical Biophysics*, 20:245–260, 1958.

[88] Manfred Eigen. Selforganization of matter and the evolution of biological macro-molecules. *Die Naturwissenschaften*, 58:465–523, 1971.

[89] Tibor Gánti. *The Principles of Life*. Oxford University Press, Oxford, UK, 2003.

[90] Athel Cornish-Bowden. How far has it come?: Systems biology. *The Biochemist*, 33:16–18, 2011.

[91] Carol E. Cleland. Is a general theory of life possible? seeking the nature of life in the context of a single example. *Biological Theory*, 7:368–379, 2013.

[92] Philip W. Anderson. *More and Different: notes from a thoughtful curmudgeon*. World Scientific, New Jersey, 2011.

[93] Alex J. Ryan. Emergence is coupled to scope, not level. *Complexity*, 13:67–77, 2007.

[94] Irun R. Cohen and David Harel. Explaining a complex living system: dynamics, multi-scaling and emergence. *Journal of The Royal Society Interface*, 4:175–182, 2007.

[95] Mark A. Bedau. Is weak emergence just in the mind? *Minds and Machines*, 18:443–459, 2008.

[96] Philip Clayton and Paul Davies, editors. *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion*. Oxford University Press, New York, 2006.

[97] Ernst Mayr. *The Growth of Biological Thought: Diversity, Evolution, and Inheritance*. Belknap Press, Cambridge, MA, 1982.

[98] Marc H. V. Van Regenmortel. Reductionism and complexity in molecular biology. *EMBO Reports*, 5:1016–1020, 2004.

[99] Yaneer Bar-Yam. A mathematical theory of strong emergence using multiscale variety. *Complexity*, 9:15–24, 2004.

[100] William P. Thurston. On proof and progress in mathematics. *Bulletin of the American Mathematical Society*, 30:161–177, 1994.

[101] Paul Cull. The mathematical biophysics of Nicolas Rashevsky. *Biosystems*, 88:178–184, 2007.

[102] Michael C. Reed. Mathematical biology is good for mathematics. *Notices of the American Mathematical Society*, 62:1172–1176, 2015.

[103] Alexandre Borovik. A mathematician's view of the unreasonable ineffectiveness of mathematics in biology. *Biosystems*, 205:104410, 2021.

[104] Jean Clairambault. Can theorems help treat cancer? *Journal of Mathematical Biology*, 66:1555–1558, 2013.

[105] Samuel M. Scheiner. Toward a conceptual framework for biology. *The Quarterly Review of Biology*, 85:293–318, 2010.

[106] Mikhail V. Blagosklonny and Arthur B. Pardee. Conceptual biology: Unearthing the gems. *Nature*, 416:373, 2002.

[107] Julie C. Barnes. Conceptual biology: a semantic issue and more. *Nature*, 417:587–588, 2002.

[108] Sydney Brenner. Nature's gift to science — Nobel lecture, december 8. https://www.nobelprize.org/prizes/medicine/2002/brenner/lecture/, 2002.

[109] Nurse. Biology must generate ideas as well as data. *Nature*, 597:305, 2021.

[110] Denis Noble. Differential and integral views of genetics in computational systems biology. *Interface Focus*, 1:7–15, 2011.

[111] Denis Noble. A theory of biological relativity: no privileged level of causation. *Interface Focus*, 2:55–64, 2012.

[112] Indika Rajapakse and Stephen Smale. Emergence of function from coordinated cells in a tissue. *Proceedings of the National Academy of Sciences*, 114:1462–1467, 2017.

[113] Nicolae Sapoval, Amirali Aghazadeh, Michael G Nute, Dinler A Antunes, Advait Balaji, Richard Baraniuk, C J Barberan, Ruth Dannenfelser, Chen Dun, Mohammadamin Edrisi, R A Leo Elworth, Bryce Kille, Anastasios Kyrillidis, Luay Nakhleh, Cameron R Wolfe, Zhi Yan, Vicky Yao, and Todd J Treangen. Current progress and open challenges for applying deep learning across the biosciences. *Nat Commun*, 13(1):1728, 04 2022.

[114] Mohammed AlQuraishi and Peter K Sorger. Differentiable biology: using deep learning for biophysics-based and data-driven modeling of molecular mechanisms. *Nat Methods*, 18(10):1169–1180, 10 2021.

[115] Thomas Naert, Özgün Çiçek, Paulina Ogar, Max Bürgi, Nikko-Ideen Shaidani, Michael M Kaminski, Yuxiao Xu, Kelli Grand, Marko Vujanovic, Daniel Prata, Friedhelm Hildebrandt, Thomas Brox, Olaf Ronneberger, Fabian F Voigt, Fritjof Helmchen, Johannes Loffing, Marko E Horb, Helen Rankin Willsey, and Soeren S Lienkamp. Deep learning is widely applicable to phenotyping embryonic development and disease. *Development*, 148(21), 11 2021.