

Reverse-engineering biological networks from large data sets

Joseph L. Natale

Department of Physics, Emory University, Atlanta, GA 30322, USA

David Hofmann and Damián G. Hernández

*Department of Physics and Initiative in Theory and Modeling of Living Systems,
Emory University, Atlanta, GA 30322, USA*

Ilya Nemenman

*Department of Physics, Department of Biology, and Initiative in Theory and
Modeling of Living Systems, Emory University, Atlanta, GA 30322, USA*

(Dated: May 24, 2017)

Much of contemporary systems biology owes its success to the abstraction of a *network*, the idea that diverse kinds of molecular, cellular, and organismal species and interactions can be modeled as relational nodes and edges in a graph of dependencies. Since the advent of high-throughput data acquisition technologies in fields such as genomics, metabolomics, and neuroscience, the automated inference and reconstruction of such interaction networks directly from large sets of activation data, commonly known as reverse-engineering, has become a routine procedure. Whereas early attempts at network reverse-engineering focused predominantly on producing maps of system architectures with minimal predictive modeling, reconstructions now play instrumental roles in answering questions about the statistics and dynamics of the underlying systems they represent. Many of these predictions have clinical relevance, suggesting novel paradigms for drug discovery and disease treatment. While other reviews focus predominantly on the details and effectiveness of individual network inference algorithms, here we examine the emerging field as a whole. We first summarize several key application areas in which inferred networks have made successful predictions. We then outline the two major classes of reverse-engineering methodologies, emphasizing that the type of prediction that one aims to make dictates the algorithms one should employ. We conclude by discussing whether recent breakthroughs justify the computational costs of large-scale reverse-engineering sufficiently to admit it as a mainstay in the quantitative analysis of living systems.

I. LAY OF THE LAND

Biological systems on all levels of organization, from cells to brains and to populations, are comprised of ensembles of interactions among smaller constitutive components [1–3]. These interactions are typically very specific, and highly coordinated spatially and temporally [4–8]. Involving not just pairs, but also larger groups of components acting in concert [9–14], they are responsible for the rich diversity of complex phenomena and behaviors that make living systems work. Although often prohibitively numerous to model individually (though see [15]), these components and their corresponding interactions can be represented formally as graphs [16], known colloquially as *biological networks* [3, 17–23].

The variables in such networks (also called nodes) typically represent biochemical or ecological species, cells, or even amino acid residues when one is interested in the biophysics of proteins. The links among the nodes represent interactions, such as chemical transformations, catalysis, and binding; cooperative or predator-prey relations among species; electrical and chemical communication among cells; or geometric proximity among amino acid residues (Fig. 1).

To answer many questions in modern data-rich biology, an intermediate step often involves the reconstruction of such networks from empirical data. The data typically

consist of joint samples of activities (often referred to as expressions, frequencies, abundances, or population sizes, depending on the context) of a large number of components measured in different biological contexts. Problems of this kind pervade the quantitative life sciences on all physical scales, even if they take different forms and use different languages across scientific disciplines.

At the smallest scale is the problem of inference of physical contacts for amino acids in a protein fold [25–27], which is a network representation of the 3D structure of the protein. Predicting such networks from the co-occurrence of amino acids promises the ability to design proteins with specific functional properties. At the cellular level, different genes activate or suppress the activities of other genes, forming networks of genetic regulatory interactions [28, 29]. Similarly, metabolites transform into each other, catalyzed by various enzymes; these form metabolic networks [22, 30, 31], as well as networks that combine both protein and metabolic modalities. Protein signaling networks characterize the structure of decision-making and information processing in individual cells [32–36]. The accurate reconstruction of different types of these cellular networks is expected to lead to successful interventions that cure some of the most debilitating diseases [37].

On the scale of the nervous system, one often reverse-engineers neural circuits [38–42] and, on a larger scale, functional connectivity networks between brain re-

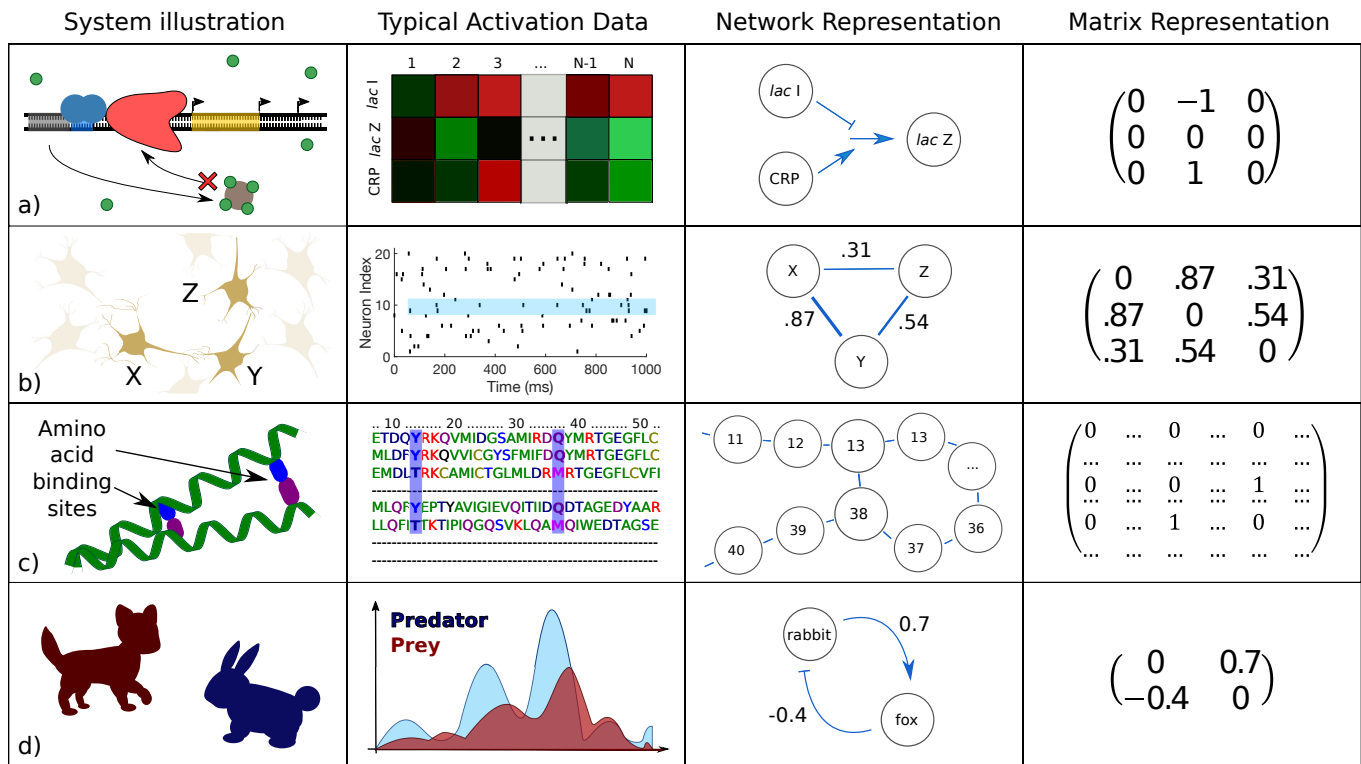


FIG. 1. Examples of biological systems whose constituent interactions have been modeled using networks. (a) The regulation of gene transcription by transcription factors and other enzymes. For example, in the classic *lac* repressor circuit [24], when lactose concentration is high and glucose concentration is low, genes for metabolizing lactose are strongly activated. Activities in this case might consist of microarray data for each mRNA species; the network shows the logic (AND) of the system. (b) Neuronal co-activation networks can be measured by computing correlations between spike patterns. In this case, the network graphs are weighted, and weights may represent correlations. (c) Spatially proximal amino acids tend to co-evolve, as they often participate in bonds that are vital to the structure and function of the protein they form. Here activity values are discrete, assuming one of 20 values to identify the amino acid at each site; the network represents bonds that are inferred to exist by noting which site pairs are highly correlated across similar proteins in different organisms. (d) Complex predator-prey interaction dynamics can be cast in a network form as well. Here activation data represent the populations of each species, and connections are labeled with inferred values for the parameters in the governing population dynamics equations.

gions [43–47]. The structure of the latter has been shown to be valuable as a diagnostic tool for some psychiatric diseases [48], and there is mounting evidence that the former can be “reprogrammed” via external interventions to repair damaged circuits [49]. Finally, on the largest scales, one can reconstruct networks of interactions among members of a particular species [50–52], or different species in an ecosystem [53–58]. This knowledge may help in forecasting ecological catastrophes [59, 60] and addressing the spread of infectious disease [61] (or other epidemics [62]).

In all of these fields, data share similar properties, and data sets often have similar sizes. This imposes uniformity not only on the question of network inference itself, but also on the obstacles and algorithmic approaches that underlie reconstruction efforts across multiple biological domains. Inference methods designed for one system type ([63], [64], and [65]) can often be adapted to accommodate others ([66–68], [32, 69, 70], and [39, 71], respectively). Moreover, morally equivalent methods have

been developed in nominally unrelated fields [72] – or else borrowed explicitly from established disciplines, such as systems identification techniques migrating to network biology from engineering [73, 74]).

An additional reason for the cross-pollination among the subfields of biological networks inference is that, like in other parts of bioinformatics, the field has benefited from advances in machine learning and related Big Data computational tools. In their turn, as is often true of mathematical approaches, these tools are applicable across multiple traditional biological subdisciplines, and hence provide for natural theoretical bridges not only among life-sciences subfields, but also to a “network” of other quantitative disciplines (physics, statistics, and computer science) [75].

However, one cannot embrace the unembraceable. Thus in this review, we will focus almost exclusively on applications of networks inference to the systems biology of the cell [76–78], and will mention bridges to other fields only briefly and haphazardly, leaving the reader

certainly thirsty for more. Starting with a few of the references that we mention, as well as using Google Scholar (another network, this time of citations), is an easy way to quench this thirst!

Before proceeding any further, it is certainly worth warning the reader that the explosive growth of the field of biological network inference has covered with a thick blanket of journal articles some treacherous rocks. A few of them are very dangerous, and can, in principle, sink the field if not addressed thoughtfully. Specifically, while fully automated network inference has become a routine procedure, it is not immediately clear that the large-scale reconstruction of entire networks from high-throughput data will necessarily result in tangible insights or actionable understanding about biological systems. One reason is that most reconstructions are not experimentally verified, remaining in the literature as collections of information (or misinformation) of dubious quality. Another comes from the fact that it is still not clear what new knowledge entire-network inference yields, besides proposing potential interactions for experimental verification. If a goal of the field is to predict response of biological systems to yet-unseen exogenous perturbations, then the bridge between a network graph and such predictive knowledge will have to be built eventually, but it is not there yet in most practical applications. Most importantly, it is usually unclear what insights are delivered by large-scale networks, or how to interpret the typical product of the reconstruction enterprise – Lander’s infamous “hairball” of decontextualized interactions [79]. One can even argue that exhaustively enumerating interactions is not inherently more insightful than cataloging the original experimental data, and both should give way to studying the system’s emergent properties [80]. Having now warned the reader, we leave these important, foundational questions aside for the remainder of the review (save the *Discussion*).

A. Scale of the biological network inference problem

Network reverse-engineering is typically done in the “low-hanging fruit,” Big Data regime. Here the data sets are large, but the number of unknowns is even larger: not all the unknowns can be learned reliably.

While reconstructions can be performed using different data types [76, 81], here we are concerned with approaches that are based exclusively on biological *activity measurements*. Suppose we have a network consisting of p nodes (e. g., a group of p interacting genes or neurons), and n simultaneous measurements of some activity variable for each of these nodes (which for our purposes fully characterizes the biological states of the nodes at a given moment in time). The activity variables can be binary (as in the characterization of whether a gene is on or off, or whether a neuron is spiking or not at a given time) or

real-valued (gene expression levels, or firing rates for neurons). In other words, the total amount of available data is $\sim np$. The goal is to identify links between pairs of the p nodes (or more generally, higher order interaction structures) from patterns in their activities. If we focus on pairwise interactions among the nodes only, then the number of unknowns is $\sim p^2$. Thus the amount of data per unknown is $\alpha \sim np/p^2 = n/p$.

In the classical statistics regime, the amount of data is typically asymptotically large compared to the number of unknowns, $\alpha \gg 1$. In contrast, network inference usually proceeds in the regime where $p \gg 1$, with typical $p \sim 10^2 \dots 10^3$. For gene expression and other high throughput cellular data, in particular, it is not uncommon to have $p \sim 10^4$. Other fields are catching up [82, 83]. The number of measurements is also typically large, $n \gg 1$. We can consider $n < p$, as in most genetic data, or $n > p$ (but not $n \gg p$), as in many neuroscience applications. More generally, $n \sim p$, so that $\alpha \sim 1$, representing a qualitative departure from the classical statistics regime.

The situation gets even worse when we remember that the total number of parameters characterizing all (higher-order) interactions in a network scales as the total number of states that the network can be in (i. e., 2^p for binary nodes, or 2^{p^S} for continuous ones, where S is the entropy of each node measured at the experimental resolution). Thus in the most general case, for biological network inference, $\alpha \ll 1$. It is clear then that, just like in most other Big Data applications, the problem cannot be solved completely, with all interactions identified. Thus networks inference necessarily is a “low-hanging fruit” problem, where the limited data allows us to focus only on the most salient features of the studied systems. This also means that, in any quantitative assessment of the quality of network reconstruction methods, we should focus a lot more on the precision (absence of false positives) of a method, rather than on its sensitivity (absence of false negatives), since the sensitivity of essentially any method on realistic data would be tiny.

B. Different ideologies for inference

In biological network inference, one can think of reconstructing actual physical interactions among the nodes or coarse-grained, phenomenological surrogates. We focus exclusively on the latter.

The notion of network inference may evoke the idea of reconstructing actual physical interactions among network nodes. For example, a regulatory interaction between two genes might mean the direct binding of a transcription factor protein, translated from one of these genes, to a specific part of the DNA sequence that controls the expression of the other gene [67]. We refer to the reconstruction of such physical, microscopically accurate interactions as the inference of *mechanistic* net-

works. In contrast, the majority of reconstruction methods focus (explicitly or not) on the inference of *effective* interaction networks [84], which keep track of purely phenomenological interactions. These may or may not be mechanistically accurate, but are sufficient to reproduce various statistics of the observed variables. Such effective interactions may correspond to subsets of the interactions in mechanistic networks. They may be compact, coarse-grained averages of some microscopic quantities. Or they may be entirely macroscopic properties that have remote and complicated relationships with the microscopic, mechanistic interactions.

One can focus on effective network inference for purely pragmatic reasons: as discussed above, even high-throughput data is insufficient to infer *all* the contributing actors in a complex system, and effective interactions may simply be the low-hanging, accessible fruit. In contrast to this pessimistic view, one may argue that every level of description requires its own proper degrees of freedom for efficient representation [80, 85, 86], and that the distinction between mechanistic and effective networks is not that clear-cut.

To wit, even mechanistic biophysical interactions are themselves effective interactions, just at a different scale. For example, the bonds between amino acids that form at protein-protein interfaces consist of electrostatic forces between constituent molecules. These forces can then be broken down in terms of quantum interactions between elementary particles, at which point the notion of an amino acid has long since disappeared. Likewise, the fact that communication between synapsing neurons requires the diffusion of neurotransmitters across the synapse undermines the notion that neurons can ever truly be in a direct, mechanistic contact. We are most sympathetic with this viewpoint, which treats the distinction between mechanistic and effective networks less as a dichotomy than as a spectrum. In what follows, we cast the issue in terms of modeling assumptions: what is the appropriate set of nodes and interactions to answer *the specific questions being asked* while working at the *desired scale*?

Our perspective notwithstanding, a few authors have distinguished explicitly between these two ideologies (see [87] as the originator of the “physical” vs. “influence” network terminology, and [46] for a more fine-grained distinction among different types of effective networks in the brain). Many other sources refrain from making such explicit distinctions, presumably either for expedience in exposition or because they take seriously the aforementioned notion of pursuing the most efficient or useful description at a given level of study, regardless of the biological implementation details at other levels. While we remain agnostic to the particular reasons for the tendency of reverse-engineering literature to avoid making this distinction at the outset, we lament the absence of explicit declarations of the intended level of description when elaborating a new algorithm by the majority of publications. By default, in this Chapter, we focus on effective inference methods, for which authors

do *not* make an effort to understand whether there is a mechanistic basis for inferred interactions, stating any exceptions at the outset when they appear.

C. Goals of this Chapter

We are now in a position to state our intended goals for this Chapter. In the following sections, we review relatively recent (within the last two decades) attempts at network inference, contending:

1. The aptness and success of a given inference method depend on the ultimate purpose of performing network reconstruction. One must first establish what kinds of predictions are desired (i. e., what does one seek to *learn* [19] using the network?), and only then decide which algorithm to use.
2. Large-scale network reverse-engineering has many fruitful applications, but it is not always the necessary – or not necessarily the best – approach for making certain kinds of predictions.

Note that we deal exclusively with inference methods that produce networks containing at most pairwise interactions. While the joint probability distribution for p discrete or continuous stochastic activation variables in a stationary state $\{g_i\}$ can be expanded [65] most generally as

$$P(\{g_i\}) \propto \exp \left[- \sum_i^p h_i(g_i) - \sum_{i,j}^p J_{ij}(g_i, g_j) \dots - \sum_{i,j,k}^p \phi_{ijk}(g_i, g_j, g_k) - \dots \right], \quad (1)$$

where functions h_i , J_{ij} , and ϕ_{ijk} denote first-, second-, and third-order interactions, respectively, it is clear from the considerations of Section I A that reliable estimation for terms of higher order than J_{ij} is prohibitively difficult. In addition, we review only the algorithms that attempt to infer *static* values for J_{ij} under the assumption that the system is in (near-)stationary conditions, although some authors have attempted to estimate networks whose topologies are dynamically evolving [88, 89].

The progression of the Chapter is as follows. First, we examine highlights of the many places where network inference has been used to advance new knowledge in contemporary systems biology and establish novel paradigms in modern medicine. Then we proceed to delineate and explicate several types of inference methods, briefly describing the operation of several representative algorithms for each of the classes we name. We conclude with a brief outlook of where the field might be headed. However, these concluding comments should be

taken with a lot of caution, since “it is difficult to make predictions, especially about the future.”

II. ROLES FOR REVERSE-ENGINEERING IN SYSTEMS BIOLOGY RESEARCH

The reverse-engineering of large-scale networks by means of automated algorithms has become such a routine procedure that it has spawned a research field of its own. Why is the task of learning networks from data considered so important?

The modern imperative to generate comprehensive parts lists for large biological systems [22] is epitomized in what one author somewhat flippantly calls “the giant maps of metabolic pathways that many molecular biologists pin to their walls” [90]. Such diagrams encode and illustrate visually the entirety of observable interactions of a particular type in a specific system. Since the mid-2000s, attempts to generate such maps have been pursued vigorously by researchers in various disciplines, but the most prominent and systematic efforts have come from the network inference Challenges of the Dialogue on Reverse-Engineering Assessment and Methods (DREAM) initiative [91, 92]. Contestants participating in these ongoing Challenges submit network reconstructions, inferred by original algorithms operating on standardized data, for comparison against (experimentally) established sets of interactions in benchmark networks.

The top-scoring networks in early competitions achieved respectable accuracies, despite the difficulties associated with defining “gold standard” benchmarks and evaluation metrics [91, 93]. However, they also lacked the ability to provide intuition (beyond structural insights) about the systems they described. As static pictures of interaction architectures, they had limited ability to *predict* a system’s behavior. The pattern of assembling a large, intricate network as the end goal, with no intention to use it as a tool for prediction – as in the iconic but largely uninformative hairball of Ref. [79] – thematized DREAM competitions roughly until 2014, nearly a decade after one reviewer declared the field to be “still in [its] ‘natural history’ phase” [2].

The emphasis of DREAM competitions has since shifted, mirroring changes in the attitude of the reverse-engineering community as a whole. Recent competitions have more strongly favored *predictive modeling*, with inferred networks serving not as ends in themselves, but as coarse summaries of high-dimensional data – a special type of statistic – to aid in projecting how the behavior or components of a system will change (as a function of time, due to changes in its environment, etc.).

This movement away from using learned topologies as a signal that the “work is done,” and instead toward viewing the entire process of network inference as an intermediate step in a fully-fledged research pipeline [94], is also

supported by theoretical work. In particular, it has been argued that structure alone provides insufficient information to achieve an adequate degree of control over the underlying system’s dynamics [95]. In fact, the object of interest is not always a network’s structural complexity (density of connections), but its *dynamical* complexity (the number of fixed points it can accommodate), which depends on other parameters beyond structure, such as its connection strengths. Indeed, only the latter is closely tied to the viability of a network architecture in the context of evolution [96].

The field’s transition – from descriptive to predictive – is a natural one, and indeed reminiscent of the progression in other branches of science. While it is not completely clear why there was this prolonged period of exploration without modeling, it is plausible that reverse-engineers first needed to convince themselves that (1) networks can, indeed, be accurately reconstructed from activity data alone, and (2) the achieved reconstructions are statistically significant and reproducible. Furthermore, experimental tools for administering systematic perturbations to the networks under study took a while to develop, so that the need to predict dynamical responses to perturbations had not emerged for a while. As confidences in the statistical power of reverse-engineering grew, and new experimental tools were developed, the next level of questions naturally emerged. It is in answering this next level of questions that network reconstructions have found their broad spectrum of highly nontrivial, often unique, and even central roles in modern systems biology. For the remainder of this section, we survey several key application areas, focusing on the most impactful types of predictions that reconstructions are capable of generating.

A. Predictions regarding individual nodes or interactions

Reconstructions can help identify intervention targets or functionally similar cohorts of biological species.

The advent of modern, high-throughput data acquisition techniques transformed the enterprise of network reconstruction from a painstaking, often collaborative process into an exercise in algorithmic design. Verifying the existence of a single interaction no longer demands corroboration by multiple independent research efforts, and connections can now be inferred in parallel directly from a single set of data. An oft-cited consequence of this change of pace contends that modern reverse-engineering dramatically increases the rate at which hypotheses about potential interactions can be generated. To this end, whole-network reconstructions allow us to rapidly elucidate both the presence and nature of individual interactions, as well as predict the function of individual nodes from knowledge about their neighbors [97, 98].

Inference methods designed for the express purpose

of proposing novel interactions for experimental verification [65, 99] have confirmed previously established gene targets [74] and identified novel targets for known transcription factors and drugs [100, 101]. Known broadly as statistical or *association* methods (see “Who talks to whom,” Section III A), algorithms in this class have also discovered entirely new interactions [10, 11, 100, 102–105], with previously unknown gene interactions often being verified experimentally [106, 107]. In a multi-algorithm litmus test, several of these methods were capable of inferring missing links in artificially corrupted, incomplete versions of established pathways [108].

Network-based strategies for the prediction of protein function [109] generalize more traditional approaches, such as clustering analyses [97], that have been used to classify genes and proteins according to their role at either the physiological or the network level. Individual gene clusters correspond to distinct functional groups in some systems [110]. They can be used to infer roles for unclassified elements according to the guilt-by-association (GBA) heuristic (i. e., assigning functions similar to those of nearby neighbors in the interaction space).

Clustering alone cannot produce a full interaction map, and its applicability is limited because its underlying assumptions are not universal among biological system types [111, 112] (GBA may be more valid for protein-protein interactions than gene-gene interactions, since the latter entail more latent or intervening steps). Nevertheless, clustering is still useful in modern reverse-engineering, predominantly in the data-processing phase that often precedes the inference of full interaction architectures [113]. Clustering the data prior to inference greatly restricts the search space by providing an effective prior to bias the set of candidate interactions. On the other hand, the same idea can also be used to *coarsen* inferred networks: “module-based” inference techniques [114] have identified entire groups of genes that are functionally related [115]. We will return to this idea of identifying coarse functional and conceptual (as opposed to simply structural) units in the *Discussion*.

B. Insights from the statistical properties of network ensembles

Certain structural statistics differentiate real biological systems from other kinds of complex networks.

While the rapid verification of microscopic interactions undeniably constitutes an improvement in the pace of discovery, it does not by itself generate categorically new kinds of knowledge. Systems biology is “more than an accelerated program of molecular biology” [79], and the relatively new tools of reverse-engineering must prove their worth by helping to play a part in that grander enterprise. This is reflected in the possibility of using reconstructions to make predictions not only about single nodes and in-

dividual connections, but about the statistical properties of network *ensembles*.

Work in this direction has produced various insights about what distinguishes biological systems and endows them with their unique characteristics among complex networks. For instance, it has been shown that the most highly connected nodes in protein networks are likely to be essential [116] for survival [117, 118]. Moreover, nodes with an exceptionally high degree (i. e., number of connections), called *hubs*, attach preferentially to nodes with low degree while tending to avoid one another [119]. This property, in part, underlies the widely observed *modular* organization of cellular systems: an efficient coding scheme in which network partitions include only components involved in related processes. This discourages overlap and ensures that (on average) no single node participates in too many processes [35]. This forms the basis for one type of biological robustness [120].

Certain modular structures recur with disproportionately high frequencies in biological systems (with respect to their chance rate of appearance in a random graph [16]). Known as *motifs* [21, 121], they can endow the network with vital control and design features, such as positive or negative feedback, and are often conserved throughout evolution [122–124]. Studying the appearance rates of different motifs across different networks can help clarify the functional “purpose” they satisfy within a given network.

While a node’s degree is its most fundamental attribute, studying other network parameters has also led to key insights. The betweenness centrality [16] for nodes in protein interaction networks has been observed to be even more highly correlated with protein essentiality than their degree [125]. Moving beyond individual nodes, it has been argued that the full degree *distribution* is approximately scale-free [20] for many systems, providing deep architectural support for the robustness of biological systems to noise and perturbations, at both environmental and genetic levels [126] (yet see [127] for a cautionary note about the associated power-law distributions).

In *network medicine* [128], clinically relevant predictions can often be made from such high-level statistics, irrespective of whether interactions can be enumerated exhaustively or determined at a fine-grained level. For instance, the aforementioned correlation between a node’s degree and its essentiality for survival begets the notion that candidate drug targets can often be ruled out immediately if they are too highly connected, such that using them risks compromising the rest of the network [129].

While one should not focus exclusively on the architectural aspects of dynamically engaged networks [96], even microscopic statistics can sometimes go beyond structure to tell rich stories about the behavior of the underlying system. Maximum Entropy [130, 131] methods [72] (see Section III A) have been used to learn the effective coupling constants that connect neurons in the retina [38, 132], where the inferred values suggest that these networks naturally reside in the neighborhood of a

critical point in their parameter spaces [133]. This might afford such networks an essentially optimal capacity for stimulus representation, as well as information storage and transmission [134, 135] (though see [136] for an alternative viewpoint). For the amino acid interaction networks that keep track of where bonds form during protein folding, the same methods corroborate the idea that geometrically proximal residues tend to coevolve [27] by showing that bond locations can be identified using a simple statistic on the ensemble of viable protein sequences (in this case, correlations between the activations of site pairs).

C. Using statistics to characterize or classify individual networks

Ensemble statistics can help identify defective or emergent properties in a network.

Sometimes, statistical surmises can be used to make statements about the typicality of a particular network. An approach known as *differential networking* (so named to contrast with *differential expression*, a popular type of approach to activation data in gene networks) has been increasingly used for this purpose.

For example, Refs. [137, 138] discuss the idea of using topological characteristics to solve supervised classification problems, such as determining whether a given network comes from a healthy or a pathological organism. This possibility is explored explicitly in [48], which nominates several criteria (reduced clustering and “small-worldness,” reduced probability of high-degree hubs, and increased robustness) as those which are markedly altered in patients with schizophrenia. The reconstruction method developed in [139] was able to identify genes that are either known tumor drivers, associated with biological processes relevant to disease, or correlated with patient prognosis for various types of cancer by examining how pathological networks differ from their counterparts in “normal” tissue. Changes in hub structure have also been used to forecast the survival outcome for breast cancer patients [140].

It is worth pointing out that the aforementioned Maximum Entropy methods [72] provide, in some sense, a complementary approach to ensemble statistics. Rather than addressing only aspects that networks have in common (or can be averaged over), these approaches are predicated on exploiting the intrinsic *variability* at the micro-scale in an attempt to reproduce what is conserved at the macro-scale. This is especially useful wherever diverse microscopic network connectivity structures are known to produce indistinguishable behavior at coarser resolutions, as in protein folding: there is no one-to-one mapping between amino acid sequence and tertiary structure, but an entire distribution of microscopic parameters – a wide variety of equally viable amino acids sequences – that code for roughly the same protein shape [25, 141–

143]. Knowing this, one can easily imagine how running Maximum Entropy methods in reverse can help determine, for example, whether a given amino acid interaction network represents a viable protein. The same might be said for evaluating the typicality of an inferred retinal network, by measuring properties like criticality [144, 145] (NB: for a selection of competing viewpoints on the criticality of neuronal networks, consult the aforementioned [136], as well as [146–149]).

D. Predicting how a given network will respond to perturbations

Reconstructions help identify and quantify response patterns in novel conditions.

Network models capture and summarize complex dependencies among states of biological components, often allowing one to predict how a system will change its state or behavior with changes in the biological environment (i. e., modifications affecting the state of one or more nodes or interactions). Commonly studied perturbations can be local [150] (e. g., knockout of a single gene, as in the simulation of deleterious mutations), multifactorial (affecting many elements) [151], or fully global [152] (applying a drug to slightly suppress the firing of all neurons in a circuit), and the system’s responses can be investigated at local or global levels as well. For instance, one might inquire about the effect of a drug or a mutation on the expression of a single gene, or the success or failure of signal propagation from start to end through a perturbed pathway.

The types of responses that are interesting to researchers vary widely, and range across a spectrum of detail. The simplest and the coarsest entail qualitative predictions: for example, is the activation state of a given node affected by a specific perturbation? Progressing to a more quantitative picture, one can try to predict the actual post-perturbation values for affected nodes, as in the prediction of gene expression levels following a knockout event [153]. At the finest granularity, models incorporating time-series measurements can be used to forecast the transient behavior for such a gene as it approaches a new steady-state expression level.

Recent DREAM Challenges have provided a testing ground for algorithms aiming to make these types of predictions. The DREAM4 Predictive Signaling Network Modeling Challenge [154] instructed contestants to predict phosphoprotein measurements “using an interpretable, predictive network” [155], and the bonus round of that year’s *in-silico* Challenge [150, 156, 157] asked competitors explicitly to predict the system’s responses to “novel” perturbations that were not encountered in the training data. The DREAM7 Network Topology and Parameter Inference Challenge [108] specified the prediction of perturbation outcomes using gene regulatory network models as a separate step from inferring their

topologies.

As we discuss later, prediction of time-course trajectories requires directed networks, but the converse is not true: directional links can sometimes be inferred from static data. On the level of qualitative predictions, the linear dynamical systems approach of [74] was able to deduce the targets of novel perturbations in a system of nine genes using only steady-state values of their expression levels, following a series of highly controlled perturbations (and the knowledge of which genes were targeted during the perturbations). We consider this result to be particularly important, for two reasons. First, it challenged previously expressed (and still later-held [158]) ideas by successfully determining a directed network, despite the fact that the applied perturbations elicited statistically significant changes in the activations of all nodes. Second, later improvements extended the abilities of the algorithm therein to determine which species were “hit” by applied perturbations even *without* specifying as inputs which genes were targeted during the data acquisition phase [159], reinforcing the idea that M static, independent, but carefully selected perturbation measurements can substitute for a series of time-course measurements taken at M intervals [160].

E. Representing the joint probability distribution for observables

Networks models can be interpreted as shorthands for joint probability distributions.

Activation values for each node depend on those of many others, rendering graphical models particularly convenient representations of their joint activities. Graphs can explicitly encode the statistical dependencies among different activation variables as connection *weights*, with the states of connected nodes given not by a stochastic transfer function, but by conditional probabilities.

A type of directed acyclic graph (DAG) known as a Bayesian network is a weighted construction whose connection strengths are typically learned [161] via Bayesian inference (i. e., computing the posterior probabilities for a set of candidate DAGs, and selecting the member with the highest value, etc.) Undirected variants, which communicate only binary dependency information via the presence or absence of symmetric links are popular in different applications. When activities are assumed to deviate normally from baseline values (an assumption that greatly simplifies the inference process), they are known as Gaussian graphical models [162].

Bayesian networks weights can be scaled so as to represent a proper, normalized probability distribution. Adjusted to match that of the observed data, the weights in such a dependency graph become an explicit encoding of the system’s joint statistics. Bayesian networks satisfy a *Markov property*, such that the activity value distribution

for a given node depends only on the values of its immediate predecessors (these activities are often discretized as binary variables for mathematical convenience, so the resulting graph neatly keeps track of the probability that a downstream node in the inferred network will be active if its predecessors are active). This directed conditional dependency structural arrangement offers a conceptually accessible and intuitive view of the system, although the presence of directed connections between two nodes does not mean there is a direct physical (i. e., mechanistic) or causal link between the corresponding species [163].

One of the most important and unique applications of network inference, this compact representation of probability distributions permits the quantitative prediction for nodal activity values, in both static and dynamic contexts. Probabilistic graphical models are particularly useful in putting numbers on answers to questions like “What is the probability of this protein being active, given that a particular stimulant is present?” or its converse: “What is the probability of the stimulant having been present, given that the expression level of this gene is high?” [164]. We discuss methods for inferring both types of probabilistic graphical models named here, and their limitations (including their ability to infer causality), in Section III.

F. Reconstructions as a part of the Big Picture

Inferred network models can be combined with existing and new methods as one part of a larger repertoire for investigating many facets of living systems.

Reconstructions are increasingly combined with other tools and prior biological knowledge to form integrated frameworks for discovery. Some reverse-engineering approaches attempt to incorporate prior knowledge explicitly into the inference process for individual networks [165–170], including one study which advocates the use of undirected gene networks (gleaned from functional association databases) as *priors* to enhance the inference of mechanistic, causal gene regulatory networks [171].

Other applications use networks to cross-reference, corroborate, or pre-screen evidence for predictions about specific systems. For example, the “network approach” to genome-wide association studies (GWAS) and disease gene prioritization is reviewed in [97], and the use of networks for the prediction of protein functions (in the general sense, not restricted to physical binding), evolutionary studies of pathogenic and non-pathogenic strains, and the bidirectional interactions between host and pathogen are reviewed for the specific context of infectious disease in [98].

We have already mentioned the work [94], which uses Bayesian networks in tandem with support vector machines to predict the toxicity of various chemicals in a supervised setting. Yet we believe the most pivotal roles to be played by reconstructed networks are those which

completely change the way we think about biological phenomena, specifically by offering new ways to predict system-wide behaviors. Such a revolution is already underway in medicine: the treatment of various diseases is no longer unilaterally viewed from within the “one-gene, one-drug” paradigm, and it is gradually becoming the new standard to view related autoimmune disorders as emanating from a network of maladies with the same root causes [172–174].

III. TWO DIFFERENT MEANINGS OF PHENOMENOLOGICAL “RECONSTRUCTION”

We distinguish two principal categories for phenomenological network inference, accounting for methods that produce undirected and directed graphical models.

Algorithms in our **first category** define an inferred interaction as an *irreducible statistical dependency* among nodes, typically quantified by some measure of the similarity among the activation profiles of different nodes. This is a structure-only approach, and should be used when it is only necessary to reconstruct the overall network topology – in other words, for applications for which it is sufficient to know “who talks to whom.” In some cases, topological maps can be augmented with weights that ascribe an effective strength or confidence level to the inferred interactions [143, 175].

Algorithms in our **second category** define interactions in terms of *asymmetric relations* capable of describing not only which nodes participate in an interaction, but also “who controls whom.” Previous classification schemes have considered the inference of unweighted, directed links as a separate endeavor from discovering quantitative input-output relationships between nodal activities [176], or further distinguish algorithms that detect the sign of interactions without an explicit direction [177, 178]. However, since both the types of data and the processing techniques needed to infer all these kinds of graphs are similar, we treat them on equal footing.

A. Who talks to whom? *Presence, absence of undirected links*

The most basic question that one can answer in the course of network reconstruction is whether a given subset of nodes can be characterized as interacting – in other words, *who talks to whom?* Since our focus here is on the unsupervised inference of interaction networks directly from activation data, any notion of “interaction” that we consider must depend on these activations alone. A natural definition for the existence of an interaction among species is the presence of statistically significant correlations among their respective activation states. Such a choice results in an undirected network with symmetric (though possibly weighted) connections.

In practice, pairwise statistical dependencies are typically quantified by introducing a *similarity* metric, such as the first-order Pearson correlation. The Pearson correlation coefficient is a normalized, pairwise dependency measure bounded by the interval $[-1, 1]$. Positive (negative) values indicate an increasing (decreasing) linear relationship. While its value is always zero for statistically independent variables, a vanishing Pearson correlation cannot rule out nonlinear correlations. Conversely, in the absence of nonlinear effects, finite sampling can cause independent variables to appear correlated, so that connections can be inferred where no otherwise discernible interaction exists. To avoid inferring such spurious interactions, one must apply a threshold to filter raw correlation values.

When nonlinear effects cannot be ignored, one can quantify statistical dependencies using information-theoretic measures [179–181], which generalize the notion of correlation to such nonlinear cases. D’Haeseleer et al. [182] were the first to employ the mutual information to uncover gene-to-gene dependencies, while Butte et al. applied mutual information “relevance networks” [183] to propose single-gene determinants of anticancer agent susceptibility [184] for experimental verification. Mutual information-based methods must still contend with the same sampling and bias problems faced by linear correlation coefficients, and therefore require thresholding as well.

Even under conditions of perfect sampling, neither Pearson correlations nor the mutual information can disambiguate so-called *direct* interactions from *indirect* interactions – statistical dependencies that are already accounted for by links involving other species. Note that this notion of “indirect” is distinct from its usage in the context of mechanistic networks. There, “direct” typically refers to physical contact, which often occurs between nodes whose activations are not included in the network model (unobserved, latent, or marginalized degrees of freedom in the system). Here instead we are concerned with statistical redundancies within the set of *observed* activation variables. For example, consider the case of three genes in a regulatory cascade: $X \rightarrow Y \rightarrow Z$. Inference methods based on measuring correlations between the associated activation variables would find a link between X and Z , which is *indirect*, in the sense that it is not actually needed to account for the joint statistics of X , Y , and Z .

While sometimes inconvenient, indirect links are not always superfluous. They are useful when probing the network at the single-node level, as when trying to discover a previously unknown member in an established pathway, propose a novel interaction for experimental verification, or predict the overall effect on the activation state of one node by perturbing another. On the other hand, in applications for which inferred networks must be treated as whole entities (e. g., when they encode normalized probability distributions; see MaxEnt methods described below), this sort of redundancy can be mini-

mized by examining *conditional* dependency structures.

There exist several approaches to studying conditional dependencies. The most intuitive is to work explicitly with either partial correlation coefficients [99] or the conditional mutual information [10, 11, 185–187] between two activation variables X and Y , given another variable (or set of variables) Z :

$$I(X; Y|Z) = I(X; Y, Z) - I(X; Z), \quad (2)$$

where $I(X, Z)$ is the mutual information between X and Z . In principle, one can refine a reconstruction by removing links between any pair of species X and Y that are associated with statistically insignificant values of $I(X; Y|Z)$. However, reliable estimation of this quantity is much more difficult than it is for the pairwise quantities, such as $I(X, Y)$, since it requires sufficiently dense concurrent sampling of at least three variables.

In order to dispose of indirect links without incurring the aforementioned estimation problems, some algorithms make additional assumptions and thus append ancillary filtering steps to the basic mutual information-based procedure. For instance, the Algorithm for the Reconstruction of Accurate Cellular Networks (ARACNe) [65, 100] invokes the Data Processing Inequality [180] to delete the weakest link in every closed triplet of nodes (this would be an exact step if the studied network was a tree). The Context Likelihood of Relatedness (CLR) method [102] determines the presence or absence of a link by assessing its strength against all other mutual information scores computed for that graph, as a background significance threshold. MRNET [188] builds a network iteratively, including a link between two variables if one is both a good predictor of the other and yields information that is non-redundant with that provided by the previously inferred links.

An alternative approach to solving the conditional independence problem is to use full probabilistic models that allow conditioning on the complete set of marginals, rather than requiring the progressive computation of higher-order partial correlations [187]. In particular, if a set of continuous, real-valued activation variables are (assumed to be) normally distributed, one can condition a single interaction on the full set of remaining variables. In this case the statistical independence of any two nodes can be ascertained by examining the elements of the inverse of the covariance matrix: $\Sigma_{ij}^{-1} = 0$ if and only if i and j are conditionally independent, given all other variables. An important facet of such multivariate Gaussian distributions is that they correspond to the least constrained, *maximum-entropy* models that satisfy the full set of first and second-order marginals for continuous variables [72, 131]. These first two moments correspond to the individual means and the pairwise correlations, which are usually well measured even in sparsely sampled data sets.

Beyond Gaussian variables, the Maximum Entropy principle has been a successful modeling approach in neuroscience [38, 189–191], natural images [192], the in-

ference of gene networks (from expression data) [193] and signal transduction networks (from phosphorylation proteomics data) [194], and the prediction of amino acid contacts in proteins [141, 142, 195, 196], multidrug effects [197], protein structural attributes [26], antibody diversity [198], and even the dynamics of flocking birds [199]. The joint probability distribution for a maximum entropy model has a particular form, known in statistical mechanics as the Boltzmann distribution. If we ask to match only the empirical means $\langle x_i \rangle$ and pairwise correlations $\langle x_i x_j \rangle$ to those of the observed data, the distribution with maximal entropy is

$$P(\vec{x}) = \frac{1}{\mathcal{Z}} \exp \left(\sum_i h_i x_i + \sum_{ij} J_{ij} x_i x_j \right). \quad (3)$$

Here parameters h_i and J_{ij} are known as the fields and the couplings, respectively, and \mathcal{Z} is the partition function (compare to the full expansion in Section IC).

For discrete variables, the Maximum Entropy model retains the form of Eq. (3), but is known as the Ising model (for binary variables) or Potts model (for categorical variables with more than two accessible states). In the discrete case, fitting the parameters $\{h_i, J_{ij}\}$ is highly nontrivial. Many methods exist, but their effectiveness depends on the system size and the density of its interactions, as well as on other properties [200–204]. One algorithm worth mentioning is the *adaptive cluster expansion*, which was developed in the context of the MaxEnt problem [204, 205]. It is closely related to information-theoretic approaches, being equivalent to relevance networks [183] for clusters of size two, and similar to conditional mutual information methods for clusters of size three.

Due to the limitations of finite sampling, both solving for the inverse of the covariance matrix and learning the parameters of an Ising model can constitute ill-posed problems. One way to avoid this is to impose a regularization [206], which invokes additional constraints on the interaction coefficients to ensure that the inference problem is well-defined – and moreover, that the inferred network generalizes well to unseen data. Regularization is often done in one of two common ways: either the interactions coefficients are assumed to be small (for example, using an L_2 norm) [205] or the interaction structure of the system is presumed to be sparse, so that the overall number of the interactions is small (this may be done explicitly by specifying the number of non-zero coefficients [74] or by invoking an L_1 norm [203]).

Frequently cited as the rationale behind these regularization procedures is the inherent sparsity of natural networks [117, 207, 208]. Indeed, for protein studies, the nodes in networks used to describe tertiary protein structure represent real amino acids in the three-dimensional space; they can therefore be connected to only a small subset of all possible neighbors. Similarly, the number of transcription factors that can influence a given gene's

expression levels is limited by the physical extent and arrangement of its promoter sequence. While the general ubiquity of sparseness in biological systems is debated [69], the enforcement of sparsity constraints can be justified as a purely pragmatic measure in the “low-hanging fruit” inference regime.

B. Who controls whom? *Causal relations and directed links*

Directed network inference differs in an important way from that of undirected, symmetric, mutual-influence graphs: since questions of causality (or, more generally, the flow of information) are built not upon a single, universally agreeable concept like statistical correlation – but rather on more subtle, less straightforward notions like *control* – there exist many diverse criteria for establishing directed connections. Each method has its own operating definition of what counts as an interaction, and how to infer its direction.

Though disparate, the aforementioned definitions can be conveniently divided into approximately two subclasses, depending on the intended application of the inference procedure. In certain cases, it is enough to know the direction or causal *sense* of an inferred interaction. For example, will silencing a certain gene or disabling a particular neuron result in a collapse of the entire system? Can the intracellular concentration of a reactant be increased by introducing more of the product? Answers to questions like these do not require numbers, entailing purely qualitative predictions. On the other hand, if the goal is to use a reconstructed network to predict the amount by which one gene’s expression level increases when two other genes are suppressed, directed connections must be weighted by quantitative values representing the effective *strengths* of interactions. We describe methods of both types, leaving it as an exercise for the reader to think about when a directed topology suffices, and when it is necessary to infer fully signed and weighted graphs.

Before we delve into specific methods, we advise the reader to tread with caution. The particular definitions of directed influence we explore in the following methods do not always correspond to our intuitive and/or mathematically formal notions of causality. As a result, producing a graph with directed links does not automatically satisfy a reverse-engineer’s desire to uncover system-wide causes in an ontological sense, and should not be treated as such despite one’s instincts. Instead, great care needs to be taken with each method in order to ensure that all idiosyncratic constraints are met, and to avoid generalizing or extrapolating beyond the predictive power of each algorithm.

To expound on this point, it is worth asking at the outset whether it is even possible to infer causal information from passive observations of activation variables [209, 210]. It has long been understood [211] that

proximal causal relations can be inferred reliably when the observer is able to *interact* with the system in accordance with a principled protocol (as is done in many controlled experimental interventions [212, 213], including genetic knockouts [150, 153] and multifactorial perturbations [151, 214]). While this is old news to engineering audiences, it has also been shown that causal information (or at least a lower-bound estimate of causal effects) can be extracted from purely observational data when the equivalence class for the fully directed graph can be ascertained first [215, 216][217].

We mention again a surprising corollary of this result that directed influence (a less stringent condition, and slightly less nebulous concept, than causal influence) can often be established without time series data, using only static measurements. Where there was once a prevalent belief in the reverse-engineering community that the inference of directed edges required temporal data [65], there is now a tradition of algorithms which accept static data as inputs [69, 73, 74, 212, 218–220]. However, we focus for coherence predominantly on methods that operate on time-series data.

We organize this subsection as follows. We first make a few general remarks about the inference of directed interaction patterns. We then explore a class of methods which presume that the measured activities can be treated as deterministic variables that change smoothly in accordance with a particular, predetermined quantitative law. Afterwards we switch to model-free deterministic methods, for which there is no need to specify a mathematical form or law in advance in order to detect interactions. We then treat the more general situation, in which activations are regarded as stochastic variables. Again we start with methods requiring a parameterized model and conclude with a discussion of stochastic model-free methods.

A naïve but conceptually intuitive approach to inferring directed connections is to take the presence of strong temporal correlations between the trajectories of different activity variables as evidence for a (casual) interaction between the corresponding species. It is common for changes in one activity variable to succeed that of another in time (consider a gene whose expression level is observed to increase consistently in response to the elevation of another), but the proxy of temporal precedence is not robust as a criterion for declaring control relations [57] because it also appears in the absence of causal influence. Despite its limitations, this strategy, combined with a projection method known as multidimensional scaling [63] in an algorithm entitled “Correlation Metric Construction,” was originally used to infer the first steps of the glycolytic pathway [221] and more recently applied to study the pharmacokinetics of the anti-cancer drug Gemcitabine [222].

In physics and engineering, signed and directed connections are often used to encode the weighted coupling constants that appear in systems of differential equa-

tions [223]. To write down such a system, one needs to first have in mind a particular quantitative form for a dynamical law, according to which activations will be presumed to interact. One then fits the model parameters, typically with some optimization or statistical learning technique that takes time series data as input, and reports the learned values as the weights for the corresponding connections, sometimes adding additional, unobserved, hidden variables in the process [86].

The inherent directionality of this method, which works best for small systems ($p \sim 10$), can be understood immediately by examining the matrix J_{ij} of pairwise interactions in Eq. (4) below: since this matrix is not constrained to be symmetric, couplings between two species can differ in the forward and backward directions. For continuous activation variables $\{x_i\}$, many popular models can be subsumed as special cases of the general form (though see [86, 224] for alternative forms):

$$\frac{dx_i(t)}{dt} = f_i \left(x_i + \sum_j^p J_{ij} x_j + u_i + \xi_i \right), \quad (4)$$

which includes at most pairwise interactions of strengths $\{J_{ij}\}$ between all element pairs i and j . Here the functions $\{f_i\}$ can be chosen according to the desired level of computational complexity (controlled by the amount of data available) or biochemical detail, or both. In the reverse-engineering of biological networks, many early applications were linear activation models [225–228], for which $f_i(x) \propto x$. The sum determines the net (excitatory and inhibitory) effect on the activation of node i at time t , given its interactions with all other elements j . The next term accounts for external driving of the node, (i. e., any extrinsic perturbation that increases or decreases its activation value by an amount $u_i(t)$), and $\xi_i(t)$ represents noise.

Linear, “additive” regulatory models are based on the assumption that dynamical systems can be *linearized* about their steady-states. They are relatively easy to fit in sparsely sampled conditions, especially when the terms in Eq. (4) are discretized to form a linear difference equation [227, 229]. Early work countered under-sampling by augmenting the number of data points for multilinear regression via nonlinear interpolation [226], or imposing sparsity constraints on singular decomposition algorithms [73]. Another approach to decreasing the number of interactions that must be inferred is to first cluster the nodes [113]. In any case, data are typically taken during the system’s approach to steady-state conditions (whether its natural equilibrium or another fixed point of its dynamics) after a perturbation.

A straightforward modification of the basic linear model, realized by overlaying the sum in Eq. (4) with a sigmoidal threshold function, leads to one version of the *artificial neural network* construction. Early methods based on neural networks were used to infer interactions between individual [230] and aggregate “genes” which encompass multiple degrees of freedom at the biological

level [228]. Modern improvements use multilayer perceptrons [231]. Early neural-inspired architectures known as gene circuits [232] have also been used to infer mechanistic interactions [233].

Nonlinear models are attractive because they can capture more sophisticated dynamical behaviors than their linear counterparts (e. g., oscillations and multistability). Nonlinear reverse-engineering schemes based on mass-action kinetic laws like Michaelis-Menten or Hill equations [21] are also used in reconstruction [234, 235].

An important causal inference method based on the assumption of an underlying deterministic system, but which does not require the definition of an explicit dynamical model, is the convergent cross-mapping (CCM) approach [57]. As noted in [237], an essentially identical method had been developed earlier to study synchronization in chaotic dynamical systems [238]. The method draws from Takens’ theorems [239], which provide both the conceptual framework and mathematical justification for a brand of state space reconstruction – reverse-engineering of the phase-space portrait for a dynamical system – known as *delay embedding*. Consider a multidimensional dynamical system, a special case of the general form (4) whose parameters are fixed, and whose temporal evolution $\mathbf{x}(t)$ is confined to a subspace determined by a d -dimensional attractor [240]. Under very general conditions, the attractor’s state space can be reconstructed [239] from measurements of a single time series $\{x_t, x_{t+\tau}, x_{t+2\tau}, \dots\}$, sampled at an interval τ . The number of consecutive time points needed to span the reconstruction space is given by the attractor dimension d ; both τ and d are often found using Ragwitz’ criterion [241], but alternative methods have been proposed as well [242, 243].

Delay embedding refers to the entire process of defining these two parameters and arriving at a reconstruction space onto which the time series can be mapped. It provides the substrate for causal inference via CCM as follows. For any two measured times series $\{x_t\}$ and $\{y_t\}$, the variables x and y are said to be causally linked if they belong to the same underlying dynamic system (i. e., the time series they represent are samples from the same attractor [57, 239, 240]). The direction of an interaction between x and y variables can be estimated by 1) using delay embedding to obtain reconstruction manifolds \mathcal{M}_x and \mathcal{M}_y for x_t and y_t , respectively [240]; 2) projecting one of the variables, say x , onto the other manifold – hence the name cross-mapping – and using the resulting, projected values to predict the values taken by the original time series (which *converge* to the measured values for a large enough number of samples); and 3) measuring (with any suitable measure, e. g. RMSE or correlation function) the deviation of the postdicted values $\{\hat{x}_t\}$ from the actual values $\{x_t\}$. A causal interaction is declared if the prediction quality does not decay to zero for a growing number of samples.

In the original work, Sugihara et al. [57] did not analyze thoroughly the influence of noise on reconstruction.

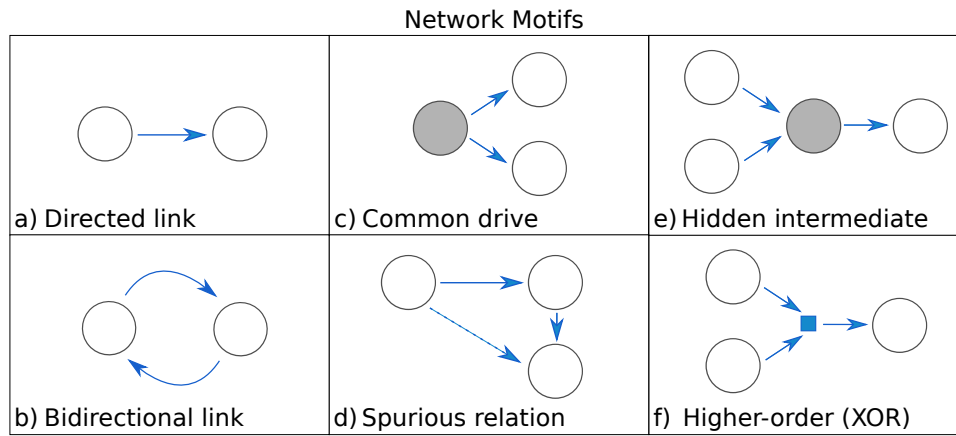


FIG. 2. Simple directed network motifs help illustrate basic problems in directed network reconstruction. This list is not intended to be comprehensive, but to address some simple yet important scenarios. Links can represent mechanistic interactions or effective relations (i. e., information transfer). Nodes represent stochastic or deterministic activation variables, which can be either continuous or discrete. Here, dashed links represent spurious (erroneously inferred) interactions, dark nodes represent unobserved (hidden) variables, and the small square in f) refers to a computation that involves more than two nodes (in this case, a third-order interaction). **a)** The simplest scenario: a directed link between two nodes. **b)** A bidirectional coupling models a simple system with feedback (e. g., the predator-prey system of Fig. 1). **c)** A hidden common drive (dark node) to two observed nodes results in a correlative relation between those nodes. If care is not taken, this can be confounded with a direct causal interaction. **d)** A situation similar to that of c), with the difference that measurements of all three nodes are accessible. Naïve pairwise methods infer a spurious link between the initial and final nodes in the feedforward chain. Multivariate methods are required in this scenario to infer the correct links. **e)** In the case of a hidden node relaying the causal interaction, network reconstruction methods may infer the correct direction of interaction, but the inferred links will be effective rather than strictly causative since an intervention at the hidden node can disrupt the interaction. **f)** The logic gate XOR entails a higher-order interaction. The output is 0 if both input nodes carry the same value, and 1 if they are different: simultaneous knowledge of the states of both nodes is required to determine how each of the inputs affects the output. This is a classical example of a scenario where X and Y carry synergistic (as opposed to unique, or redundant) information [236].

Indeed, Takens’ original theorems allow for noise in the measurement procedure only (i. e., intrinsic stochasticity is prohibited; the breakdown of inference based on CCM in the presence of intrinsic noise has been demonstrated explicitly [244–246], and a thorough analysis of state space reconstruction in the presence of noise can be found in [241]). Nevertheless, artificially added measurement noise can actually improve the detection of causality [247].

Several other considerations must be taken into account when inferring causal relations with CCM. First, it seems that the outcome is quite sensitive to the sampling methods used to obtain training data (for example, eliminating nonstationarity on the way to the attractor is key) [237]. Second, CCM fails to infer the accurate coupling strengths and even the direction of causal interaction when time series are synchronous [246]. Third, it has been shown that the predictions made by CCM do not always conform to our intuitive notions of causality, even for certain rudimentary systems like a simple resistor-inductor (R-L) circuit with a sinusoidal driving voltage, where CCM does not unequivocally determine the causal dependence of the current on the voltage [244]. Finally, Cobey and Baskerville [245] provide a thorough numerical analysis of the limits of CCM, suggesting that the standard approach is generally prone to failure if the system dynamics are oscillatory and proposing a modifica-

tion in the algorithm to alleviate this shortcoming [245].

For stochastic activations, early attempts to reconstruct the directionality of interactions included autoregressive models [248, 249], but autoregression by itself makes no assertions about causality. However, a method due to Granger [250] combines autoregression with the aforementioned notion of temporal precedence to infer quantify a robust stand-in for causality – namely, Weiners predictability [251]. The framework for Granger Causality (GC) is built upon two central assumptions [252]:

1. The cause x occurs before the effect y .
2. The causal series $\{x_t\}$ contains unique information about the time series being caused $\{y_t\}$ that is not available in any other series $\{w_t\}$.

More generally, $\{w_t\}$ represents the entirety of processes that can influence $\{x_t\}$ and $\{y_t\}$. In the ideal scenario, for which these three variables together contain “all the information available in the universe at time t ” [252] (i. e., in the closed system under investigation), GC guarantees that one can reconstruct the direction of the causal relationship between x and y . By definition, a variable x “Granger-causes” variable y if knowledge of past values of both x and y reduces the variance of the prediction error for y , in comparison with the history of y alone.

Typically, these predictions are carried out via linear regression, and the direction of causality is decided by statistical tests on the variances of the respective residuals (prediction errors). However, this implicitly assumes (at most) linear relations between variables. Nonlinear extensions of GC exist, but these extensions can be more difficult to use in practice and their statistical properties are less well understood [253–257].

Granger causality can be extended to multivariate scenarios [258] as well, although finding Granger-causal links among all possible candidate interactions then becomes a combinatorially hard problem. For the particular case of inferring causal relations between the activity of distinct brain areas (using electroencephalograms or local field potential time series), it has been found to be of crucial importance to employ a multivariate approach rather than bivariate techniques [259].

A more general approach to the reverse-engineering of directed links between stochastic variables is to learn an explicit model for the joint probability distribution of the observed activities. This approach, based on *probabilistic graphical models*, was discussed earlier for undirected networks. For the directed case, one can define a class of models known as Bayesian networks [260–263] which decompose the joint distribution into separate factors representing conditional probabilities. Edges are drawn starting from the nodes corresponding to variables being conditioned on (called the “parents”) and ending on the conditioned variables (the “children”) [211, 263]. Since the joint distribution of a Bayesian network is an exact product of conditional probabilities, the resulting graphical structure is a *directed acyclic graph* (DAG). Thus in order to be eligible for representation by a Bayesian network, systems need to satisfy the necessary criteria for forming a DAG. If the phenomenon in question is known to encompass cyclic dependencies (e. g., autoregulation pathways in gene regulatory networks, or autapses in neural networks), the only recourse is to “unroll” the cyclic dynamics in time, forming a *dynamic Bayesian network* [163, 264–267]. The performance of dynamic Bayesian nets has been compared directly against that of Granger causality [268], and favorably so when the observed time series are shorter than a certain length (NB: In general, findings like these should be taken with a grain of salt, since 1) they could be artifactual results that depend on idiosyncratic features of the data, and 2) notions of error and accuracy tend to rest on the existence of a reference network containing only the “correct” edges, which is in our opinion a dubious concept; see comments on evaluation metrics in the *Discussion*. In [268], the authors are clear in their admission that “the causal relationship derived from these two approaches could be different, in particular when we face the data obtained from experiments,” in accordance with our introductory statements about the nonuniform definitions of causality that are assumed by different methods.).

With the conditional probability framework in place, one needs to select 1) a quantitative form for the under-

lying model that parameterizes the conditional probabilities, 2) a scoring or objective function that quantifies the quality of fit, and 3) an optimization or search routine by which to learn the parameters values that extremize the objective function. An example of such a parameterization, used quite frequently in the literature, is again that of linear regression [161, 263]. The choice of a specific parametric representation of conditional probabilities is often dictated by our knowledge or assumptions about the domain (prior knowledge) [269], or pragmatic principles favoring computationally simple models (Occam’s razor). Standard objectives are the maximization of the likelihood function [264] or posterior probability distribution [161], as well as the Bayesian Information Criterion (BIC) [266], which penalizes for large numbers of parameters. Since the optimization search is an NP-hard problem [261, 263], exact methods are often computationally infeasible, so one often reverts to heuristics like greedy hill-climbing (which adds, deletes, or reverses edges to encourage maximal ascent in the objective score [270]), stochastic hill-climbing, or Monte Carlo methods [271].

An impressively comprehensive and thorough body of work regarding the concept of causality and its formal description via Bayesian nets has been provided by their originator, Judea Pearl [211]. Pearl introduces a conceptual framework called the *do-formalism* (known variously as the do-calculus, the intervention-calculus, etc.), which formally describes the use of experimental interventions to ascertain a causal structure. In the do-formalism, $p(y|do(x))$ denotes conditioning on a variable x that is experimentally controlled rather than simply measured (i. e., observed passively). In other words, this notation distinguishes the more familiar observational conditioning $p(y|x)$ from “interventional conditioning” [152, 272].

While correlation does not in general imply causal influence, Pearl reveals specific cases for which the conditional probability distribution – reflecting associative dependencies – is equivalent to that which denotes the corresponding mechanistic dependencies: in such situations, interventions which manipulate the values of parent nodes are clearly and unambiguously seen to have direct effects on the children, and the Bayesian graph is therefore also the correct casual graph.

It is often difficult to satisfy all the criteria for modeling a causal system with DAGs. In certain circumstances, it is easier to work with model-free stochastic frameworks, such as that of the transfer entropy (TE). TE was introduced twice independently, by the physicists Schreiber [273] and Paluš [274], and has since proven to be a versatile and useful tool for inferring the direction of information transfer in neuroscience [275–277], physiology [243, 278], climatology [279–281] and economics [282, 283]. TE is simply the conditional mutual information (2) between a target variable Y and the entire history of values assumed by a source variable X , given the history of the target:

$$\mathcal{T}(X \rightarrow Y) = I(\mathbf{X}_{t-}; Y_t | \mathbf{Y}_{t-}). \quad (5)$$

Here the arrow denotes the direction of information transfer (i. e., X informs Y) and \mathbf{X}_{t-} and \mathbf{Y}_{t-} respectively denote the histories of the corresponding stochastic processes up to, but not including, t ; Y_t denotes the value taken by the target variable at time t . Conditioning on the history of the target ensures that only those bits of information that are unique (in the sense discussed earlier for Granger Causality; for a formal treatment see [236, 284]) to the source variable are considered.

Like all information-theoretic measures, TE and its surrogates [57] suffer from the curse of dimensionality because of the need to estimate entire probability distributions (discrete variables) or probability densities (continuous variables) for long time series and many variables. For discrete variables, the simplest estimation procedure entails simply counting frequencies to produce a histogram that approximates the desired distribution. A substantially more accurate estimation of information-theoretic quantities for discrete variables (especially if the data set is small) can be obtained by computing entropies directly with the NSB estimator [285, 286]. In the continuous case, a standard approach is to bin the data, rendering the distribution effectively discrete and therefore amenable to histogram methods. While less “data hungry” alternatives exist for continuous variables (such as kernel estimators [287]), they suffer from the same systematic estimation biases that are associated with histogram methods [288], and may even reverse the inferred direction of information flow [289]. Nearest neighbor estimators [277, 288] are some of the most commonly used in practice. In all cases, statistical testing against surrogate data or empirical control data [290] is recommended to help ameliorate the bias problem.

An approach to dimensionality reduction based on the concept of Markov chains has been proposed for the estimation of TE [280]. This approach is particularly useful in the case of delayed coupling between variables [291]: estimation of the delay time can prevent the inclusion of unnecessary time steps when tracking the history of the source variable (i. e., X_{t-} in Eq. (5)), which can clearly reduce the dimensionality of the latent representation. Finally, the curse of dimensionality can also be alleviated by first constructing an explicit, low-dimensional model of the time series (and hence, parameterizing the probability distribution). For the simplest case – linear dependence between X and Y with additive Gaussian noise – it has been shown analytically that TE will always recover the same network as Granger Causality, up to a constant factor [292].

Since some authors speak loosely about inferring causality when computing the TE or related quantities like the directed information [105], we reiterate that, although causal interaction is a necessity for information transfer, the converse is not true: information transfer, as quantified by TE and other information-theoretic functionals, does not imply underlying causal interactions. In fact, we caution readers that some methods for the detection of causal or directed influence have been rou-

tinely applied in ways that differ markedly from the intentions of their originators. For instance, the directed information was initially designed to infer achievable information rates on a known communication channel with feedback [293], rather than the inference of directed networks (for a thorough discussion, see [277]). However, TE specifically has been extended using the aforementioned do-formalism in a new procedure known as *information flow* [211], a more appropriate measure for inferring causality under certain constraints [272, 294]. Notably, this measure can correctly resolve the connectivities of an XOR circuit (see Fig. 2f) even in special scenarios where the conditional mutual information fails [272], a fact overlooked by authors who have contended that conditional mutual information is sufficient for this purpose (see, for instance, the argument in [186]). Finally, we note that TE and similar methods have not achieved widespread implementation for large systems ($p \gg 1$) due to the aforementioned, intrinsic difficulty of estimating information theoretic measures in high dimensional spaces. Multivariate approaches to TE estimation and related methods are a subject of ongoing research.

IV. DISCUSSION

Since the year 2000, some thirty review articles that we know of have been published on the inference of gene networks alone (in addition to those referenced or mentioned throughout, see [295–322]), and an increasing number have begun to specialize on the unique challenges faced by network reverse-engineers rather than merely listing different algorithms [29, 97, 323–328]. One DREAM report [92] notes that the number of PubMed articles on reverse-engineering had doubled each year for over a decade through 2009, and “novel” algorithms (new twists on the same foundational principles we outline above) continue to emerge even as we write [329].

Has this explosive growth in the number of reverse-engineering algorithms and studies helped carve out a niche for large-scale reverse-engineering in contemporary systems biology repertoires? Or has a staunch directive on the reconstruction of entire microscopic networks actually encumbered and obfuscated our understanding of the working principles that underlie these complex systems?

One major impediment to assessing the promise of reverse-engineering algorithms stems from the way in which they are assessed: we observe a rampant, pervasive, and potentially counterproductive tendency to draw direct, quantitative comparisons between reconstructions produced by different algorithms. In other words, despite the commoditization of network inference tools, there is still no consensus on the correct way to *evaluate* reconstruction results [91] – and perhaps for good reason! In the context of effective network inference, the notion that reconstructions can be *checked for accuracy* contradicts our very premise, that algorithms both among and within

each of the classes we have described make diverse assumptions about what should count as an interaction. Recent work [93, 330] notwithstanding, we believe this issue continues to be confounded by a repeated mismatch between algorithms and metrics (as in the use of the area under receiver-operator characteristic curves [331], a measure that presupposes the existence of a valid confusion matrix, to give an overall rank or “score” to effective reconstructions [324, 332]).

The methods in different classes also differ in more concrete ways: they vary in the extent to which they can infer strengths, signs, and directions for the interactions they detect. This might be thought of as a “feature, not a bug” of reverse-engineering technologies: having a selection of versatile algorithms, each tailored to particular situations or designed with different inference goals in mind, increases the chances that researchers can make use of reverse-engineering algorithms. Yet the question of whether systems biologists should persist in pursuing whole-network reconstruction as a go-to modality or learning tool hinges not solely on whether the inference goals are achieved by the time the smoke clears, but on the attainment of a reasonable *tradeoff* between the computational effort consumed by inference algorithms and the (ideally, unique) benefits they afford to researchers.

Do the spectrum and short history of network inference successes live up to such high hopes? Along these lines, we have argued that reverse-engineering over the past two decades has played at least five distinct research roles – the acceleration of hypothesis generation and verification at the single-node/single-interaction level, the illumination of statistical properties that render biological networks unique among complex systems, the diagnosis of individual networks as either typical or perturbed (paralleled by the use of within-class variation to make theoretical statements about the system), the prediction of how the activities in a given network will respond to exogenous perturbations, and the compact encoding of joint probability distributions – that go far beyond the trivial task of piecing together which of a set of observed system elements engage in physical contacts or the transfer of biologically relevant information. The roles we have identified represent a far cry from the (three) uses of effective influence networks – identification of functional modules, probing the response to perturbations, and helping determine the underlying mechanistic interactions – named by the authors of Ref. [28] ten years ago.

While it is impossible to say which of recent attempts to use networks as compressed “statistics” to help make (quantitative or qualitative) predictions will have the biggest impact down the road, it is clear that new precedents for the prediction of drug targets and systemic responses in network medicine [101] point to a significant departure from the more traditional, reductionist ways of thinking. The consequences here will almost certainly include dramatic impacts on the ways medicine is practiced in the lifetime of the reader. With this example in

mind, we reiterate our assertion that reverse-engineering yields its most succulent fruit when it is used to augment other methods of expanding our understanding of how living systems work, rather than employed disposable as an end goal in itself. Indeed, changes in the ways network inference has been used over time seem to be in accordance with this sentiment: whereas in 2003 the field was still firmly entrenched in its “pattern-detection phase” [333] (to better understand the state of the art at that time, we recommend [19]), it was around the time of publication of [92] in 2009 that the DREAM4 Challenges first introduced predictive modeling tasks as part of the main annual competition.

Indeed, the DREAM competitions play a unique part in the reverse-engineering culture. They not only echo changes in the field’s priorities but also inform them: they have helped set the precedent in establishing inferred networks as tools for making predictions (as in the DREAM8 prompt to anticipate the responses of cellular signals to yet-unseen perturbations [334]). More radically, some of the most recent Challenges go as far as skipping the hitherto-canonical intermediate step of network inference entirely, asking competitors to infer macroscopic properties or outcomes using wholly different types of data [335]. While we clearly do not advocate for the complete abandonment of automated, network-scale reverse-engineering from large data sets, we do view the foundation’s decreasing reliance on methods which require the construction of a detailed microscopic model prior to making inference about the macroscopic system as a progressive step. In fact, we contend that, given suitable alternatives, whole-network reverse-engineering may not be justified in every case.

If the reverse-engineering of entire microscopic networks is not always the right tool for the job, what might be done instead? As a starting point, we suggest asking:

- Given a reverse-engineered network, can we find any further compressions of that network that still preserve information about (i. e., are equally good at predicting) the macroscopic properties and observables it encodes?
- Can we identify any coarse *functional units* (perhaps with their own set of interaction rules and dynamics) that might supplant individual nodes and edges as the elements of a common parlance for the study of living systems?

For instance, might more appropriate “parts lists” for biological systems consist not of individual species’ activations, but of larger physical or conceptual elements (e. g., negative feedback loops and operons) with their own dynamical interaction laws? Alternatively, attractors of the dynamics of biological networks may be a more laconic descriptions of the networks than the interactions among the nodes themselves [14, 336]. This possibility may be motivated via a historical analogy: *renormalization group* theory in physics [337] offers a systematic way

to deduce an appropriate new vocabulary (and the corresponding syntax) when one changes the physical scale at which a system is to be observed. The effective interaction rules which emerge (say, the interactions between groups of Ising spins) are not always easily reducible to the familiar dynamics of microscopic activation variables (the nearest-neighbor interactions associated with individual spins), but which nonetheless account accurately for their effects at the new scale.

A recent line of work, inspired directly by statistical physics, formalizes the argument that only a small subset of parameter combinations are easily learnable from data, and therefore that only certain (combinations of) microscopic parameters can be *relevant* in determining a complex system’s macroscopic or emergent properties [338–341]. By systematically integrating out “sloppy” parameters or parameter combinations, whose values remain relatively unconstrained, one can assemble coarse, parsimonious models in terms of the remaining “stiff” parameters that serve as effective, low-dimensional compressions of a system’s microscopic statistics.

Answers to the second question – that of finding higher-level explanatory structures in terms of which system’s behavior can be understood – have been explored since the inception of “module-based” inference [119, 122]. In fact, newer and more powerful tools have sparked a resurgence [114, 342–344] of this approach. Around the same time, it was demonstrated that the flow of information in development, from promoter sequence to expression, can be reliably understood in terms of coarse, multiple-sequence patterns called graph-mers [345] that encompass entire sequence motifs. Ultimately, we believe that it will be work in directions such as these, which involve gross reconceptualizations regarding the fundamental actors in the biological dynamics, that will supersede whole-network reverse-engineering.

If the end goal of emulating physics-style modeling is prediction, the penultimate is certainly intuition and conceptual understanding. We entertain *phenomenological* approaches like renormalization because they promise to yield interpretable models, not intractably large sets of detailed equations. Yet we still stress that, while searching for modularity and simple descriptions entails an invocation of the engineering *mindset* that has informed systems biology since its inception, the principles of good biological design often differ markedly from what works in that context; an open mind is necessary to dream up fitting new constructs. Whatever the case, we are confident that it is only by focusing on phenomenological (rather than microscopic) accuracy that we can deliver a satisfying confutational blow to famous Rutherford’s quip that “all sciences are either physics or stamp collecting [346] and begin removing the major impediments to the advancement of formal theories in biology [347].

V. ACKNOWLEDGMENTS

This work was funded in part by the James S. McDonnell foundation Grant JSMF/ 220020321 and the National Science Foundation Grant NSF/PoLS/1410978, and by the Laney Graduate School Fellowship (LGSF) at Emory University.

VI. TRY ON YOUR OWN: BECOME A REVERSE-ENGINEER

By now we hope to have made a convincing case for our contention that different reverse-engineering methodologies are, in general, best-suited for answering different types of questions. We have reviewed the most prominent such questions, and illustrated how the “goals” fulfilled by specific algorithms are really manifestations of their underlying assumptions about what should count as an interaction.

Since no one definition of biological interaction can be considered more “correct” than the others in all contexts (different algorithms merely capture different aspects of the same system), a diversity of goals and operational idiosyncracies might be viewed as a blessing rather than a curse. Yet choices should be made at the outset regarding what one wishes to learn by doing reverse-engineering, because these choices inform which algorithms are best suited for the job.

In this section, we simulate the conditions under which the need for such choices arises. Imagine that you have just been handed a set of high-throughput data, for a system whose interaction architectures have not yet been fully mapped. Follow the series of prompts in the box to embark on an exploratory challenge with a representative set of actual experimental data.

Consider a set of multi-electrode recordings from the retina of a salamander (we thank M. Berry for providing us with data from [191]; download link at https://figshare.com/articles/bint_fishmovie32_100_mat/5009840). As explained in detail in the `README.txt` file, the data consists of the responses from $p = 160$ ganglion cells to the presentation of a naturalistic stimulus – in this case, a short (~ 20 sec) movie of a fish tank, repeated $n = 297$ times. The activity of each neuron is binarized as 0 (when the neuron is not firing an action potential) and 1 (when it is firing an action potential) within discrete time bins of length 20 ms.

1. Of the methods discussed in this Chapter, which are clearly applicable to this particular set of data? Are there any which are not?
2. What kinds of predictions might a researcher want to make using this data?

Consider multiple levels of analysis, from single nodes in the neuronal network (Will removing a

single node cause the network to collapse? Can we predict a future value for a given neuron, given the values of certain others?) to multiple nodes (Are there any functional groups that seem to be operating as a unit? Are there hub structures present?) to the entire system as an emergent whole (What can we say about the percentage of time the system is silent, versus when it is spiking? What other information would we need to say something about the “typicality” of the recorded networks, with respect to their structural and dynamical properties?).

3. Crowdsourcing [348] – the idea that conglomerate predictions, made by combining the wisdom of many independent thinkers, are more accurate than those of any individual – is a popular strategy in DREAM competitions [349, 350] (for recent examples, see the closed Sage Bionetworks-DREAM Breast Cancer Prognosis (DREAM7,

2012), NIEHS-NCATS-UNC DREAM Toxicogenetics (DREAM8, 2013) and ICGC-TCGA DREAM Somatic Mutation Calling (DREAM 8.5-9, 2013-2014) Challenges). Yet we have seen that different reverse-engineering methods often yield disparate – even antagonistic or contradictory – predictions. For which combination of the following algorithms would you feel comfortable following the “wisdom of crowds” (say, averaging the results, or taking majority rules)?

Think about ARACNe, CLR, Bayesian networks (static and dynamics), MaxEnt approaches, and possibly other methods. Given the assumptions these methods make, would you take the union or intersection of the set of results produced by Bayesian methods and ARACNe? MaxEnt and CLR? Other combinations? When do you think crowdsourcing in general is a good strategy?

-
- [1] L. H. Hartwell, J. J. Hopfield, S. Leibler, and A. W. Murray, *Nature* **402**, C47 (1999).
 - [2] S. R. Proulx, D. E. L. Promislow, and P. C. Phillips, *Trends Ecol Evol SPECIAL ISSUE: BUMPER BOOK REVIEW*, **20**, 345 (2005).
 - [3] X. Zhu, M. Gerstein, and M. Snyder, *Genes Development* **21**, 1010 (2007).
 - [4] B. N. Kholodenko, *Nature Rev Mol Cell Biol* **7**, 165 (2006).
 - [5] S. B.-T. de Leon and E. H. Davidson, *Annual Rev Biophys Biomol Struct* **36**, 191 (2007).
 - [6] B. N. Kholodenko, J. F. Hancock, and W. Kolch, *Nature Rev Mol Cell Biol* **11**, 414 (2010).
 - [7] J. O. Dubuis, G. Tkačik, E. F. Wieschaus, T. Gregor, and W. Bialek, *Proc Natl Acad Sci (USA)* **110**, 16301 (2013).
 - [8] X. Wang, X. Wei, B. Thijssen, J. Das, S. M. Lipkin, and H. Yu, *Nature Biotechnol* **30**, 159 (2012).
 - [9] N. E. Buchler, U. Gerland, and T. Hwa, *Proc Natl Acad Sci (USA)* **100**, 5136 (2003).
 - [10] K. Wang, M. Saito, B. C. Bisikirska, M. J. Alvarez, W. K. Lim, P. Rajbhandari, Q. Shen, I. Nemenman, K. Basso, A. A. Margolin, U. Klein, R. Dalla-Favera, and A. Califano, *Nature Biotechnol* **27**, 829 (2009).
 - [11] A. A. Margolin, K. Wang, A. Califano, and I. Nemenman, *IET Syst Biol* **4**, 428 (2010).
 - [12] E. Ganmor, R. Segev, and E. Schneidman, *Proc Natl Acad Sci (USA)* **108**, 9679 (2011).
 - [13] D.-Y. Cho, Y.-A. Kim, and T. M. Przytycka, *PLOS Comp Biol* **8** (2012).
 - [14] L. Merchan and I. Nemenman, *J Stat Phys* **162**, 1294 (2016).
 - [15] W. S. Hlavacek, *Mol Syst Biol* **5**, 240 (2009).
 - [16] E. Estrada, *The Structure of Complex Networks: Theory and Applications* (Oxford University Press, Oxford, UK, 2011).
 - [17] H. Kitano, *Foundations of Systems Biology* (MIT Press, 2001).
 - [18] U. Alon, *Science* **301**, 1866 (2003).
 - [19] E. Alm and A. Arkin, *Curr Opin Struct Biol* **13**, 193 (2003).
 - [20] A.-L. Barabási and Z. N. Oltvai, *Nature Rev Genet* **5**, 101 (2004).
 - [21] U. Alon, *An Introduction to Systems Biology: Design Principles of Biological Circuits* (CRC Press, 2006).
 - [22] B. Ø. Palsson, *Systems Biology: Properties of Reconstructed Networks* (Cambridge University Press, 2006).
 - [23] N. Subramanian, P. Torabi-Parizi, R. A. Gottschalk, R. N. Germain, and B. Dutta, *Wiley Interdiscip Rev: Syst Biol Med* **7**, 13 (2015).
 - [24] F. Jacob and J. Monod, *J Mol Biol* **3**, 318 (1961).
 - [25] F. Morcos, T. Hwa, J. N. Onuchic, and M. Weigt, in *Protein Structure Prediction*, Methods in Molecular Biology No. 1137, edited by D. Kihara (Springer New York, 2014) pp. 55–70.
 - [26] N. Halabi, O. Rivoire, S. Leibler, and R. Ranganathan, *Cell* **138**, 774 (2009).
 - [27] C. Feinauer, H. Szurmant, M. Weigt, and A. Pagnani, *PLOS ONE* **11**, e0149166 (2016).
 - [28] M. Bansal, V. Belcastro, A. Ambesi-Impiombato, and D. di Bernardo, *Mol Syst Biol* **3** (2007).
 - [29] J. Linde, S. Schulze, S. G. Henkel, and R. Guthke, *EXCLI Journal* **14**, 346 (2015).
 - [30] T. Çakır, M. M. W. B. Hendriks, J. A. Westerhuis, and A. K. Smilde, *Metabolomics* **5**, 318 (2009).
 - [31] T. Çakır and M. J. Khatibipour, *Frontiers Bioeng Biotechnol* **2** (2014).
 - [32] B. N. Kholodenko, A. Kiyatkin, F. J. Bruggeman, E. Sontag, H. V. Westerhoff, and J. B. Hoek, *Proc Natl Acad Sci (USA)* **99**, 12841 (2002).
 - [33] J. Stark, R. Callard, and M. Hubank, *Trends Biotechnol* **21**, 290 (2003).
 - [34] R. J. Prill, J. Saez-Rodriguez, L. G. Alexopoulos, P. K. Sorger, and G. Stolovitzky, *Sci Sign* **4**, mr7 (2011).
 - [35] R. Cheong, A. Rhee, C. J. Wang, I. Nemenman, and A. Levchenko, *Science* **334**, 354 (2011).

- [36] Q. Ni, A. Ganesan, N.-N. Aye-Han, X. Gao, M. D. Allen, A. Levchenko, and J. Zhang, *Nature Chem Biol* **7**, 34 (2011).
- [37] S. J. Little, S. L. K. Pond, C. M. Anderson, J. A. Young, J. O. Wertheim, S. R. Mehta, S. May, and D. M. Smith, *PLOS ONE* **9**, e98443 (2014).
- [38] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, *Nature* **440**, 1007 (2006).
- [39] L. M. Bettencourt, G. J. Stephens, M. I. Ham, and G. W. Gross, *Phys Rev E* **75**, 021915 (2007).
- [40] R. Friedrich, C. Genoud, and A. A. Wanner, *Frontiers Neur Circ* **7**, 71 (2013).
- [41] K. Chung and K. Deisseroth, *Nature Meth* **10**, 508 (2013).
- [42] J. R. Epp, Y. Niibori, H.-L. Liz Hsiang, V. Mercaldo, K. Deisseroth, S. A. Josselyn, and P. W. Frankland, *eNeuro* **2**, ENEURO.0022 (2015).
- [43] I. H. Stevenson, J. M. Rebesco, L. E. Miller, and K. P. Körding, *Curr Opin Neurobiol* **18**, 582 (2008).
- [44] E. T. Bullmore and O. Sporns, *Nature Rev Neurosci* **10**, 186 (2009).
- [45] E. T. Bullmore and D. S. Bassett, *Annual Rev Clin Psych* **7**, 113 (2011).
- [46] K. J. Friston, *Brain Connectivity* **1**, 13 (2011).
- [47] C. Capone, C. Filosa, G. Gigante, F. Ricci-Tersenghi, and P. D. Giudice, *PLOS ONE* **10**, e0118412 (2015).
- [48] M.-E. Lynall, D. S. Bassett, R. Kerwin, P. J. McKenna, M. Kitzbichler, U. Muller, and E. T. Bullmore, *J Neurosci* **30**, 9477 (2010).
- [49] R. G. Canter, J. Penney, and L.-H. Tsai, *Nature* **539**, 187 (2016).
- [50] N. Eagle, A. S. Pentland, and D. Lazer, *Proc Natl Acad Sci (USA)* **106**, 15274 (2009).
- [51] D. Angluin, J. Aspnes, and L. Reyzin, in *Proc 21st Intern Conf on Algorithmic Learning Theory*, ALT'10, edited by M. Hutter, F. Stephan, V. Vovk, and T. Zeugmann (Springer-Verlag, Berlin, Heidelberg, 2010) pp. 104–118.
- [52] I. Psorakis, S. J. Roberts, I. Rezek, and B. C. Sheldon, *J R Soc Interf* **9**, 3055 (2012).
- [53] R. M. May, *Math Biosci* **12**, 59 (1971).
- [54] R. M. May, *Nature* **238**, 413 (1972).
- [55] R. M. May, *Trends Ecol Evol* Twenty years of TREE - part 2, **21**, 394 (2006).
- [56] M. W. McCoy, B. M. Bolker, K. M. Warkentin, and J. R. Vonesh, *Amer Naturalist* **177**, 752 (2011).
- [57] G. Sugihara, R. M. May, H. Ye, C.-h. Hsieh, E. R. Deyle, M. Fogarty, and S. Munch, *Science* **338**, 496 (2012).
- [58] E. R. Deyle, R. M. May, S. B. Munch, and G. Sugihara, *Proc R Soc B* **283**, 20152258 (2016).
- [59] H. Moon and T.-C. Lu, *Scientific Rep* **5**, 9450 (2015).
- [60] R. Compton, H. Moon, and T.-C. Lu, “Catastrophe prediction via estimated network autocorrelation,” (2015), uS Patent 9,020,875.
- [61] G. K. Nikolopoulos, E. Pavlitina, S. Q. Muth, J. Schneider, L. D. Mina Psychogiou and Williams, D. Paraskevis, V. Sypsa, G. Magiorkinis, P. Smyrnov, A. Korobchuk, T. I. Vasylyeva, B. Skaathun, M. Malliori, E. Kafetzopoulos, A. Hatzakis, and S. R. Friedman, *Scientific Rep* **6**, 38100 (2016).
- [62] D. B. Bahr, R. C. Browning, H. R. Wyatt, and J. O. Hill, *Obesity* **17**, 723 (2009).
- [63] A. Arkin and J. Ross, *The Journal of Physical Chemistry* **99**, 970 (1995).
- [64] J.-H. S. Hofmeyr and A. Cornish-Bowden, *J Theor Biol* **182**, 371 (1996).
- [65] A. A. Margolin, I. Nemenman, K. Basso, C. Wiggins, G. Stolovitzky, R. D. Faveira, and A. Califano, *BMC Bioinf* **7**, S7 (2006).
- [66] W. A. Schmitt, R. M. Raab, and G. Stephanopoulos, *Genome Res* **14**, 1654 (2004).
- [67] M. J. Dunlop, R. S. Cox, J. H. Levine, R. M. Murray, and M. B. Elowitz, *Nature Gen* **40**, 1493 (2008).
- [68] A. F. Villaverde, J. Ross, F. Morán, and J. R. Banga, *PLOS ONE* **9**, e96732 (2014).
- [69] A. de la Fuente, P. Brazhnik, and P. Mendes, *Trends Genet* **18**, 395 (2002).
- [70] S. D. Santos, P. J. Vermeer, and P. I. Bastiaens, *Nature Cell Biol* **9**, 324 (2007).
- [71] P. Bandaru, M. Bansal, and I. Nemenman, *J Comp Biol* **18**, 147 (2011).
- [72] R. R. Stein, D. S. Marks, and C. Sander, *PLOS Comp Biol* **11**, e1004182 (2015).
- [73] M. K. S. Yeung, J. Tegnér, and J. J. Collins, *Proc Natl Acad Sci (USA)* **99**, 6163 (2002).
- [74] T. S. Gardner, D. di Bernardo, D. Lorenz, and J. J. Collins, *Science* **301**, 102 (2003).
- [75] A. F. Villaverde and J. R. Banga, *J R Soc Interf* **11**, 20130505 (2014).
- [76] C. Christensen, J. Thakar, and R. Albert, *IET Syst Biol* **1**, 61 (2007).
- [77] M. Vidal, M. E. Cusick, and A.-L. Barabási, *Cell* **144**, 986 (2011).
- [78] C. F. Quo, C. Kaddi, J. H. Phan, A. Zollanvari, M. Xu, M. D. Wang, and G. Alterovitz, *Brief Bioinf* **13**, 430 (2012).
- [79] A. D. Lander, *BMC Biol* **8**, 40 (2010).
- [80] P. W. Anderson, *Science* **177**, 393 (1972).
- [81] M. Hecker, S. Lambeck, S. Toepfer, E. P. van Someren, and R. Guthke, *Biosyst* **96**, 86 (2009).
- [82] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller, *Nature Meth* **10**, 413 (2013).
- [83] D. A. Schwarz, M. A. Lebedev, T. L. Hanson, D. F. Dimitrov, G. Lehew, J. Meloy, S. Rajangam, V. Subramanian, P. J. Ifft, Z. Li, A. Ramakrishnan, A. Tate, K. Z. Zhuang, and M. A. L. Nicolelis, *Nature Meth* **11**, 670 (2014).
- [84] M. Timme and J. Casadiego, *J Phys A* **47**, 343001 (2014).
- [85] N. Goldenfeld and L. P. Kadanoff, *Science* **284**, 87 (1999).
- [86] B. C. Daniels and I. Nemenman, *Nature Comm* **6** (2015).
- [87] T. S. Gardner and J. J. Faith, *Phys Life Rev* **2**, 65 (2005).
- [88] A. P. Parikh, W. Wu, R. E. Curtis, and E. P. Xing, *Bioinf* **27**, i196 (2011).
- [89] L. Song, M. Kolar, and E. P. Xing, *Bioinf* **25**, i128 (2009).
- [90] M. Newman, *SIAM Rev* **45**, 167 (2003).
- [91] G. Stolovitzky, D. Monroe, and A. Califano, *Ann New York Acad Sci* **1115**, 1 (2007).
- [92] G. Stolovitzky, R. J. Prill, and A. Califano, *Ann New York Acad Sci* **1158**, 159 (2009).
- [93] P. C. Boutros, A. A. Margolin, J. M. Stuart, A. Califano, and G. Stolovitzky, *Genome Biol* **15**, 462 (2014).
- [94] J. Yamane, S. Aburatani, S. Imanishi, H. Akanuma, R. Nagano, T. Kato, H. Sone, S. Ohsako, and W. Fu-

- jibuchi, *Nucl Acids Res* **44**, 5515 (2016).
- [95] A. J. Gates and L. M. Rocha, *Scientific Rep* **6** (2016).
- [96] M. Tikhonov and W. Bialek, *Phys Biol* **13** (2016).
- [97] D. Yu, M. Kim, G. Xiao, and T. H. Hwang, *Genomics and Informatics* **11**, 200 (2013).
- [98] N. J. Mulder, R. O. Akinola, G. K. Mazandu, and H. Rapanoel, *Comput Struct Biotechn J* **11**, 1 (2014).
- [99] A. de la Fuente, N. Bing, I. Hoeschele, and P. Mendes, *Bioinf* **20**, 3565 (2004).
- [100] K. Basso, A. A. Margolin, G. Stolovitzky, U. Klein, R. Dalla-Favera, and A. Califano, *Nature Gen* **37**, 382 (2005).
- [101] P. Csermely, T. Korcsmáros, H. J. Kiss, G. London, and R. Nussinov, *Pharmacology and Therapeutics* **138**, 333 (2013).
- [102] J. J. Faith, B. Hayete, J. T. Thaden, I. Mogno, J. Wierzbowski, G. Cottarel, S. Kasif, J. J. Collins, and T. S. Gardner, *PLOS Biol* **5**, e8 (2007).
- [103] A. Rao, A. O. Hero, D. J. States, and J. D. Engel, *Comput Syst Bioinf Conf* **6**, 145 (2007).
- [104] G. Altay and F. Emmert-Streib, *BMC Syst Biol* **4**, 132 (2010).
- [105] C. Kaleta, A. Göhler, S. Schuster, K. Jahreis, R. Guthke, and S. Nikolajewa, *BMC Syst Biol* **4**, 116 (2010).
- [106] S. Kumari, J. Nie, H.-S. Chen, H. Ma, R. Stewart, X. Li, M.-Z. Lu, W. M. Taylor, and H. Wei, *PLOS ONE* **7**, e50411 (2012).
- [107] S. Ma, P. Kemmeren, D. Gresham, and A. Statnikov, *PLOS ONE* **9**, e106479 (2014).
- [108] P. Meyer, T. Cokelaer, D. Chandran, K. H. Kim, P.-R. Loh, G. Tucker, M. Lipson, B. Berger, C. Kreutz, A. Raue, B. Steiert, J. Timmer, E. Bilal, H. M. Sauro, G. Stolovitzky, and J. Saez-Rodriguez, *BMC Syst Biol* **8**, 13 (2014).
- [109] R. Sharan, I. Ulitsky, and R. Shamir, *Mol Syst Biol* **3**, 88 (2007).
- [110] X. Wen, S. Fuhrman, G. S. Michaels, D. B. Carr, S. Smith, J. L. Barker, and R. Somogyi, *Proc Natl Acad Sci (USA)* **95**, 334 (1998).
- [111] C. J. Wolfe, I. S. Kohane, and A. J. Butte, *BMC Bioinf* **6**, 227 (2005).
- [112] J. Gillis and P. Pavlidis, *PLOS Comp Biol* **8**, e1002444 (2012).
- [113] R. Bonneau, D. J. Reiss, P. Shannon, M. Facciotti, L. Hood, N. S. Baliga, and V. Thorsson, *Genome Biol* **7**, R36 (2006).
- [114] R. De Smet and K. Marchal, *Nature Rev Microbiol* **8**, 717 (2010).
- [115] E. Segal, M. Shapira, A. Regev, D. Pe'er, D. Botstein, D. Koller, and N. Friedman, *Nature Gen* **34**, 166 (2003).
- [116] N. Kravchenko-Balasha, A. Levitzki, A. Goldstein, V. Rotter, A. Gross, F. Remacle, and R. Levine, *Proc Natl Acad Sci (USA)* **109**, 4702 (2012).
- [117] H. Jeong, S. P. Mason, A.-L. Barabási, and Z. N. Oltvai, *Nature* **411**, 41 (2001).
- [118] X. He and J. Zhang, *PLOS Genet* **2**, e88 (2006).
- [119] S. Maslov and K. Sneppen, *Science* **296**, 910 (2002).
- [120] N. Lemke, F. Herédia, C. K. Barcellos, A. N. d. Reis, and J. C. M. Mombach, *Bioinf* **20**, 115 (2004).
- [121] S. S. Shen-Orr, R. Milo, S. Mangan, and U. Alon, *Nature Gen* **31**, 64 (2002).
- [122] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, *Science* **298**, 824 (2002).
- [123] R. Sharan, S. Suthram, R. M. Kelley, T. Kuhn, S. McCuine, P. Uetz, T. Sittler, R. M. Karp, and T. Ideker, *Proc Natl Acad Sci (USA)* **102**, 1974 (2005).
- [124] J. M. Stuart, E. Segal, D. Koller, and S. K. Kim, *Science* **302**, 249 (2003).
- [125] M. W. Hahn and A. D. Kern, *Mol Biol Evol* **22**, 803 (2005).
- [126] H. Kitano, *Nature Rev Genet* **5**, 826 (2004).
- [127] M. P. Stumpf and M. A. Porter, *Science* **335**, 665 (2012).
- [128] A.-L. Barabási, N. Gulbahce, and J. Loscalzo, *Nature Rev Genet* **12**, 56 (2011).
- [129] P. Lecca and C. Priami, *Drug Discovery Today* **18**, 256 (2013).
- [130] E. T. Jaynes, *Phys Rev* **106**, 620 (1957).
- [131] E. T. Jaynes, *Proc IEEE* **70**, 939 (1982).
- [132] G. Tkačik, O. Marre, T. Mora, D. Amodei, M. J. Berry II, and W. Bialek, *J Stat Mech: Thy Exp* **2013**, P03011 (2013).
- [133] J. M. Beggs and D. Plenz, *J Neurosci* **23**, 11167 (2003).
- [134] J. M. Beggs, *Philosophical Trans R Soc London A: Math, Phys, Eng Sci* **366**, 329 (2008).
- [135] W. L. Shew and D. Plenz, *Neuroscientist* **19**, 88 (2013).
- [136] D. J. Schwab, I. Nemenman, and P. Mehta, *Phys Rev Lett* **113**, 068102 (2014).
- [137] A. de la Fuente, *Trends Genet* **26**, 326 (2010).
- [138] T. Ideker and N. J. Krogan, *Mol Syst Biol* **8**, 565 (2012).
- [139] M. Grechkin, B. A. Logsdon, A. J. Gentles, and S.-I. Lee, *PLOS Comp Biol* **12**, e1004888 (2016).
- [140] I. W. Taylor, R. Linding, D. Warde-Farley, Y. Liu, C. Pesquita, D. Faria, S. Bull, T. Pawson, Q. Morris, and J. L. Wrana, *Nature Biotechnol* **27**, 199 (2009).
- [141] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, *Proc Natl Acad Sci (USA)* **106**, 67 (2009).
- [142] D. S. Marks, L. J. Colwell, R. Sheridan, T. A. Hopf, A. Pagnani, R. Zecchina, and C. Sander, *PLOS ONE* **6**, e28766 (2011).
- [143] L. J. Colwell, M. P. Brenner, and A. W. Murray, *PLOS ONE* **9**, e107723 (2014).
- [144] W. L. Shew, H. Yang, T. Petermann, R. Roy, and D. Plenz, *J Neurosci* **29**, 15595 (2009).
- [145] G. Tkačik, T. Mora, O. Marre, D. Amodei, S. E. Palmer, M. J. Berry, and W. Bialek, *Proc Natl Acad Sci (USA)* **112**, 11508 (2015).
- [146] D. R. Chialvo, *Nature Phys* **6**, 744 (2010).
- [147] I. Mastromatteo and M. Marsili, *J Stat Mech: Thy Exp* **2011**, P10012 (2011).
- [148] T. Mora and W. Bialek, *J Stat Phys* **144**, 268 (2011).
- [149] J. M. Beggs and N. Timme, *Frontiers Physiol* **3**, 163 (2012).
- [150] A. Pinna, N. Soranzo, and A. de la Fuente, *PLOS ONE* **5**, e12912 (2010).
- [151] R. C. Jansen, *Nature Rev Genet* **4**, 145 (2003).
- [152] F. Emmert-Streib, *PeerJ* **1**, e10 (2013).
- [153] S. M. M. Ud-Dean and R. Gunawan, *Bioinf* **32**, 875 (2016).
- [154] F. Eduati, A. Corradin, B. Di Camillo, and G. Toffolo, *PLOS ONE* **5**, e12789 (2010).
- [155] <http://dreamchallenges.org/project/dream4-predictive-signaling-network-modeling>. The solution presented in [154] infers a network using Boolean truth tables, one of the most popular approaches during the early stages of automated network inference [351]. This approach has since fallen out of favor, yielding to the more sophisticated methods we

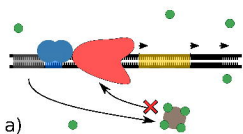
- discuss in Section III, but Bayesian networks are often still discretized to Boolean values for convenience.
- [156] A. Greenfield, A. Madar, H. Ostrer, and R. Bonneau, *PLOS ONE* **5**, e13397 (2010).
- [157] V. A. Huynh-Thu, A. Irrthum, L. Wehenkel, and P. Geurts, *PLOS ONE* **5**, e12776 (2010).
- [158] A. Wagner, *J Comp Biol* **11**, 53 (2004).
- [159] D. di Bernardo, M. J. Thompson, T. S. Gardner, S. E. Chobot, E. L. Eastwood, A. P. Wojtovich, S. J. Elliott, S. E. Schaus, and J. J. Collins, *Nature Biotechnol* **23**, 377 (2005).
- [160] M. Bansal, G. Della Gatta, and D. di Bernardo, *Bioinf* **22**, 815 (2006).
- [161] N. Friedman, M. Linial, I. Nachman, and D. Pe'er, *J Comp Biol* **7**, 601 (2000).
- [162] J. Schäfer, K. Strimmer, J. F. Mendes, S. Dorogovtsev, A. Povolotsky, F. Abreu, and J. Oliveira, in *AIP Conf Proc*, Vol. 776 (AIP, La Herradura, Spain, 2005) pp. 263–276.
- [163] A. J. Hartemink, D. K. Gifford, T. S. Jaakkola, and R. A. Young, in *Pac Symp Biocomp*, Vol. 6 (World Scientific Pub Co Inc, Stanford, CA, 2001) pp. 437–449.
- [164] C. J. Needham, J. R. Bradford, A. J. Bulpitt, and D. R. Westhead, *Nature Biotechnol* **24**, 51 (2006).
- [165] A. V. Werhli and D. Husmeier, *Stat Appl Genet Mol Biol* **6**, 15 (2007).
- [166] F. Geier, J. Timmer, and C. Fleck, *BMC Syst Biol* **1**, 11 (2007).
- [167] S. Mukherjee and T. P. Speed, *Proc Natl Acad Sci (USA)* **105**, 14313 (2008).
- [168] A. Greenfield, C. Hafemeister, and R. Bonneau, *Bioinf* **29**, 1060 (2013).
- [169] Y. Li and S. A. Jackson, *G3: Genes, Genomes, Genetics* **5**, 1075 (2015).
- [170] M. Ghanbari, J. Lasserre, and M. Vingron, *BMC Syst Biol* **9**, 84 (2015).
- [171] M. E. Studham, A. Tjärnberg, T. E. Nordling, S. Nelander, and E. L. L. Sonnhammer, *Bioinf* **30**, i130 (2014).
- [172] K.-I. Goh, M. E. Cusick, D. Valle, B. Childs, M. Vidal, and A.-L. Barabási, *Proc Natl Acad Sci (USA)* **104**, 8685 (2007).
- [173] J. Loscalzo and A.-L. Barabási, *Wiley Interdiscip Rev: Syst Biol Med* **3**, 619 (2011).
- [174] X. Zhou, J. Menche, A.-L. Barabási, and A. Sharma, *Nature Comm* **5** (2014).
- [175] S. Horvath, *Weighted network analysis: applications in genomics and systems biology* (Springer Science & Business Media, 2011).
- [176] J. J. Rice, Y. Tu, and G. Stolovitzky, *Bioinf* **21**, 765 (2005).
- [177] B. Zhang and S. Horvath, *Stat Appl Genet Mol Biol* **4**, 1128 (2005).
- [178] P. Khosravi, V. H. Gazestani, L. Pirhaji, B. Law, M. Sadeghi, B. Goliaei, and G. D. Bader, *Algorithms for Molecular Biology* **10**, 23 (2015).
- [179] C. Shannon, *Bell Syst Techn J* **27**, 379 (1948).
- [180] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, 2012).
- [181] A. Levchenko and I. Nemenman, *Curr Opin Biotechnol* **28**, 156 (2014).
- [182] P. Dhaeseleer, X. Wen, S. Fuhrman, and R. Somogyi, in *Information Processing in Cells and Tissues*, edited by M. Holcombe and R. Paton (1998) pp. 203–212.
- [183] A. J. Butte and I. S. Kohane, in *Pac Symp Biocomp* (World Scientific Pub Co Inc, Stanford, CA, 2000) pp. 418–429.
- [184] A. J. Butte, P. Tamayo, D. Slonim, T. R. Golub, and I. S. Kohane, *Proc Natl Acad Sci (USA)* **97**, 12182 (2000).
- [185] K. Wang, I. Nemenman, N. Banerjee, A. A. Margolin, and A. Califano, in *Res Comput Mol Biol (RECOMB)*, Lecture Notes in Computer Science No. 3909, edited by A. Apostolico, C. Guerra, S. Istrail, P. A. Pevzner, and M. Waterman (Springer, Berlin-Heidelberg, 2006) pp. 348–362.
- [186] K.-C. Liang and X. Wang, *EURASIP J Bioinf Syst Biol* **2008**, 253894 (2008).
- [187] X. Zhang, X.-M. Zhao, K. He, L. Lu, Y. Cao, J. Liu, J.-K. Hao, Z.-P. Liu, and L. Chen, *Bioinf* **28**, 98 (2012).
- [188] P. E. Meyer, K. Kontos, F. Lafitte, and G. Bontempi, *EURASIP J Bioinf Syst Biol* **2007**, 79879 (2007).
- [189] I. E. Ohiorhenuan, F. Mechler, K. P. Purpura, A. M. Schmid, Q. Hu, and J. D. Victor, *Nature* **466**, 617 (2010).
- [190] G. D. Field, J. L. Gauthier, A. Sher, M. Greschner, T. A. Machado, L. H. Jepson, J. Shlens, D. E. Gunning, K. Mathieson, W. Dabrowski, L. Paninski, A. M. Litke, and E. J. Chichilnisky, *Nature* **467**, 673 (2010).
- [191] G. Tkačik, O. Marre, D. Amodei, E. Schneidman, W. Bialek, and M. J. Berry, *PLOS Comp Biol* **10**, e1003408 (2014).
- [192] M. Bethge and P. Berens, in *NIPS 20* (MIT Press, Vancouver, Canada, 2007) pp. 97–104.
- [193] T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff, *Proc Natl Acad Sci (USA)* **103**, 19033 (2006).
- [194] J. W. Locasale and A. Wolf-Yadlin, *PLOS ONE* **4**, e6522 (2009).
- [195] F. Morcos, A. Pagnani, B. Lunt, A. Bertolino, D. S. Marks, C. Sander, R. Zecchina, J. N. Onuchic, T. Hwa, and M. Weigt, *Proc Natl Acad Sci (USA)* **108**, E1293 (2011).
- [196] D. T. Jones, D. W. Buchan, D. Cozzetto, and M. Pontil, *Bioinf* **28**, 184 (2012).
- [197] K. Wood, S. Nishida, E. Sontag, and P. Cluzel, *Proc Natl Acad Sci (USA)* **109**, 12254 (2012).
- [198] T. Mora, A. M. Walczak, W. Bialek, and C. G. Callan, *Proc Natl Acad Sci (USA)* **107**, 5405 (2010).
- [199] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, *Proc Natl Acad Sci (USA)* **109**, 4786 (2012).
- [200] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cogn Sci* **9**, 147 (1985).
- [201] Y. Roudi, J. Tyrcha, and J. Hertz, *Phys Rev E* **79**, 051915 (2009).
- [202] T. Hastie, R. Tibshirani, and J. Friedman, *Unsupervised Learning* (Springer, 2009).
- [203] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty, *Ann Stat* **38**, 1287 (2010).
- [204] S. Cocco and R. Monasson, *Phys Rev Lett* **106**, 090601 (2011).
- [205] S. Cocco and R. Monasson, *J Stat Phys* **147**, 252 (2012).
- [206] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, 2006).
- [207] H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A.-L. Barabási, *Nature* **407**, 651 (2000).
- [208] R. D. Leclerc, *Mol Syst Biol* **4**, 213 (2008).

- [209] J. Karvanen, arXiv preprint arXiv:1403.1124 (2014).
- [210] M. Nitzan, J. Casadiego, and M. Timme, *Sci Adv* **3**, e1600396 (2017).
- [211] J. Pearl, *Causality*, 2nd ed. (Cambridge University Press, 2009).
- [212] A. Wagner, *Bioinf* **17**, 1183 (2001).
- [213] E. J. Molinelli, A. Korkut, W. Wang, M. L. Miller, N. P. Gauthier, X. Jing, P. Kaushik, Q. He, G. Mills, D. B. Solit, C. A. Pratilas, M. Weigt, A. Braunstein, A. Pagnani, R. Zecchina, and C. Sander, *PLOS Comp Biol* **9** (2013).
- [214] J. Tegnér and J. Björkegren, *Trends in genetics: TIG* **23**, 34 (2007).
- [215] M. H. Maathuis, M. Kalisch, and P. Bühlmann, *Ann Stat* **37**, 3133 (2009).
- [216] M. H. Maathuis, D. Colombo, M. Kalisch, and P. Bühlmann, *Nature Meth* **7**, 247 (2010).
- [217] Once the equivalence class is determined, formal causality detection methods can be subsequently applied to estimate the full causal graph. We refer curious readers to [211, 352] for a wealth of both philosophical explanations and more rigorous treatments of algorithms designed explicitly to detect causality in its many guises.
- [218] J. Tegnér, M. K. S. Yeung, J. Hasty, and J. J. Collins, *Proc Natl Acad Sci (USA)* **100**, 5944 (2003).
- [219] E. Sontag, A. Kiyatkin, and B. N. Kholodenko, *Bioinf* **20**, 1877 (2004).
- [220] D. Djordjevic, A. Yang, A. Zadorian, K. Rungruechareon, and J. W. K. Ho, *PLOS ONE* **9**, e111661 (2014).
- [221] A. Arkin, P. Shen, and J. Ross, *Science* **277**, 1275 (1997).
- [222] P. Lecca, D. Morpurgo, G. Fantaccini, A. Casagrande, and C. Priami, *BMC Syst Biol* **6**, 51 (2012).
- [223] L. Ljung, *System identification* (Wiley Online Library, 1999).
- [224] M. Schmidt and H. Lipson, *Science* **324**, 81 (2009).
- [225] T. Chen, H. L. He, and G. M. Church, in *Pac Symp Biocomp* (World Scientific Pub Co Inc, Stanford, CA, 1999) pp. 29–40.
- [226] P. D’haeseleer, X. Wen, S. Fuhrman, and R. Somogyi, in *Pac Symp Biocomp* (World Scientific Pub Co Inc, Stanford, CA, 1999) pp. 41–52.
- [227] D. C. Weaver, C. T. Workman, and G. D. Stormo, in *Pac Symp Biocomp* (World Scientific Pub Co Inc, Stanford, CA, 1999) pp. 112–123.
- [228] E. Mjolsness, T. Mann, R. Castao, and B. J. Wold, in *Adv Neural Inf Proc Syst 12*, edited by S. A. Solla, T. K. Leen, and K. Mller (MIT Press, 2000) pp. 928–934.
- [229] E. P. van Someren, L. F. Wessels, and M. J. Reinders, in *Ismb* (Eighth International Conference on Intelligent Systems for Molecular Biology (ISMB-2000), San Diego, CA, 2000) pp. 355–366.
- [230] M. Wahde and J. Hertz, *Biosyst* **55** (2000).
- [231] M. Grimaldi, R. Visintainer, and G. Jurman, *PLOS ONE* **6**, e28646 (2011).
- [232] E. Mjolsness, D. H. Sharp, and J. Reinitz, *J Theor Biol* **152**, 429 (1991).
- [233] A. Crombach, K. R. Wotton, D. Cicin-Sain, M. Ashyraliyev, and J. Jaeger, *PLOS Comp Biol* **8**, e1002589 (2012).
- [234] C. J. Oates, F. Dondelinger, N. Bayani, J. Korkola, J. W. Gray, and S. Mukherjee, *Bioinf* **30**, i468 (2014).
- [235] N. M. Mangano, S. L. Brunton, J. L. Proctor, and J. N. Kutz, *IEEE Trans Mol Biol Multi-Scale Commun* **2**, 52 (2016).
- [236] M. Wibrál, J. T. Lizier, and V. Priesemann, *Frontiers Robotics AI* **2** (2015).
- [237] M. Luo, H. Kantz, N.-C. Lau, W. Huang, and Y. Zhou, *Proc Natl Acad Sci (USA)*, 201510571 (2015).
- [238] N. F. Rulkov, M. M. Sushchik, L. S. Tsimring, and H. D. Abarbanel, *Phys Rev E* **51**, 15 (1995).
- [239] F. Takens, in *Symposium on Dynamical Systems and Turbulence*, edited by D. A. Rand and L. S. Young (Springer, 1981) pp. 366–381.
- [240] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, 2nd ed. (Westview Press, Boulder, CO, 2015).
- [241] M. Ragwitz and H. Kantz, *Phys Rev E* **65**, 12 (2002).
- [242] M. Small and C. K. Tse, *Physica D: Nonlinear Phenomena* **194**, 283 (2004).
- [243] L. Faes, G. Nollo, and A. Porta, *Computers Biol Med Computing complexity in cardiovascular oscillations Selected Papers from the 6th ESGCO conference*, **42**, 290 (2012).
- [244] J. M. McCracken and R. S. Weigel, *Phys Rev E* **90**, 062903 (2014).
- [245] S. Cobey and E. B. Baskerville, *PLOS ONE* **11**, e0169050 (2016).
- [246] D. Mønster, R. Fusaroli, K. Tylén, A. Roepstorff, and J. F. Sherson, in *Proceedings of the 1st International Conference on Complex Information Systems (COMPLEXIS 2016)*, edited by V. Muñoz, O. Gusikhin, and V. Chang (SCITEPRESS, Funchal, Madeira, Portugal, 2016) pp. 48–56.
- [247] J.-J. Jiang, Z.-G. Huang, L. Huang, H. Liu, and Y.-C. Lai, *Scientific Rep* **6** (2016).
- [248] C. Wiggins and I. Nemenman, *Exp Mech* **43**, 361 (2003).
- [249] R. Opgen-Rhein and K. Strimmer, *BMC Bioinf* **8**, S3 (2007).
- [250] C. W. J. Granger, *Econometrica* **37**, 424 (1969).
- [251] N. Wiener, *Modern Math for Engineers* **1**, 125 (1956).
- [252] C. W. J. Granger, *J Econometrics* **39**, 199 (1988).
- [253] W. A. Freiwald, P. Valdes, J. Bosch, R. Biscay, J. C. Jimenez, L. M. Rodriguez, V. Rodriguez, A. K. Kreiter, and W. Singer, *J Neurosci Meth* **94**, 105 (1999).
- [254] N. Ancona, D. Marinazzo, and S. Stramaglia, *Phys Rev E* **70**, 056221 (2004).
- [255] Y. Chen, G. Rangarajan, J. Feng, and M. Ding, *Phys Lett A* **324**, 26 (2004).
- [256] B. Gourévitch, R. L. Bouquin-Jeannès, and G. Faucon, *Biological Cybern* **95**, 349 (2006).
- [257] D. Marinazzo, M. Pellicoro, and S. Stramaglia, *Phys Rev Lett* **100**, 144103 (2008).
- [258] M. Ding, Y. Chen, and S. L. Bressler, in *Handbook of Time Series Analysis: Recent Theoretical Developments and Applications* (Wiley, Weinheim, 2006).
- [259] K. J. Blinowska, R. Kuś, and M. Kamiński, *Phys Rev E* **70**, 050902 (2004).
- [260] J. Pearl, in *Proceedings, Cognitive Science Society, UC Irvine* (UCLA Computer Science Department Technical Report 850021 (R-43), Los Angeles, CA, 1985) pp. 329–334.
- [261] E. Charniak, *AI Magazine* **12**, 50 (1991).
- [262] Z. Ghahramani, *Int J Pattern Recogn Artificial Intelligence* **15**, 9 (2001).
- [263] N. Friedman, *Science* **303**, 799 (2004).
- [264] N. Friedman, K. Murphy, and S. Russell, in *Proc Four-*

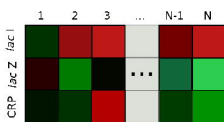
- teenth Conf Uncertainty in Artificial Intelligence (UAI-98)*, edited by G. F. C. Cooper and S. Moral (Morgan Kaufmann Publishers Inc., San Francisco, CA, 1998) pp. 139–147.
- [265] V. A. Smith, E. D. Jarvis, and A. J. Hartemink, *Bioinf* **18**, S216 (2002).
- [266] I. Nachman, A. Regev, and N. Friedman, *Bioinf* **20**, i248 (2004).
- [267] E. Acerbi, T. Zelante, V. Narang, and F. Stella, *BMC Bioinf* **15**, 387 (2014).
- [268] C. Zou and J. Feng, *BMC Bioinf* **10**, 122 (2009).
- [269] J. Zhu, B. Zhang, E. N. Smith, B. Drees, R. B. Brem, L. Kruglyak, R. E. Bumgarner, and E. E. Schadt, *Nature Gen* **40**, 854 (2008).
- [270] I. Pournara and L. Wernisch, *Bioinf* **20**, 2934 (2004).
- [271] N. Friedman and D. Koller, in *Proc Sixteenth Conf Uncertainty in Artificial Intelligence (UAI-00)*, edited by C. Boutilier and M. Goldszmidt (Morgan Kaufmann Publishers Inc., San Francisco, CA, 2000) pp. 201–210.
- [272] N. Ay and D. Polani, *Adv Compl Syst* **11**, 17 (2008).
- [273] T. Schreiber, *Phys Rev Lett* **85**, 461 (2000).
- [274] M. Paluš, V. Komárek, Z. Hrnčíř, and K. Štěrbová, *Phys Rev E* **63**, 6 (2001).
- [275] C. J. Honey, R. Kttr, M. Breakspear, and O. Sporns, *Proc Natl Acad Sci (USA)* **104**, 10240 (2007).
- [276] O. Stetter, D. Battaglia, J. Soriano, and T. Geisel, *PLOS Comp Biol* **8**, 1 (2012).
- [277] M. Wibral, R. Vicente, and M. Lindner, in *Directed Information Measures in Neuroscience*, Understanding Complex Systems, edited by M. Wibral, R. Vicente, and J. T. Lizier (Springer Berlin Heidelberg, 2014) pp. 3–36.
- [278] L. Faes, G. Nollo, and A. Porta, *Entropy* **15**, 198 (2013).
- [279] B. Pompe and J. Runge, *Phys Rev E* **83**, 051122 (2011).
- [280] J. Runge, J. Heitzig, V. Petoukhov, and J. Kurths, *Phys Rev Lett* **108**, 5 (2012).
- [281] J. Runge, J. Heitzig, N. Marwan, and J. Kurths, *Phys Rev E* **86**, 061121 (2012).
- [282] O. Kwon and J.-S. Yang, *EPL (Europhys Lett)* **82**, 68003 (2008).
- [283] J. Kim, G. Kim, S. An, Y.-K. Kwon, and S. Yoon, *PLOS ONE* **8**, e51986 (2013).
- [284] N. Bertschinger, J. Rauh, E. Olbrich, J. Jost, and N. Ay, *Entropy* **16**, 2161 (2014).
- [285] I. Nemenman, F. Shafee, and W. Bialek, in *Advances in Neural Information Processing Systems 14*, edited by T. G. Dietterich, S. Becker, and Z. Ghahramani (MIT Press, 2002).
- [286] I. Nemenman, *Entropy* **13**, 2013 (2011).
- [287] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning*, 2nd ed. (Springer series in statistics, Springer Berlin, 2001).
- [288] A. Kraskov, H. Stögbauer, and P. Grassberger, *Phys Rev E* **69**, 066138 (2004).
- [289] D. W. Hahs and S. D. Pethel, *Phys Rev Lett* **107**, 128701 (2011).
- [290] M. Lindner, R. Vicente, V. Priesemann, and M. Wibral, *BMC Neuroscience* **12**, 119 (2011).
- [291] M. Wibral, N. Pampu, V. Priesemann, F. Siebenhner, H. Seiwert, M. Lindner, J. T. Lizier, and R. Vicente, *PLOS ONE* **8**, 19 (2013).
- [292] L. Barnett, A. B. Barrett, and A. K. Seth, *Phys Rev Lett* **103**, 4 (2009).
- [293] J. Massey, in *Proc. Int. Symp. Inf. Theory Applic. (ISITA-90)* (1990) pp. 303–305.
- [294] J. T. Lizier and M. Prokopenko, *Eur Physical J B* **73**, 605 (2010).
- [295] P. Dhaeseleer, S. Liang, and R. Somogyi, *Bioinf* **16**, 707 (2000).
- [296] L. F. Wessels and M. J. Reinders, in *Pac Symp Biocomp*, Vol. 6 (World Scientific Pub Co Inc, Stanford, CA, 2001) pp. 508–519.
- [297] P. Brazhnik, A. de la Fuente, and P. Mendes, *Trends Biotechn* **20**, 467 (2002).
- [298] H. De Jong, *J Comp Biol* **9**, 67 (2002).
- [299] E. P. van Someren, L. F. Wessels, E. Backer, and M. J. Reinders, *Pharmacogenomics* **3**, 507 (2002).
- [300] R. Albert, in *Complex networks* (Springer Berlin Heidelberg, 2004) pp. 459–481.
- [301] N. A. van Riel, *Brief Bioinf* **7**, 364 (2006).
- [302] R. Albert, *The Plant Cell* **19**, 3327 (2007).
- [303] K.-H. Cho, S.-M. Choo, S. Jung, J.-R. Kim, H.-S. Choi, and J. Kim, *IET Syst Biol* **1**, 149 (2007).
- [304] J. Goutsias and N. Lee, *Curr Pharmaceutical Design* **13**, 1415 (2007).
- [305] F. Markowetz and R. Spang, *BMC Bioinf* **8** (2007).
- [306] L. Kaderali and N. Radde, in *Computational Intelligence in Bioinformatics* (Springer, 2008) pp. 33–74.
- [307] G. Karlebach and R. Shamir, *Nature Rev Mol Cell Biol* **9**, 770 (2008).
- [308] R. Bonneau, *Nature Chemical Biology* **4**, 658 (2008).
- [309] W.-P. Lee and W.-S. Tzou, *Brief Bioinf* **10**, 408 (2009).
- [310] I. Shmulevich and J. D. Aitchison, *Meth Enzymol* **467**, 335 (2009).
- [311] C. Sima, J. Hua, and S. Jung, *Curr Genomics* **10**, 416 (2009).
- [312] A. Ay and D. N. Arnosti, *Critical Rev Biochem Mol Biol* **46**, 137 (2011).
- [313] C. A. Penfold and D. L. Wild, *Interface Focus* **1**, 857 (2011).
- [314] D. F. Veiga, B. Dutta, and G. Balázs, *Mol bioSyst* **6**, 469 (2010).
- [315] E. Hernández-Lemus and J. M. Siqueiros-García, *Complex Adaptive Systems Modeling* **1**, 8 (2013).
- [316] B. Ristevski, *Nonlinear Anal Model Control* **18**, 444 (2013).
- [317] N. Vijesh, S. K. Chakrabarti, and J. Sreekumar, *J Biomed Sci Eng* **6**, 223 (2013).
- [318] A. F. Villaverde, J. Ross, and J. R. Banga, *Cells* **2**, 306 (2013).
- [319] F. Emmert-Streib, M. Dehmer, and B. Haibe-Kains, *Frontiers Cell Devel Biol* **2**, 38 (2014).
- [320] Y. X. R. Wang and H. Huang, *Journal of Theoretical Biology Network-based biomarkers for complex diseases*, **362**, 53 (2014).
- [321] X. Dong, A. Yambartsev, S. A. Ramsey, L. D. Thomas, N. Shulzhenko, and A. Morgun, *Bioinf and Biol Insights* **9**, 61 (2015).
- [322] C. Liseron-Monfils and D. Ware, *Curr Plant Biol* **3**, 30 (2015).
- [323] D. Marbach, R. J. Prill, T. Schaffter, C. Mattiussi, D. Floreano, and G. Stolovitzky, *Proc Natl Acad Sci (USA)* **107**, 6286 (2010).
- [324] C. J. Oates and S. Mukherjee, *Ann Stat* **6**, 1209 (2012).
- [325] A. Noor, E. Serpedin, M. Nounou, H. Nounou, N. Mohamed, and L. Chouchane, *Advances in Bioinformatics* **2013**, e953814 (2013).

- [326] J. Omony, *Ann Res Rev Biol* **4**, 577 (2013).
- [327] S. A. Thomas and Y. Jin, *Evolutionary Intelligence* **7**, 29 (2014).
- [328] R. Bonneau and T. Aijo, *bioRxiv*, 051847 (2016).
- [329] D. M. Budden and E. J. Crampin, *BMC Syst Biol* **10**, 89 (2016).
- [330] C. Siegenthaler and R. Gunawan, *PLOS ONE* **9**, e90481 (2014).
- [331] D. J. Hand, *Machine Learning* **77**, 103 (2009).
- [332] S. R. Maetschke, P. B. Madhamshettiwar, M. J. Davis, and M. A. Ragan, *Brief Bioinf* **15**, 195 (2013).
- [333] L. Goodman, *Genome Res* **9**, 673 (1999).
- [334] F. Zhu and Y. Guan, *Bioinf* **30**, 2772 (2014).
- [335] D. P. Noren, B. L. Long, R. Norel, K. Rhissorakrai, K. Hess, C. W. Hu, A. J. Bisberg, A. Schultz, E. Engquist, L. Liu, X. Lin, G. M. Chen, H. Xie, G. A. M. Hunter, P. C. Boutros, O. Stepanov, DREAM 9 AML-OPC Consortium, T. Norman, S. H. Friend, G. Stolovitzky, S. Kornblau, and A. A. Qutub, *PLOS Comp Biol* **12**, e1004890 (2016).
- [336] A. H. Lang, H. Li, J. J. Collins, and P. Mehta, *PLOS Comp Biol* **10**, e1003734 (2014).
- [337] K. G. Wilson, *Scient Amer* **241**, 158 (1979).
- [338] B. B. Machta, R. Chachra, M. K. Transtrum, and J. P. Sethna, *Science* **342**, 604 (2013).
- [339] P. Mehta and D. J. Schwab, *arXiv preprint arXiv:1410.3831* (2014).
- [340] M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna, *J Chem Phys* **143**, 07B201_1 (2015).
- [341] S. Bradde and W. Bialek, *arXiv preprint arXiv:1610.09733* (2016).
- [342] A. Joshi, R. D. Smet, K. Marchal, Y. V. d. Peer, and T. Michoel, *Bioinf* **25**, 490 (2009).
- [343] T. Michoel, R. De Smet, A. Joshi, Y. Van de Peer, and K. Marchal, *BMC Syst Biol* **3**, 49 (2009).
- [344] E. Bonnet, M. Tatari, A. Joshi, T. Michoel, K. Marchal, G. Berx, and Y. V. d. Peer, *PLOS ONE* **5**, e10162 (2010).
- [345] X. Li, C. Panea, C. H. Wiggins, V. Reinke, and C. Leslie, *PLOS Comp Biol* **6**, e1000761 (2010).
- [346] J. B. Birks, *Rutherford at Manchester* (Heywood, London, 1962).
- [347] W. Bialek, *Biophysics: Searching for principles* (Princeton University Press, Princeton, 2012).
- [348] J. Surowiecki, *The wisdom of crowds* (Knopf Doubleday Publishing Group, 2005).
- [349] D. Marbach, J. C. Costello, R. Küffner, N. M. Vega, R. J. Prill, D. M. Camacho, K. R. Allison, M. Kellis, J. J. Collins, G. Stolovitzky, *et al.*, *Nature Meth* **9**, 796 (2012).
- [350] J. Saez-Rodriguez, J. C. Costello, S. H. Friend, M. R. Kellen, L. Mangravite, P. Meyer, T. Norman, and G. Stolovitzky, *Nature Rev Genet* **17**, 470 (2016).
- [351] S. Liang, S. Fuhrman, and R. Somogyi, (1998).
- [352] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search* (Springer Science & Business Media, 2012).

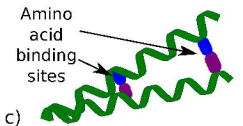
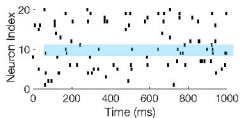
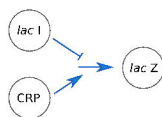
System illustration



Typical Activation Data



Network Representation



... 10 20 30 40 50 ..

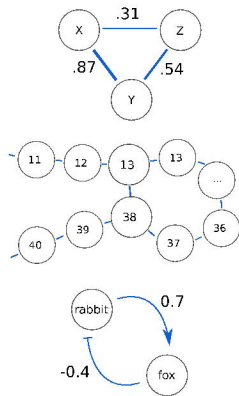
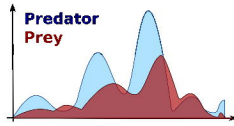
ETDQYRKQVMDIGSAMIRDQYMRGTGEGFLC

MLDFYRKQVVICGYSFMIFDQYMRGTGEGFLC

EMDLTRKCAMICTGLMLDRMRRTGEGFLCVFI

MLQFYEPYAVIGIEVQITIDODTAGEDYAAR

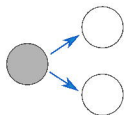
LLQFRTTKTIPIQGQSVKLAQMIQIWE DTAGSE



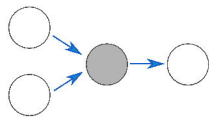
Network Motifs



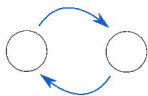
a) Directed link



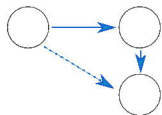
c) Common drive



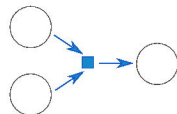
e) Hidden intermediate



b) Bidirectional link



d) Spurious relation



f) Higher-order (XOR)