Network Structure Inference, A Survey: Motivations, Methods, and Applications

IVAN BRUGERE, University of Illinois at Chicago, USA

BRIAN GALLAGHER, Lawrence Livermore National Laboratory

TANYA Y. BERGER-WOLF, University of Illinois at Chicago

Networks represent relationships between entities in many complex systems, spanning from online social interactions to biological cell development and brain connectivity. In many cases, relationships between entities are unambiguously known: are two users “friends” in a social network? Do two researchers collaborate on a published article? Do two road segments in a transportation system intersect? These are directly observable in the system in question. In most cases, relationships between nodes are not directly observable and must be inferred: Does one gene regulate the expression of another? Do two animals who physically co-locate have a social bond? Who infected whom in a disease outbreak in a population?

Existing approaches for inferring networks from data are found across many application domains and use specialized knowledge to infer and measure the quality of inferred network for a specific task or hypothesis. However, current research lacks a rigorous methodology that employs standard statistical validation on inferred models. In this survey, we examine (1) how network representations are constructed from underlying data, (2) the variety of questions and tasks on these representations over several domains, and (3) validation strategies for measuring the inferred network’s capability of answering questions on the system of interest.

CCS Concepts: • **Information systems** → **Network data models**; **Data mining**; *Temporal data*; • **Computing methodologies** → **Network science**; Unsupervised learning; Learning in probabilistic graphical models;

Additional Key Words and Phrases: Graph mining, relational learning, computational science

ACM Reference format:

Ivan Brugere, Brian Gallagher, and Tanya Y. Berger-Wolf. 2018. Network Structure Inference, A Survey: Motivations, Methods, and Applications. *ACM Comput. Surv.* 51, 2, Article 24 (April 2018), 39 pages.

<https://doi.org/10.1145/3154524>

1 INTRODUCTION

Networks represent relationships between entities in many complex systems, spanning from online social interactions to biological cell development and brain connectivity. In many cases, relationships between entities are unambiguously known: Are two users “friends” in a social

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344 (LLNL-JRNL-703477), and with support of National Science Foundation Grants No. III-1514126 (Berger-Wolf), No. CNS-1248080 (Berger-Wolf), and No. IGERT CNS-1069311 (Brugere).

Authors’ addresses: I. Brugere, University of Illinois at Chicago, 1200 W Harrison St, Chicago, IL 60607; email: ibruge2@uic.edu; B. Gallagher, Lawrence Livermore National Laboratory, 7000 East Ave, Livermore, CA 94550; T. Y. Berger-Wolf, University of Illinois at Chicago, 1200 W Harrison St, Chicago, IL 60607; email: tanyabw@uic.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 0360-0300/2018/04-ART24 \$15.00

<https://doi.org/10.1145/3154524>

network? Do two researchers collaborate on a published article? Do two road segments in a transportation system intersect? These are directly observable in the system in question [104]. In *most* cases, however, relationships between entities are not directly observable and must be inferred: Does one gene regulate the expression of another? Do two animals who physically co-locate have a social bond? Who infected whom in a disease outbreak in a population?

Networks are mathematical representations (i.e., models) used to answer these types of questions about data collected on individual entities. There is a broad range of the questions over application domains and a variety of ways in which networks are used to answer these questions. How do we know whether a particular network on data is the most *useful* representation for answering a given question? What is the “right” network representation and how do we *compare* the utility of many possible representations for our particular question? Finally, how can we measure whether networks are the *correct* models to answer a particular question on the system of interest?

Various approaches that use specialized knowledge to infer and measure the quality of inferred network for a specific task or hypothesis are found across many application domains. The rigor of these methodologies varies across domains and currently there are no common best-practices for evaluating the quality of networks inferred from data. Such practices would measure some combination of general statistical properties, including model significance-testing and uniqueness, sensitivity (i.e., the change in quality relative to underlying data measurements and model parameters), and generalizability (i.e., network quality measured over multiple hypotheses or tasks).

In this survey, we examine: (1) how network representations are constructed from underlying data; (2) the variety of questions and tasks on these representations over several domains; and (3) validation strategies for measuring the inferred network’s capability for answering questions on the system of interest.

1.1 Motivation: Networks Represent Complex Relationships

Networks are a natural choice of data representation across many domains. First, networks by design represent *higher-order structure* emerging from dyadic relationships, which serve as units of further analysis. These structures include neighborhoods, ego-nets, communities/modules, and connected components. For example, in the computational biology domain, clusters and motifs often represent shared biological function. The mutual connectivity provides stronger evidence for the common function than individual pairwise relationships. Individual, dyadic, neighborhood, or aggregate population analysis may be the most appropriate depending on the question of interest.

Second, networks represent heterogeneity among entities in the local network topology. Rather than analysis on population aggregates (e.g., histograms), networks allow local querying of complex, non-metric attribute and feature spaces. *Homophily* yields correlation in node attributes among entities close in the network. However, these correlations also tend to be non-monotonic: the *most similar* node to a query entity may be arbitrarily distant in the network. Due to this correlation, network topology often represents local clusters as overlapping, heterogeneous relationships. For example, a user’s “friends” in an online social network often cluster into functional units: friends from work, school, from the user’s hometown, and so on, where nodes within each cluster are correlated in some—often unknown—attribute. The effectiveness of simple heuristics such as counting common neighbors in the link prediction problem [117] shows the latent information within social networks.

Third, networks are *interpretable* models for further analysis and hypothesis generation. Researchers can visually explore small networks and examine relationships between nodes to compare against their knowledge-base and intuition from the domain. Furthermore, a shared vocabulary of descriptive network measures enables researchers to compare networks according to density, degree distribution, clustering coefficient, centralities, diameter, average path length,

triangle counts [90, 181], and graphlet distributions [142]. Many higher-level network measures have also recently been developed including robustness [37, 144], local information efficiency [7], and routing efficiency [114, 189]. Using these shared measures, researchers can compare properties across classes of networks and generate hypotheses explaining similarities or differences.

Finally, networks are *common* models for data and can be re-used in multiple studies. The breadth of tools and support for network analysis allows researchers of various disciplines to apply sophisticated off-the-shelf analysis and visualization techniques, as well as easier storage, querying, and portability in graph databases. Researchers also have a common vernacular developed in the area of “network science,” across biology, physics, and computer science [24].

A cautionary note, however: when inferring networks from data, researchers ought to consider whether higher-order structures are *meaningful* and which descriptive measures are appropriate on the inferred network. Available tools and convenience can motivate researchers to translate their problem into a network formulation, whether or not a network is the best model for the question of interest. This survey aims to clarify *when* network models are appropriate and how they are inferred.

1.2 Inferring Network Models from Data

Inferred networks are a class of networks where the node and/or edge definitions are inferred from non-explicitly relational data. Work in network science often focuses on applications of *explicit networks*: network representations where the meaning of nodes and edges is unambiguous and categorical. For example, in the Facebook network, two adjacent users are categorically “friends.” In weighted networks, such as a road transportation network, nodes represent road intersections and edges represent the road segments between them. These edges are unambiguous, change with low frequency, and have strict spatial constraints for new edge appearance. Measuring weights (e.g., travel time) is a matter of sensing/measuring traffic over the network, but these measurements are constrained to adjacent nodes on the known topology. In general, estimating weights in explicit network applications is a matter of accurately measuring a *known* relationship on a *known* topology, rather than learning a *hypothesized* relationship on an *unknown* topology.

Seemingly explicit networks often also have biases and arbitrary choices (e.g., thresholding, see: Section 2.5.1) in the collection or construction of the network. These choices are out of the control of researchers using these networks. Starting from underlying data and translating data to a network model enables researchers to measure and account for these biases to ensure the network model actually represents the appropriate characteristics of the system.

1.3 Challenges

Inferring networks from data presents several challenges:

- **Noisy measurements:** The underlying relationships of interest within real datasets are often noisy and confounded by overlapping relationships at varying scales (e.g., temporal, spatial). Furthermore, such errors are propagated when measuring higher-order properties of the inferred network (e.g., path lengths, triangle counts, network centralities, communities).
- **Lack of ground truth and/or model assumptions:** Determining whether a particular method accurately encodes the relationship of interest in the network requires: (1) ground-truth data, (2) model assumptions (e.g., Exponential Random Graph Model), or (3) some stability assumption (e.g., predictability over time). In many instances, no such assumptions or data are available and researchers are left with an under-determined problem and ad-hoc tuning of network model parameters [161].

- **Large parameter spaces and sensitivity:** Many network inference methods have thresholds or other parameters that define a large space of possible networks. Often this leads to ad-hoc selection or parameter-space sampling (see Section 2.5.1). Although sensitivity analysis is routine for method parameters of the subsequent task on the network (e.g., prediction, classification), most evaluation methodologies do not incorporate inferred network structure(s) into the sensitivity analysis.
- **Varying model complexity:** Descriptive interpretation of edges, paths, and communities is increasingly challenging with additional model complexity. For example, a social network defined via thresholding on who-calls-whom call record data is more interpretable than a linear regression on node feature vectors. Determining the validity and interpretability of these higher-order features is crucial to the many analyses that use them.

1.4 Outline of This Survey

This survey organizes recent work focusing on inferring networks from data. We draw from several domains where these methods are applied and general method development from machine learning. The study of inferring networks from data is largely unstructured. In Section 2, we first organize the work under a common problem formulation, taxonomy, and terminology. In Section 3, we examine various application domains to highlight unique challenges found in these areas.

1.5 Audience

Our survey is best suited for domain researchers using network inference in applications, as well as machine-learning researchers interested in methods and methodology on networks. For novices in the networks space, Section 1.6 structures many introductory surveys related to our work.

Our survey is also relevant to the broader network science community working on traditional *explicit* networks. These networks often represent an aggregation of longitudinal *interactions*, which are not informative to explain the *current* behavior of the population. Closer scrutiny of edges in the network “as given” may yield more appropriate representations and affect all aspects of downstream analysis [29]. Researchers routinely construct networks from “behavior” (i.e., node features) or other data to evaluate against their explicit network. Much of the work in scientific domains (including biology, ecology, chemistry) learn dynamic *interaction* networks in the absence of ground truth. This work has developed varied strategies to evaluate the quality of networks, in general, for prediction, classification, or the discovery of robust edge relationships over time or question(s) of interest.

1.6 Meta-Review: Comparison to Existing Surveys

In this section, we provide a brief meta-review of existing surveys and work that is related but distinct from network structure inference. See Section 2.3 for how these areas relate to our proposed taxonomy of related work within our problem.

Figure 1 provides a map of research in related but distinct network science problems. Kolaczyk [103] is the most relevant work to our survey. The author organizes related work in three categories: (1) the link prediction problem, where some edges and all nodes are known and the task is to infer new edges, (2) interaction networks where all nodes are known and the task is to infer edge relationships (e.g., by correlation), and (3) network tomography where some edges and nodes are known, and the task is to infer unobserved “interior” nodes and edges [76, 133, 206]. Of these, category (2) is primarily within the scope of what we define as the network structure inference; we focus with much greater depth on the network inference methodology and applications, where no edge definition is known *a priori*.

Fig. 1. An overview of related areas in machine learning and network science for this survey, several surveyed domains, and principle surveys and introductory work for these sub-areas.

We define network structure inference as distinct from the large body of work in relational learning. One branch contains attribute inference and prediction *on* networks. Given a network, these methods infer missing attributes using local estimates [159, 195], or link prediction: predicting edges at a later time step or by node attribute similarity [79, 117, 118]. These are two fundamental tasks that serve as evaluation methodologies for many network inference methods.

A second branch in relational learning focuses on inferring probabilistic relational models from data [65, 97]. While graphical models are one strategy applied to our problem, generally these models treat attributes or variables as nodes. Previous work in probabilistic relational models learns relationships between *explicit* input and output graphs, using node attribute inference, entity resolution, and link prediction tasks in a *supervised* setting [130].

Other recent surveys cover broad statistical network modeling [66] and multi-layer networks [22, 98]. Our survey draws on parametric network models, which is one class of model inferring parameters according to a particular distribution (e.g., attribute-edge joint distributions [137]). Research in network synthesis on multi-layer networks is one case of network structure inference for a particular task, where input data is a collection of networks.

Many network structure inference applications define edges by association measures over time (e.g., correlation in time series), so dynamics and temporal networks are an important aspect of network models for prediction tasks [2, 87, 88]. Finally, “network representation learning” is typically used to describe graph embedding methods. This is a rapidly growing area that aims to learn a low-dimensional representation of the network structure of a *given* network [40, 70, 78], which preserves desired structural properties. These models are used for comparison of nodes or networks, or predictive tasks in the embedded representation.

Our survey draws on several application areas. For a more exhaustive and general survey of complex networks in different application domains, da Fontoura Costa et al. [41], Havlin et al. [81] provide more breadth. However, we focus specifically on comparing motivations, evaluation methodologies, and challenges in several of these areas. Recent domain-focused meta-studies [120] and surveys in computational biology [83, 116, 162, 197], ecology [60, 106, 139, 141], neuroscience [19, 31, 152, 163], political science [112], and epidemiology [43, 95, 190] all have significant discussion of network topology inference specific to the domain. However, few have

network structure inference as a methodological focus and are limited to discussion of the single domain. Our survey focuses on challenges *across* each of these areas. We provide value to domain researchers both within and across fields, as well as to researchers in machine learning interested in predictive model development on networks.

2 PROBLEM DESCRIPTION

We give a high-level formulation of network structure inference to clarify the methodologies often encountered for this problem. This definition is very broad and not meant to be novel or exhaustive. We use this formalization for added precision of defining the constituent elements we will discuss.

2.1 Preliminaries

We define a network $G = \langle \mathbf{V}, \mathbf{E}, \mathbf{A} \rangle$ as a tuple containing a set $\mathbf{V}$ of n nodes, $v_i \in \mathbf{V}$, a set $\mathbf{E}$ of m node pairs, $e_{ij} \in \mathbf{E}$, and set $\mathbf{A}$ containing node and/or edge attribute sets. A particular attribute, the weight of an edge w_{ij} is a scalar value, $|w_{ij}| \leq 1$, where $w_{ij} = 0$ denotes the absence of an edge. An unweighted network is a special case of a weighted network where $w_{ij} \in \{0, 1\}$. Edge and node *features* are a particular type of attribute, derived by a function measuring some local edge or node property (e.g., node degree). Time-varying networks are defined as a t -length sequence of static network snapshots: $\mathbf{G} = (\mathbf{G}_1, \dots, \mathbf{G}_k, \dots, \mathbf{G}_t)$, with time-varying attributes and/or edges.¹

2.2 Data Science Motivations for Network Structure Inference

We organize network structure inference through the perspective of hypothesis-driven data science [35, 84]. Under this perspective, the value of a network can be stated simply: For a question of interest on the underlying system and its relevant data, are networks an informative and useful *data model* for *better* answering the question?

Figure 2 gives an overview of a general data science methodology and introduces all the relevant terms to our taxonomy. We start with a broad question of interest and try to (1) locate or collect the data relevant to computationally modeling our question. There are several possible data sources, collection processes, and modalities for analysis, coupled with several possible representations (e.g., networks, tensors). Our choice of data representation informs and constrains our hypotheses about the question of interest. (2) Domain science (orange, top) and machine learning (purple, bottom) generate complimentary analyses for hypothesis generation. (3) Novel computational methods are developed to test these hypotheses. These tested models serve as a solution approaching an answer, (4) closing the gap in understanding for many questions in novel ways. For example, data-driven science has developed novel, large-scale methodologies relevant across scientific domains, including experimental design [8, 74, 102] and randomization techniques [56, 99].

Figure 2 highlights several levels of modeling that can impact the final capability for answering a question of interest. Often different individuals will be responsible for generating and collecting the underlying data, inferring or defining the network, and developing the predictive models. For example, machine-learning researchers will rarely control the underlying sampling rate of specialized data collection processes in bioinformatics, neuroscience, or climate science. One novel approach in bioinformatics uses feedback from the network with active learning to guide the often costly and labor intensive data collection and experimentation process [177].

In the machine-learning community, there is a great deal of sensitivity analysis for the subsequent downstream predictive models on a network, without testing the parameters of the

¹While there are other representations of time-varying (or temporal, dynamic) networks [88], the time series of static networks is by far the most common.

Fig. 2. A common methodology for data science, with a focus on networks as models of underlying data. A researcher typically (1) uses *relevant* data and *appropriate* data models to formalize a particular question. (2) Traditional domain science (top, orange) typically combines direct observation and simple statistical models. Data-driven science with machine learning (bottom, purple) augments this intuition, hypotheses, and statistical models with sophisticated features and patterns, yielding additional challenges for visualization, hypothesis testing, and generation. (3) The output of this analysis is a plausible explanation or some supported hypothesis. (4) The result of data-driven science are considered to be shrinking this “gap” in understanding, toward robust and verifiable answers to the original question (see Cao [35], Hey et al. [84]).

underlying network representation. This strategy is effective at producing sophisticated models “given a network,” but is less suited at answering our scientific questions on the original system. Previous work demonstrates that the incorrect network representation for a particular question can have a large impact on performance [30]. Therefore, the correct network representation may be *more* important than developing the best predictive model.

2.3 A Taxonomy for Network Structure Inference

Below, we taxonomize work for the network inference problem according to:

- (1) The varying *difficulty* of defining edges and nodes, given the nature of the *underlying data*.
- (2) The *type of method* (e.g., mutual information, correlation) used for network inference.
- (3) The types of hypotheses and *tasks* used to evaluate the quality of the network representation.

2.3.1 Explicit, Discrete, and Inferred Edge Definitions. Networks model three broad types of relationships (i.e., edges), typically based on a measure of interaction between entities. Categorical, *explicit* relationships are unambiguously known in the system—such as the “friend” or “follower” relations in online social networks. *Discrete* interactions denote unambiguous transactions occurring between two entities—such as phone calls or text messages in mobile device data. The primary task of defining edges on these interactions is to select an appropriate threshold to measure the strength of the relationship. *Inferred* interactions denote some statistical measure of similarity, beyond simple transaction counts. For example, the definition of a spatiotemporal *co-location* interaction between two entities can be inferred by specifying “how close” for “what duration,” both of which are continuous thresholds. This is a more challenging to measure relative to discrete interactions.

2.3.2 Explicit and Inferred Node Definitions. Similarly, inferring nodes is of varying difficulty in different domains. In most applications, nodes are explicit in two ways: (1) the node definition is unambiguous (e.g., a “user,” “animal,” or “gene” is represented by a node) and (2) there is a node *correspondence* between time steps in the data (e.g., this particular node corresponds to a user “ivan” at each time step). Functional brain networks and climate networks are two examples of domains where the node definition is neither explicitly given nor is consistent over time. Here, nodes correspond to groups of two-dimensional earth surface pixels or three-dimensional brain voxels. The *scale* of node definition can dramatically change the performance of predictive models and the stability of descriptive statistics derived from the network [33].

2.3.3 Hypotheses and Tasks on Inferred Networks. Figure 2 (2–3) illustrates the core novelty in machine learning and data science for hypothesis generation, model development, and testing. In machine learning, there is typically a bias toward evaluating models using a novel prediction tasks. However, typically these models do little to explain the underlying system and are more useful for generating new hypotheses. In the statistics domain and many natural sciences, there is a bias toward descriptive modeling, where the “task” is parameter inference on an assumed parametric (or non-parametric, but still with parameters [188]) statistical model on the data. These parameters are useful at interpreting the underlying system, such as fitting the exponent of degree distributions, or reporting other aggregate statistics on a network. Both of these aspects have a role for understanding the mechanisms of the underlying system.

Applications driven by descriptive models aim to reconstruct and describe the “hidden” relational structure with the greatest fidelity, relative to the domain knowledge base. A typical example is gene regulatory networks (GRNs). These networks are inferred on data measuring individual genetic expression over different experimental settings. In these networks, nodes represent genes or functional gene families and edges represent inferred positive or negative gene expression relationships (e.g., “gene A reduces the expression of gene B under some context”). Inference of these networks typically identifies new, high quality *candidate* regulation relationships given high accuracy of inferring known edges. These interactions can be experimentally tested to build greater understanding of cellular processes and to develop potential personalized medical treatments.

In contrast, predictive methods aim to discover network structure that maximizes predictive performance and may not reconstruct the underlying process with the highest fidelity. Instead, they may focus on those aspects or modalities that are most predictive. Modeling the predictive aspects of the data allows researchers to learn regular relationships between modalities (e.g., call, text, and location in mobile phone data) or over time (e.g., periodicities). This informs further hypothesis generation to understand the *mechanism* underlying the predictive relationships. However, often highly predictive relationships are also uninteresting and can drive the structure of the inferred network. Several domains such as climate science tend to remove known periodic dynamics as preprocessing of the underlying data, prior to inferring the network structure.

2.3.4 Models for Network Structure Inference. We organize related work broadly along the type of network structure inference *model* used, including parametric, non-parametric, and thresholded interaction/correlation measures. In these groups, we categorize the type of task performed on the network, including edge and attribute prediction, descriptive analysis, and model selection.

Table 1 summarizes work across several domains, introducing an example scientific question driving the analysis, as well as the network structure inference model used to realize the network. In Table 1 (Column “Model”), we label these models under two broad categories.

First, *parametric models* make assumptions on the *distribution* of edges in the network. We identify graphical models (**GM**) and other network models using maximum likelihood methods (**ML**), relative to the assumption on the input data (e.g., the information flow between nodes [68]). Causal

Table 1. A Summary of Related Work, Across Domains

Domain	Example	Sec.	Model	Task	Citations
Computational Biology	Discover interactions between genes in cellular processes	3.1	I	DH, DR	Zhang and Horvath [204]
			IE		Butte and Kohane [32], Faith et al. [57], Meyer et al. [125]
			GM		Allen and Liu [4], Lebre et al. [113], Mani and Cooper [119], Mukherjee and Speed [128], Toh and Horimoto [178]
			R		Haury et al. [80], Yuan and Lin [199]
			CM	PE	Barzel and Barabasi [14], Feizi et al. [61]
Climate	Describe relationships in environmental system dynamics	3.2	IC	DN, PA	Kawale et al. [92], Paluš et al. [134], Steinhäuser et al. [170], Tsonis and Roebber [179], Tupikina et al. [182], Yamasaki et al. [196]
			IE		Donges et al. [50], Hlinka et al. [85]
			CM	DN	Ebert-Uphoff and Deng [55], Kretschmer et al. [107], Runge et al. [153]
			GM		Zerener et al. [202]
			R		Zhou et al. [205]
Neuroscience	Model relationships between brain regions, physiological structure, and function	3.3	IC	DN, MS	Bialonski et al. [18], Zalesky et al. [200]
			IF		Lachaux et al. [109], Pfurtscheller and Andrew [138], Ponten et al. [140], Zhan et al. [203]
			CM		David et al. [44], Dhamala et al. [48], Friston et al. [64], Ramsey et al. [145], Roebroek et al. [150], Rosa et al. [151]
			ML	PA	Cao et al. [34], Papalexakis et al. [135]
Epidemiology	Model hidden networks from observed infections	3.4	I	PA, MS	Adar and Adamic [1], Haydon et al. [82]
			GM		Britton and O'Neill [26], Groendyke et al. [73], Stack et al. [168]
			ML		Du et al. [51], Gomez-Rodriguez et al. [67, 68], Myers and Leskovec [129], Netrapalli and Sanghavi [132]
Ecology	Describe and predict animal behavior	3.5	I	DH, MS	Aplin et al. [6], Haddadi et al. [75], Psorakis et al. [143], Spiegel et al. [165]
			IE		Barrett et al. [12]
			R	DN	Farine [58], Whitehead and James [193]
				PE	Farine et al. [59], Fletcher et al. [63]
Mobile	2cmPredict social influence on individual mobility	3.6	I	PE, DN	De Choudhury et al. [45], Eagle and Pentland [53], Mastrandrea et al. [121], Sekara and Lehmann [158]

Interaction, IC: Correlation, IE: Entropy, IF: Frequency, I: Novel measures

CM: Causal model GM: Graphical model, ML: Maximum likelihood, R: Regression

Prediction, PA: Attributes, PE: Edges

Descriptive Analysis, DN: Nodes, DR: Roles, DH: Other high-order, MS: Model Selection

models (**CM**) typically generate causal networks—a special case of graphical models—use Granger causality [71] or other causal definitions [119, 124]. These networks represent strong relationships between nodes that control for confounding factors of possible adjacent nodes.

Second, *non-parametric models* tend to directly measure interactions between nodes and use statistical tests to determine appropriate edge weights (see Section 2.5.1). We categorize work related to novel and “ad-hoc” interaction measures (**I**) between the data associated with pairs of nodes, correlation networks (**IC**) that measure linear, cross, or some other correlation, entropy (**IE**), frequency domain measures (**IF**), and regression (**R**). The *edge definition* may be defined by an underlying parametric model (e.g., linear correlation); however, we refer to these models as “non-parametric,” because they do not have a model for the edge structure as a whole.

Table 1 (Column “Task”) categorizes rows within each domain by the type of task performed, under the caveat that one study may use several evaluation strategies, or that the actual task may be loosely described as the canonical task (e.g., edge prediction).

First, we denote predictive tasks, including edge prediction (**PE**) and attribute prediction (**PA**). Attribute prediction can also describe prediction of the original data, or tasks such as change-point prediction [62] and rank prediction, which are higher-order attribute tasks. Predictive models are relatively rare across application domains, because researchers are interested in descriptive and interpretable models that give insight into the underlying system. We observe some specialization in both information networks and epidemiology, which aims to predict the extent or timing of an epidemic over a population in varying contact models.

Second, descriptive analysis is broken by node-oriented statistics (**DN**), which includes reporting distributions of simple node statistics, such as degree distribution, clustering coefficient, correlation distributions, and so on. This often constitutes the base-level exploratory analysis and qualitative evaluation of the inferred network. Role-oriented analysis (**DR**) aims to characterize nodes using network features by the structural roles they play in the system (e.g., bridges between social communities). This may also include factor analysis and interpretation. Other high-order analysis (**DH**) examines communities and larger subgraph structures beyond node and edge-based descriptive statistics.

In descriptive and predictive cases, we find a sizable collection of work in model selection (**MS**). This work uses varying network representations with some evaluation criteria over models. These varying models correspond to different hypotheses to how the network may have been generated.

2.3.5 Evaluating Networks and Tasks. Evaluation strategies over the entire methodology of Figure 2 measure the performance of the *final task* (prediction or statistical inference). In this context, the network serves as a model of the data. However, the fidelity of this model in terms of fitting error, or reconstruction against a partial ground-truth network does not measure the network’s usefulness for answering questions on the original system. It is more appropriate to think of the space of networks as *possible* representations with some utility for answering a specific question. “The network” is typically seen as uncovering the true relational structure of the data with some error [187], rather than one of many representations for a particular purpose or question. Most scientific domains use an evaluation strategy focused on a final task (a network *for* location prediction, brain activity response, etc.), because there is no evaluation framework for comparing the inferred network structure. For example, we simply cannot survey baboons [59] as we can humans [53] to discover their real friendships. Furthermore, uncovering a general-purpose, robust network from complex data may not appropriately model the overlapping modalities of the data. For example, the existence of many *contexts* may mean that functional brain networks are inherently probabilistic. Discovering the most *general* network model will be less informative in any particular question, because it does not account for this complexity [30].

2.4 Network Structure Inference

The network structure inference problem represents some input data as a network and *validates* this network relative to performance on some task(s) or hypothesis. Often, model selection is done offline, and we only see the result of the final network representation evaluated on the task. Currently, there exists no standardized methodology for evaluation of networks inferred from data. Our work structures this problem, particularly in the absence of network validation data.

The network topology inference problem combines two basic models: first, a *network model* $\mathcal{R}(\mathbf{D}, \alpha) \rightarrow \mathbf{G}$ constructs network $\mathbf{G}$ on input data $\mathbf{D}$ under some parameters α . This model may be parametric statistical models (e.g., Exponential Random Graph Models) or non-parametric and threshold-based interaction networks. Second, the problem uses a *task model* $\mathcal{T}(\mathbf{G}, \beta) \rightarrow p_1, p_2 \dots$ on input $\mathbf{G}$ under some parameters, with task outputs (e.g., prediction “ p_i ”). These approximate the unknown ideal function $\mathcal{T}^*(\mathbf{G})$ of a network task (e.g., classification, prediction) with error $e()$.

This formulation may seem overly simplistic. However, it succinctly clarifies the relationships between input data, the network model, and the task model. It explicitly formulates network $\mathbf{G}$ as a *model* on data $\mathbf{D}$ for task $\mathcal{T}^*$, approximated by $\mathcal{T}$. This formulation captures simple *interaction network* methodologies (see Section 2.5.1), which separately infer the network (often by hand-tuned thresholds) and validate task performance, as well as parametric inference methods that learn the network model parameters. To our knowledge, all network structure inference methodologies can be broadly formulated in this pattern and all network inference models *should* be formulated relative to a particular task or hypothesis. In much of the existing work, the network or task models will not be formulated explicitly, or possible model combinations may be sparsely-explored.

We can instantiate several tasks within this formulation. In the context of network prediction tasks, our predictive model can output predictions of (1) edges, (2) attributes, or (3) the original data. For one instantiation, on a validation edge-set $\mathbf{E}^*$, we can evaluate:

$$\underset{\mathbf{G}}{\operatorname{argmin}} e(\mathcal{T}(\mathcal{R}(\mathbf{D}, \alpha), \beta), \mathbf{E}^*). \quad (1)$$

This infers $\mathbf{G}$ over network model parameters α and task model parameters β . Practically, this might be done by a model selection over a grid search of parameter pairs [30], or an iterative re-training [122]. In this context, the suitability of both the network model, task model, and the appropriate error function will determine the performance of the inferred network. Network model selection methodologies mitigate some of the biases from offline, hand-tuned network construction by exploring many possible model combinations.

This formulation highlights two key challenges in network structure inference. First, the parameter space of possible $\mathbf{G}$ from $\mathcal{R}(\mathbf{D}, \alpha)$ may be large and parameter search will likely be non-convex with respect to the performance of the task of interest. Second, well-performing models in the parameter-space may yield very different network topologies. Summarizing and reconciling these differences may be important for interpreting which network model is most suitable. This proliferation of plausible task and network models makes further hypothesis generation and testing more challenging for the researcher; more plausible models means additional interpretations of the mechanisms involved in the underlying behavior of the system.

2.5 Network Models and Inference Methodologies

Our problem description is flexible enough to incorporate varying types of network models (e.g., regression, correlation, parametric network models) and tasks (e.g., prediction, descriptive statistics and hypothesis testing). Below, we describe several formulations for network structure inference according to their network model and task model.

2.5.1 Interaction Networks. The most prevalent class of network topology inference is measuring pairwise interactions (e.g., correlations) between nodes and choosing a threshold to define a sufficient degree of interaction. This threshold may be chosen by some statistical test, by tuning on some desired criteria (e.g., a desired network density), or ad-hoc offline testing. This has been discussed in the context of *discrete* interactions as described above:

“Inferring networks from pairwise interactions of cell-phone call or email records simply reduces down to selecting the right threshold τ such that an edge (u, v) is included in the network if u and v interacted more than τ times in the dataset. Similarly, inferring networks of interactions between proteins in a cell usually reduces to determining the right threshold.” [129]

Researchers often make several application-specific decisions around these thresholds:

“From this complete correlation graph, only the edges with significant correlation (> 0.5) were retained. But using the same threshold for positive and negative correlations is not appropriate as negative correlations are usually weaker and many nearby locations have high positive correlation.” [93]

“We let δ as a user-controlled parameter, where larger δ values correspond to less predicted regulations, and only focus on designing a significance score $s(t, g)$ that leads to ‘good’ prediction for some values of δ .” [80]

These methods typically produce a fixed network model and explore networks under varying threshold settings in an offline trial-and-error fashion. Unfortunately, these negative results are rarely reported, so we often only have a description of the final network definition.

We use our formulation for this process. Assume the interaction threshold τ is given by hand-tuning or domain knowledge; we have some feature matrix $\mathbf{D}$ that has some similarity between features, measured by $\mathcal{R}()$, and $\mathcal{T}()$ is an edge prediction task evaluated on $\mathbf{E}^*$. This method is expressed as

$$\mathcal{R}(\mathbf{D}, \tau) \rightarrow \mathbf{G}; \underset{\beta}{\operatorname{argmin}} e(\mathcal{T}(\mathbf{G}, \beta), \mathbf{E}^*). \quad (2)$$

When these interaction networks are evaluated in the absence of ground truth, the network may be measured through autocorrelation. In this case, the same network inference is applied, $\mathcal{R}(\mathbf{D}^*, \alpha) \rightarrow \mathbf{G}^*$ for some hold-out data $\mathbf{D}^*$.

2.5.2 Parametric Network Models and Parameter Inference. Maximum-likelihood methods assume some parametric model family to represent relationships between nodes, such as time between interaction, likelihood of information transmission over time. For clarity, we work through one specific application, in epidemiology and information networks such as blogs, although the pattern is similar in other applications. Structure inference methods on disease transmission networks share the assumption that we are observing the “arrival” of infection or information/attribute value at nodes (i.e., computers, blog pages, individuals) over time, but are unable to observe the topology that transmitted the information. For an edge e_{ij} , the likelihood of information transmission (or infection, in the epidemiology context) time difference $t_j - t_i$ is given by an information transmission/infection model, yielding the likelihood that v_i transmitted information to v_j [129].

Where input data $\mathbf{D}$ are infection times of each node, we can formulate these methods as

$$\underset{\alpha}{\operatorname{argmin}} \mathcal{R}(\mathbf{D}, \alpha) \rightarrow \mathbf{G}; \underset{\beta}{\operatorname{argmin}} e(\mathcal{T}(\mathbf{G}, \beta), \mathbf{E}^*). \quad (3)$$

In information network applications, E^* is typically provided by a known network. Processes are simulated on this network to generate input data for the maximum-likelihood relational data model $\mathcal{R}()$. This method is used to “reconstruct” E^* only from input data D .

2.5.3 Learning Network Structure and Task Model Parameters. Previous work in statistical relational learning on *explicit* networks has focused on learning relationships between (categorical) attributes and a predictive task, such as link prediction [69, 130], and distinguishing correlated effects between these processes [108].

Previous work has also used a methodology for the network structure inference problem, which maximizes performance of particular task(s) on the inferred network [45, 59]. Consider interaction networks over varying thresholds τ . A naive solution for this type of approach is to explore the parameter space of $\mathcal{R}(D, \tau)$ and evaluate the task performance on each inferred network. However, this is very costly, especially as $\mathcal{R}()$ requires more parameters.

Learning network structure paired with task parameters $\mathcal{T}()$ yields both the network and task definition according to some optimization or model selection criteria. For example, previous work has learned network models that iteratively discover features in a supervised LDA model, and the logistic regression weights that perform well for an edge prediction task [122]. Other work uses model selection to select task and network models by performance [30, 183], or minimum description length (MDL) *encoding cost* [28]. A related area of work treats dynamic networks as a sampling process over varying sampling *rates* [148], where network models are evaluated according to the most predictive representation for the given task [62, 174].

In Equation (1), we have already formulated paired evaluation of network and task parameters:

$$\underset{G}{\operatorname{argmin}} \ e(\mathcal{T}(\mathcal{R}(D, \alpha), \beta), E^*). \quad (4)$$

With more context, we now see $\mathcal{R}()$ can construct networks of varying interaction measures, sampling rate criteria, and so on, and this methodology selects the $G()$ corresponding to minimum error on $\mathcal{T}()$. Therefore, this simple formulation defines a very general model selection methodology.

3 APPLICATIONS OF NETWORK STRUCTURE INFERENCE

Applications of network structure inference are too numerous to meaningfully cover. Instead, we look closely at a few applications where inferring networks from data is illustrative of the challenges and expertise in these domains.

3.1 Computational Biology: Discovering New Genetic Regulation

Networks are constructed in computational biology to model many different processes. Protein-protein interaction (PPI) is the most common type of network in this domain, constructed mostly experimentally through, for example, yeast two-hybrid experiments others that physically test for binding of one protein to another or other types of physical interactions. Metabolic networks are process networks that model relationships between enzymes, metabolites (typically nodes) on processes (typically edges) such as reactions or pathways. We focus only on gene regulation networks (GRNs) and gene co-expression networks (GCNs), which infer relationships between genes under different experimental scenarios (for further reading, see Reference [162]).

3.1.1 Underlying Data in Gene Regulatory Networks. Next-generation high-throughput microarray technologies allow the sequencing of genomes and measuring the *expression* of particular genes at a large scale and low cost [160]. “Expression” is the measurement of how groups of genes produce different phenotypic specializations through the production of different proteins.

Fig. 3. A consensus gene regulatory network (GRN) of *Staphylococcus aureus* generated from microarray data in the DREAM5 challenge ensemble methodology (see Figure 4). Network modules in agreement with Gene Ontology (GO) database are colored and labeled according to their shared function (grey are not coherent in GO). The detail (right) shows inferred gene regulation related to pathogenesis. (From Marbach et al. [120].)

In cellular development, different gene co-expression can be responsible for RNA translation or nucleotide metabolism, yielding many complex functions from gene interaction (see Reference [13]). Figure 3 illustrates a small network with annotated functional clusters.²

The output of microarray analysis (with notable simplification) is a 2D data matrix of numeric values measuring the expression of a gene (row), on a particular experimental design, subject, or time step (column) [10]. Defining edges between genes simplifies to comparing expression profiles across the different columns of the data.

3.1.2 Discovering New Gene Interactions From Data. The high-level “task” for gene regulatory networks is *link prediction* on the network learned from data, to discover unknown gene regulation candidates that can be experimentally tested. The network model inferred from data should agree with databases of biologically known interactions and function, while providing few verified false-positive regulations.

Table 1 illustrates that the inference of GRNs is very mature relative to other domains, since the inferred network is clearly verifiable according to current domain knowledge, and directs future investigation and hypothesis generation. One unique challenge in gene regulation is the issue of confounding factors, including indirect and transitive associations, resulting in many spurious edges. Recent work has measured these “direct” (e.g., causal) edges in noisy expression datasets [14, 61].

3.1.3 Translating Gene Expression Data to Interaction Discovery. Marbach et al. [120] introduce an ensemble approach associated with the DREAM5 network inference challenge. The authors present 29 different network inference methods across different model types, including regression, mutual information, correlation (a.k.a. “relevance networks” in this domain), Bayesian networks, ensembles, and other novel approaches (e.g., random forests, neural networks, Gaussian mixture models, etc.) This list demonstrates the maturity and variety of methods applied to this problem. These methods use pairwise (e.g., gene-to-gene correlation) or group-wise (e.g., many-to-one group LASSO) measures of interaction intensity, which yield a directed, unweighted network signifying “gene A regulates gene B.”

² All figures reprinted with permission and attribution.

Fig. 4. The general methodology design for the DREAM5 network inference challenge. This methodology (1-2) generates one simulation dataset and three experimental microarray datasets from three well-studied species. The 29 participating inference methods all (3) generate inferred network output, (4) a consensus network is constructed for each dataset, and then (5) validated against known edges in synthetic networks and against experimentally known edges in two of the real datasets. (From Marbach et al. [120].)

Figure 4 illustrates an ensemble methodology for the DREAM5 network structure inference challenge. The authors (1) generate a ground-truth network on three species using experimental trials as well as synthetic network data. In (2), these experiments yield four different datasets of biologically-tested networks of gene regulation, as well as the raw (non-relational) gene expression. The collection of inference methods produce (3) inferred networks on each of the four datasets. The authors (4) integrate these 29 different inferred networks to produce an ensemble network. This network is (5) validated against the ground-truth networks generated in step (1). Finally, the authors show that the ensemble method discovered 59 potential interactions, of which 29 show some support and 20 show strong support for being biologically significant.

We examine this process in such detail to demonstrate the “complete” methodology for network structure inference, from input data, to network representation, to a final output (learning the regulatory network). Within our formulation, this methodology does not have an explicit *task model*, as the network itself is the object of interest. Therefore, this application is typically *descriptive* modeling against evaluation data. Although these models do “predict” new potential gene regulation via unknown edges, these interactions are usually checked manually through experimentation.

3.2 Environmental Science: Discovering Climate Relationships and Predicting Outcomes

Networks inferred to understand climate dynamics are among the most difficult to model of any domain, with much of the work to formalize and validate these networks still in early development. Within this domain, researchers want to discover robust, *causal* relationships between climatic variables, over different spatial regions of earth. This modeling can improve prediction of changing hydrological processes, land-cover, ecosystem productivity, and polar or sea ice cover, which are key aspects for climate change mitigation. Two unique challenges exist for inferring climate networks:

- (1) Input data is typically noisy, highly spatially autocorrelated, multivariate time series of climatic variables collected over varying regimes and sensor (e.g., satellite) quality. Domain scientists produce a consistent “reanalysis” data products that attempt to mitigate problems of varying sensors and errors in data. However, inferring accurate networks from these data requires significant understanding of the data ingestion process [115, 155] and its introduced biases and variability.

- (2) The structure of climate networks is not well understood aside from a handful of climate indices—coarse spatial locations on earth where dynamics are well-studied and regulate or correlate with other environmental processes (e.g., El Niño and La Niña oscillation cycles). Therefore, validating correctness of the inferred network is suitable for unsupervised strategies such as relational or predictive modeling of the original data.

3.2.1 Network Inference Methods in Climate and Environmental Science. Nearly all studies constructing climate networks use some time series similarity as an underlying relational measurement. Previous work has used linear correlation [49, 170, 179, 196] or mutual information [49, 85], either a hand-picked [49, 179] or a simple statistical test [196] to set similarity threshold τ —where similarity greater than τ is considered a binary edge in the network.

There is considerable focus on formulating these simple pairwise comparison methods and often the “recipe” of the network, according to parameter settings and preprocessing choices, varies greatly from study to study. These networks are typically binary rather than weighted, because the final output of interest is a binary decision on the similarity distribution: (e.g., [“significant,” “not significant”]). However, typically these measures will have no “natural” threshold, which gives this binary classification. Instead, these networks can be gradually densified or sparsified by weakening or strengthening the similarity threshold.

Descriptive statistical work has been very popular *downstream* from these varying network “recipes.” Donges et al. [49] reports clustering coefficient, betweenness centrality, closeness centrality on correlation and mutual information networks (Equations (5) and (6)). Tsonis et al. [180] reports community structure that tends to cluster in spatially contiguous locations, on account of the autocorrelation present in these networks. Little work focuses on evaluating these networks as predictive models for the input data. Steinhäuser et al. [170] use both descriptive statistics and predictive performance to evaluate the inferred network.

As work utilizing network models grows in this climate and environmental science, researchers have developed more sophisticated techniques for determining edge significance [92], or conditional dependencies using causality [55, 107, 153].

3.2.2 A Methodology for Translating Environmental Sensing Data to Interaction Network. We will step through a concrete example of constructing a climate network from spatially-gridded time series data of global surface air temperature (SAT) [49]. The authors measure the similar dynamics of pairwise earth locations (corresponding to nodes n_i, n_j) and construct edges between locations with “significant” similarity. The authors use two standard measures, linear correlation and mutual information between time series X_i and X_j :

$$P_{ij} = \frac{\sum_{t=1}^{|X_i|} (X_{i,t} - \bar{X}_i)(X_{j,t} - \bar{X}_j)}{\sqrt{\sum_{t=1}^{|X_i|} (X_{i,t} - \bar{X}_i)^2} \sqrt{\sum_{t=1}^{|X_j|} (X_{j,t} - \bar{X}_j)^2}}, \quad (5)$$

$$M_{ij} = \sum_{b=1}^{|B|} p_b(X_i, X_j) \log \frac{p_b(X_i, X_j)}{p_b(X_i)p_b(X_j)}. \quad (6)$$

Equation (5) is the sample Pearson correlation between two time series, where $\bar{X}_i$ is the sample mean of time series X_i . The denominator represents the product of the sample standard deviations of X_i and X_j . This measures linear relationship of X_i and X_j over the length of the time series. Equation (6) is the discrete mutual information estimate between two time series, where $p_b(X_i, X_j)$ is the joint cumulative distribution of the b th discretization window, and $p_b(X)$ is the marginal

Fig. 5. On the input data of averaged global surface air temperature (SAT) at different spatial regions, (a) shows the distribution of pairwise linear correlation measures (P_{ij}) vs. geographic distance between nodes, on a logarithmic color bar scale. This illustrates a strong spatial autocorrelation between nearby points. (b) Shows the pairwise distribution for mutual information pairwise calculation vs. geographic distance, showing less spatial autocorrelation. The horizontal bars indicate thresholds on the P_{ij} and M_{ij} scales, which produce the same network density ($\rho = 0.005$). (c) Shows the linear correlation vs. mutual information. The starred quadrant (top-left) denotes edges defined by mutual information but not by correlation. (From Donges et al. [49].)

cumulative distribution of the b th discretization window. This measure compares the shape of the joint and marginal probability density functions. When the joint distribution is equal to the product of marginal distributions: $\log(1) = 0$ yields no mutual information (e.g., X_i and X_j are independent).

Varying similarity threshold τ produces networks of varying edge densities ρ and other network measures. The authors select thresholds for correlation and mutual information ($\tau_{corr} = 0.682$ and $\tau_{MI} = 0.398$) such that they produce the same network density ($\rho = 0.005$).

Figure 5(a) shows the density of pairwise linear correlation measures (P_{ij}) vs. geographic distance between nodes, on a logarithmic color bar scale. This illustrates a strong spatial autocorrelation between nearby points. Figure 5(b) shows the pairwise distribution for mutual information pairwise calculation vs. geographic distance, showing less spatial autocorrelation. The horizontal bars indicate thresholds on the P_{ij} and M_{ij} scales, which produce the same network density ($\rho = 0.005$). Figure 5(c) shows the linear correlation vs. mutual information. The starred quadrant (top-left) denotes edges defined by mutual information but not by correlation. This demonstrates that under a fixed network density, the chosen similarity measure can yield a very different network structure.

While several methods have been developed to test edge significance, little work has focused on the validity of higher-order structures such as paths or communities. While the graph *definition* of paths or communities are valid on these networks, no known work in this domain measures the *interpretation* of these relationships with respect to the original time series data.

Figure 6 explores the sub-spaces of different node measures on the inferred network. Figure 6(a) shows the distribution of betweenness centrality per node vs. the degree per node. Similarly, 6(b) shows betweenness centrality and closeness centrality. The authors demonstrate that degree-preserving edge re-wiring randomization indeed destroys the rank-order correlation between the marginal distributions of the node measures.

There are two drawbacks of this analysis that re-occur across domains. First, while this methodology tests some global relational structure of the network, we are unable to interpret the relationship between any two nodes at a high geodesic distance (≥ 2). This means that we cannot measure properties we associate with networks, such as flow or routing. Second, significance

Fig. 6. (a) The distribution of betweenness centrality of nodes vs. degree of nodes. (b) The distribution of betweenness centrality of nodes vs. closeness centrality of nodes. (From Donges et al. [49].)

analysis is done at a particular threshold setting, without a sensitivity analysis on the original threshold choice. In the machine-learning settings, the parameter sensitivity will often be on the *prediction model* parameters at a particular network definition threshold.

3.3 Neuroscience: Describing Functional Brain Structure and Its Connections

Biological research suggests that the brain activates interconnected, often spatially distant, regions along neuronal pathways [166]. This interconnected complexity makes networks a very natural model to study the brain. These studies are broadly in two areas: structural and functional brain networks. *Structural networks* (also known as “effective connectivity,” tractography, or the brain connectome) map the physical axon pathways between neurons, which may be relatively long and spatially distant. With some simplification, these networks are analogous to the physical layer in communication networks, where nodes are *explicitly* connected by cables and routers. *Functional networks* are analogous to the logical layer in communication networks. These networks model *how* neuronal signals (i.e., “traffic”) flows over this physical layer to activate other neurons (i.e., “resources”) to perform different types of behavior such as auditory, visual, or motor behaviors. Researchers are only starting to understand the underlying routing and information-seeking on this physical network, and the complex contexts that change how the behavior is realized within the functional layer. Researchers aim to better understand and predict this routing, as well as the collective activation dynamics in different areas of the brain.

Much work compares the topology of structural and functional networks using descriptive network statistics and higher-order structures, such as cluster and communities [147, 152, 167], especially under different experimental conditions that may affect these structures, such as spinal cord injuries, epilepsy, or schizophrenia. However, all of these studies infer the network models differently, therefore it is an open challenge to rigorously synthesize these results.

3.3.1 Underlying Data in Brain Networks. The underlying data for structural or function networks are primarily derived from two sources. First, bio-medical imaging technologies, including Magnetic Resonance Imaging (MRI), functional MRI (fMRI), and Diffusion Tensor Imaging (DTI), detect structure of varying densities and water content. These procedures produce a flat two-dimensional image of pixels, or a three-dimensional space of voxels (often as a series of imaging over time). For example, DTI is used to construct structural networks. These images can accurately trace axon tissue connectivity by measuring flow vector orientation through the voxel space. fMRI similarly measures blood flow to voxels, a surrogate for “activity” at this location. Inferring a functional network on fMRI data then amounts to measuring statistical interactions (e.g., correlation) between activations in different brain areas. Second, non-invasive sensors,

Fig. 7. A methodology for constructing brain networks. (1) For a subject or population of subjects, (2) sensing techniques such as Diffusion Tensor Imaging (DTI) or functional Magnetic Resonance Imaging (fMRI) are used to sense connectivity or activity of the brain, respectively. (3) Given sensed data at recording sites (nodes), edges are inferred by different measures on the underlying data (see Section 3.3.2). (4) Subsequent scientific study is conducted, using the network as a data model. (From Sporns [166].)

including electroencephalography (EEG) and magnetoencephalography (MEG), measure and localize electrical current at a probed location. Typically, these probes yield fewer and spatially coarser nodes than those defined from fMRI voxel data. These techniques produce time-series estimating electrical current (a measure for “activity”) at reference locations.

Figure 7 illustrates defining both structural and functional networks, from (1) the data collection on individual subjects to (4) the final analysis task. The left path of Figure 7 illustrates constructing structural networks. (2) *Anatomical parcellation* techniques use DTI or similar imaging to determine physical connectivity in the brain. These techniques are very accurate in recovering tracts of connectivity, unambiguously. (3) These tracts are translated into nodes and edges, where previous work shows significant effects of node definition on descriptive measures such as average path length and clustering coefficient [201]. Finally, (4) researchers use the networks as models to ask questions about the brain of the original subject or population.

The right path of Figure 7 illustrates inferring functional networks from sensed neuronal activity. This activity can be for a range of stimulus, such as music preferences [194], image/language associations [135], or for experimental conditions such as an Alzheimer’s patient cohort [175]. (2) fMRI, EEG, or MEG sensors measure activity at different *recording sites* (contact locations, pixel, or voxel locations). As in structural network construction, pixel aggregation or node definition mapping may be applied. (3) These activity response signals are compared between recording sites (nodes) with time series similarity measures (e.g., cross-correlation). “Sufficiently” similar time series are interpreted as latent connections between nodes, yielding the final functional network.

One class of experiments focuses on functional networks from fMRI, coupled with a particular experimental user task (e.g., speaking, listening, motor). Another class focuses on inferring

Fig. 8. A study methodology comparing structural and functional brain networks. (a) A schematic of the rich-club structural network. Blue nodes indicate the k -core decomposition of the structural network ($k > 10$), grey nodes indicate non-rich nodes. Red (“rich-club”) edges connect two rich-club nodes, orange (“feeder”) edges connect exactly one rich-club and one non-rich node, yellow (“local”) edges connect two non-rich nodes. (b) The network with colored nodes and edges, visualized in a brain coordinate system. (c) An adjacency matrix comparing structural topology sensed from Diffusion Tensor Imaging with rich-club edge coloring (left) to three thresholded values of functional connections for resting brain state inferred from fMRI for the same node-set (right). These nodes are ordered according to brain function in different regions of the brain. (d) The distribution of rich-club nodes within these different labeled regions. (From Sporns [166].)

resting-state networks (RSNs) of the brain [72]. These networks are constructed in much the same way as other functional networks, except this resting connectivity is informative of very robust functional clusters. Functional networks for different user tasks can then be characterized at a higher level (e.g., cognitively “difficult” tasks) by comparing to the resting-state network (RSN).

Another network of particular research interest is the “rich-club” structural sub-networks [184–186]. This network is essentially a k -core decomposition of the structural network, which indicates the global “backbone” of connectivity (where $k > 10$ is set in comparison to degree-preserving randomized networks). Nodes within the rich-club network are used to characterize the broader network into “rich-club edges” connecting two rich-club nodes, “feeder edges” connecting exactly one rich-club node, and “local edges,” which connect two non rich-club nodes (see Figure 8). Analogous to routing in communication networks, information can flow locally within one local region for a particular behavior or be routed through backbones to physically distant regions.

3.3.2 Methods for Inferring Networks. Neuroimaging time series are the dominant underlying data in neuroscience, therefore methods for constructing functional brain networks are almost exclusively in the domain of thresholded pairwise similarity measures, with some exceptions of parametric network modeling [100]. Sakkalis [156] provides an in-depth review of these different measures, including cross-correlation [18, 200], frequency domain analysis, such as discrete Fourier transform (DFT), and discrete wavelet transform (DWT) and domain-driven “coherence” measures [109, 138, 140, 203]. Finally, significant work focuses on causal models [145], including Granger causality [48, 150] and dynamic causal models (DCM) [44, 64, 151].

Where these methods have threshold parameters, τ , they are often validated by measuring *robustness* of network statistics across varying thresholds [105], then using these thresholds for

distinguishing patient cohorts by label or network statistic distribution. For example, previous work uses a paired t -test or other simple statistical test [176]. Kramer et al. [104] proposes a bootstrapping [56] method in the frequency domain, providing more general p -values without model assumptions.

3.3.3 Dynamic Functional Brain Networks. Since the underlying data of functional brain networks is often time series, a time series of networks (dynamic networks or time-evolving networks) are a natural extension in this domain [89]. Network construction using time-series similarity measures (e.g., cross-correlation) generalize to the dynamic setting, computing on time-series subsequences. The advantage of introducing the complexity of dynamics is discovering distinct connectivity “states” over the course of the experiment. Because fMRI response can change very quickly as activity occurs over the brain, these states are lost under global time series measures [42, 149, 198]. Challenges of network validation generalize to this dynamic setting, with the added challenge of appropriate *temporal scale* [62, 174].

3.3.4 Comparing Functional and Structural Brain Networks. Figure 8 illustrates a complete case study summarizing many of the topics discussed above. This work integrates structural networks across 75 individuals, sensed from Diffusion Tensor Imaging (DTI) with functional networks sensed from fMRI in resting state using Pearson correlation. These different views of the network enable researchers to study how function and physical connectivity are correlated. (a) Illustrates a schematic layout of *rich-club* nodes, feeder, and local nodes. (b) Shows the layout of these nodes in a brain coordinate system, with the same node and edge coloring. (c) Illustrates an adjacency matrix comparing structural connections (left) with three thresholded values of functional connections within the same node-set (right). Furthermore, nodes in different spatial regions of the brain (e.g., default mode network, motor, auditory, frontal) are labeled according to primary function, showing structural and functional edges between these regions. (d) Shows the distribution of rich-club nodes within these different labeled regions.

3.4 Epidemiology, Blogs, Information Networks: Modeling Virus Spread and Information Flow

3.4.1 Networks in Epidemiology. Networks are used in epidemiology to simulate the spread of disease over a family of parametric network models (e.g., random, small-world, exponential random graphs) representing contact between entities over time [95]. Network structure inference in epidemiology and information networks aims to discover an unobservable network (e.g., physical contact networks, sexual networks, malware transmission) over which information or infection is spread. We observe the *effects* of transmission on the infected node (e.g., symptoms), but not the edges over which they spread. The focus in this domain is therefore parametric inference for these models from data.

Modeling *contact networks* allows researchers to simulate different outbreaks on these networks. Figure 9 illustrates types of network data and models used in the epidemiology domain [43]. Figure 9(a) visualizes a contact network from survey data, where squares are respondents and edges identify the transmission risk for HIV through contact by drug use or sex [16]. Figure 9(b) shows a snowball sample of respondents (as squares) and their partners. In this example, edges between circle nodes are unknown. Subsequent modeling for edges between circle nodes can test the spread over this population under different unobserved contact assumptions. Figure 9(c) illustrates a model of households (cliques), sparsely connected to others. This is intended to model contact networks and potential spread through family-unit environments. Finally, Figure 9(d) illustrates a lattice network with sparse edges outside of the local neighborhood. This model was previously used to capture the “small-world” property of information and disease spread [23].

Fig. 9. Examples of networks inferred from data, or modeled in the epidemiology domain. (a) Depicts a contact network where squares are respondents and edges are sexual or drug contact that might transmit HIV [16]. (b) A sexual network derived from snowball sampling of respondents (squares), where the edges between non-respondents (circle nodes) are unknown. (c) A network of households (cliques) and their interconnections for modeling infection in realistic social contact networks. (d) The “small-world” network property modeled through a lattice with sparse edges connecting distant nodes. (From Danon et al. [43].)

Historically, inferring networks in this area focuses largely on parametric graphical modeling using Markov Chain Monte Carlo (MCMC) [27] and maximum likelihood methods, incorporating modeling assumptions in transmission and interaction rates.

Given a compartmental transmission model [25] (e.g., the susceptible-infected-recovered SIR model or susceptible-exposed-infected-recovered SEIR model), these methods measure the likelihood of possible *sequences*, or trees of infection, where infection times from a root are monotonically increasing. Let t_i and t_j denote the infection times of nodes n_i and n_j , then the transmission model yields the likelihood $P(\text{“}j \text{ infected i”} | t_j - t_i)$.

3.4.2 Contact Network Inference. Early work in the epidemiology domain focused on the inference of either spread parameters (e.g., infection rate) or network model parameters on random graphs [26], Poisson, and power-law networks [126], as well as fitting of real-world data to a contact network model [9]. In these analytical and simulation results, the network model is known and no structural inference is necessary. These studies generally model the spread of an epidemic under possible individuals in contact (called “contact tracing”) [136] and discovering the root individual(s) of the infection (called “transmission tracing”). Early work also formulates association network heuristics based on time and distance of potential contacts [82].

3.4.3 Infection-Time Cascades. Previous work in machine learning uses network structure inference to represent the spread of information between nodes, where the edges of transmission are unobservable. Maximum-likelihood methods have focused on learning a network under assumed transmission rate models, using statistical inference for these parameters [129].

Figure 10(a) illustrates the intuition of network construction by information propagation through unobservable edges. To recover the unobservable true network G^* , each sequence of non-decreasing infection times (e.g., “cascades”) supports the possible transmission between nodes with sufficiently close infection times. The key step of this work is learning the parameters of a transmission model that measures the likelihood of a cascade according to differences in “adjacent” infection times in the network.

Myers and Leskovec [129] use convex programming to learn a maximum-likelihood network under a fixed transmission time probability distributions $w(t)$ and recovery-time distribution $r(t)$. To learn the transmission weight matrix A , the authors use an Independent Cascade model [96] where an uninfected node n_i is exposed to infection by adjacent infected nodes n_j at each time step using a Bernoulli process with a probability $A_{j,i}$. The authors present a convex optimization formulation of their likelihood function, with regularization. This model, and most

Fig. 10. (a) Illustrates constructing network G^* on a collection of cascades $c_1, c_2, \dots, c_k$. For a source node s , these methods estimate the likelihood of infecting i given the observation of information arrival (“infection”) at node i . (b) An estimated network inferred from hyperlink arrival times at nodes. (From Gomez-Rodriguez et al. [68].)

Fig. 11. Transmission histograms, showing poor fit of exponential and Raleigh models. (From Du et al. [51].)

of the subsequent work, is evaluated on synthetic network models where the underlying network is explicitly known. The “task” is the accurate reconstruction of the network that simulated these infection times. Similar to gene regulatory networks, the final evaluation of the network is the network itself, rather than a subsequent task on the inferred network. These methods then also typically present qualitative results on real-world datasets.

Gomez-Rodriguez et al. [68] scales this previous work by fixing a global edge transmission probability β and solving this simplified problem. For many applications, this fixed transmission probability assumption can be made. The primary insight under this assumption is that we can simply use the most likely propagation tree over a set of nodes in a cascade c (see Figure 10). Given a cascade set C , the authors marginalize their likelihood function relative to edge selection and prove this function is monotonic and submodular. Therefore, edges can be greedily selected with an approximation factor of $(1 - 1/e)$ [131].

The transmission model of [129] assumes repeated Bernoulli trials of fixed probability. Gomez-Rodriguez et al. [67] introduce transmission likelihood functions that vary over time. Given a set of cascades C , the method infers the parameters of transmission rate models for each edge. The model uses a hazard function $H(\tau_i|\tau_j; \alpha_{j,i})$ that measures the instantaneous infection rate on node n_i from infected node n_j , given the parametric function on α . The authors present three different transmission model functions: Exponential, Power-law, and Rayleigh, and prove that the optimization for transmission rates matrix A is convex.

In many real-world applications, information can propagate in a multitude of ways. Older information might be promoted by an influential node, causing multi-modal spikes in propagation after some delays. Therefore, these transmission rates may not decay monotonically, nor according to any simple parametric function.

Figure 11 illustrates this intuition on real transmission data between blog sites. The figure reports the probability density function (y -axis) over varying time delays between post creations on different nodes (x -axis). The data poorly fits any single exponential or Rayleigh transmission

model. Du et al. [51] extends previous work to address this limitation. The key addition of this work is to kernelize the hazard function $H(\tau_*|\tau_j, \alpha_{j,i})$ over m different kernels. These kernels serve as a piecewise approximation of the time-lag distribution, which can then be used to estimate the likelihood of transmission between nodes, given observed “infection” data.

We have dug deeper into the technical details in this section to demonstrate the added modeling complexity accounting for richer assumptions of the underlying transmission process. However, there is still an important gap. Previous work does not typically account for null modeling. In this instance the underlying process may not be a network, but a confounding process that might individually “infect” nodes.

3.5 Ecology: Inferring Animal Social Networks to Explain Social Behavior

3.5.1 Networks in Ecology. Networks in ecology model systems at differing levels of abstraction. In systems ecology, traditional graphical models are used to model an ecosystem at a high level, with relationships between species, environmental variables, services, and other processes [127]. For example, food webs model who-eats-whom within an ecosystem [141]. A second application of networks in behavioral ecology model social systems of animal populations [60, 191]. These networks are the focus of this section.

Analogous to sociology and political science, traditional fieldwork data in ecology are collected from direct observation and “surveys,” measuring interactions or other relationships among individuals in the population [12, 173], particularly over time [47, 123, 139]. Networks serves as models to summarize and describe these social or behavioral relationships within these populations.

Direct observation of populations allows researchers to incorporate their own intuition and experience into the definition of these networks. In practice, much of the work in this area uses ad-hoc, intuitive network definitions with some sensitivity analysis. These networks derived from direct observation are typically categorical (e.g., kinship relations) or discrete (e.g., thresholding on number of interactions, where interactions are implicitly decided by observers). Methodological consideration are well established in the field on these data, including edge strength thresholding [38], sampling, hypothesis testing [39], and randomization strategies [75, 91]. Each of these provide several different choices for validating the definition of these networks. These networks are tightly coupled with a particular hypothesis and experimental cohorts in the population.

Two networks of interest in animal social networks measure *affiliations* and *associations* between individuals [193]. Affiliations describe *intentional* social relationships between individuals [39, 192] (e.g., grooming pairs of primates), while associations describe a broader set of interactions, which might be driven by structural factors rather than social affinity (e.g., environmental resources, sex, age, and other individual attributes) [15]. Whitehead and James [193] introduce a generalized affiliation index using a linear regression model, using it to remove sets of predictive structural features. The resulting model by subtraction is the affiliative network model.

3.5.2 Instrumentation and Sensing of Animal Populations. Recent instrumentations of individuals and the environment allow the observation of ecosystems and populations at an unprecedented scale using geo-location sensors such as GPS, proximity sensors, radio-telemetry, and Passive Integrated Transponder (PIT) tags [94, 106, 154], as well as individual identification by photographs [17]. This instrumentation allows the study of detailed individual behavior and social dynamics that are outside the direct observation of researchers. This abundance of data requires novel statistical techniques for inferring networks from *implicit* interactions.

Defining meaningful “interactions” directly from data is non-trivial. No known work compares the biases of interaction and/or affiliation sampling via traditional direct observation fieldwork, against the interaction or affiliation inference from underlying data. Presently, these sensors are

most effective at recording simple co-location or trajectories. Challenging independent problems such as activity recognition (e.g., grooming, conflict) are much more easily identified by direct observation by researchers than by data, but these direct observations do not scale. Future research will likely integrate the strengths of these modalities to improve fieldwork data collection.

3.5.3 Studies and Network Inference Methods on Instrumented Data. The key difference between data from traditional fieldwork and from instrumented technologies is that the former tend to be discrete counts (e.g., number of co-locations or grooming events), while the latter are continuous data without these higher-level labels (e.g., relative distances between individuals). To translate to discrete co-location events—and subsequently a network—requires defining “how close” for “which duration” constitutes a co-location edge, or “how correlated” for “which duration” constitutes a higher-order edge in the network.

The simplest method for setting these closeness and persistence thresholds for co-location is by domain knowledge, or by sampling the parameter-space in some way. Haddadi et al. [75] use this strategy in GPS data from sheep, ranging from individuals co-locating for 1min at 1m, to 5min at 3.5m. The authors have some known “affiliations” (as described above), which are used to validate network accuracy at these different thresholds when the individuals are mixed into a larger population. Aplin et al. [6] collected data from passive integrated transponder (PIT) tags of individuals sensed by radio-frequency identification (RFID) antennae at feeder sites. This work defines associations as two individuals co-occurring at the site within 30 seconds before or after the other on a sliding 75s window. Co-occurrence is categorical due to the physical design of feeders, so only the “persistence” of interaction need be defined. This threshold generates a stream of pairwise associations, which is then thresholded again ($\tau \geq 0.02$) to produce an aggregated association network.

Psorakis et al. [143] define edges using Gaussian mixture models (GMM) on co-occurrence data for a similar feeder system. This approach mitigates the “persistence” threshold by fitting Gaussian distributions to a one-dimensional space of occurrence counts (and generalized to continuous two-dimensional geographic space [59]). These distributions capture *events* of co-occurrence among several individuals. Hamede et al. [77] used a randomization approach to define non-random associations on proximity sensors on a population of wild Tasmanian devils (*Sarcophilus harrisii*). Internal thresholds on these sensors detect co-location within 30cm. This work studies disease transmission through physical contact of the animals, so this thresholding is appropriate.

We enumerate several instances to demonstrate that these network constructions are typically “ad-hoc” from domain knowledge. In each case, the sensitivity of downstream hypotheses is not tested against the choice of network threshold parameters. This may result in inadvertent fine-tuning of parameters and incorrect conclusions.

3.5.4 Translating Co-Location Data to Innovation Spread in Networks. Aplin et al. [5] proposed a network inference task to measure the learning of a feeding behavior in a population of great tits (*Parus major*). In this experiment, feeders were instrumented with RFID antennae and recorded the visitation of each unique bird using PIT tags. The feeders had a sliding door that needed to be open to the left or right to access the food. The feeder recorded the bird’s puzzle solution. The authors investigated whether birds learn by example from others at the feeding sites. Figure 12 shows a thresholded, aggregated network over individuals, weighted by the frequency of co-location events at any feeder, using the Gaussian mixture model (GMM) method described above for interaction “events.” Yellow nodes represent trained individuals, red nodes represent individuals that have learned the correct behavior to solve the feeder by the end of the study.

These networks visually show strong network modules between red and black individuals. Figure 12(b) shows a strong network separation between the two behaviors, where trained

Fig. 12. Networks inferred from co-location at feeding stations sensed using RFID over three different populations. These edges are colored showing learned behavior (red edges) of obtaining food through an instrumented puzzle mechanism, the trained individuals (yellow nodes), and affiliated individuals using the default strategy (black edges). (From Aplin et al. [5].)

Fig. 13. A “co-sitting” network inferred from geo-location data on a population of baboons. Edge weights are proportional to the time two individuals are within 1.5m for at least 1min. (From Farine et al. [59].)

individuals are within the red cluster. Figure 12(c) shows two strong red clusters around both trained individuals, but also that the correct behavior has spread across a component of untrained individuals.

3.5.5 Testing Biological Hypotheses Using Networks as Models. Farine et al. [59] evaluates networks inferred from geo-location data of individual olive baboons (*Papio anubis*) within a troop. The authors evaluate varying networks for their ability to be used as models for predicting the future location of individuals.

Figure 13 shows an affiliation network inferred on a large timescale. Individual baboons are fitted with high-resolution GPS collars and are represented as nodes. Edge weights indicate the proportion of time two individuals are within 1.5m for at least 1min, over much of the course of data collection. The authors also generate an alternative k -nearest neighbor network model over varying k at shorter timescales. For a given node, the location centroid of *neighbors* in each of these networks is used as a prediction of the individual’s location.

This collection of networks represents different *hypotheses* for the mechanism of predictable interactions between baboons: how many individuals do baboons pay attention to at a time? How long do these relationships persist? And do individuals have fixed preferences that are more predictive than interactions over time?

This study finds that a k -nearest neighbor network of $k = 4, 5, 6$ are most predictive and statistically significant. This is biologically meaningful and agrees with previous work showing grooming cliques in baboons tend to be size 4–8 [52]. This size indicates the biological capability of baboons to manage attention and social relationships at a given time. The study also finds

that at longer timescales on the order of 10min, the network of long-term affiliations becomes more predictive. So, different networks are more appropriate models of behavior on different timescales.

3.6 Mobile Social Networks: Studying Human Mobility Through Social Relationships

3.6.1 Underlying Data in Mobile Networks. Phones and other mobile devices are among the most versatile and informative sensors of personal and social activity [110]. The modeling of mobile phone data as networks is motivated by the complex, overlapping, and dynamic modalities that are sensed by these devices. Mobile devices collect physical proximity (bluetooth, WiFi), physical location (GPS), direct communication (SMS, voice), interactions with applications, interactions through other online social networks and email [101]. Integrating these modalities promises to give a rich picture of large-scale human mobility, dynamics and scale [157], geography and communication [21, 146], and offline face-to-face social networks.

While much of this underlying data is similar to data collected for animal social networks (proximity, location, discrete interactions), there are notable trade-offs between privacy and experimental design within these domains. While animals are not due rights to data privacy, they are also unable to comply with instructions or be surveyed for ground-truth network edges.

In human experiments, contact diaries [121] or Facebook friends [158] have been collected to validate networks inferred from proximity sensors. Experiments on mobile users are necessarily less invasive, while topics such as disease spread and sexual contact networks are often sensed in animal populations. Data privacy requires careful, informed consent and secure storage [171]; location has been shown very effective to uniquely identify users using only a few data points [46].

Several mobile datasets have been collected for the purposes of social research [20]. The “Reality Mining” dataset is the first large-scale collection, on 100 participants (faculty and students) in the MIT Media Laboratory [54]. This anonymized dataset contains call logs, Bluetooth device proximity, cell tower ID (a proxy for location), and other fields. Similar mobile data collection projects followed, including the Lausanne Data Collection Campaign on 170 student participants [111], the Social fMRI study on 130 participants [3], and the SensibleDTU project of 1,000 participants [172]. These subsequent studies collected more detailed user activity, surveys, online social network activity, and detailed user demographics, addressing the limitations of previous efforts.

Finally, the SocioPatterns platform [11, 36] uses a specialized proximity sensor design to record face-to-face interactions. These sensors have been deployed in an academic conference setting [164], elementary schools [169], high schools [121], and several other environments. The specificity of these sensors for detecting individual *interactions* between users addresses the challenges of using general proximity sensors (e.g., Bluetooth) for population studies.

3.6.2 Studies and Methods on Mobile Data. The primary task in inferring networks from mobile data is related to comparison across modalities, for edge or attribute prediction. Previous work focuses on predicting Facebook friends from Bluetooth co-location [158], as well as survey-reported friends from proximity and call record data [53].

Figure 14 examines this latter task. The network in Figure 14(a) is inferred using principle components analysis of Bluetooth proximity counts across different times and locations (e.g., work, off-campus, weekday, weekend). The authors assign edge weights according to the PCA factor that by inspection corresponded to “non-work” hours (e.g., “close friends are those co-located outside of work”). Figure 14(b) reports the ground-truth social network, self-reported from a user survey, accurately reconstructed by the inferred network.

Fig. 14. (a) A weighted latent network inferred by Bluetooth proximity and call records on the population of Reality Mining users. (b) The self-reported ground-truth friendship network. (From Eagle and Pentland [53].)

Fig. 15. Two networks inferred over varying network thresholds τ , creating sparser networks from left to right as the threshold criteria becomes more strict (e.g., more than τ email interactions for an edge between individuals v_i and v_j). (From De Choudhury et al. [45].)

While the discovery that friends meet or call after work is not particularly surprising, this demonstrates the integration approach of these modalities for the simple edge prediction task. The principle components measure to infer the network edges also incorporates domain knowledge of work schedules. Previous work shows that human mobility in urban environments is highly periodic between a small set of locations (e.g., home and work) [53]. Therefore, incorporating these periodicities explicitly is a key aspect of this domain.

De Choudhury et al. [45] revisit the discussion of setting similarity threshold τ for an interaction measure calculated on data. Even under a fixed interaction measure, varying τ realizes many possible networks. In Figure 15, from left to right the number of required emails exchanged increases to define an edge, thus reducing the resulting network density. Any predictive task on this network balances novelty against the task difficulty: at a low threshold, a dense graph is realized and edge prediction may not perform better than random, because the definition of the edge is simply noise. However, a very high threshold may infer a very sparse network, where edges are trivially easy to predict (but uninteresting).

The authors generate one-year aggregated networks in a university email dataset and the Enron email dataset (note that the timescale is arbitrarily set). The authors tune the global threshold τ according to performance across several different classification tasks on sets of node-level

features from each inferred network. These tasks include classification of node class (e.g., undergraduate, graduate, faculty, staff), gender, “community” (with class labels provided by stochastic block modeling [86]), where each of these tasks may be independently of interest. If these node demographics and communities were separable as a set of behaviors at some “natural” threshold, then this analysis would discover the threshold yielding the maximal classification accuracy. The authors also predict future communication activity using simple linear regression, reporting the accuracy at these same τ . Each of these tasks yield a similar range of high performing τ thresholds, suggesting that classification and prediction agree across multiple views of the network.

4 CONCLUSION

This survey aims to provide a vocabulary and structure to the problem of inferring networks from data. Typically, this problem is addressed in data preprocessing, often with several *artful* steps of parameter tuning or feature selection. We propose investigation of this problem using a more general and rigorous methodology for building networks appropriate for data science questions.

We survey several domains to illustrate their varying questions, challenges, and how the nature of the data drives the methodological specializations in the areas. For example—with some simplification—we observe that gene regulatory networks are methodologically very mature, with a breadth of interaction measures appropriate for multivariate, *matrix* data (e.g., microarray), including regression and graphical models. Climate networks and brain networks are mature in *time series* interaction measures, including causal and frequency-based analysis, respectively. The problems in each of these areas are still exploratory, focusing on integrating and validating networks from different data (e.g., structural and functional brain networks) to develop data science tools downstream from these robust network models. Animal social networks inherit traditional *direct observation* data in relatively simple formats, with instrumented sensing now becoming more feasible. This yields relatively simple network models over straightforward parameters (e.g., distance and persistence), with a focus on experimental design. Epidemiology historically studies observed infection data *spreading* across a hidden contact network. Therefore, modeling these transmission functions is a key to this area. We hope that this *data-driven* summary might help locate models and expertise on networks derived from different underlying data modalities.

Previous work often assumes that the objective of network inference is uncovering “the network” representation that is obscured by noise. Often in this context, the network inference method tries to reconstruct known ground-truth networks from data. In contrast, our work treats a network as a model to perform a particular *task*, where we often cannot access the ground-truth network, or assume its parametric form. Analogous to clustering for a classification *task*, there are many possible clusterings that are only as valuable as the accuracy improvement they provide of a downstream task (e.g., classification) or question. Conceptualizing network inference within the complete data science methodology—from data (to network) to task models for particular questions—focuses on a tighter coupling of network models constructed from data and task models.

There are several methodologies across domains that use randomization, causality, and significance testing strategies to rigorously learn the network model under some assumptions. While these networks are appropriate according to their structural assumptions, they may not be the most informative for the question/task(s) of interest. Currently, no general, statistically rigorous methodology exists to learn and evaluate networks over particular task(s). Furthermore, there is little understanding of the criteria for network models and predictive models that would make

them appropriate for this paired evaluation. We anticipate this will be an exciting area of future research and look forward to the followup publications.

REFERENCES

- [1] Eytan Adar and Lada A. Adamic. 2005. Tracking information epidemics in blogspace. In *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence*. 207–214. <https://doi.org/10.1109/WI.2005.151>
- [2] Charu Aggarwal and Karthik Subbian. 2014. Evolutionary network analysis: A survey. *ACM Comput. Surv.* 47, 1 (May 2014), 10:1–10:36. <https://doi.org/10.1145/2601412>
- [3] Nadav Aharony, Wei Pan, Cory Ip, Inas Khayal, and Alex Pentland. 2011. Social fMRI: Investigating and shaping social mechanisms in the real world. *Pervas. Mobil. Comput.* 7, 6 (Dec. 2011), 643–659. <https://doi.org/10.1016/j.pmcj.2011.09.004>
- [4] Genevera I. Allen and Zhandong Liu. 2012. A log-linear graphical model for inferring genetic networks from high-throughput sequencing data. In *Proceedings of the 2012 IEEE International Conference on Bioinformatics and Biomedicine (BIBM'12)*. IEEE Computer Society, 1–6. <https://doi.org/10.1109/BIBM.2012.6392619>
- [5] Lucy M. Aplin, Damien R. Farine, Julie Morand-Ferron, Andrew Cockburn, Alex Thornton, and Ben C. Sheldon. 2015. Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature* 518, 7540 (Feb. 2015), 538–541. <https://doi.org/10.1038/nature13998>
- [6] Lucy M. Aplin, Damien R. Farine, Julie Morand-Ferron, and Ben C. Sheldon. 2012. Social networks predict patch discovery in a wild population of songbirds. *Proceedings of the Royal Society of London B: Biological Sciences* (2012). <https://doi.org/10.1098/rspb.2012.1591>
- [7] Mahmoudreza Babaei, Przemyslaw Grabowicz, Isabel Valera, Krishna P. Gummadi, and Manuel Gomez-Rodriguez. 2016. On the efficiency of the information networks in social media. In *Proceedings of the 9th ACM International Conference on Web Search and Data Mining (WSDM'16)*. ACM, 83–92. <https://doi.org/10.1145/2835776.2835826>
- [8] Lars Backstrom and Jon Kleinberg. 2011. Network bucket testing. In *Proceedings of the 20th International Conference on World Wide Web (WWW'11)*. ACM, 615–624. <https://doi.org/10.1145/1963405.1963492>
- [9] Shweta Bansal, Bryan T. Grenfell, and Lauren Ancel Meyers. 2007. When individual behaviour matters: Homogeneous and network models in epidemiology. *J. Roy. Soc. Interface* 4, 16 (2007), 879–891. <https://doi.org/10.1098/rsif.2007.1100>
- [10] Ziv Bar-Joseph, Anthony Gitter, and Itamar Simon. 2012. Studying and modelling dynamic biological processes using time-series gene expression data. *Nature Rev. Genet.* 13, 8 (Aug. 2012), 552–564. <https://doi.org/10.1038/nrg3244>
- [11] Alain Barrat, Ciro Cattuto, Vittoria Colizza, Jean-François Pinton, Wouter Van den Broeck, and Alessandro Vespignani. 2008. High resolution dynamical mapping of social interactions with active RFID. *ArXiv e-prints* (Nov. 2008). <http://arxiv.org/abs/0811.4170v2>
- [12] Louise Barrett, Peter Henzi, and David Lusseau. 2012. Taking sociality seriously: The structure of multi-dimensional social networks as a source of information for individuals. *Philos. Trans. Roy. Soc. London B: Biol. Sci.* 367, 1599 (2012), 2108–2118. <https://doi.org/10.1098/rstb.2012.0113>
- [13] Tanya Barrett, Stephen E. Wilhite, Pierre Ledoux, Carlos Evangelista, Irene F. Kim, Maxim Tomashevsky, Kimberly A. Marshall, Katherine H. Phillippy, Patti M. Sherman, Michelle Holko, Andrey Yefanov, Hyeseung Lee, Naigong Zhang, Cynthia L. Robertson, Nadezhda Serova, Sean Davis, and Alexandra Soboleva. 2013. NCBI GEO: Archive for functional genomics data sets—update. *Nucleic Acids Res.* 41, D1 (2013), D991–D995. <https://doi.org/10.1093/nar/gks1193>
- [14] Baruch Barzel and Albert-Laszlo Barabasi. 2013. Network link prediction by global silencing of indirect correlations. *Nature Biotechnol.* 31, 8 (Aug. 2013), 720–725. <https://doi.org/10.1038/nbt.2601>
- [15] Lars Bejder, David Fletcher, and Stefan Bräger. 1998. A method for testing association patterns of social animals. *Animal Behav.* 56, 3 (1998), 719–725. <https://doi.org/10.1006/anbe.1998.0802>
- [16] David C. Bell, John S. Atkinson, and Jerry W. Carlson. 1999. Centrality measures for disease transmission networks. *Social Networks* 21, 1 (1999), 1–21. [https://doi.org/10.1016/S0378-8733\(98\)00010-0](https://doi.org/10.1016/S0378-8733(98)00010-0)
- [17] Tanya Y. Berger-Wolf, Daniel I. Rubenstein, Charles V. Stewart, Jason A. Holmberg, Jason Parham, Sreejith Menon, Jonathan Crall, Jon Van Oast, Emre Kiciman, and Lucas Joppa. 2017. Wildbook: Crowdsourcing, computer vision, and data science for conservation. *ArXiv e-prints* (Oct. 2017). <https://arxiv.org/abs/1710.08880>
- [18] Stephan Bialonski, Martin Wendler, and Klaus Lehnertz. 2011. Unraveling spurious properties of interaction networks with tailored random networks. *PLoS ONE* 6, 8 (Aug. 2011), e22826. <https://doi.org/10.1371/journal.pone.0022826>
- [19] Concha Bielza and Pedro Larrañaga. 2014. Bayesian networks in neuroscience: A survey. *Front. Comput. Neurosci.* 8 (Oct. 2014), 131. <https://doi.org/10.3389/fncom.2014.00131>

- [20] Vincent D. Blondel, Adeline Decuyper, and Gautier Krings. 2015. A survey of results on mobile phone datasets analysis. *EPJ Data Sci.* 4, 1 (2015), 1–55. <https://doi.org/10.1140/epjds/s13688-015-0046-0>
- [21] Vincent D. Blondel, Gautier Krings, and Isabelle Thomas. 2010. Regions and borders of mobile telephony in Belgium and in the Brussels metropolitan zone. *Brussels Studies* (2010). <https://doi.org/2078.1/95261>
- [22] Stefano Boccaletti, Ginestra Bianconi, Regino Criado Herrero, Charo I. del Genio, Jesus Gómez-Gardeñes, Miguel Romance, Irene Sendiña-Nadal, Zhen Wang, and Massimiliano Zanin. 2014. The structure and dynamics of multi-layer networks. *Phys. Rep.* 544, 1 (2014), 1–122. <https://doi.org/10.1016/j.physrep.2014.07.001>
- [23] Michael Boots and Akira Sasaki. 1999. “Small worlds” and the evolution of virulence: Infection occurs locally and at a distance. *Philos. Trans. Roy. Soc. London B: Biol. Sci.* 266, 1432 (1999), 1933–1938. <https://doi.org/10.1098/rspb.1999.0869>
- [24] Ulrik Brandes, Garry Robins, Ann Mcranie, and Stanley Wasserman. 2013. What is network science? *Network Sci.* 1, 1 (2013), 1–15. <https://doi.org/10.1017/nws.2013.2>
- [25] Fred Brauer. 2008. *Compartmental Models in Epidemiology*. Springer, Berlin, 19–79. https://doi.org/10.1007/978-3-540-78911-6_2
- [26] Tom Britton and Philip D. O’Neill. 2002. Bayesian inference for stochastic epidemics in populations with random social structure. *Scand. J. Stat.* 29, 3 (Sept. 2002), 375–390. <https://doi.org/10.1111/1467-9469.00296>
- [27] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. 2011. *Handbook of Markov Chain Monte Carlo*. CRC Press. 2011021891 <https://books.google.com/books?id=qfRsAIKZ4rIC>.
- [28] Ivan Brugere and Tanya Y. Berger-Wolf. 2017. Network model selection using task-focused minimum description length. *ArXiv e-prints* (Oct. 2017). <https://arxiv.org/abs/1710.05207>
- [29] Ivan Brugere, Chris Kanich, and Tanya Y. Berger-Wolf. 2017. Evaluating social networks using task-focused network inference. In *Proceedings of the 13th International Workshop on Mining and Learning with Graphs (MLG’17)*. <https://arxiv.org/abs/1707.02385>
- [30] Ivan Brugere, Chris Kanich, and Tanya Y. Berger-Wolf. 2017. Network model selection for task-focused attributed network inference. In *Proceedings of the 2017 IEEE International Conference on Data Mining Workshop (ICDMW’17)*. <https://arxiv.org/abs/1708.06303>
- [31] Ed Bullmore and Olaf Sporns. 2009. Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Rev. Neurosci.* 10, 3 (Mar. 2009), 186–198. <https://doi.org/10.1038/nrn2575>
- [32] Atul J. Butte and Isaac S. Kohane. 2000. Mutual information relevance networks: Functional genomic clustering using pairwise entropy measurements. *Pacific Symp. Biocomput.* 5 (2000), 415–426. https://doi.org/10.1142/9789814447331_0040
- [33] Leila Cammoun, Xavier Gigandet, Djalel Meskaldji, Jean Philippe Thiran, Olaf Sporns, Kim Q. Do, Philippe Maeder, Reto Meuli, and Patric Hagmann. 2012. Mapping the human connectome at multiple scales with diffusion spectrum MRI. *J. Neurosci. Methods* 203, 2 (2012), 386–397. <https://doi.org/10.1016/j.jneumeth.2011.09.031>
- [34] Bokai Cao, Lifang He, Xiaokai Wei, Mengqi Xing, Philip S. Yu, Heide Klumpp, and Alex D. Leow. 2017. *t-BNE: Tensor-based Brain Network Embedding*. 189–197. <https://doi.org/10.1137/1.9781611974973.22>
- [35] Longbing Cao. 2017. Data science: A comprehensive overview. *ACM Comput. Surv.* 50, 3, Article 43 (June 2017), 42 pages. <https://doi.org/10.1145/3076253>
- [36] Ciro Cattuto, Wouter Van den Broeck, Alain Barrat, Vittoria Colizza, Jean-François Pinton, and Alessandro Vespignani. 2010. Dynamics of person-to-person interactions from distributed RFID sensor networks. *PLOS ONE* 5, 7 (July 2010), e11596. <https://doi.org/10.1371/journal.pone.0011596>
- [37] Hau Chan, Hanghang Tong, and Leman Akoglu. 2014. Make it or break it: Manipulating robustness in large networks. In *Proceedings of the 2014 SIAM International Conference on Data Mining*. 325–333. <https://doi.org/10.1137/1.9781611973440.37>
- [38] Darren P. Croft, Jens Krause, Safi K. Darden, Indar W. Ramnarine, Jolyon J. Faria, and Richard James. 2009. Behavioural trait assortment in a social network: Patterns and implications. *Behav. Ecol. Sociobiol.* 63, 10 (2009), 1495–1503. <https://doi.org/10.1007/s00265-009-0802-x>
- [39] Darren P. Croft, Joah R. Madden, Daniel W. Franks, and Richard James. 2011. Hypothesis testing in animal social networks. *Trends Ecol. Evol.* 26, 10 (2011), 502–507. <https://doi.org/10.1016/j.tree.2011.05.012>
- [40] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. 2017. A survey on network embedding. *ArXiv e-prints* (Nov. 2017). <https://arxiv.org/abs/1711.08752>
- [41] Luciano da Fontoura Costa, Osvaldo N. Oliveira Jr., Gonzalo Travieso, Francisco Aparecido Rodrigues, Paulino Ribeiro Villas Boas, Lucas Antiquiera, Matheus Palhares Viana, and Luis Enrique Correa Rocha. 2011. Analyzing and modeling real-world phenomena with complex networks: A survey of applications. *Adv. Phys.* 60, 3 (2011), 329–412. <https://doi.org/10.1080/00018732.2011.572452>

- [42] E. Damaraju, E. A. Allen, A. Belger, J. M. Ford, S. McEwen, D. H. Mathalon, B. A. Mueller, G. D. Pearlson, S. G. Potkin, A. Preda, J. A. Turner, J. G. Vaidya, T. G. van Erp, and V. D. Calhoun. 2014. Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage: Clin.* 5 (2014), 298–308. <https://doi.org/10.1016/j.nicl.2014.07.003>
- [43] Leon Danon, Ashley P. Ford, Thomas House, Chris P. Jewell, Matt J. Keeling, Gareth O. Roberts, Joshua V. Ross, and Matthew C. Vernon. 2011. Networks and the epidemiology of infectious disease. In *Interdisciplinary Perspectives on Infectious Diseases*, Vol. 2011. <https://doi.org/10.1155/2011/284909>
- [44] Olivier David, Isabelle Guillemain, Sandrine Saillet, Sebastien Rey, Colin Deransart, Christoph Segebarth, and Antoine Depaulis. 2008. Identifying neural drivers with functional MRI: An electrophysiological validation. *PLoS Biol.* 6, 12 (Dec. 2008), e315. <https://doi.org/10.1371/journal.pbio.0060315>
- [45] Munmun De Choudhury, Winter A. Mason, Jake M. Hofman, and Duncan J. Watts. 2010. Inferring relevant social networks from interpersonal communication. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. ACM, 301–310. <https://doi.org/10.1145/1772690.1772722>
- [46] Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Nature Sci. Rep.* 3 (Mar 2013), 1376. <https://doi.org/10.1038/srep01376>
- [47] Shermin de Silva, Ashoka D. G. Ranjeewa, and Sergey Kryazhimskiy. 2011. The dynamics of social networks among female Asian elephants. *BMC Ecol.* 11, 1 (2011), 1–16. <https://doi.org/10.1186/1472-6785-11-17>
- [48] Mukeshwar Dhamala, Govindan Rangarajan, and Mingzhou Ding. 2008. Analyzing information flow in brain networks with nonparametric granger causality. *NeuroImage* 41, 2 (2008), 354–362. <https://doi.org/10.1016/j.neuroimage.2008.02.020>
- [49] Jonathan F. Donges, Yong Zou, Norbert Marwan, and Juergen Kurths. 2009. Complex networks in climate dynamics. *Eur. Phys. J. Special Topics* 174, 1 (2009), 157–179. <https://doi.org/10.1140/epjst/e2009-01098-2>
- [50] Jonathan F. Donges, Yong Zou, Norbert Marwan, and Juergen Kurths. 2009. The backbone of the climate network. *Europhys. Lett.* 87, 4 (2009), 48007. <https://doi.org/10.1209/0295-5075/87/48007>
- [51] Nan Du, Le Song, Ming Yuan, and Alex J. Smola. 2012. Learning networks of heterogeneous influence. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). 2780–2788. <http://papers.nips.cc/paper/4582-learning-networks-of-heterogeneous-influence>.
- [52] Robin M. Dunbar. 1992. Time: A hidden constraint on the behavioural ecology of baboons. *Behav. Ecol. Sociobiol.* 31, 1 (01 Jul 1992), 35–49. <https://doi.org/10.1007/BF00167814>
- [53] Nathan Eagle and Alex Sandy Pentland. 2009. Eigenbehaviors: Identifying structure in routine. *Behav. Ecol. Sociobiol.* 63, 7 (2009), 1057–1066. <https://doi.org/10.1007/s00265-009-0739-0>
- [54] Nathan Eagle and Alex (Sandy) Pentland. 2006. Reality mining: Sensing complex social systems. *Personal Ubiqu. Comput.* 10, 4 (Mar. 2006), 255–268. <https://doi.org/10.1007/s00779-005-0046-3>
- [55] Imme Ebert-Uphoff and Yi Deng. 2012. Causal discovery for climate research using graphical models. *J. Climate* 25, 17 (Feb. 2012), 5648–5665. <https://doi.org/10.1175/JCLI-D-11-00387.1>
- [56] Bradley Efron and R. J. Tibshirani. 1993. *An Introduction to the Bootstrap*. Chapman & Hall.
- [57] Jeremiah J. Faith, Boris Hayete, Joshua T. Thaden, Ilaria Mogno, Jamey Wierzbowski, Guillaume Cottarel, Simon Kasif, James J. Collins, and Timothy S. Gardner. 2007. Large-scale mapping and validation of *Escherichia coli* transcriptional regulation from a compendium of expression profiles. *PLoS Biol.* 5, 1 (Jan. 2007), e8. <https://doi.org/10.1371/journal.pbio.0050008>
- [58] Damien R. Farine. 2017. A guide to null models for animal social network analysis. *Methods Ecol. Evol.* 8, 10 (2017), 1309–1320. <https://doi.org/10.1111/2041-210X.12772>
- [59] Damien R. Farine, Ariana Strandburg-Peshkin, Tanya Y. Berger-Wolf, Brian Ziebart, Ivan Brugere, Jia Li, and Margaret C. Crofoot. 2016. Both nearest neighbours and long-term affiliates predict individual locations during collective movement in wild baboons. In *Nature Scientific Reports. Sci. Rep.* 6, 27704. <https://doi.org/10.1038/srep27704>
- [60] Damien R. Farine and Hal Whitehead. 2015. Constructing, conducting and interpreting animal social network analysis. *J. Animal Ecol.* 84, 5 (2015), 1144–1163. <https://doi.org/10.1111/1365-2656.12418>
- [61] Soheil Feizi, Daniel Marbach, Muriel Medard, and Manolis Kellis. 2013. Network deconvolution as a general method to distinguish direct dependencies in networks. *Nature Biotech.* 31, 8 (Aug. 2013), 726–733. <https://doi.org/10.1038/nbt.2635>
- [62] Benjamin Fish and Rajmonda S. Caceres. 2017. A task-driven approach to timescale detection in dynamic networks. In *Proceedings of the 13th International Workshop on Mining and Learning with Graphs (MLG'17)*. http://www.mlworkshop.org/2017/paper/MLG2017_paper_17.pdf.
- [63] Robert J. Fletcher, Miguel A. Acevedo, Brian E. Reichert, Kyle E. Pias, and Wiley M. Kitchens. 2011. Social network models predict movement and connectivity in ecological landscapes. *Proc. Natl. Acad. Sci. U.S.A.* 108, 48 (2011), 19282–19287. <https://doi.org/10.1073/pnas.1107549108>

- [64] Karl J. Friston, Baojuan Li, Jean Daunizeau, and Klaas E. Stephan. 2011. Network discovery with DCM. *NeuroImage* 56, 3 (Jun. 2011), 1202–1221. <https://doi.org/10.1016/j.neuroimage.2010.12.039>
- [65] Lise Getoor, Nir Friedman, Daphne Koller, Avi Pfeffer, and Ben Taskar. 2007. Probabilistic relational models. *Introduction to Statistical Relational Learning* (2007), 129. <https://mitpress.mit.edu/books/introduction-statistical-relational-learning>.
- [66] Anna Goldenberg, Alice X. Zheng, Stephen E. Fienberg, and Edoardo M. Airolidi. 2010. A survey of statistical network models. *Found. Trends Mach. Learn.* 2, 2 (Feb. 2010), 129–233. <https://doi.org/10.1561/2200000005>
- [67] Manuel Gomez-Rodriguez, Jure Leskovec, David Baldazzi, and Bernhard Schölkopf. 2014. Uncovering the structure and temporal dynamics of information propagation. *Network Sci.* 2, 1, 26–65. <https://doi.org/10.1017/nws.2014.3>
- [68] Manuel Gomez-Rodriguez, Jure Leskovec, and Andreas Krause. 2012. Inferring networks of diffusion and influence. *ACM Trans. Knowl. Discov. Data* 5, 4 (Feb. 2012), 21:1–21:37. <https://doi.org/10.1145/2086737.2086741>
- [69] Neil Zhenqiang Gong, Ameet Talwalkar, Lester Mackey, Ling Huang, Eui Chul Richard Shin, Emil Stefanov, Elaine (Runting) Shi, and Dawn Song. 2014. Joint link prediction and attribute inference using a social-attribute network. *ACM Trans. Intell. Syst. Technol.* 5, 2 (Apr. 2014), 27:1–27:20. <https://doi.org/10.1145/2594455>
- [70] Palash Goyal and Emilio Ferrara. 2017. Graph embedding techniques, applications, and performance: A survey. *CoRR* abs/1705.02801 (2017). <http://arxiv.org/abs/1705.02801>
- [71] Clive W. J. Granger. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37, 3 (1969), 424–438. <https://doi.org/10.2307/1912791>
- [72] Michael D. Greicius, Ben Krasnow, Allan L. Reiss, and Vinod Menon. 2003. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proc. Natl. Acad. Sci. U.S.A.* 100, 1 (2003), 253–258. <https://doi.org/10.1073/pnas.0135058100>
- [73] Chris Groendyke, David Welch, and David R. Hunter. 2011. Bayesian inference for contact networks given epidemic data. *Scand. J. Stat.* 38, 3 (Sept. 2011), 600–616. <https://doi.org/10.1111/j.1467-9469.2010.00721.x>
- [74] Huan Gui, Ya Xu, Anmol Bhasin, and Jiawei Han. 2015. Network A/B testing: From sampling to estimation. In *Proceedings of the 24th International Conference on World Wide Web (WWW'15)*. ACM, 399–409. <https://doi.org/10.1145/2736277.2741081>
- [75] Hamed Haddadi, Andrew J. King, Alison P. Wills, Damien Fay, John Lowe, A. Jennifer Morton, Stephen Hailes, and Alan M. Wilson. 2011. Determining association networks in social animals: Choosing spatial-temporal criteria and sampling rates. *Behav. Ecol. Sociobiol.* 65, 8 (2011), 1659–1668. <https://doi.org/10.1007/s00265-011-1193-3>
- [76] Hamed Haddadi, Miguel Rio, Gianluca Iannaccone, Andrew Moore, and Richard Mortier. 2008. Network topologies: Inference, modeling, and generation. *IEEE Commun. Surveys Tutor* 10, 2 (2008), 48–69. <https://doi.org/10.1109/COMST.2008.4564479>
- [77] Rodrigo K. Hamede, Jim Bashford, Hamish McCallum, and Menna Jones. 2009. Contact networks in a wild Tasmanian devil (*Sarcophilus harrisii*) population: Using social network analysis to reveal seasonal variability in social behaviour and its implications for transmission of devil facial tumour disease. *Ecol. Lett.* 12, 11 (2009), 1147–1157. <https://doi.org/10.1111/j.1461-0248.2009.01370.x>
- [78] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Representation learning on graphs: Methods and applications. *ArXiv e-prints* (Sept. 2017). <https://arxiv.org/abs/1709.05584>
- [79] Mohammad Al Hasan and Mohammed J. Zaki. 2011. A survey of link prediction in social networks. In *Social Network Data Analytics SE - 9*, Charu C. Aggarwal (Ed.). Springer US, 243–275. https://doi.org/10.1007/978-1-4419-8462-3_9
- [80] Anne-Claire Haury, Fantine Mordelet, Paola Vera-Licona, and Jean-Philippe Vert. 2012. TIGRESS: Trustful inference of gene regulation using stability selection. *BMC Syst. Biol.* 6, 1 (2012), 145. <https://doi.org/10.1186/1752-0509-6-145>
- [81] S. Havlin, D. Y. Kenett, E. Ben-Jacob, A. Bunde, R. Cohen, H. Hermann, J. W. Kantelhardt, J. Kertész, S. Kirkpatrick, J. Kurths, J. Portugali, and S. Solomon. 2012. Challenges in network science: Applications to infrastructures, climate, social systems and economics. *Eur. Phys. J. Special Topics* 214, 1 (2012), 273–293. <https://doi.org/10.1140/epjst/e2012-01695-x>
- [82] D. T. Haydon, M. Chase-Topping, D. J. Shaw, L. Matthews, J. K. Friar, J. Wilesmith, and M. E. J. Woolhouse. 2003. The construction and analysis of epidemic trees with reference to the 2001 UK foot-and-mouth outbreak. *Proc. Roy. Soc. London B: Biol. Sci.* 270, 1511 (Jan. 2003), 121–127. <https://doi.org/10.1098/rspb.2002.2191>
- [83] Michael Hecker, Sandro Lambeck, Susanne Toepfer, Eugene van Someren, and Reinhard Guthke. 2009. Gene regulatory network inference: Data integration in dynamic models—A review. *Biosystems* 96, 1 (2009), 86–103. <https://doi.org/10.1016/j.biosystems.2008.12.004>
- [84] Anthony Hey, Stewart Tansley, and Kristin Tolle. 2009. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Microsoft Research.
- [85] Jaroslav Hlinka, David Hartman, Martin Vejmelka, Jakob Runge, Norbert Marwan, Jürgen Kurths, and Milan Paluš. 2013. Reliability of inference of directed climate networks using conditional mutual information. *Entropy* 15, 6 (2013), 2023. <https://doi.org/10.3390/e15062023>

- [86] Jake M. Hofman and Chris H. Wiggins. 2008. Bayesian approach to network modularity. *Phys. Rev. Lett.* 100, 25 (Jun. 2008), 258701. <https://doi.org/10.1103/PhysRevLett.100.258701>
- [87] Petter Holme. 2015. Modern temporal network theory: A colloquium. *Eur. Phys. J. B.* 88, 9 (2015), 1–30. <https://doi.org/10.1140/epjb/e2015-60657-4>
- [88] Petter Holme and Jari Saramäki. 2012. Temporal networks. *Phys. Rep.* 519, 3 (2012), 97–125. <https://doi.org/10.1016/j.physrep.2012.03.001>
- [89] R. Matthew Hutchison, Thilo Womelsdorf, Elena A. Allen, Peter A. Bandettini, Vince D. Calhoun, Maurizio Corbetta, Stefania Della Penna, Jeff H. Duyn, Gary H. Glover, Javier Gonzalez-Castillo, Daniel A. Handwerker, Shella Keilholz, Vesa Kiviniemi, David A. Leopold, Francesco de Pasquale, Olaf Sporns, Martin Walter, and Catie Chang. 2013. Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage* 80 (Oct. 2013), 360–378. <https://doi.org/10.1016/j.neuroimage.2013.05.079>
- [90] Alon Itai and Michael Rodeh. 1978. Finding a minimum circuit in a graph. *SIAM J. Comput.* 7, 4 (Nov. 1978), 413–423. <https://doi.org/10.1137/0207033>
- [91] Richard James, Darren P. Croft, and Jens Krause. 2009. Potential banana skins in animal social network analysis. *Behav. Ecol. Sociobiol.* 63, 7 (2009), 989–997. <https://doi.org/10.1007/s00265-009-0742-5>
- [92] Jaya Kawale, Snigdhasu Chatterjee, Dominick Ormsby, Karsten Steinhaeuser, Stefan Liess, and Vipin Kumar. 2012. Testing the significance of spatio-temporal teleconnection patterns. In *Proceedings of ACM SIGKDD 2012 International Conference on Knowledge Discovery and Data Mining (KDD'12)*. ACM, 642–650. <https://doi.org/10.1145/2339530.2339634>
- [93] Jaya Kawale, Stefan Liess, Arjun Kumar, Michael Steinbach, Peter Snyder, Vipin Kumar, Auroop R. Ganguly, Nagiza F. Samatova, and Fredrick Semazzi. 2013. A graph-based approach to find teleconnections in climate data. *Stat. Anal. Data Min.* 6, 3 (2013), 158–179. <https://doi.org/10.1002/sam.11181>
- [94] Roland Kays, Margaret C. Crofoot, Walter Jetz, and Martin Wikelski. 2015. Terrestrial animal tracking as an eye on life and planet. *Science* 348, 6240 (2015). <https://doi.org/10.1126/science.aaa2478>
- [95] Matt J. Keeling and Ken T. D. Eames. 2005. Networks and epidemic models. *J. Roy. Soc. Interface* 2, 4 (Sept. 2005), 295–307. <https://doi.org/10.1098/rsif.2005.0051>
- [96] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of ACM SIGKDD 2003 International Conference on Knowledge Discovery and Data Mining (KDD'03)*. ACM, 137–146. <https://doi.org/10.1145/956750.956769>
- [97] Angelika Kimmig, Lilyana Mihalkova, and Lise Getoor. 2014. Lifted graphical models: A survey. *Mach. Learn.* 99, 1 (2014), 1–45. <https://doi.org/10.1007/s10994-014-5443-2>
- [98] Mikko Kivelä, Alex Arenas, Marc Barthélemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. 2014. Multi-layer networks. *J. Complex Networks* 2, 3 (2014), 203–271. <https://doi.org/10.1093/comnet/cnu016>
- [99] Ariel Kleiner, Ameet Talwalkar, Purnamrita Sarkar, and Michael I. Jordan. 2014. A scalable bootstrap for massive data. *J. Roy. Stat. Soc.: Series B (Statistical Methodology)* 76, 4 (2014), 795–816. <https://doi.org/10.1111/rssb.12050>
- [100] Florian Klimm, Danielle S. Bassett, Jean M. Carlson, and Peter J. Mucha. 2014. Resolving structural variability in network models and the brain. *PLoS Comput. Biol.* 10, 3 (Mar. 2014), e1003491. <https://doi.org/10.1371/journal.pcbi.1003491>
- [101] Bryan Klimt and Yiming Yang. 2004. The enron corpus: A new dataset for email classification research. In *Proceedings of the Conference on Machine Learning (ECML'04)*. Springer, 217–226. https://doi.org/10.1007/978-3-540-30115-8_22
- [102] Ron Kohavi, Alex Deng, Brian Frasca, Toby Walker, Ya Xu, and Nils Pohlmann. 2013. Online controlled experiments at large scale. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'13)*. ACM, New York, NY, 1168–1176. <https://doi.org/10.1145/2487575.2488217>
- [103] Eric D. Kolaczyk. 2009. Network topology inference. In *Statistical Analysis of Network Data SE-7*. Springer, New York, 1–48. https://doi.org/10.1007/978-0-387-88146-1_7
- [104] Mark A. Kramer, Uri T. Eden, Sydney S. Cash, and Eric D. Kolaczyk. 2009. Network inference with confidence from multivariate time series. *Phys. Rev. E* 79, 6 (Jun. 2009), 61916. <https://doi.org/10.1103/PhysRevE.79.061916>
- [105] Mark A. Kramer, Eric D. Kolaczyk, and Heidi E. Kirsch. 2008. Emergent network topology at seizure onset in humans. *Epilepsy Res.* 79, 2–3 (2008), 173–186. <https://doi.org/10.1016/j.epilepsyres.2008.02.002>
- [106] Jens Krause, Stefan Krause, Robert Arlinghaus, Ioannis Psorakis, Stephen Roberts, and Christian Rutz. 2013. Reality mining of animal social systems. *Trends Ecol. Evol.* 28, 9 (Sept. 2013), 541–551. <https://doi.org/10.1016/j.tree.2013.06.002>
- [107] Marlene Kretschmer, Dim Coumou, Jonathan F. Donges, and Jakob Runge. 2016. Using causal effect networks to analyze different arctic drivers of mid-latitude winter circulation. *J. Climate* (Mar 2016). <https://doi.org/10.1175/JCLI-D-15-0654.1>
- [108] Timothy La Fond and Jennifer Neville. 2010. Randomization tests for distinguishing social influence and homophily effects. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. ACM, 601–610. <https://doi.org/10.1145/1772690.1772752>

- [109] Jean-Philippe Lachaux, Antoine Lutz, David Rudrauf, Diego Cosmelli, Michel Le Van Quyen, Jacques Martinerie, and Francisco Varela. 2002. Estimating the time-course of coherence between single-trial brain signals: An introduction to wavelet coherence. *Neurophys. Clin./Clin. Neurophysiol.* 32, 3 (Jun. 2002), 157–174. [https://doi.org/10.1016/S0987-7053\(02\)00301-5](https://doi.org/10.1016/S0987-7053(02)00301-5)
- [110] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell. 2010. A survey of mobile phone sensing. *IEEE Commun. Mag.* 48, 9 (Sept. 2010), 140–150. <https://doi.org/10.1109/MCOM.2010.5560598>
- [111] Juha K. Laurila, Daniel Gatica-Perez, Imad Aad, Jan Blom, Olivier Bornet, Trinh Minh Tri Do, Olivier Dousse, Julien Eberle, and Markus Miettinen. 2013. From big smartphone data to worldwide research: The mobile data challenge. *Pervas. Mobile Comput.* 9, 6 (Dec. 2013), 752–771. <https://doi.org/10.1016/j.pmcj.2013.07.014>
- [112] David Lazer. 2011. Networks in political science: Back to the future. *Polit. Sci. Pol.* 44, 01 (2011), 61–68. <https://doi.org/10.1017/S1049096510001873>
- [113] Sophie Lebre, Jennifer Becq, Frederic Devaux, Michael Stumpf, and Gaelle Lelandaïs. 2010. Statistical inference of the time-varying structure of gene-regulation networks. *BMC Syst. Biol.* 4, 1 (2010), 130. <https://doi.org/10.1186/1752-0509-4-130>
- [114] Jure Leskovec and Eric Horvitz. 2014. Geospatial structure of a planetary-scale social network. *IEEE Trans. Comput. Social Syst.* 1, 3 (Sept. 2014), 156–163. <https://doi.org/10.1109/TCSS.2014.2377789>
- [115] S. Levitus, J. I. Antonov, O. K. Baranova, T. P. Boyer, C. L. Coleman, H. E. Garcia, A. I. Grodsky, D. R. Johnson, R. A. Locarnini, A. V. Mishonov, J. R. Reagan, C. L. Sazama, D. Seidov, I. Smolyar, E. S. Yarosh, and M. M. Zweng. 2013. The world ocean database. *Data Sci. J.* 12 (2013), WDS229–WDS234. <https://doi.org/10.2481/dsj.WDS-041>
- [116] Huai Li, Jianhua Xuan, Yue Wang, and Ming Zhan. 2008. Inferring regulatory networks. *Front. Biosci.* 13 (2008), 263–275. <https://doi.org/10.2741/2677>
- [117] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. *J. Amer. Soc. Info. Sci. Technol.* 58, 7 (May 2007), 1019–1031. <https://doi.org/10.1002/asi.20591>
- [118] Linyuan Lü and Tao Zhou. 2011. Link prediction in complex networks: A survey. *Physica A: Stat. Mech. Appl.* 390, 6 (2011), 1150–1170. <https://doi.org/10.1016/j.physa.2010.11.027>
- [119] Subramani Mani and Gregory F. Cooper. 2004. Causal discovery using a bayesian local causal discovery algorithm. *Studies Health Technol. Info.* 107, Pt 1 (2004), 731–735. <https://doi.org/10.3233/978-1-60750-949-3-731>
- [120] Daniel Marbach, James C. Costello, Robert Kuffner, Nicole M. Vega, Robert J. Prill, Diogo M. Camacho, Kyle R. Allison, Manolis Kellis, James J. Collins, and Gustavo Stolovitzky. 2012. Wisdom of crowds for robust gene network inference. *Nature Methods* 9, 8 (Aug. 2012), 796–804. <https://doi.org/10.1038/nmeth.2016>
- [121] Rossana Mastrandrea, Julie Fournet, and Alain Barrat. 2015. Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PLoS ONE* 10, 9 (Sept. 2015), e0136497. <https://doi.org/10.1371/journal.pone.0136497>
- [122] Julian McAuley, Rahul Pandey, and Jure Leskovec. 2015. Inferring networks of substitutable and complementary products. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'15)*. ACM, 785–794. <https://doi.org/10.1145/2783258.2783381>
- [123] David B. McDonald. 2007. Predicting fate from early connectivity in a social network. *Proc. Natl. Acad. Sci. U.S.A.* 104, 26 (Jun. 2007), 10910–10914. <https://doi.org/10.1073/pnas.0701159104>
- [124] Christopher Meek. 1995. Causal inference and causal explanation with background knowledge. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence (UAI'95)*. Morgan Kaufmann Publishers Inc., 403–410. <http://dl.acm.org/citation.cfm?id=2074204>
- [125] Patrick Meyer, Kevin Kontos, Frederic Lafitte, and Gianluca Bontempi. 2007. Information-theoretic inference of large transcriptional regulatory networks. *EURASIP J. Bioinfo. Syst. Biol.* 2007, 1 (2007), 79879. <https://doi.org/10.1155/2007/79879>
- [126] Lauren Ancel Meyers, Babak Pourbohloul, M. E. J. Newman, Danuta M. Skowronski, and Robert C. Brunham. 2005. Network theory and SARS: Predicting outbreak diversity. *J. Theoret. Biol.* 232, 1 (2005), 71–81. <https://doi.org/10.1016/j.jtbi.2004.07.026>
- [127] Isobel Milns, Colin M. Beale, and V. Anne Smith. 2010. Revealing ecological networks using Bayesian network inference algorithms. *Ecology* 91, 7 (Feb. 2010), 1892–1899. <https://doi.org/10.1890/09-0731.1>
- [128] Sach Mukherjee and Terence P. Speed. 2008. Network inference using informative priors. *Proc. Natl. Acad. Sci. U.S.A.* 105, 38 (Sept. 2008), 14313–14318. <https://doi.org/10.1073/pnas.0802272105>
- [129] Seth Myers and Jure Leskovec. 2010. On the convexity of latent social network inference. In *Advances in Neural Information Processing Systems* 23. 1741–1749. <http://papers.nips.cc/paper/4113-on-the-convexity-of-latent-social-network-inference>.
- [130] Galileo Mark Namata, Ben London, and Lise Getoor. 2015. Collective graph identification. *ACM Transactions on Knowledge Discovery from Data*. <https://doi.org/10.1145/2818378>
- [131] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Math. Program.* 14, 1 (1978), 265–294. <https://doi.org/10.1007/BF01588971>

- [132] Praneeth Netrapalli and Sujay Sanghavi. 2012. Learning the graph of epidemic cascades. *SIGMETRICS Perform. Eval. Rev.* 40, 1 (Jun. 2012), 211–222. <https://doi.org/10.1145/2318857.2254783>
- [133] Jian Ni, Haiyong Xie, Sekhar Tatikonda, and Yang Richard Yang. 2010. Efficient and dynamic routing topology inference from end-to-end measurements. *IEEE/ACM Trans. Network.* 18, 1 (Feb. 2010), 123–135. <https://doi.org/10.1109/TNET.2009.2022538>
- [134] M. Paluš, D. Hartman, J. Hlinka, and M. Vejmelka. 2011. Discerning connectivity from dynamics in climate networks. *Nonlin. Process. Geophys.* 18, 5 (2011), 751–763. <https://doi.org/10.5194/npg-18-751-2011>
- [135] Evangelos E. Papalexakis, Alona Fyshe, Nicholas D. Sidiropoulos, Partha Pratim Talukdar, Tom M. Mitchell, and Christos Faloutsos. 2014. Good-enough brain model: Challenges, algorithms and discoveries in multi-subject experiments. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'14)*. ACM, 95–104. <https://doi.org/10.1145/2623330.2623639>
- [136] D. M. Patrick, M. L. Rekart, A. Jolly, S. Mak, M. Tyndall, J. Maginley, E. Wong, T. Wong, H. Jones, C. Montgomery, and R. C. Brunham. 2002. Heterosexual outbreak of infectious syphilis: Epidemiological and ethnographic analysis and implications for control. *Sex. Transm. Infect.* 78, suppl 1 (2002), i164–i169. https://doi.org/10.1136/sti.78.suppl_1.i164
- [137] Joseph J. Pfeiffer III, Sebastian Moreno, Timothy La Fond, Jennifer Neville, Jennifer Fewell, and David B. McDonald. 2014. Attributed graph models: Modeling network structure with correlated attributes. In *Proceedings of the 23rd International Conference on World Wide Web (WWW'14)*. ACM, 831–842. <https://doi.org/10.1145/2566486.2567993>
- [138] Gert Pfurtscheller and Colin Andrew. 1999. Event-related changes of band power and coherence: Methodology and interpretation. *J. Clin. Neurophysiol.* 16, 6 (Nov. 1999), 512–519. <http://dx.doi.org/10.1097/00004691-199911000-00003>
- [139] Noa Pinter-Wollman, Elizabeth A. Hobson, Jennifer E. Smith, Andrew J. Edelman, Daizaburo Shizuka, Shermin de Silva, James S. Waters, Steven D. Prager, Takao Sasaki, George Wittemyer, Jennifer Fewell, and David B. McDonald. 2013. The dynamics of animal social networks: Analytical, conceptual, and theoretical advances. *Behavioral Ecology*. Oxford University Press. <https://doi.org/10.1093/beheco/art047>
- [140] S. C. Ponten, F. Bartolomei, and C. J. Stam. 2016. Small-world networks and epilepsy: Graph theoretical analysis of intracerebrally recorded mesial temporal lobe seizures. *Clin. Neurophysiol.* 118, 4 (Apr. 2016), 918–927. <https://doi.org/10.1016/j.clinph.2006.12.002>
- [141] Stephen R. Proulx, Daniel E. L. Promislow, and Patrick C. Phillips. 2005. Network thinking in ecology and evolution. *Trends Ecol. Evol.* 20, 6 (2005), 345–353. <https://doi.org/10.1016/j.tree.2005.04.004>
- [142] Nataša Pržulj, Derek G. Corneil, and Igor Jurisica. 2004. Modeling interactome: Scale-free or geometric? *Bioinformatics* 20, 18 (2004), 3508–3515. <https://doi.org/10.1093/bioinformatics/bth436>
- [143] Ioannis Psorakis, Stephen J. Roberts, Iead Rezek, and Ben C. Sheldon. 2012. Inferring social network structure in ecological systems from spatio-temporal data streams. *J. Roy. Soc. Interface* 9, 76 (2012), 3055–3066. <https://doi.org/10.1098/rsif.2012.0223>
- [144] Manish Purohit, B. Aditya Prakash, Chanhyun Kang, Yao Zhang, and V. S. Subrahmanian. 2014. Fast influence-based coarsening for large networks. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'14)*. ACM, 1296–1305. <https://doi.org/10.1145/2623330.2623701>
- [145] J. D. Ramsey, S. J. Hanson, C. Hanson, Y. O. Halchenko, R. A. Poldrack, and C. Glymour. 2010. Six problems for causal inference from fMRI. *NeuroImage* 49, 2 (2010), 1545–1558. <https://doi.org/10.1016/j.neuroimage.2009.08.065>
- [146] Carlo Ratti, Stanislav Sobolevsky, Francesco Calabrese, Clío Andris, Jonathan Reades, Mauro Martino, Rob Claxton, and Steven H. Strogatz. 2010. Redrawing the map of great britain from a network of human interactions. *PLoS ONE* 5, 12 (Dec. 2010), e14248. <https://doi.org/10.1371/journal.pone.0014248>
- [147] Jaap C. Reijneveld, Sophie C. Ponten, Henk W. Berendse, and Cornelis J. Stam. 2007. The application of graph theoretical analysis to complex networks in the brain. *Clin. Neurophysiol.* 118, 11 (Dec. 2007), 2317–2331. <https://doi.org/10.1016/j.clinph.2007.08.010>
- [148] Bruno Ribeiro, Nicola Perra, and Andrea Baronchelli. 2013. Quantifying the effect of temporal resolution on time-varying networks. *Sci. Rep.* 3 (Oct. 2013), 3006. <http://dx.doi.org/10.1038/srep03006>
- [149] Lucy F. Robinson, Lauren Y. Atlas, and Tor D. Wager. 2015. Dynamic functional connectivity using state-based dynamic community structure: Method and application to opioid analgesia. *NeuroImage* 108 (2015), 274–291. <https://doi.org/10.1016/j.neuroimage.2014.12.034>
- [150] Alard Roebroeck, Elia Formisano, and Rainer Goebel. 2005. Mapping directed influence over the brain using granger causality and fMRI. *NeuroImage* 25, 1 (2005), 230–242. <https://doi.org/10.1016/j.neuroimage.2004.11.017>
- [151] M. J. Rosa, K. Friston, and W. Penny. 2012. Post-hoc selection of dynamic causal models. *J. Neurosci. Methods* 208, 1 (2012), 66–78. <https://doi.org/10.1016/j.jneumeth.2012.04.013>
- [152] Mikail Rubinov and Olaf Sporns. 2010. Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* 52, 3 (Sept. 2010), 1059–1069. <https://doi.org/10.1016/j.neuroimage.2009.10.003>

- [153] Jakob Runge, Vladimir Petoukhov, and Jürgen Kurths. 2013. Quantifying the strength and delay of climatic interactions: The ambiguities of cross correlation and a novel measure based on graphical models. *J. Climate* 27, 2 (Sept. 2013), 720–739. <https://doi.org/10.1175/JCLI-D-13-00159.1>
- [154] Christian Rutz, Zackory T. Burns, Richard James, Stefanie M. H. Ismar, John Burt, Brian Otis, Jayson Bowen, and James J. H. St Clair. 2012. Automated mapping of social networks in wild birds. *Curr. Biol.* 22, 17 (2012), R669–R671. <https://doi.org/10.1016/j.cub.2012.06.037>
- [155] Suranjana Saha, Shrinivas Moorthi, Xingren Wu, Jiande Wang, Sudhir Nadiga, Patrick Tripp, David Behringer, Yu-Tai Hou, Hui-ya Chuang, Mark Iredell, Michael Ek, Jesse Meng, Rongqian Yang, Malaquías Peña Mendez, Huug van den Dool, Qin Zhang, Wanqiu Wang, Mingyue Chen, and Emily Becker. 2014. The NCEP climate forecast system version 2. *J. Climate* 27, 6 (2014), 2185–2208. <https://doi.org/10.1175/JCLI-D-12-00823.1>
- [156] Vangelis Sakkalis. 2011. Review of advanced techniques for the estimation of brain connectivity measured with EEG/MEG. *Comp. Biol. Med.* 41, 12 (2011), 1110–1117. <https://doi.org/10.1016/j.compbiomed.2011.06.020>
- [157] Jari Saramäki and Esteban Moro. 2015. From seconds to months: An overview of multi-scale dynamics of mobile telephone calls. *Eur. Phys. J. B.* 88, 6 (2015), 1–10. <https://doi.org/10.1140/epjb/e2015-60106-6>
- [158] Vedran Sekara and Sune Lehmann. 2014. The strength of friendship ties in proximity sensor data. *PLoS ONE* 9, 7 (July 2014), e100915. <https://doi.org/10.1371/journal.pone.0100915>
- [159] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI Mag.* 29, 3 (2008), 93. <http://hdl.handle.net/1903/7546>
- [160] Jay Shendure and Hanlee Ji. 2008. Next-generation DNA sequencing. *Nature Biotechnol.* 26, 10 (Oct. 2008), 1135–1145. <https://doi.org/10.1038/nbt1486>
- [161] Caroline Siegenthaler and Rudiyanto Gunawan. 2014. Assessment of network inference methods: How to cope with an underdetermined problem. *PLoS ONE* 9, 3 (03 2014), 1–7. <https://doi.org/10.1371/journal.pone.0090481>
- [162] Chao Sima, Jianping Hua, and Sungwon Jung. 2009. Inference of gene regulatory networks using time-series data: A survey. *Curr. Genom.* 10, 6 (Sept. 2009), 416–429. <https://doi.org/10.2174/138920209789177610>
- [163] Sean L. Simpson, F. DuBois Bowman, and Paul J. Laurienti. 2013. Analyzing complex functional brain networks: Fusing statistics and network science to understand the brain. *Stat. Surveys* 7 (2013), 1–36. <https://doi.org/10.1214/13-SS103>
- [164] Timo Smieszek, Stefanie Castell, Alain Barrat, Ciro Cattuto, Peter J. White, and Gérard Krause. 2016. Contact diaries versus wearable proximity sensors in measuring contact patterns at a conference: Method comparison and participants’ attitudes. *BMC Infect. Dis.* 16, 1 (2016), 1–14. <https://doi.org/10.1186/s12879-016-1676-y>
- [165] Orr Spiegel, Stephan T. Leu, Andrew Sih, and C. Michael Bull. 2016. Socially interacting or indifferent neighbours? Randomization of movement paths to tease apart social preference and spatial constraints. *Methods Ecol. Evol.* 7, 8 (2016), 971–979. <https://doi.org/10.1111/2041-210X.12553>
- [166] Olaf Sporns. 2014. Contributions and challenges for network models in cognitive neuroscience. *Nature Neurosci.* 17, 5 (May 2014), 652–660. <https://doi.org/10.1038/nn.3690>
- [167] Olaf Sporns and Richard F. Betzel. 2016. Modular brain networks. *Annu. Rev. Psychol.* 67, 1 (Jan. 2016), 613–640. <https://doi.org/10.1146/annurev-psych-122414-033634>
- [168] J. Conrad Stack, Shweta Bansal, V. S. Anil Kumar, and Bryan Grenfell. 2013. Inferring population-level contact heterogeneity from common epidemic data. *J. Roy. Soc. Interface* 10, 78 (2013). <https://doi.org/10.1098/rsif.2012.0578>
- [169] Juliette Stehlé, François Charbonnier, Tristan Picard, Ciro Cattuto, and Alain Barrat. 2013. Gender homophily from spatial behavior in a primary school: A sociometric study. *Social Networks* 35, 4 (2013), 604–613. <https://doi.org/10.1016/j.socnet.2013.08.003>
- [170] Karsten Steinhäuser, Nitesh V. Chawla, and Auroop R. Ganguly. 2011. Complex networks as a unified framework for descriptive analysis and predictive modeling in climate science. *Stat. Anal. Data Min.* 4, 5 (2011), 497–511. <https://doi.org/10.1002/sam.10100>
- [171] Arkadiusz Stopczynski, Riccardo Pietri, Alex Pentland, David Lazer, and Sune Lehmann. 2014. Privacy in sensor-driven human data collection: A guide for practitioners. *ArXiv e-prints* (Mar. 2014). <https://arxiv.org/abs/1403.5299>
- [172] Arkadiusz Stopczynski, Vedran Sekara, Piotr Sapiezynski, Andrea Cuttone, Mette My Madsen, Jakob Eg Larsen, and Sune Lehmann. 2014. Measuring large-scale social networks with high resolution. *PLoS ONE* 9, 4 (Apr. 2014), e95978. <https://doi.org/10.1371/journal.pone.0095978>
- [173] Cédric Sueur, Armand Jacobs, Frédéric Amblard, Odile Petit, and Andrew J. King. 2011. How can social network analysis improve the study of primate behavior? *Amer. J. Primatol.* 73, 8 (Aug. 2011), 703–19. <https://doi.org/10.1002/ajp.20915>
- [174] Rajmonda Sulo, Tanya Berger-Wolf, and Robert Grossman. 2010. Meaningful selection of temporal resolution for dynamic networks. In *Proceedings of the 8th Workshop on Mining and Learning with Graphs (MLG’10)*. ACM, New York, NY, 127–136. <https://doi.org/10.1145/1830252.1830269>

- [175] Kaustubh Supekar, Vinod Menon, Daniel Rubin, Mark Musen, and Michael D. Greicius. 2008. Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. *PLoS Comput. Biol.* 4, 6 (2008), 1–11. <https://doi.org/10.1371/journal.pcbi.1000100>
- [176] Kaustubh Supekar, Vinod Menon, Daniel Rubin, Mark Musen, and Michael D. Greicius. 2008. Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. *PLoS Comput. Biol.* 4, 6 (Jun. 2008), e1000100. <https://doi.org/10.1371/journal.pcbi.1000100>
- [177] Yuriy Sverchkov and Mark Craven. 2017. A review of active learning approaches to experimental design for uncovering biological networks. *PLOS Comput. Biol.* 13, 6 (2017), 1–26. <https://doi.org/10.1371/journal.pcbi.1005466>
- [178] Hiroyuki Toh and Katsuhisa Horimoto. 2002. Inference of a genetic network by a combined approach of cluster analysis and graphical Gaussian modeling. *Bioinformatics* 18, 2 (2002), 287–297. <https://doi.org/10.1093/bioinformatics/18.2.287>
- [179] Anastasios A. Tsonis and Paul J. Roebber. 2004. The architecture of the climate network. *Physica A: Stat. Mech. Appl.* 333, 0 (2004), 497–504. <https://doi.org/10.1016/j.physa.2003.10.045>
- [180] Anastasios A. Tsonis, Geli Wang, Kyle L. Swanson, Francisco A. Rodrigues, and Luciano da Fontura Costa. 2011. Community structure and dynamics in climate networks. *Climate Dynam.* 37, 5–6 (2011), 933–940. <https://doi.org/10.1007/s00382-010-0874-3>
- [181] Charalampos E. Tsourakakis, U. Kang, Gary L. Miller, and Christos Faloutsos. 2009. DOULION: Counting triangles in massive graphs with a coin. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'09)*. ACM, 837–846. <https://doi.org/10.1145/1557019.1557111>
- [182] Liubov Tupikina, Nora Molkenthin, Cristóbal López, Emilio Hernández-García, Norbert Marwan, and Jürgen Kurths. 2016. Correlation networks from flows. The case of forced and time-dependent advection-diffusion dynamics. *PLoS ONE* 11, 4 (04 2016), 1–12. <https://doi.org/10.1371/journal.pone.0153703>
- [183] Toni Vallès-Català, Tiago P. Peixoto, Roger Guimerà, and Marta Sales-Pardo. 2017. On the consistency between model selection and link prediction in networks. *ArXiv e-prints* (May 2017). <https://arxiv.org/abs/1705.07967>
- [184] Martijn P. van den Heuvel, René S. Kahn, Joaquín Goñi, and Olaf Sporns. 2012. High-cost, high-capacity backbone for global brain communication. *Proc. Natl. Acad. Sci. U.S.A.* 109, 28 (July 2012), 11372–11377. <https://doi.org/10.1073/pnas.1203593109>
- [185] Martijn P. van den Heuvel and Olaf Sporns. 2011. Rich-club organization of the human connectome. *J. Neurosci.* 31, 44 (Nov. 2011), 15775–15786. <https://doi.org/10.1523/jneurosci.3539-11.2011>
- [186] Martijn P. van den Heuvel and Olaf Sporns. 2015. Network hubs in the human brain. *Trends Cogn. Sci.* 17, 12 (Dec. 2015), 683–696. <https://doi.org/10.1016/j.tics.2013.09.012>
- [187] Dan J. Wang, Xiaolin Shi, Daniel A. McFarland, and Jure Leskovec. 2012. Measurement error in network data: A re-classification. *Soc. Networks* 34, 4 (2012), 396–409. <https://doi.org/10.1016/j.socnet.2012.01.003>
- [188] Larry Wasserman. 2006. *All of Nonparametric Statistics (Springer Texts in Statistics)*. Springer, New York. <https://doi.org/10.1007/0-387-30623-4>
- [189] Duncan J. Watts and Steven H. Strogatz. 1998. Collective dynamics of “small-world” networks. *Nature* 393, 6684 (Jun. 1998), 440–442. <https://doi.org/10.1038/30918>
- [190] David Welch, Shweta Bansal, and David R. Hunter. 2011. Statistical inference to advance network models in epidemiology. *Epidemics* 3, 1 (Mar. 2011), 38–45. <https://doi.org/10.1016/j.epidem.2011.01.002>
- [191] Tina Wey, Daniel T. Blumstein, Weiwei Shen, and Ferenc Jordán. 2008. Social network analysis of animal behaviour: A promising tool for the study of sociality. *Animal Behav.* 75, 2 (2008), 333–344. <https://doi.org/10.1016/j.anbehav.2007.06.020>
- [192] Hal Whitehead, Lars Bejder, and C. Andrea Ottensmeyer. 2005. Testing association patterns: Issues arising and extensions. *Animal Behav.* 69, 5 (2005). <https://doi.org/10.1016/j.anbehav.2004.11.004>
- [193] Hal Whitehead and Richard James. 2015. Generalized affiliation indices extract affiliations from social network data. *Methods Ecol. Evol.* 6, 7 (2015), 836–844. <https://doi.org/10.1111/2041-210X.12383>
- [194] R. W. Wilkins, D. A. Hodges, P. J. Laurienti, M. Steen, and J. H. Burdette. 2014. Network science and the effects of music preference on functional brain connectivity: From beethoven to eminem. *Sci. Rep.* 4 (Aug. 2014), 6130. <https://doi.org/10.1038/srep06130>
- [195] Rongjing Xiang and Jennifer Neville. 2008. Pseudolikelihood EM for within-network relational learning. In *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM'08)*. 1103–1108. <https://doi.org/10.1109/ICDM.2008.148>
- [196] K. Yamasaki, A. Gozolchiani, and S. Havlin. 2008. Climate networks around the globe are significantly affected by El Niño. *Phys. Rev. Lett.* 100, 22 (Jun. 2008), 228501. <https://doi.org/10.1103/PhysRevLett.100.228501>
- [197] Donghyeon Yu, MinSoo Kim, Guanghua Xiao, and Tae Hyun Hwang. 2013. Review of biological network data and its applications. *Genom. Informat.* 11, 4 (Dec. 2013), 200–210. <https://doi.org/10.5808/GI.2013.11.4.200>

- [198] Qingbao Yu, Erik B. Erhardt, Jing Sui, Yuhui Du, Hao He, Devon Hjelm, Mustafa S. Cetin, Srinivas Rachakonda, Robyn L. Miller, Godfrey Pearlson, and Vince D. Calhoun. 2015. Assessing dynamic brain graphs of time-varying connectivity in fMRI data: Application to healthy controls and patients with schizophrenia. *NeuroImage* 107 (2015), 345–355. <https://doi.org/10.1016/j.neuroimage.2014.12.020>
- [199] Ming Yuan and Yi Lin. 2006. Model selection and estimation in regression with grouped variables. *J. Roy. Stat. Soc: Series B (Stat. Methodol.)* 68, 1 (Feb. 2006), 49–67. <https://doi.org/10.1111/j.1467-9868.2005.00532.x>
- [200] Andrew Zalesky, Alex Fornito, and Ed Bullmore. 2012. On the use of correlation as a measure of network connectivity. *NeuroImage* 60, 4 (2012), 2096–2106. <https://doi.org/10.1016/j.neuroimage.2012.02.001>
- [201] Andrew Zalesky, Alex Fornito, Ian H. Harding, Luca Cocchi, Murat Yücel, Christos Pantelis, and Edward T. Bullmore. 2010. Whole-brain anatomical networks: Does the choice of nodes matter? *NeuroImage* 50, 3 (Apr. 2010), 970–983. <https://doi.org/10.1016/j.neuroimage.2009.12.027>
- [202] Tanja Zerenner, Petra Friederichs, Klaus Lehnertz, and Andreas Hense. 2014. A gaussian graphical model approach to climate networks. *Chaos* 24, 2 (2014). <https://doi.org/10.1063/1.4870402>
- [203] Yang Zhan, David Halliday, Ping Jiang, Xuguang Liu, and Jianfeng Feng. 2006. Detecting time-dependent coherence between non-stationary electrophysiological signals—A combined statistical and time-frequency approach. *J. Neurosci. Methods* 156, 1–2 (Sept. 2006), 322–332. <https://doi.org/10.1016/j.jneumeth.2006.02.013>
- [204] Bin Zhang and Steve Horvath. 2005. A general framework for weighted gene co-expression network analysis. *Stat. Appl. Genet. Mol. Biol.* 4, 1 (2005). <https://doi.org/10.2202/1544-6115.1128>
- [205] Dong Zhou, Avi Gozolchiani, Yosef Ashkenazy, and Shlomo Havlin. 2015. Teleconnection paths via climate network direct link detection. *Phys. Rev. Lett.* 115, 26 (Dec. 2015), 268501. <https://doi.org/10.1103/PhysRevLett.115.268501>
- [206] Hui Zhou, Wencai Du, Shaochun Xu, and Qinling Xin. 2011. An empirical study of network topology inference. In *Computer and Information Science 2011 SE-17*, Roger Lee (Ed.). Studies in Computational Intelligence, Vol. 364. Springer, Berlin, 213–225. https://doi.org/10.1007/978-3-642-21378-6_17

Received October 2016; revised August 2017; accepted October 2017