

A Bayesian Active Learning Experimental Design for Inferring Signaling Networks

Robert O Ness^{1,2}, Karen Sachs³, Parag Mallick³, and Olga Vitek^{*,2}

Abstract—Machine learning methods for learning network structure, applied to quantitative proteomics experiments, reverse-engineer intracellular signal transduction networks. They provide insight into the deregulation or rewiring of signaling within the context of a disease or a phenotype. To learn the causal patterns of influence between proteins in a regulatory network, the methods require experiments that include targeted interventions. The experiments fix the activity of specific signaling proteins, and measure the downstream effects. However, the interventions are costly, have varying availability and effectiveness, and add complexity and cost.

We describe a Bayesian active learning strategy for selecting optimal interventions. Our approach takes as input pathway databases such as KEGG and historic datasets in repositories such as Cytobank, expresses them in form of prior probability distributions on the signaling network structures, and selects interventions that maximize the expected information gain for inferring the signaling network. Evaluations on simulated and real data show that the strategy reduces the detection error of validated edges as compared to an unguided choice of interventions, and avoids redundant interventions, thereby increasing the effectiveness of the experiment.

I. INTRODUCTION

Signaling networks describe chains of protein interactions that determine how cells process signals from their environment. The deregulation of signaling networks occurs under many conditions, e.g. in diseases such as cancer [29], gene knockouts, or introduction of a drug. The patterns of such deregulation can be inferred from quantitative proteomic experiments conducted under the conditions of interest, using causal inference and Bayesian networks [11], [37].

In these investigations, signaling is induced with a stimulus perturbation, and a measurement technology acquires information on the activity of signaling proteins [41]. Bulk experiments quantify aggregate signaling activity across a sample. In contrast, single cell technologies provide cell-level resolution of signaling activity. For example, in flow cytometry cells are chemically fixed, and intracellular signaling proteins are tagged with fluorescently-labeled antibodies. The cytometer then records the antibodies' fluorescence in individual cells, each reflecting the relative abundance of signaling proteins in different states of enzymatic activity [31]. Similarly, in mass cytometry (CyTOF) experiments, intracellular signaling proteins are tagged with heavy-metal

isotopes, and the mass spectrometer records the mass-to-charge ratio of the charged isotope tags [1].

Causal Bayesian networks represent signaling proteins as nodes, and regulatory relationships as directed edges. It interprets a network as a topological map of the underlying signaling network. By comparing the structure inferred under a condition to *canonical pathways* in sources such as KEGG and Reactome, we can learn the patterns of network deregulation. Data repositories, such as Cytobank [4], provide *historic data*, which can be incorporated in the analysis and interpretation of experimental results [11], [44]. This prior information is especially important for higher throughput experiments because it helps eliminate spurious correlations and false discoveries of relationships between proteins [28].

To distinguish causal relationships from statistical associations, causal network inference requires *targeted interventions* on some proteins [8], [28], e.g. using small-molecule inhibitors that block a protein's enzymatic activity. An insufficient set of interventions results in only a partially causally oriented network [30]. At the same time, increasing the number of interventions increases the complexity of the experiment and the cost. This cost is wasted when targeted interventions redundantly orient the same edges.

In this paper, we propose a strategy for optimal design of bulk or single-cell proteomic experiments aiming at causal inference. The design prioritizes targeted interventions, and provides a criterion to stop adding interventions. It combines prior knowledge in the form of canonical pathways imported from sources such as KEGG [21] with historic data. The strategy outputs a sequence of interventions that we call a “batch”, i.e. a minimal subset of candidate interventions that contributes maximal causal information given the available data. We then describe an active learning framework, that iterates between selecting interventions and acquiring data to obtain a fully inferred causal network. To the best of our knowledge, this is the first active learning approach to experimental design for inference of signaling networks.

II. BACKGROUND

A. Directed graphs as causal models of signaling

A causal Bayesian network denotes a set of p signaling proteins with p nodes $V = \{v_1, \dots, v_p\}$. The nodes are variables representing levels of signaling activity of the proteins. For example, v_1 can take discrete signaling states “active” or “inactive”, or continuous values quantifying the abundance of a protein form. The model expresses causal relations between nodes with a directed acyclic graph structure (DAG) G . The edge direction in the DAG represents the causal effect of a

* To whom correspondence should be addressed o.vitek@neu.edu

¹Dpt. of Statistics, Purdue University, West Lafayette, 47907, USA

²College of Science, College of Computer and Information Science, Northeastern University, Boston, 02115, USA

³School of Medicine, Stanford University, Palo Alto, 94305, USA




Fig. 1: Illustration of DAG equivalence classes. The DAG $\text{Raf} \rightarrow \text{Mek} \rightarrow \text{Erk}$ is the “ground truth” canonical MAPK signaling pathway, which we seek to learn by causal inference. The left box shows the equivalence class P , represented by the PDAG $\text{Raf} - \text{Mek} - \text{Erk}$. The PDAG contains three DAGs, all statistically indistinguishable in absence of interventions. The cardinality of P is 3. The middle box shows the PDAG P_{Erk} , obtained after an intervention on Erk. P_{Erk} is a subclass of P that has eliminated $\text{Raf} \leftarrow \text{Mek} \leftarrow \text{Erk}$. The cardinality of P_{Erk} is 2. The right box shows the single ground truth DAG obtained after an additional intervention on Raf. An alternative single intervention on Mek simultaneously compels the direction of both edges, and is more effective at discovering the ground truth.

change in the signaling state of a parent node on the state of the child. The DAG is best interpreted as a snapshot of a dynamic system [19]. This interpretation is strongest when the signaling response has reached some quasi-steady-state.

Each node in the DAG has a conditional probability distribution given its parents. It is a probabilistic representation of the regulatory influences of the parents on the child [30]. A key advantage of the probabilistic interpretation is that it encodes *conditional independence*, i.e. the probability that the state of a protein is independent of the state of all its upstream proteins, if we know (i.e. condition on) the states of its direct parents. This allows us to ignore the correlation between a protein and proteins more than one step upstream.

From the statistical perspective, the goal of causal inference is to infer the DAG structure representing the signaling network, using associations between proteins as input. However, statistical associations are not sufficient to orient the edges in the DAG [28]. We illustrate this with the simple 3-protein canonical MAPK signaling pathway $\text{Raf} \rightarrow \text{Mek} \rightarrow \text{Erk}$. Imagine that the structure of the pathway is unknown, and needs to be inferred. A causal inference algorithm would (1) detect pairwise statistical correlations between abundances of each pair of the three proteins, pointing to the three candidate edges, (2) test Raf and Erk for conditional independence, given the state of Mek, and (3) in presence of conditional independence, eliminate the edge between Raf and Erk. After that, additional interventions are required to orient the edges between Raf and Mek, and between Mek and Erk. The left box in Figure 1 illustrates that, in absence of interventions, the ground truth is statistically indistinguishable from the other two causally incorrect DAGs.

A set of statistically indistinguishable DAGs form a *Markov equivalence class* P , comprised of DAGs with same edges but varying orientations, which have equal statistical likelihood for the dataset [23], [30]. A Markov equivalence class is represented with a partially directed acyclic graph (PDAG). The directed edges in a PDAG have the same

orientation in all the DAGs in P . The undirected edges in the PDAG have varying directions in the DAGs, and therefore represent edges with uncertain causality. Figure 1 illustrates PDAGs and DAGs in the MAPK pathway example. The cardinality of a PDAG is defined as the number of its DAGs.

Targeted interventions proceed by fixing the state of a node, such that it does not vary with that of its parents [8], [23], [30]. Fixing the node introduces an additional constraint, and eliminates members of the equivalence class that fail to satisfy this constraint. For example, in Figure 1 a small inhibitor fixes Erk’s enzymatic activity state to “off”, such that the activity of Erk becomes independent of the state of Mek. The intervention fails to regulate the activity of Mek, and therefore eliminates the DAG with an edge $\text{Erk} \rightarrow \text{Mek}$.

The reduced equivalence class is formally defined as a *transition-sequence Markov equivalence class*, i.e. the equivalence class after a sequence of “transitions” (interventions) [42]. Each additional intervention orients more edges in a PDAG, and a sufficiently large set of interventions compels all the edges. In Figure 1, interventions on Erk and Raf eliminated all but the ground truth from the equivalence class.

As the example illustrates, a batch of interventions targeting Erk and Raf would reveal the ground truth DAG. However, so would a batch containing a single intervention on Mek. The goal of this work is to identify batches of interventions, which reveal the most causality while minimizing the number of interventions they contain, and by extension the experimental complexity and cost.

B. Bayesian inference of causal networks

This work focuses on a Bayesian approach to learning causal PDAG representations of signaling, where experimentalists (1) start with background knowledge about the signaling network, such as likely pathways or motifs, (2) use the background knowledge to construct a *prior* distribution on graph structures, (3) collect experimental measurements of the signaling states of proteins, and (4) estimate a *posterior* distribution on structures based on the prior and the experimental measurements. The Bayesian approach is advantageous, as it reveals the “rewiring” of signaling between conditions by examining differences between the prior and the posterior distributions.

More formally, the approach uses the background knowledge δ to construct a prior probability distribution of possible DAG structures $\pi(G|\delta)$ in the space of possible structures \mathbb{G} . The experimentalist collects a data set D of measurements on the proteins $V \in G$, acquired under a batch of targeted interventions S . A statistical likelihood function $p(D|S, G)$ quantifies the likelihood that the observations were generated from a graph G . Finally, inference of the DAG structure relies on a posterior probability distribution $\pi(G|D, S, \delta)$

$$\pi(G|D, S, \delta) \propto p(D|S, G)\pi(G|\delta) \quad (1)$$

In many biological applications, inferring a single DAG with high posterior probability is not of the main interest. Instead, we are interested in local features of the graph, such

as the presence of particular edges or network motifs. This is expressed through a function f on a graph that quantifies the feature of interest, and through its posterior expectation $E\{f|D, S, \delta\}$ across all graphs.

$$E\{f|D, S, \delta\} = \sum_{\mathbb{G}} f(G)\pi(G|D, S, \delta) \quad (2)$$

For example, if f is an indicator of the presence of an edge in the network, then $E\{f|D, S, \delta\}$ is the posterior probability of the presence of that edge. In this work, our feature of interest quantifies the causal insight provided by an intervention.

Bayesian inference of DAG structures relies on a *Bayesian scoring function* $Score(G, D, S, \delta)$ that takes as arguments a DAG, a dataset quantifying protein activity, a set of interventions, and the structured prior knowledge. It returns values proportional to the posterior distribution $\pi(G|D, S, \delta)$. Algorithms searching for DAGs with high $Score(G, D, S, \delta)$ must explore the combinatorially large search space of possible DAGs [24], [25], [35]. For computational tractability, some algorithms approximate the posterior probabilities $\pi(G|D, S, \delta)$ (see [23], [39] for background, and [7], [17] for an example). Since the full search space over all possible DAGs \mathbb{G} is combinatorially large, the common practice is to sample a set of DAGs from their posterior distribution through a random process such as bootstrap ([13], [20]) or MCMC [14], and keep a sample of high scoring networks. $E\{f|D, S, \delta\}$ is then approximated as

$$E\{f|D, S, \delta\} \approx \sum_{G \in \Omega} f(G) \frac{Score(G, D, S, \delta)}{\sum_{G' \in \Omega} Score(G', D, S, \delta)} \quad (3)$$

where Ω denotes a sample of high scoring DAGs.

C. PDAG representation of uncertainty in causal effects

Experiments with incomplete sets of interventions lack information to fully infer the causal orientation of the edges in a DAG. In order to characterize this uncertainty, algorithms take a single DAG and determine the space of DAGs having the same conditional independence structure, topology and likelihood $p(D|S, G)$, but different edge orientations as the input DAG [5], [42]. These DAGs form a *Markov equivalence class*. The algorithms return a PDAG, which represents the class by preserving the shared edge structure, while presenting edges with conflicting orientations as undirected.

In Bayesian inference of causal networks, we are interested in Markov equivalent graphs with not only the same likelihood, but also the same posterior probability and the same score $Score(G, D, S, \delta)$ [5]. As seen from Equation 1, Markov equivalent graphs have the same posterior only if they have the same prior probability $\pi(G|\delta)$. This last condition does not hold when the prior probabilities encode available causal information, such that for some edges orientation in one direction is more probable than the other. A contribution of this manuscript is an implementation of a DAG-to-PDAG conversion algorithm that accounts for informative prior probabilities of causal direction of edges, and outputs a PDAG representing a class of DAGs that are Markov equivalent and have the same posterior and score.

D. Active learning for the optimal design of causal network inference experiments

In the context of causal inference from experiments, active learning is the task of including targeted interventions in the design, to optimize the inference of edge orientation. For example, in Figure 1 an intervention on Mek compels both edges in the graph, and is thus more valuable than an intervention on Erk, which only orients one edge.

Previous work used active learning to distinguish members of statistically equivalent graphs [9], [16], [26], [32], [42]. However, these approaches work with either a single PDAG, or a selection of highly-likely PDAGs. The approach in this manuscript differs by working with the entire probability distribution of PDAGs, characterizing each PDAG by its posterior probability of containing the causal truth.

Alternative Bayesian approaches to active learning of causal networks also exist [27], [43]. However, they require the experimentalists to represent their background knowledge in terms of topological orderings, i.e. an ordering of nodes such that the “from” node for every edge occurs earlier in the ordering of the “to” node. In contrast, this manuscript represents background knowledge in terms of probabilities of edge presence and orientation. This more intuitive approach simplifies the experimentalists’ work with pathways.

Finally, the proposed approach is similar in spirit to other methods in the bioinformatics literature that use historic data to inform experimental design. E.g., Rossel and Muller [34] used a sequential Bayesian method to plan sample size. Guan et al. [15] used available data to find optimal orderings of high-throughput experiments. King et al. [22] constructed a “robot scientist” that applied an active-learning strategy to functional genomics. To our knowledge, this manuscript is the first to apply active learning to inferring the structure of cell signaling pathways.

III. METHODS

A. Prior knowledge for causal graph structure learning

Quantifying causal knowledge with edge probabilities.

We propose to use probabilities of edge presence and edge orientation as a means of modeling signaling events. A set of signaling proteins, represented by nodes in V , has up to $\frac{|V|(|V|-1)}{2}$ possible edges. *Presence probability* of an edge between nodes $\{u, v\}$, denoted $P(u - v)$ or $\bar{\pi}_{uv}$ as a shorthand, quantifies the confidence that the edge is present in G . *Orientation probability* for the edge from u to v , denoted $P(u \rightarrow v|u - v)$ or $\vec{\pi}_{uv}$ as a shorthand, is the conditional probability of this orientation, *given* that the edge is present. Since only two orientations are possible, $P(u \rightarrow v|u - v) = 1 - P(u \leftarrow v|u - v)$. The goal of targeted interventions is to resolve the orientations of the edges, i.e. coerce orientation probabilities towards 0 or 1.

Let δ be a set of edge probabilities $\delta_{u,v}$, where

$$\delta_{u,v} \stackrel{\text{def}}{=} \{\bar{\pi}_{uv}, \vec{\pi}_{uv}\} \quad (4)$$

In Bayesian setting, we wish to use the edge probabilities to quantify prior causal knowledge. Let $\bar{I}_{uv}(G)$ be an indicator

function for the presence an an edge between u and v in G , and $\vec{I}_{uv}(G)$ be an indicator function for the orientation of the edge from u to v . We map edge probabilities δ to a probability distribution on DAG structures using using an *edge-wise prior*

$$\pi(G|\delta) = c \prod_{uv \in G} (1 - \bar{\pi}_{uv})^{1 - \vec{I}_{uv}(G)} (\bar{\pi}_{uv} \vec{\pi}_{uv})^{\vec{I}_{uv}(G)} \vec{I}_{uv}(G) \quad (5)$$

where c is a normalizing function on the space of graphs that corrects for the acyclicity constraint [3].

When nothing is known about the presence or orientation of the edges, we specify the *uninformative* edge probabilities [39], where

$$\bar{\pi}_{uv} \approx \frac{1}{2} + \frac{1}{2(|V| - 1)}, \quad \text{and} \quad \vec{\pi}_{uv} = \frac{1}{2} \quad (6)$$

The intuition behind Equation 6 is that two nodes are more likely to be linked in a small network than in a large network. Therefore, the uninformative presence probability approaches .5 as network size increases. The uninformative orientation probability is .5 for either direction. These uninformative edge-wise priors correspond to the marginal probabilities of an edge in the case when there is a uniform probability distribution on the space of graphs [39].

Upon conducting an experiment with interventions S and collecting a dataset D , the next step is to *update* the DAG probability distribution with condition-specific information in the data using Bayes rule

$$\pi(G|D, S, \delta) \propto p(D|S, G)\pi(G|\delta) \quad (7)$$

Note that the Bayes rule is agnostic of the process that selected the interventions in S . S can be selected by any approach, such as applying the available inhibitors, or using the proposed active learning approach below.

The updated probability distribution $\pi(G|D, S, \delta)$ maps back to edge probabilities using the approach in Equation 2. Let $f(\cdot)$ in Equation 2 be an indicator function $\vec{I}_{uv}(G)$ for the presence, and an indicator function $\vec{I}_{uv}(G)$ for the orientation, of an edge between nodes u and v in a DAG. Then, the updated presence and orientation probabilities of an edge after observing data D is defined as

$$\begin{aligned} \bar{\pi}_{uv|D,S} &= \sum_{G \in \mathbb{G}} \vec{I}_{uv}(G) \pi(G|D, S, \delta) \\ \vec{\pi}_{uv|D,S} &= \frac{1}{\bar{\pi}_{uv}} \sum_{G \in \mathbb{G}} \vec{I}_{uv}(G) \pi(G|D, S, \delta) \end{aligned} \quad (8)$$

In other words, these are average frequencies of edge presence and orientation over all the DAGs, weighted by the posterior probabilities of the DAGs.

Incorporating pathway knowledge and historic data.

When prior information is available, we propose to construct informative edge probabilities $\delta_{u,v}$. In the simplest case, a directed edge between u and v in the canonical pathway is viewed as a hypothesis that an edge linking these nodes is also present under the condition of interest, and is oriented

from u to v . Denoting the set of edges in the canonical pathway as \mathbb{K} , the background knowledge $\delta_{u,v}$ is defined as in Equation 4 where

$$\begin{aligned} \bar{\pi}_{uv} &= \begin{cases} \sim 1 & \text{if } u - v \in \mathbb{K} \\ \sim 0 & \text{otherwise} \end{cases} \\ \vec{\pi}_{uv} &= \begin{cases} \sim 1 & \text{if } u - v \text{ is oriented } u \rightarrow v \in \mathbb{K} | u - v \in \mathbb{K} \\ \sim 0 & \text{if } u - v \text{ is oriented } u \leftarrow v \in \mathbb{K} | u - v \in \mathbb{K} \end{cases} \end{aligned} \quad (9)$$

The notation ~ 1 (probability near 1) and ~ 0 (probability near 0) emphasizes that the Bayesian approach avoids the boundary probabilities of 0 and 1.

In some cases, experimentalists may wish to use alternative specifications. For example, in absence of canonical pathway information it may be inappropriate to assign $\bar{\pi}_{uv} \sim 0$ to each edge, and subjective edge probabilities may be a better choice. For edges where no assessments can be made, we use the uninformative edge probabilities in Equation 6.

In addition to incorporating prior knowledge from canonical pathways, we also seek to make use of information in historic datasets. We define historic data D_0 as previous experiments, which quantified the activity of the proteins in the same network, under the same signaling conditions as the pending causal inference experiment, but lacking targeted interventions. We update the canonical knowledge $\pi(G|\delta)$ with the condition-specific information in the historic data

$$\pi(G|D_0, \delta) \propto p(D_0|G)\pi(G|\delta) \quad (10)$$

Here $p(D_0|G)$ is the likelihood that the historic data came from the graph G , and $\pi(G|D_0, \delta)$ is the updated distribution on graph structures, which now captures the full state of our prior information before the interventions.

Sampling DAGs from a DAG distribution Similarly to Equations 2 and 3, the proposed Bayesian inference on causal networks relies on sampling a set of DAGs Ω from a distribution $\pi(G|D, S, \delta)$. Our implementation uses Bayesian bootstrap sampling from a distribution of DAGs [13] and [20] with random graph starts [18] and greedy search, and the posterior distributions are derived using Bayesian Dirichlet approximation[17]. When the signal-to-noise ratio in the historic data is high and/or the edge-wise prior is informative, the sampling concentrates on a smaller set of most probable DAGs. When the signal-to-noise ratio is low, weight is distributed more evenly among graphs, and the sampling must cover a larger number of graphs with similar $\pi(G|D, S, \delta)$.

Representing uncertainty in causal effects with PDAGs.

We incorporate the informative edge orientation probabilities in δ to express the uncertainty in edge orientation using PDAGs. We view the conversion of a DAG to a PDAG as the function f in Equations 2 and 3. Applying these equations requires that DAG members of the same equivalence class have the same posterior probability, otherwise different instances of the same PDAG would have different probabilities, i.e. the same f would have different probabilities in

Equation 2 and scores in 3. Current conversion algorithms are not compatible with edge-wise prior probabilities.

Algorithm 1 is a DAG-to-PDAG conversion algorithm that incorporates informative edge-wise prior probabilities. It starts with a DAG from Ω . Next, every directed edge is converted to an undirected edge if it meets three conditions. The first two conditions are the same as in the prior literature [5], [42]. First (lines 4 and 10 in Algorithm 1), reversing edge direction should not change the number of immoral v-structures (i.e., 3-node motifs with one child and two parents, with no edge between parents). Second (line 6 of Algorithm 1), the child node of the edge should not be targeted by an intervention in S . In this manuscript we introduce a third condition (lines 8 in Algorithm 1), stating that the edge should have an uninformative prior orientation probability of .5. These conditions create an equivalence class P and its PDAG, where all the members have the same value for $\pi(G|D_0, \delta)$. If the algorithm is applied to two Markov equivalent DAGs with different causal information in their informative edge-wise priors, two different PDAGs are returned.

Algorithm 1 DAG to PDAG Algorithm

Inputs: A DAG G , an (optional) set of selected intervention targets S , a set of edge probabilities δ .

```

1: procedure PDAG( $G, S, \delta$ )
2:   for edge  $e$  in  $G$  do
3:     if  $e$  is in an immoral v-structure then
4:       Fix direction of  $e$ 
5:     if  $e$ 's child is targeted by  $S$  then
6:       Fix direction of  $e$ 
7:     if  $e$ 's orientation probability in  $\delta \neq .5$  then
8:       Fix direction of  $e$ 
9:   for edge  $e$  in  $G$  do
10:    if  $e$  is not fixed then
11:      if reversing  $e$ 's direction
          will not add a new v-structure
          or introduce a cycle then
12:        Make  $e$  undirected
13:    $P \leftarrow G$ 
14:   return ( $P$ )

```

B. Bayesian active learning with causal information gain

Overview. The strategy for selecting optimal interventions is overviewed in Figure 2. The active learning algorithm takes as input a sample of DAGs from a DAG distribution, and a set of candidate interventions. The algorithm sequentially evaluates the expected causal information gain of the interventions, and outputs a minimally-sized batch of interventions that maximizes the expected information gain. We detail the components of the algorithm below.

Defining the causal information gain of an intervention.

Suppose that the true causal DAG G were known. Let $S \subseteq V$ denote a batch of candidate interventions. As discussed in

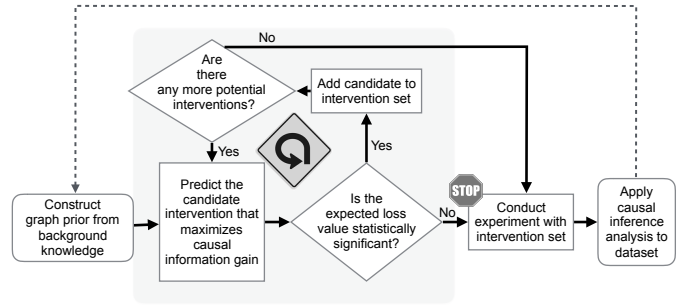


Fig. 2: Overview of the proposed method. A probability distribution of possible graph structures is constructed from canonical pathways and historic data. Interventions are iteratively added to the design until the expected causal information gain we can expect from an additional intervention becomes small. We then stop adding interventions to the design and use the newly acquired data to infer the causal network.

Section II-C, a PDAG is derived directly from a DAG's topology and a set of interventions, and can be determined before collecting data. Therefore, we could devise an algorithm $H(G, S)$ that first determines a PDAG, then counts the number of oriented edges in the PDAG, i.e. the number of oriented edges in G that could be inferred from data with interventions S .

Suppose that we consider an additional intervention on node v , which leads to $H(G, \{S, v\})$ edges correctly oriented edges. We define the causal *information gain* $IG(G, S, v)$ of the intervention on v as the increase in correctly oriented edges, i.e. $IG(G, S, v) = H(G, \{S, v\}) - H(G, S)$. $IG(G, S, v)$ is non-negative, and can be zero if v fails to orient any edges beyond those oriented by S . Note that an equivalent definition of information gain is the reduction in the number of unoriented edges. This definition parallels the information theory notation of entropy, where the information gain is viewed as entropy reduction.

Selecting interventions that maximize the expected information gain

Of course in practice the true causal DAG G is unknown. Therefore, similarly to Equation 2 we consider the *expected information gain*, which averages the information gain over all possible graphs, weighted by their prior distribution $\pi(G|D_0, \delta)$

$$EIG_{S,v} = \sum_G IG(G, S, v) \pi(G|D_0, \delta) \quad (11)$$

Moreover, similarly to Equation 3, we approximate the expected information gain as

$$EIG_{S,v} \approx \sum_{G \in \Omega} IG(G, S, v) \frac{Score(G, D_0, S, v, \delta)}{\sum_{G \in \Omega} Score(G, D_0, S, v, \delta)} \quad (12)$$

where Ω denotes a sample of high scoring DAGs, and $Score$ is a Bayesian scoring function returning a value proportional to $\pi(G|D_0, \delta)$. We then select the candidate v that maximizes the approximated expected information gain. Algorithm 2 details these steps. Note that in Bayesian

decision theory, $-IG$ is a loss function, and we select the candidate v that minimizes the expected loss [2].

Algorithm 2 *Expected Information Gain*

Inputs: A set of DAGs Ω , Bayesian scoring function $Score$, a set of pre-selected intervention targets S , a candidate for next intervention v , and a set of edge probabilities δ .

```

1: procedure EIG( $\Omega$ ,  $Score$ ,  $S$ ,  $v$ ,  $\delta$ )
2:   Initialize array  $IG$  of size  $|\Omega|$ 
3:   for  $i$  in  $1:|\Omega|$  do
4:      $P_S \leftarrow \text{PDAG}(G_i, S, \delta)$ 
5:      $P_{S,v} \leftarrow \text{PDAG}(G_i, \{S, v\}, \delta)$ 
6:      $H_S \leftarrow \text{num. of directed edges in } P_S$ 
7:      $H_{S,v} \leftarrow \text{num. of directed edges in } P_{S,v}$ 
8:      $IG_i \leftarrow H_{S,v} - H_S$ 
9:   return WeightedMean( $IG$ ,  $Score$ )

```

Algorithm 3 *Active Learning*

Inputs: A set of DAGs Ω , Bayesian scoring function $Score$, a set U of candidates for next intervention, and a set of edge probabilities δ .

Parameter: α , stopping criterion for the information gain.

```

1: procedure ACTIVELEARNING( $\Omega$ ,  $Score$ ,  $U$ ,  $\alpha$ ,  $\delta$ )
2:    $S \leftarrow \text{null}$ 
3:   while length( $U$ ) > 0 do
4:     TopCandidate  $\leftarrow \text{null}$ 
5:     MaxEIG  $\leftarrow 0$ 
6:     for node  $v \in U$  do
7:        $EIG_{S,v} \leftarrow \text{EIG}(\Omega, Score, S, v, \delta)$ 
8:       if  $EIG_{S,v} > \text{MaxEIG}$  then
9:         TopCandidate  $\leftarrow v$ 
10:        MaxEIG  $\leftarrow EIG_{S,v}$ 
11:   if  $\neg \text{Stop}(\text{TopCandidate}, \Omega, \alpha)$  then
12:      $S \leftarrow \text{cat}(S, \text{TopCandidate})$ 
13:      $U \leftarrow U[-\text{TopCandidate}]$ 
14:   else
15:     return ( $S$ )

```

Starting point, iterations and stopping criteria The proposed active learning strategy is summarized in Algorithm 3. It starts with the empty set of selected interventions $S = \emptyset$, a set of candidate interventions U , and a set of DAGs Ω . For each intervention $v \in U$ and each DAG G in Ω , the $\text{EIG}(\Omega, Score, S, v)$ algorithm calculates P_S , $P_{S,v}$ and $IG(G, S, v)$, and returns the expected information gain. The candidate with the maximum expected information gain is added to the batch S . In the next iteration, the expected information gain for the remaining candidates is evaluated, while accounting for the effect of the interventions that are already in the batch.

After a certain point, additional interventions to the batch become counterproductive. For example, in Figure 1 an intervention on Mek would orient the $Mek - Erk$ edge. Including an additional intervention on Erk would provide no

additional information. In the case of Figure 1 the true graph is known, and we stop adding interventions when the information gain is 0. Since in real-life situations the structure of the true DAG is unknown, we stop adding interventions when the probability that at least some information gain occurs is below a parameter α . Similarly to Equations 3 and 12, the probability that at least some information gain occurs is

$$q_{S,v} = \sum_{G \in \Omega} I_{\{IG(G,S,v) > 0\}} \frac{Score(G, D_0, S, v, \delta)}{\sum_{G \in \Omega} Score(G, D_0, S, v, \delta)} \quad (13)$$

where v is the candidate that maximizes EIG, and $I_{\{\cdot\}}$ is the indicator function. We add v to the set of interventions if $q_{S,v} > \alpha$, and stop if $q_{S,v} \leq \alpha$. Higher values of α will result in a smaller intervention batch. Setting probability threshold α to ~ 0 (i.e., stopping when the probability of at least some information gain is near 0) is equivalent to stopping when expected information gain is near 0.

When the signal-to-noise ratio in the historic data is low, there is greater weight on graphs where the most optimal intervention candidates have no information gain. This leads to the triggering of the stopping criteria with a smaller batch of interventions. Thus when there is more uncertainty in the data, the procedure avoids the risk of wasteful use of interventions, instead favoring running the experiment with a smaller batch and then relying on the new experimental data to evaluate unused interventions.

C. Inference of causal network from data acquired post-intervention

The next step in the investigation is to apply the selected interventions, and collect a new dataset D . The new dataset updates Equation 7 as

$$\pi(G|S, D_0, D, \delta) \propto p(D|S, G)\pi(G|D_0, \delta) \quad (14)$$

The updated edge probabilities are obtained as in Equations 8 and 9. They balance the canonical representation of the signaling with the condition-specific signaling behavior quantified under the condition of interest, after the interventions. A large deviation of the posterior edge probability from the prior in δ indicates network rewiring or deregulation.

The active learning procedure is iterative in nature, in that the results of the intervention experiment can be viewed as a new instance of historic data. They can inform the selection of new interventions by substituting $\pi(G|S, D_0, D, \delta)$ in Equations 11 and 12, and repeating the overall procedure.

D. Implementation and computational complexity

The proposed strategy is implemented in an open-source R package *bninfo*, available on Github. The edge-wise prior is constructed as a data frame in R, either manually or through an interface to the API of KEGG provided by *bninfo*. Historic data is represented as a data frame and pre-selected interventions S as an array. The package *bninfo* implements the algorithms for converting a DAG and a edge-wise to PDAG, calculating the expected gain, and selecting the optimal batch of interventions. The Bayesian network

structure learning is performed with the existing R package *bnlearn* [38].

The main scalability bottleneck in the proposed strategy is the selection of interventions in the while loop in Algorithm 3. The complexity of calculating the expected information gain for a single intervention (Algorithm 2 and line 7 in Algorithm 3) is in the order of the number of edges in the input DAG. However, the interventions in a batch are selected sequentially (i.e., the selection of the j th member of the batch depends on the $j-1$ previously selected interventions). Therefore, the selection of j candidate interventions requires up to $j!$ calculations of the expected gain. The running time can be reduced by parallelization of Algorithm 2, or by limiting the number of candidate interventions. Moreover, sampling Ω from a distribution of DAGs has well-known scalability challenges in Bayesian literature. In the worst case, the computational time scales exponentially with the number of proteins. Bayesian bootstrap sampling can be split among nodes on a cluster, and sped up by parallelization. For the datasets described below, the generation of intervention batches took 70 minutes (17 protein DREAM4 network with 14 candidate interventions and 500 sampled graphs) and 20 minutes (11 protein T-cell network with 5 candidate interventions and 500 sampled graphs) on a 16 node cluster.

E. Metrics used for performance evaluation

We evaluated the proposed strategy using datasets containing some notion of “ground truth”, i.e. situations where the true structure of the causal graph is known. We used the proposed active learning strategy to determine an optimal intervention batch S . To evaluate the performance of S , we considered the data D that would be experimentally acquired with the selected interventions. We then inferred causal networks from D , and derived posterior edge probabilities as described in Section III-C. Finally, we evaluated whether S can lead to network inference that correctly detects edges in the “ground truth” graph.

We evaluated the performance of edge detection using two metrics. The first is the *true positive rate of edge detection*, i.e. the proportion of correctly detected edges among the edges in the ground truth network [13]. Our cutoff for edge’s detection is presence and orientation probabilities greater than uninformative prior probabilities described in Equation 6. The second metric is the L_1 edge error, which quantifies the overall probability of prediction error, i.e. the probability of either discovering a false edge or missing a true edge [27], [43]. Given the set of interventions S and the ground truth network, the L_1 edge detection error is defined as

$$\begin{aligned}
 L_1(G, S) &= \sum_{u,v} \vec{I}_{uv}(G)(1 - \vec{\pi}_{uv|D,S}) \quad (15) \\
 &+ (1 - \vec{I}_{uv}(G))(\vec{\pi}_{uv|D,S}) \\
 &+ (1 - \bar{I}_{uv}(G))(\bar{\pi}_{uv|D,S})
 \end{aligned}$$

where $\vec{I}_{uv}(G)$ and $\bar{I}_{uv}(G)$ are as in Equation 9.

IV. DATASETS

There are currently no publicly available datasets that implement active learning approach to causal network inference. We therefore use two publicly available datasets, adapted to provide a measure of “ground truth”.

A. DREAM4 Network

“Ground truth” network The 17-node network in Figure 3A was used in the DREAM4 Predictive Signaling Network Challenge [33]. The network contains canonical pathways downstream of four receptors (dark grey nodes in Figure 3A): two inflammatory (TNFa, IL1a), one insulin (IGF-I), and one growth factor receptor (TGFA). We use this network to evaluate the relative advantages of active learning, and the importance of the use of prior information.

Prior information regarding the network structure We used as the background information the fact that TNFa, IL1a, IGF-I and TGFA are receptors, i.e. proteins that receive signals from the environment, and activate downstream proteins. This presents causal information, in that the remaining proteins in the pathway are downstream of the receptors.

Historic data We use the DREAM4 challenge as a historic dataset. The challenge provided antibody-based measurements (sandwich immunoassays with the Luminex xMAP platform) from hepatocellular carcinoma cell lines (HepG2), which quantified the activity levels of the signaling proteins at bulk (i.e., non-single-cell) resolution. The dataset is comprised of samples with 5 stimulus conditions, namely no stimulus, stimulus on TNFa, on IL1a, on IGF1, and on TGFA. The dataset also includes 5 intervention conditions, namely no inhibition, inhibition on ikk, on mek12, on pi3k, and on p38. In total there were 25 samples, one for each stimulus and intervention pair.

We processed the historic dataset as follows. We imputed missing values using a neural network model that predicted missing values of a protein given the values of the protein’s neighbors in the ground truth network. Several proteins from the pathway (map3k7, ras, map3k1, and mkk4) were not quantified in the historic dataset. Approaches exist for learning Bayesian network structure with hidden variables [6], [12], but these are beyond the scope of this manuscript. So to eliminate this artifact we applied a model that predicts a protein’s values given the values of its parents in the model and common biochemical assumptions on signaling dynamics [40], and used the model to predict the values for the hidden nodes. The publicly available challenge data is normalized to the 0-1 range, we then discretized the quantification values to binary on/off variables using .5 as a cutoff.

Candidate interventions All the non-receptor nodes in the network were considered as candidates for interventions. Due to the small number of samples and inhibitions, it was not possible to set aside a portion of this dataset for performance evaluation. Therefore, we fit a causal network model to the challenge data and the ground truth network, and used the

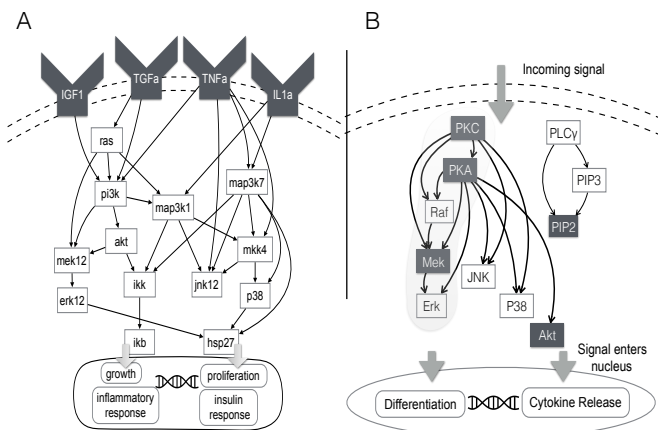


Fig. 3: The “ground truth” networks in the experimental datasets. A: The DREAM4 Predictive Signaling Network Challenge network. Dark nodes indicate receptors. B: Native T-cell signaling network in response to antigen. The edges in the PKC → Raf → Mek → Erk cascade, and the edge PKA → Raf belong to the canonical MAPK pathway. Dark nodes indicate targets of experimental interventions.

model fit to generate synthetic post-intervention datasets. The simulation mimicked the design of the challenge data, in that it contained one biological sample for each intervention. We then evaluated the performance of the selected interventions on these synthetic datasets.

B. Flow Cytometry Measurements of T-cell Signaling

“Ground truth” network The network in Figure 3B contains 11 phosphoproteins and phospholipids involved in the native CD4+ T-cell signaling response to antigen, and their canonical edges. The network was used to validate causal inference from an experimental dataset [37], and has been subsequently used as a benchmark in multiple causal inference studies [8], [10].

Prior information regarding the network structure We used as the background knowledge the edges in the canonical MAPK pathway. Although CD4+ T-cell signaling has been extensively studied, we assumed that no prior information is available regarding the remaining edges. This assumption allowed us to compare the case of a minimally informative prior to the case of a uninformative prior, and focus performance evaluation on detection edges that were not addressed in the background knowledge.

Historic data The experimental dataset in [37] contains single cell fluorescence-based quantifications of 11 phosphoproteins and phospholipids in human primary naive CD4+ T-cells. These analytes are downstream of the receptors CD3 and CD28 that provide co-stimulatory signals required for T-cell activation. We used as the “historic data” a portion of this dataset that was acquired without any targeted interventions. This portion of the dataset contained 11672 cells.

Candidate interventions The dataset in [37] also contained single cell quantifications, acquired after activating or inhibiting five signaling proteins (dark nodes in Figure 3B). These

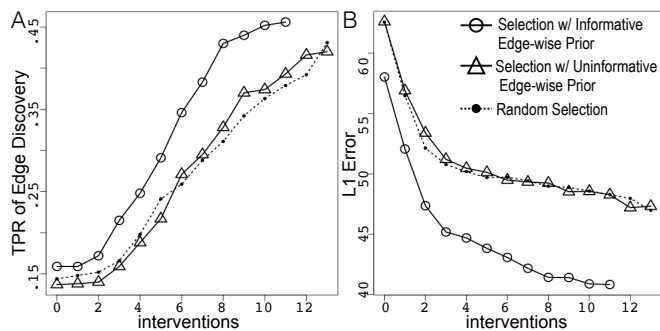


Fig. 4: Performance of the proposed strategy on the DREAM4 dataset. Dotted line: uninformative edge-wise priors, randomly selected interventions. Solid line with triangles: uninformative edge-wise priors, interventions selected with active learning. Solid line with circles: informative edge-wise priors, interventions selected active learning. A: True Positive Rate (TPR) of detecting ground truth edges. B: L_1 error of detecting ground truth edges.

five proteins were considered as candidates for interventions in this manuscript. The post-intervention experimental datasets were used to compare the information gain projected in our approach with the actual information gain after the interventions.

V. RESULTS

A. Informative prior edge probabilities reduced the required number of interventions in the DREAM4 dataset

In the DREAM4 dataset, we compared (1) random ordering of interventions, and (2) proposed active learning strategy with uninformative prior probabilities of edge presence (Equation 6), and (3) the same active learning strategy, but with informative priors. The informative priors encoded the fact that the receptor nodes have upstream positions in the network. All the receptor-originating edges were assigned the prior orientation probability of ~ 1 , all receptor-terminating edges were assigned the presence probability of ~ 0 , and the remaining edges were assigned the uniform prior presence probability in Equation 6. All the non-receptor nodes in the network were candidate targets for interventions. In order to exhaust the information in the prior and historic data, the active learning approach with both priors used the most liberal stopping criteria for growing a batch. It only stopped adding interventions when all the additional candidates has expected information gain of ~ 0 .

Figure 4 summarizes the results. With uninformative edge-wise priors, random selection of interventions performed similarly to the proposed active learning in both True Positive Rate of edge detection and L_1 loss. Both metrics depend on correct detection of edge presence as well as edge orientation. The contribution to edge detection of the added samples provided by each intervention experiment overshadowed the selection strategies prioritization of interventions that better resolve edge orientation.

At the same time, the results demonstrate the efficiency gain in selecting the interventions, brought by encoding the knowledge of receptor identities into the prior. For example, while with the uninformative prior the true positive rate

Prior	Intervention targets selected by active learning
Informative	(1) hsp27, (2) mek12, (3) map3k1, (4) jnk12, (5) pi3k, (6) mkk4, (7) ikk, (8) akt, (9) p38, (10) erk12, (11) ikb
Uninformative	(1) jnk12, (2) hsp27, (3) mkk4, (4) mek12, (5) pi3k, (6) map3k1, (7) map3k7, (8) ras, (9) ikk, (10) akt, (11) ikb, (12) p38, (13) erk12

TABLE I: Intervention targets selected by active learning in the DREAM4 dataset. The informative prior edge probabilities required a smaller intervention batch.

of more than .35 could be achieved with on average 9 interventions, the informative prior only required on average 6 interventions. Table I shows the specific interventions that were selected at a conservative cutoff ($q_{S,v} \leq 0.01$). The results indicate that informative edge-wise priors are important for such bulk experiments with a small number of replicate samples. The prior knowledge removed uncertainty in edge presence, increasing the contribution of improved detection of edge orientation to overall performance.

B. The ordering of T-cell interventions by active learning matched their contribution to causal inference

As above, for the T-cell dataset we compared (1) random ordering of interventions, (2) proposed active learning strategy with uninformative prior probabilities of edge presence (Equation 6), and (3) the same active learning strategy, but with informative priors. In the latter case we assumed the prior knowledge of the canonical MAPK pathway, and assigned a high prior probability (~ 1) to the edges in the PKC \rightarrow Raf \rightarrow Mek \rightarrow Erk cascade, and to the edge PKA \rightarrow Raf. The remaining potential edges were assigned the uninformative prior probability of both presence and orientation. We then considered the five interventions in [37] as the set of candidate interventions.

Figure 5 summarizes the results. Selection with an uninformative edge-wise prior did not outperform random selection of interventions in terms of True Positive Rate of detecting edges. However, it had a smaller L_1 error for the first three selected interventions. With this experiment, the intervention datasets served not just to resolve causality, but to improve edge detection by adding variation in signaling activity not present in preceding datasets. As with the DREAM4 data, this lead to performance gains due to improved edge detection overshadowed gains owed to the selection strategy.

In contrast, active learning with the edge-wise prior encoding the MAPK edges outperformed random selection both in terms of greater True Positive Rate, and smaller L_1 error. Table II shows the specific interventions that were selected at a conservative cutoff ($q_{S,v} \leq 0.01$). For example, while with the uninformative prior the true positive rate of .75 could be achieved with on average 5 interventions, the informative prior only required on average 3 interventions.

The network structure provides some insight into the role of informative edge-wise priors in improving the performance. Since the orientation probability of an edge depends on the orientation probability of its neighbors, the edge-wise prior reduced error by reducing uncertainty in the orientation

Prior	Intervention targets selected by active learning
Informative	(1) PKA, (2) PKC, (3) PIP2
Uninformative	(1) PKA, (2) PIP2, (3) PKC, (4) Akt, (5) Mek

TABLE II: Intervention targets selected by active learning in the T-cell dataset. The use of an informative edge-wise prior eliminates two interventions from the batch.

of edges neighboring the MAPK edges. In addition, the edge-wise prior enabled the causal inference procedure to down-weight graphs where the MAPK edges were not present or had the wrong orientation, increasing sensitivity. Finally, additional causal information encoded in the prior made interventions on Mek and Akt interventions less useful, enabling the stopping criteria to eliminate them from the batch.

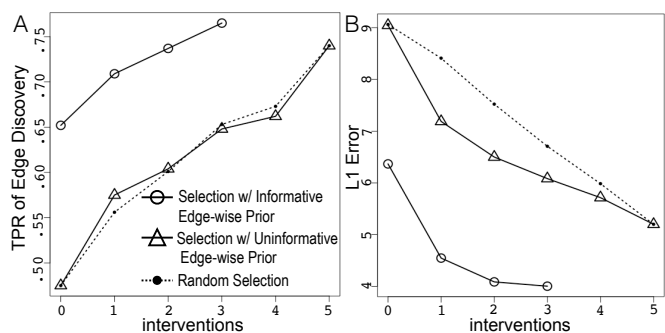


Fig. 5: Performance of the proposed strategy on the T-cell signaling dataset. Lines and panels are as in Figure 4.

VI. DISCUSSION

Our results showed that an active learning strategy, combined with informative priors, is the most effective at suggesting the smallest batch of target interventions for inference of causal networks. It optimizes the causal information in the intervention experiments, while controlling the experiment time and cost.

The background information comes in the form of prior knowledge on the presence and orientation of edges in the system available, e.g., in pathway databases such as KEGG, as well as from historic datasets available, e.g. from repositories such as Cytobank. The proposed strategy can be used with any level of prior information or uncertainty. When prior information and historic data indicate an intervention is potentially wasteful, it will tend to omit it from the batch, electing rather to reserving it for potential use in future experiments.

The active learning strategy in this manuscript suggests interventions based on the prior information and on the topology of the network. In practice, however, other factors can affect the utility of an intervention. For example, small molecule inhibitors vary both in cost and efficacy. Their inhibitory effects may only occur with some probability, and may also have off-target effects. The proposed framework can be easily extended to incorporate this type of “soft intervention” [8], as well as cost considerations into the

expected information gain. Alternatively, it can produce a batch of interventions with a fixed pre-specified number of targets, while selecting the targets with the most expected causal information gain.

The proposed methodology relies on signaling network modeling by means of direct acyclic graphs. In practice, however, cell signaling often displays feedback loops. This can be addressed by refining the biological interpretation of the graph structures. In particular, cycles in signaling often involve regulatory feedback loops that involve transcription. In this case, an activation of the signaling pathways causes transcription, which then results in the translation of new signaling proteins which then change the initial signaling pathway. Since the time scale of signaling is in seconds and minutes, and the time scale of transcription is in minutes and hours, collecting the data at an appropriate time point can help resolve the confounding between the initial causal effect of the signaling, and the feedback.

This work addressed both the inference of causal networks from bulk experiments and single cell experiments. Inferring an edges orientation depends of first detecting its presence, and since many edges between the proteins can potentially exist, bulk experiments often lack replicates to confidently detect the edges. Therefore, in bulk experiments it is often useful to add interventions not only to improve the detection of edge orientation, but also to improve the detection of edge presence through increased sample size.

In contrast, single cell experiments characterize protein signaling activity in thousands to millions of individual cells, and increase our confidence in the inferred edge. If the historic data was collected under a minimal number of conditions, intervention data may resolve both edge orientation and presence by virtue of adding variation in signaling activity. Moreover, single-cell experiments do not eliminate the need for true biological replication. The proposed approach can be extended to population level inference by modeling subjects as additional nodes in the network [36].

Overall, we believe that the proposed strategy is an important step towards an informed experimental design for inference of causal networks, and advocate its practical use.

ACKNOWLEDGEMENTS

We thank M. Scutari for guidance in using the R package *bnlearn*. This work was supported in part by the NSF CAREER award DBI-1054826, and by the Sy and Laurie Sternberg award to OV.

REFERENCES

- [1] D. R. Bandura, V. I. Baranov, O. I. Ornatsky, A. Antonov, R. Kinach, X. Lou, S. Pavlov, S. Vorobiev, J. E. Dick, and S. D. Tanner. Mass cytometry: Technique for real time single cell multitarget immunoassay based on inductively coupled plasma time-of-flight mass spectrometry. *Analytical Chemistry*, 81(16):6813–6822, 2009.
- [2] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Science & Business Media, 2013.
- [3] R. Castelo and A. Siebes. Priors on network structures. Biasing the search for Bayesian networks. *International Journal of Approximate Reasoning*, 24(1):39–57, 2000.

- [4] T. J. Chen and N. Kotecha. CytoBank: Providing an analytics platform for community cytometry data analysis and collaboration. In *High-Dimensional Single Cell Analysis*, pages 127–157. Springer, 2014.
- [5] D. M. Chickering. A transformational characterization of equivalent Bayesian network structures. In *Proceedings of the Eleventh Conference On Uncertainty In Artificial Intelligence*, pages 87–98. Morgan Kaufmann Publishers Inc., 1995.
- [6] D. M. Chickering and D. Heckerman. Efficient approximations for the marginal likelihood of Bayesian networks with hidden variables. *Machine Learning*, 29(2-3):181–212, 1997.
- [7] G. F. Cooper and C. Yoo. Causal discovery from a mixture of experimental and observational data. In *Proceedings of the Fifteenth Conference On Uncertainty in Artificial Intelligence*, pages 116–125. Morgan Kaufmann Publishers Inc., 1999.
- [8] D. Eaton and K. P. Murphy. Exact Bayesian structure learning from uncertain interventions. In *International Conference On Artificial Intelligence and Statistics*, pages 107–114, 2007.
- [9] F. Eberhardt, C. Glymour, and R. Scheines. On the number of experiments sufficient and in the worst case necessary to identify all causal relations among n variables. *arXiv Preprint arXiv:1207.1389*, 2012.
- [10] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the Graphical LASSO. *Biostatistics*, 9(3):432–441, 2008.
- [11] N. Friedman. Inferring cellular networks using probabilistic graphical models. *Science*, 303(5659):799–805, 2004.
- [12] N. Friedman et al. Learning belief networks in the presence of missing values and hidden variables. In *ICML*, volume 97, pages 125–133, 1997.
- [13] N. Friedman, M. Goldszmidt, and A. Wyner. Data analysis with Bayesian networks: A bootstrap approach. In *Proceedings of the Fifteenth Conference On Uncertainty In Artificial Intelligence*, pages 196–205. Morgan Kaufmann Publishers Inc., 1999.
- [14] N. Friedman and D. Koller. Being Bayesian about network structure: A Bayesian approach to structure discovery in Bayesian networks. *Machine Learning*, 50(1-2):95–125, 2003.
- [15] Y. Guan, M. Dunham, a. Caudy, and O. Troyanskaya. Systematic planning of genome-scale experiments in poorly studied species. *PLoS Comput Biol*, 6(3):e1000698, 2010.
- [16] Y.-B. He and Z. Geng. Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research*, 9(11), 2008.
- [17] D. Heckerman, D. Geiger, and D. M. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20(3):197–243, 1995.
- [18] J. S. Ide and F. G. Cozman. Random generation of Bayesian networks. In *Advances In Artificial Intelligence*, pages 366–376. Springer, 2002.
- [19] T. Ideker and N. J. Krogan. Differential network biology. *Molecular Systems Biology*, 8(1):565, 2012.
- [20] S. Imoto, S. Y. Kim, H. Shimodaira, S. aburatani, K. Tashiro, S. Kuhara, and S. Miyano. Bootstrap analysis of gene networks based on Bayesian networks and nonparametric regression. *Genome Informatics*, 13:369–370, 2002.
- [21] M. Kanehisa, Y. Sato, M. Kawashima, M. Furumichi, and M. Tanabe. Kegg as a reference resource for gene and protein annotation. *Nucleic Acids Research*, 44(D1):D457–D462, 2016.
- [22] R. D. King, K. E. Whelan, F. M. Jones, P. G. Reiser, C. H. Bryant, S. H. Muggleton, D. B. Kell, and S. G. Oliver. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427(6971):247–252, 2004.
- [23] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- [24] K. B. Korb and a. E. Nicholson. *Bayesian Artificial Intelligence*. CRC Press, 2010.
- [25] D. Margaritis. *Learning Bayesian Network Model Structure from Data*. PhD thesis, US Army, 2003.
- [26] S. Meganck, P. Leray, and B. Manderick. Learning causal Bayesian networks from observations and experiments: A decision theoretic approach. In *Modeling Decisions for Artificial Intelligence*, pages 58–69. Springer, 2006.
- [27] K. P. Murphy. Active learning of causal Bayes net structure. 2001.
- [28] R. O. Ness, K. Sachs, and O. Vittek. From correlation to causality: Statistical approaches to learning regulatory relationships in large-scale biomolecular investigations. *Journal of Proteome Research*, 2016.

- [29] T. Pawson and N. Warner. Oncogenic re-wiring of cellular signaling pathways. *Oncogene*, 26(9):1268–1275, 2007.
- [30] J. Pearl. *Causality: Models, Reasoning and Inference*, volume 29. Cambridge Univ Press, 2000.
- [31] O. D. Perez and G. P. Nolan. Simultaneous measurement of multiple active kinase states using polychromatic flow cytometry. *Nature Biotechnology*, 20(2):155–162, 2002.
- [32] I. Pournara and L. Wernisch. Reconstruction of gene networks using Bayesian learning and manipulation experiments. *Bioinformatics*, 20(17):2934–2942, 2004.
- [33] R. J. Prill, J. Saez-Rodriguez, L. G. Alexopoulos, P. K. Sorger, and G. Stolovitzky. Crowdsourcing network inference: The DREAM predictive signaling network challenge. *Science Signaling*, 4(189):mr7, 2011.
- [34] D. Rossell and P. Müller. Sequential stopping for high-throughput experiments. *Biostatistics*, 14(1):75–86, 2013.
- [35] S. J. Russell, P. Norvig, J. F. Canny, J. M. Malik, and D. D. Edwards. *Artificial Intelligence: A Modern Approach*, volume 2. Prentice Hall Upper Saddle River, 2003.
- [36] K. Sachs, A. J. Gentles, R. Youland, S. Itani, J. Irish, G. P. Nolan, and S. K. Plevritis. Characterization of patient specific signaling via augmentation of bayesian networks with disease and patient state nodes. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 6624–6627. IEEE, 2009.
- [37] K. Sachs, O. Perez, D. Pe’er, D. a. Lauffenburger, and G. P. Nolan. Causal Protein-Signaling Networks Derived From Multiparameter Single-cell Data. *Science (New York, N.Y.)*, 308(5721):523–9, Apr. 2005.
- [38] M. Scutari. Learning Bayesian networks with the bnlearn R package. *Journal of Statistical Software*, 35(3):1–22, 2010.
- [39] M. Scutari. On the prior and posterior distributions used in graphical modelling. *Bayesian Analysis*, 8(3):505–532, 2013.
- [40] C. Terfve, T. Cokelaer, D. Henriques, A. MacNamara, E. Goncalves, M. K. Morris, M. van Iersel, D. A. Lauffenburger, and J. Saez-Rodriguez. Cellnoptr: A flexible toolkit to train protein signaling networks to data using multiple logic formalisms. *BMC systems biology*, 6(1):1, 2012.
- [41] C. Terfve and J. Saez-Rodriguez. Modeling signaling networks using high-throughput phospho-proteomics. In *Advances In Systems Biology*, pages 19–57. Springer, 2012.
- [42] J. Tian and J. Pearl. Causal discovery from changes. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 512–521. Morgan Kaufmann Publishers Inc., 2001.
- [43] S. Tong and D. Koller. Active learning for structure in Bayesian networks. In *International Joint Conference On Artificial Intelligence*, volume 17, pages 863–869. Lawrence Erlbaum Associates Ltd, 2001.
- [44] A. V. Werhli and D. Husmeier. Reconstructing gene regulatory networks with Bayesian networks by combining expression data with multiple sources of prior knowledge. *Statistical Applications In Genetics and Molecular Biology*, 6(1), 2007.