

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

11-1-2021

NOTMAD: estimating bayesian networks with sample-specific structures and parameters

Benjamin Lengerich

Massachusetts Institute of Technology & Broad Institute of MIT and Harvard

Caleb Ellington

Carnegie Mellon University, United States

Bryon Aragam

University of Chicago

Eric P. Xing

Carnegie Mellon University, United States & Mohamed bin Zayed University of Artificial Intelligence

Manolis Kellis

Massachusetts Institute of Technology & Broad Institute of MIT and Harvard

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Artificial Intelligence and Robotics Commons](#)

Preprint: arXiv;

Archived with thanks to arXiv;

Preprint License: CC by NC-ND 4.0;

Uploaded 05 May 2022

Recommended Citation

B. Lengerich, C. Ellington, B. Aragam, E. Xing, and M. Kellis, "NOTMAD: estimating bayesian networks with sample-specific structures and parameters", arXiv, Nov 2021, doi: 10.48550/arXiv.2111.01104

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Machine Learning Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

NOTMAD: Estimating Bayesian Networks with Sample-Specific Structures and Parameters

Ben Lengerich^{*1, 2}, Caleb Ellington³, Bryon Aragam⁴, Eric P. Xing^{3, 5}, and Manolis Kellis^{1, 2}

¹Massachusetts Institute of Technology

²Broad Institute of MIT and Harvard

³Carnegie Mellon University

⁴University of Chicago

⁵Mohamed bin Zayed University of Artificial Intelligence

November 2, 2021

Abstract

Context-specific Bayesian networks (i.e. directed acyclic graphs, DAGs) identify context-dependent relationships between variables, but the non-convexity induced by the acyclicity requirement makes it difficult to share information between context-specific estimators (e.g. with graph generator functions). For this reason, existing methods for inferring context-specific Bayesian networks have favored breaking datasets into subsamples, limiting statistical power and resolution, and preventing the use of multi-dimensional and latent contexts. To overcome this challenge, we propose NOTEARS-optimized Mixtures of Archetypal DAGs (NOTMAD). NOTMAD models context-specific Bayesian networks as the output of a function which learns to mix archetypal networks according to sample context. The archetypal networks are estimated jointly with the context-specific networks and do not require any prior knowledge. We encode the acyclicity constraint as a smooth regularization loss which is back-propagated to the mixing function; in this way, NOTMAD shares information between context-specific acyclic graphs, enabling the estimation of Bayesian network structures and parameters at even single-sample resolution. We demonstrate the utility of NOTMAD and sample-specific network inference through analysis and experiments, including patient-specific gene expression networks which correspond to morphological variation in cancer.

1 Introduction

Estimating the structure and parameters of Bayesian networks (BNs) is a classic problem in which we seek to summarize relationships between variables by factorizing the joint distribution as a probabilistic graphical model. This approach has been applied to many problems in biology [34], genetics [45], machine learning [22], and causal inference [38]. However, the relationships between variables are often non-stationary and context-specific. Previous attempts to model these context-specific relationships, such as time-varying networks [3, 31, 37], have enjoyed success in estimating network connections which vary according to a single observed axis of variation (e.g. time); however, when dealing with complex processes which vary according to many processes simultaneously, we often do not know *a priori* how observed covariates relate to network structure. Thus, we would like a method which can estimate BNs with structures and coefficients that vary according to several, and possibly latent, continuous factors. Such a method would enable us to understand large, heterogeneous datasets as directed relationships which rewire in continuously-changing contexts.

Toward this end, we propose NOTEARS-optimized Mixtures of Archetypal DAGs (NOTMAD), a framework for network estimation which views BNs as the output of a learnable function of sample context. There are two main challenges in designing such a method: how can we enforce acyclicity (traditionally expressed as

*blengeri@mit.edu

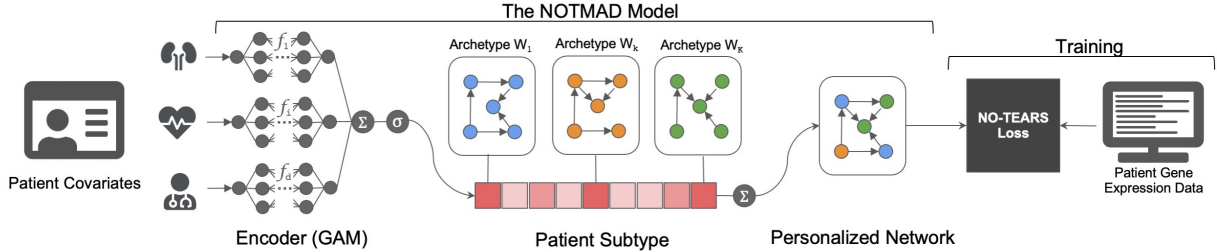


Figure 1: Our motivating example: patient-specific network inference for understanding cancer at high resolution. We know that each patient’s tumor is uniquely shaped by genetic and environmental factors. We would like to summarize the gene regulatory relationships in each tumor as a BN to identify mutations and therapeutic targets. However, the assays which measure gene expression destroy the tissue sample, so traditional methods of network inference which treat networks in isolation are unable to estimate patient-specific networks. Our framework instead views each network as the output of a function of patient-specific covariates, which can be used to predict the patient’s gene network as a personalized mixing of archetypal networks. In this view, the generative function of regulatory networks is shared between many patients and can be learned efficiently, even when we have only a single sample from each patient.

a combinatorial constraint) on the *output* of a learned function, and how should we encode latent network structure in a learnable function? Our key insight is that acyclicity function can be imposed as a smooth regularization loss on the graph generator, allowing context-specific networks to be modeled as mixtures of archetypal DAGs. As a result, NOTMAD is able to generate context-specific BNs at single-sample resolution, allowing us to ask new questions regarding sample heterogeneity and latent structure.

Motivating Example: Gene Regulatory Network Inference for Personalized Cancer Analysis

A common task in bioinformatics is to infer regulatory networks from gene expression data. That is, we seek a graphical model which factorizes the distribution of gene expression, from which we can infer regulatory relationships. This graphical model representation is especially useful for cancer analysis: cancer is a highly individualized disease [28], for which distinct biological processes are important for understanding and treating different tumors; moreover, cluster-based analyses are limited because tumors may not form discrete clusters [42]. Thus, we seek patient-specific BNs to identify dysregulated pathways and therapeutic targets at patient-specific granularity. However, bulk RNA sequencing assays which measure gene expression destroy the physical sample; as a result, we are given only a single instance from which to view each tumor. Thus, traditional methods to estimate BNs by clustering samples are incapable of providing patient-specific inferences. In this paper, we develop a method (Figure 1) which views BNs as the output of a learnable function of contextual patient information. This enables us to estimate personalized gene regulatory networks, and perform personalized analysis of cancer patients while sharing statistical power between many samples.

Contributions This paper contributes NOTMAD, a method for estimating context-specific BNs which enables sample-specific network inference, enabling analyses of network heterogeneity. Underlying NOTMAD is a meta-view of structured models in which parameters are outputs of a shared, learnable graph generator function. Previously, it was thought that the structured aspect of these models prohibited graph generator functions; we show that these structured models (e.g. context-specific BNs) can be efficiently modeled as interpolations between a small number of latent archetypes, and the non-convexity of the model domain can be enforced with regularization. This perspective unites two areas of recent interest in machine learning: structure learning and contextualization. Python code of NOTMAD is available at <https://github.com/cnellington/SampleSpecificDAGs>.

2 Preliminaries

DAG Structure Learning The basic DAG learning problem is formulated as: Let $X \in \mathbb{R}^{n \times p}$ be a data matrix consisting of n IID observations of the random vector $X = (X_1, \dots, X_p)$. Let \mathbb{D} denote the space of DAGs $G = (V, E)$ of p nodes. Given X , we seek weighted adjacency matrix $W \in \mathbb{R}^{p \times p}$ which defines a DAG $G \in \mathbb{D}$ with structure defined by the binarized adjacency matrix $A(W)$ where $A(W)_{ij} = 1 \iff W_{ij} \neq 0$, and describes the joint distribution $\mathbb{P}(X)$ via a structural equation model [22, 38].

Structural Equation Models To translate the weighted adjacency matrix $W \in \mathbb{R}^{p \times p}$ into a joint distribution on X , we use a structural equation model (SEM). Linear SEM models $X_j = W_j^T X + \epsilon_j$, where $X = (X_1, \dots, X_p)$ is a random vector and $\epsilon = (\epsilon_1, \dots, \epsilon_p)$ is a random noise vector. This form can be modified to use more flexible transition functions with $\mathbb{E}(X_j | X_{pa(X_j)}) = f(W_j, X)$ (e.g., for Boolean X_j , we can use a logistic link function). For the remainder of this paper, we will assume linear SEM with ϵ Gaussian and drop this specification; however, we note that NOTMAD easily extends to other SEM by replacing this portion of loss function. Thus, a DAG W defines a Bayesian Network (BN) distribution $\text{BN}(X|W)$ which provides the likelihood of observing X under structure W and imposes least-squares loss $\ell(W; X) = \frac{1}{2n} \|X - XW\|_F^2$.

3 Inferring Context-Specific Bayesian Networks

3.1 Problem Definition

We are interested in learning BNs which are *context-specific*; i.e., the parameters and/or structure of the BNs may vary according to context. In this problem, for each observation $X \in \mathbb{R}^p$, we also observe contextual data $C \in \mathbb{R}^m$. We can describe this as factorizing

$$P(X, C) = \int_W dW P(X|W) P(W|C) P(C),$$

where $P(X|W) = \text{BN}(X|W)$ (recalling $\text{BN}(\cdot|W)$ was defined previously as the distribution implied by BN structure W). Drawing back to the motivating example discussed in Section 1, we can view X as gene expression, W as a gene regulatory network, and C as patient covariates which contextualize the network. This is similar to a varying-coefficient (VC) description of BNs, which would model $W = f(C)$ with some learnable f . However, it is not tractable in practice to estimate context-specific BNs using this VC form because DAGs are high-dimensional and DAG space is non-convex, so defining f to generate DAGs is unsolved.

We are particularly interested in the extreme case where there is at most one observation per context, i.e. sample-specific BNs. In such cases, there is limited power to infer each BN. This differs from other applications in we many samples per context (e.g. C is discrete and takes on finitely-many values), and cluster-based models may succeed. In the most challenging setting of sample-specific inference, we require a tool to share power between sample-specific estimators.

3.2 Our Approach: Smooth Network Inference and Contextualization

As described above, the main challenge in inference of context-specific BNs is that W is a high-dimensional, structured latent variable which must satisfy acyclicity. To overcome this challenge, we assume that context-specific BNs lie on a subspace measured by a latent variable $Z \in \mathbb{R}^K$ such that $C \perp (X, W) \mid Z$. A graphical diagram of this meta-model is shown in Fig. 2. Drawing back to the motivating example, we can imagine that tumors correspond to latent molecular subtypes [24, 44] and mutations Z which govern the underlying gene regulatory functions W ; this conditional independence implies that knowing the patient covariates C would not give any additional insight into the gene expression values X if either the molecular subtype or the graph structure W were already known.

With this meta-model, $\mathbb{P}(W|X, C) \propto \mathbb{P}(X|W) \int_Z dZ \mathbb{P}(W|Z) \mathbb{P}(Z|C)$, where $\mathbb{P}(X|W) = \text{BN}(X|W)$ and our challenge reduces to defining tractable $\mathbb{P}(W|Z)$ and $\mathbb{P}(Z|C)$. In NOTMAD, we use a smooth DAG criterion to define $\mathbb{P}(W|Z)$ and use contextual parameter generation to estimate latent Z to generate W . In this way, our approach to context-specific network inference builds on two developing themes in machine learning: smooth DAG inference and contextual parameter generation.

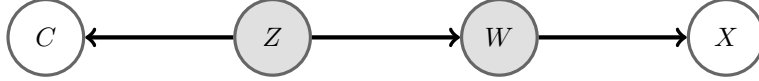


Figure 2: Graphical Model. Contextual covariates C and observations X are observed, while subtype Z and BN parameters W are latent.

Smooth DAG Inference via NOTEARS Estimating DAG structure, even for a single population model, from data is NP-hard [9], largely due to the acyclicity constraint which has traditionally been formulated as a combinatorial optimization. Recent work [46], however, has translated this combinatorial optimization into a continuous program through the matrix exponential: $h(W) = \text{tr}(e^{W \cdot W}) - p = 0 \iff W \in \mathbb{D}$, where \cdot is the Hadamard product and e^A is the matrix exponential of A . This expression of acyclicity is useful because $h(W)$ has a simple gradient: $\nabla h(W) = (e^{W \cdot W})^T \cdot 2W$. In NOTMAD, we relax $h(W)$ into a regularizer which is backpropagated for gradient-based training.

Contextual Parameter Generation As datasets continue to increase in size and complexity, sample-specific inference has driven interest in many application areas [2, 6, 12, 17, 29]. Several frameworks have been proposed to estimate meta-models which use contextual information C to generate model parameters θ . In general, we can write the operation of these frameworks as $\theta = Z^T Q$, where $Z = g_\phi(C)$. This form encompasses several well-known framework of contextualized models: if g is a linear model and Q is an identity matrix, we have the varying-coefficients model [18], if g indexes the k -nearest neighbors in Q by a learned distance metric on C , we have personalized regression [26], if g is a deep neural network and $\|Z\| = 1$, we have Contextual Explanation Networks [4], if g is a few steps of gradient-descent beginning at the initialization ϕ for the instances grouped by C , we have meta-learning [11, 36]. In NOTMAD, Q is a learned dictionary of DAG archetypes and g is a differentiable context encoder (details in Sec.4).

3.3 Related Work

This problem of estimating sample-specific BNs has its roots in context-specific independence (CSI) [5], which permitted that the independence relations captured in BN sparsity patterns may hold only for certain contexts. In general, CSI transforms the traditional BN DAG into a multi-graph; a long line of work have invested into developing alternative representations and approaches for estimating CSIs [10, 14, 30, 32], and more recent approaches have used causality rules to perform instance-specific independence tests [19–21]. However, these approaches rely on the combinatorial view of acyclicity and hence suffer scalability challenges.

There have also been approaches to design estimators of multiple BNs without sample-specific covariates. These frameworks include priors governing network rewiring [15, 25] and sample-specific networks as deviations from a population network [23]. For example, LIONESS [23] estimates linear network for sample i as the weighted difference between the network estimated by all n samples and the network estimated by $n - 1$ samples, leaving out the i th sample. While this approach showcases the clever idea of the exchangeability of these estimators and differences in linear SEM, it does not share statistical power between sample estimators (the statistical power is only shared for the population model). A similar approach was used to identify sample-specific dysregulated pathways in a gene regulatory network by examining the residual of neighborhood regressions [7], although this approach was more robust than the above because it sought only to estimate dysregulation rather than the underlying network. There are fundamental limits to the estimation of populations of parameters without any covariate information [43]; for this reason, we concern ourselves only with the setting where we observe covariate information to contextualize the BN.

Finally, we are not the first to study mixtures of archetypal DAGs. Past methods have focused on estimating the latent mixture subtype of a population [40, 41], or of individual samples [35]; however, these methods use localized structure-search algorithms, which creates a computational bottleneck that constrains learning to only mixing a pre-defined set of archetypal DAGs. This severely hinders the adaptability of these models because it requires a perfect prior knowledge about the causal structures underlying DAG mixtures. In contrast, NOTMAD uses smooth optimization to learn archetypal DAGs without any prior knowledge.

4 NOTMAD: NOTEARS-optimized Mixtures of Archetypal DAGs

To estimate sample-specific BNs, we propose NOTMAD, which represents context-specific distributions as:

$$\mathbb{P}(X|C) = \text{BN}(X|\phi_\theta(C)), \quad (1)$$

which equates to $X = \phi_\theta(C)^T X + \epsilon$ under the linear SEM, where ϵ is a noise vector. Context-specific DAGs $W = \phi_\theta(C)$ are generated by ϕ_θ as a convex combination of a fixed number (K) of archetypal DAGs ($W_{1:K}$):

$$W = \phi_\theta(C) = \sum_{k=1}^K \sigma(f_\theta(C))_k [W_k \cdot (1 - I)] \quad (2)$$

where $\sigma(\cdot)$ is the softmax function and $f_\theta(\cdot)$ is the context encoder, a learned function parameterized by θ . The softmax constraint ensures that the weighting of each archetype is positive and sums to 1, permitting these values to be interpreted as the probability of class membership. The dictionary of archetypal DAGs, $W_{1:K}$, is estimated jointly with the context encoder f , and the entire architecture is trained end-to-end. Finally, we clamp the diagonal of each W_k to 0 to avoid the cyclic $W = I$.

To ensure that each W is acyclic, we balance the estimation loss with the NOTEARS DAG-ness loss $h(W)$. Thus, NOTMAD can be summarized as the following minimizer:

$$\underset{\theta, W_{1:K}}{\operatorname{argmin}} \sum_{i=1}^n (X^i - X^i \phi_\theta(C^i))^2 + \alpha \operatorname{tr}(e^{\phi_\theta(C^i) \cdot \phi_\theta(C^i)}) + \sum_{k=1}^K \beta \|W_k\|_1 + \gamma \operatorname{tr}(e^{W_k \cdot W_k}) \quad (3)$$

where α , β and γ are hyperparameters that trade off predictive loss against DAG-ness and sparsity, and we recall that $h(W) = \operatorname{tr}(e^{W \cdot W})$ is the acyclicity regularizer discussed in Section 3.2, here applied to both the context-specific and the archetypal DAGs. As introduced in Sec. 3.2, NOTMAD has a natural probabilistic interpretation. We have defined $\mathbb{E}[\mathbb{P}(W|Z)] = \sum_k Z_k [W_k \cdot (1 - I)]$, and constrained $Z \in \{q \in \mathbb{R}_{\geq 0}^K : \|q\| = 1\}$ by $Z = \sigma(f_\theta(C))$ such that these weightings are a convex combination. Python code of NOTMAD is available at <https://github.com/cnellington/SampleSpecificDAGs>.

Context Encoder This definition of NOTMAD is flexible with respect to the form of the context encoder f – any differentiable function can be used. We have found best success and report all experimental results using a Neural Additive Model [1], which rivals the accuracy of fully-connected neural networks while retaining interpretability. For other types of contextual data, other encoder models may be more appropriate; e.g. a convolutional neural network for images.

Fine-Tuning The NOTEARS regularizer encourages the context encoder to produce acyclic graphs; however, we are not guaranteed that the graph predicted by NOTMAD will be acyclic for all unseen test contexts. Nevertheless, any directed graph can be projected to a DAG by removing edges which violate acyclicity: for fixed W , we perform a binary search over the magnitudes in W to find the minimum value t for which $W \cdot (|W| > t) = W' \in \mathbb{D}$. This produces a DAG W' , to which we attempt to add back excluded edges from W in order of decreasing magnitude of edge weights, checking that each edge addition does not form a cycle. In practice, predictions from NOTMAD do not require fine-tuning beyond thresholding out weak edges.

Benefits of Modeling Networks as Mixtures of Archetypes In some sense, NOTMAD is comparable to a multi-output varying-coefficients model with an extra loss to encourage acyclicity. However, NOTMAD uses archetype mixing to create a bottleneck and reduce the effective dimensionality of the outputs (borrowing the dictionary-learning framework of [4]). This bottleneck has statistical benefits: constraining the sample-specific DAGs to lie in a K -dimensional polytope defined with the K archetypes as extreme points. In addition, this has computational benefits: it is not tractable to store many $p \times p$ matrices in memory; the archetype view allows efficient computation and optimization. Finally, archetypes aid interpretability: we can understand latent membership patterns by visualizing the context-subtype mapping in the additive context encoder f_θ .

Non-Convexity Is it reasonable to define sample-specific DAGs as convex combinations of archetypal DAGs? The DAG constraint is non-convex, so it is not true that the convex combination of any set of DAG archetypes would be acyclic. If DAG space is highly disconnected, then we should not expect optimization to efficiently learn to generate DAGs by combining archetypal DAGs. Is NOTMAD hopeless?

Inspecting the topology of DAG space gives hope: the compatibility of DAGs is driven by the sparsity structure of summands, not the weighting. Let $A(W)$ be the binarized structure of W (i.e. $A(W)_{ij} = 1 \iff W_{ij} \neq 0$). Then $A(W_1) + A(W_2) \in \mathbb{D} \implies aW_1 + W_2 \in \mathbb{D} \forall a$. Conversely, if $A(W_1) + A(W_2) \notin \mathbb{D}$, then $aW_1 + W_2 \in \mathbb{D}$ for at most p^2 values of a . As a result, there is minimal confusion over allowable archetype combinations: either the combination is acyclic for all weightings, or there is a small, finite set of weightings which produce DAGs and all other weightings produce cyclic graphs. Thus, the context encoder f needs only to learn compatible archetypes, not compatible weightings.

5 Experiments

BNs factorize joint distributions into tractable conditional dependencies; thus inference of context-specific BNs provides insight into the changing dependencies in heterogeneous data. Traditionally, structure learning has required multiple observations for each network; in contrast, NOTMAD enables us to estimate sample-specific networks for unclustered individual observations. Here, we evaluate the success and limitations of NOTMAD to identify the accuracy and utility of sample-specific network inference.

Baselines We compare the performance of NOTMAD against estimators of BNs at various resolutions: state-of-the-art population estimator (NOTEARS [46]), a cluster-specific estimator which first clusters the samples according to the context and then applies the population estimator to each cluster, and sample-specific BN inference by LIONESS [23]. Finally, when we have domain knowledge regarding latent sample structure, we construct “oracle”-clustered estimators by clustering the samples according to domain knowledge and applying the population estimator to each cluster. On real data, we do not have access to the underlying BN; thus, we evaluate the mean-squared error (MSE) of $\|X - XW\|_2$ on held-out samples which corresponds to the likelihood of the observed data under the probabilistic interpretation of the linear SEM.

5.1 Finance

Finally, we explore heterogeneity in a dataset of financial asset prices and news headlines from [26]. This dataset was proposed to show the utility of allowing regression parameters to change over time and industry; here we examine how the inter-relatedness of the samples drifts over time by estimating contextualized BNs.

For each date, we have as context C a summary of the news headlines as a 50-dimensional vector, while the predictors X are the trading price of 24 assets. We split the 14,555 samples into train/test sets at the 75th percentile date, which is approximately the beginning of 2011, to produce a training set of 10,916 samples. From these data, we estimate daily BNs of asset price relations contextualized by history and news.

In addition to giving the best MSE loss on held-out samples (Table 1), the estimated BNs show interpretable patterns. As shown in Fig. 3, the sample-specific networks accurately translate the changing environment into interpretable networks of financial relationships.

Dataset	Population	Clustered	Oracle Clustered	LIONESS	NOTMAD
Finance	0.58 ± 0.01	0.54 ± 0.01	-*	0.67 ± 1.94	0.43 ± 0.008
Football	0.47 ± 0.006	0.44 ± 0.006	0.41 ± 0.004	0.52 ± 0.19	0.35 ± 0.005
Cancer	$0.49 \pm 2.3e-5$	$0.37 \pm 6.5e-4$	$0.24 \pm 1.1e-4$	0.65 ± 0.14	$0.14 \pm 2.7e-5$
Single-Cell	0.56 ± 0.02	0.45 ± 0.01	0.37 ± 0.02	0.60 ± 0.08	0.28 ± 0.02

Table 1: Mean-squared error (MSE) loss of BNs applied to held-out samples. Values displayed are the mean \pm variance of the MSE over 10 bootstrap samples. *No oracle for clustering.



(a) Embeddings of context-specific BNs by date. Context-specific BNs recover trends leading to and the broad dislocation after the 2008 financial crisis. (b) Cosine similarities between context-specific BNs by date. Context-specific BNs recover macro trends including 1970s stagflation, 2000s growth, and 2008 financial crisis.

Figure 3: Representations of context-specific BNs estimated for financial asset prices.

5.2 Football

A goal in sports analytics is to isolate the influence of individual players from the abilities of teammates. Professional players are often rated by independent rating agencies; we can use these ratings to identify influences by estimating a BN to factorize player ratings. In such a BN, each node represents a player rating while edges represent how the rating of a player is influenced by the ratings of teammates. Context is critical: team strategies dictate player expectations and analyses. We use ratings¹ of 1000 professional football players from the 2020 National Football League season. We group players into 7 position categories, and construct simulated observations of lineups by team. For context, we use the team identity and a PCA-reduced representation of the ratings; for oracle clustering, we use team identity. In this dataset, $n = 1478$ (with 634 test samples), $p = 7$, and $m = 4$. As shown in Table 1, context-specific BNs capture the distributions of player ratings more accurately than population or cluster-based BNs.

5.3 Gene Regulatory Network Inference for Personalized Analysis of Cancer

We now return to the motivating example discussed in Section 1.

Dataset We use bulk RNA-seq gene expression data from The Cancer Genome Atlas² (TCGA), which measures the average expression of each gene in a patient’s tumor. We select 784 RNA-seq transcripts (features) based on annotation in the COSMIC gene census [13] and large variance over the sample. These gene expressions levels are measured for 9715 patients, which each have 12 contextual covariates: Age, Gender, Race, Sample Type, Primary Site, % stromal cells, % tumor Cells, % Normal Cells, % Neutrophil Infiltration, % Lymphocyte Infiltration, % Monocyte infiltration, and a single principal component of gene expression which captures batch effects in processing. The cancers span 20 tissues and 26 types of tumors.

Results Patient-specific BNs estimated by NOTMAD represent gene expression better than do population, cluster-specific, or tissue-specific BNs. Not only do patient-specific BNs improve average MSE, (Table 1) the patient-specific BNs also improve estimation in every tissue (Fig. 4). The largest out-performance is seen for cancers of the lymph nodes (MSE of tissue-specific BNs is 4.9x that of NOTMAD), brain (2.4x), thyroid (2.0x), prostate (1.9x), colon (1.8x), and head/neck (1.8x).

The high resolution of these patient-specific BNs unlocks new investigations of network heterogeneity. For example, are the patient-specific distributions of gene expression best viewed as a patient-specific *selection* of a distribution, or a patient-specific *mixture* of archetypal distributions (i.e. do gene regulatory networks in cancer rewire according to a small number of fixed processes, or do they interpolate between processes)? With NOTMAD, we can begin to answer this question. Here, we examine this question through an ablation study: constraining each patient-specific BNs to use only a single archetype (“One Archetype” in Fig. 4) or to use only a single BN for all tumors of a particular tissue (“Tissue-average” in Fig. 4). According to MSE,

¹<https://www.lineups.com/nfl/player-ratings>

²<https://portal.gdc.cancer.gov/>

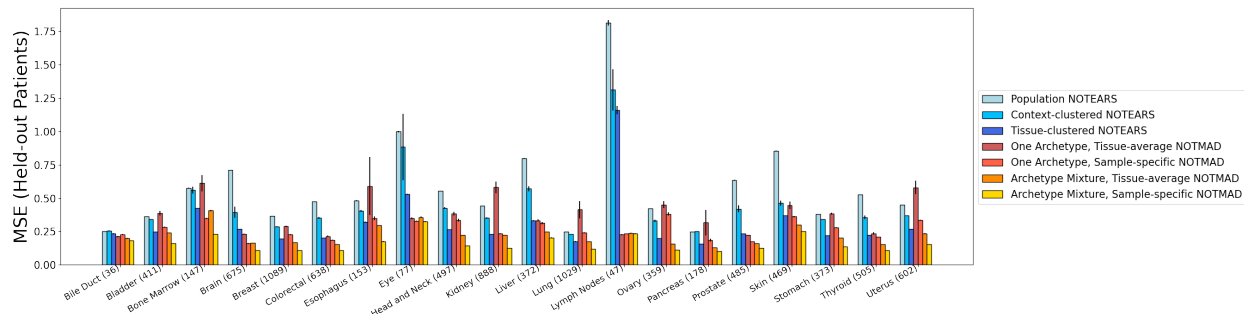


Figure 4: MSE of BNs of gene expression in cancer, grouped by tissue. The number of training samples for each tissue type is provided in parentheses. In addition to the population, cluster-specific, and tissue-specific baselines, we also include ablations of the NOTMAD model: “One Archetype, Tissue average” uses the power sharing of NOTMAD across tissues, but only estimates a single BN for each tissue; “Archetype Mixture, Tissue average” similarly estimates a single BN for each tissue but allows this tissue-specific BN to be a mixture of archetypal DAGs; “One Archetype, Sample-Specific” estimates patient-specific BNs that are constrained to use only a single archetype; “Archetype Mixture, Sample-Specific” is the entire NOTMAD model.

the full patient-specific heterogeneity of NOTMAD is required to achieve optimal summarization of gene expression distributions, suggesting that these cancers exhibit a diversity of expression patterns that cannot be resolved by a “cancer” BN (i.e. a population model), tissue-specific BNs, or patient-specific BNs which are selections of clustered networks; instead, tumors are best modeled as interpolations between latent processes.

In addition, the patient-specific BNs recover medically-significant heterogeneity. By inspecting network embeddings (Fig. 5), we see that the BNs estimated by NOTMAD tend to cluster according to tissue. Nevertheless, the patient-specific BNs separate lung adenocarcinoma (LUAD) samples from lung squamous cell carcinomas (LUSC). These subtypes of lung cancer have important distinctions in clinicopathology (e.g. LUSC is associated with a history of smoking, while LUAD is the most common form of lung cancer in non-smokers [33]) and treatment (e.g. epidermal growth factor receptor and anaplastic lymphoma kinase are two commonly mutated genes in LUAD and are targets for current therapies, while they are infrequently mutated in LUSC [27]). Since NOTMAD did *not* have access to this LUAD/LUSC distinction as contextual data, this distinction must be a result of optimizing MSE loss, indicating different regulatory networks underlying these biomedically-distinct subtypes of lung cancers.

In addition to recovering known heterogeneity of disease types, NOTMAD also identifies intra-disease heterogeneity which correspond to morphological distinctions. For example, the BNs for colon cancers tend to be interpolations between two archetypal networks, and similarly head and neck cancers have two archetypal networks underlying the patient-specific BNs. For colon cancer, this heterogeneity may correspond to the four consensus molecular subtypes [16], while for head and neck cancers, this split may correspond to differences in head and neck tumors arising from either tobacco use and/or HPV [39]. Whole-slide imaging of these tumors (Fig. 5) shows the morphological distinctions between the two latent BN archetypes, suggesting a correspondence between gene regulatory network heterogeneity and tumor morphology.

5.4 Single-Cell Gene Regulatory Network Inference

The gene regulatory network of each brain cell is unique, shaped by its cell type, cell state, history of activation, interconnections, environment, and by the age, sex, and disease status of each person. However, these networks are highly related across cells and highly structured, as they are shaped by circuitry constraints, a limited set of regulatory programs, and a finite set of active regulators. Here, we exploit this structure by using NOTMAD to simultaneously estimate network archetypes and the cell-specific networks using paired epigenetic and transcriptomic data. Importantly, we learn both membership vectors and archetypal networks *de novo* in an end-to-end joint optimization, in essence inferring a functional transformation that predicts a cell’s gene regulatory network from its DNA accessibility profile.

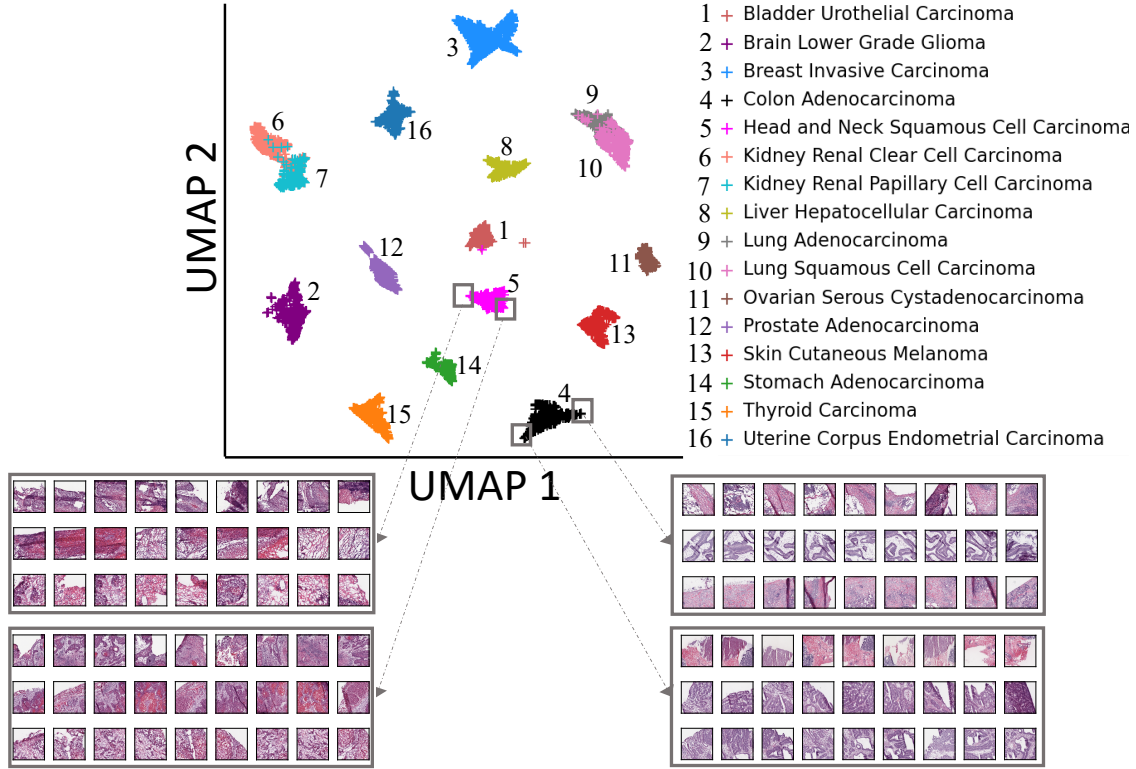
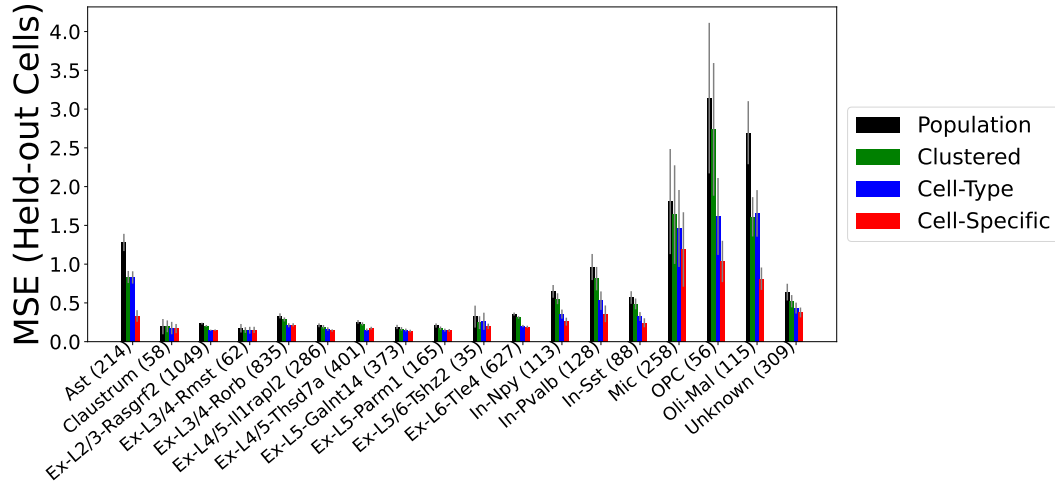


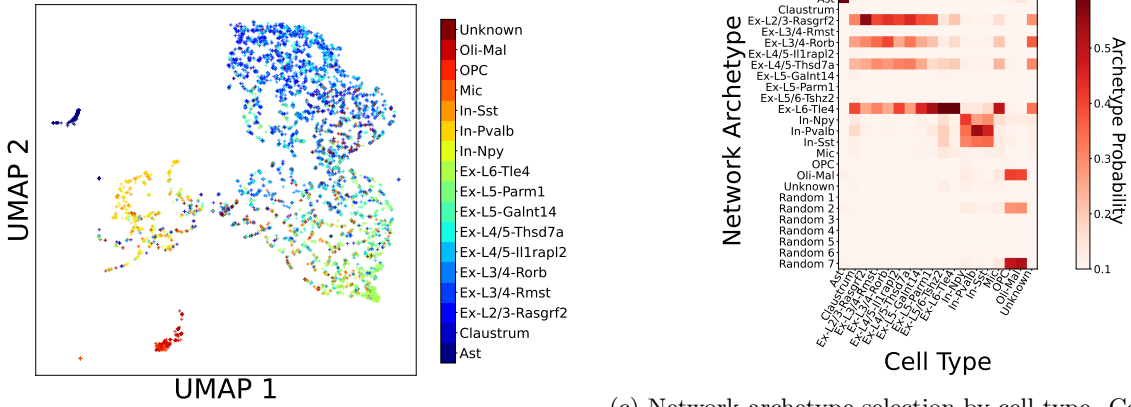
Figure 5: Patient-specific BNs embedded by UMAP. (Top) Without access to disease type during training, the patient-specific BNs accurately recover heterogeneity between cancers of disparate tissues, and cancers of different subtypes affecting the same tissue. (Bottom) Network archetype selection suggests network heterogeneity in both head and neck squamous cell carcinomas (left) and colon adenocarcinomas (right). For each archetype, we show image patches from the 3 cases (by row) which maximize NOTMAD activation of the archetype (10 patches each) and draw arrows to the corresponding sample group containing these cases.

Dataset To estimate single-cell regulatory networks, we use recently-developed dual assay technology, which allows simultaneous profiling of transcriptomic (RNA-seq) and epigenetic (ATAC-seq) markers in matched single cells. In this example, we use the SNARE-seq dataset of 10,309 cells from adult mouse brain cells [8]. To preprocess the scRNA-seq, we initially perform quality control with poor quality cell removal and doublet removal. Next, we remove empty droplets using `DropletUtils` and doublets using `scds` with a threshold of 1.0. Finally, we normalize reads with `linnorm`. Finally, we perform a PCA transform of both the scATAC-seq and scRNA-seq data to compress each to 5 covariates. The pre-processed dataset is available at <https://github.com/blengerich/SnareSeq>.

Results As shown in Figure 6, cell-specific BNs more accurately capture the heterogeneity of gene expression than do cell type-specific BNs, cluster-specific BNs, or a single population BN. This benefit is observed for all cell types, with most benefit for cell types which are less common in the training set or more uniquely differentiated due to their biological function, suggesting that NOTMAD improves performance through both sharing power between samples and capturing heterogeneity where appropriate. In addition, the cell-specific networks cluster according to the cell’s biological function; networks for all subtypes of inhibitory neurons cluster together, while networks for all subtypes of excitatory neurons also loosely cluster together. In contrast, networks for OPC, Oli-Mal, Mic, and Ast cells do not cluster with either the excitatory or inhibitory neurons.



(a) Mean-squared error (MSE) of BNs of single-cell gene expression. We compare four models: (1) a population model which estimates a single BN, (2) a clustered model which estimates a small number of discrete clusters of samples and cluster-specific BNs, (3) cell-type specific networks which a BN for each cell type, and (4) cell-specific networks estimated with NOTMAD. In all cell types, the cell-specific networks estimated by NOTMAD are the most accurate.



(b) Network embeddings, colored by cell type. Each point represents the BN for a single cell, embedded with UMAP. The networks for inhibitory and excitatory neurons tend to cluster according to their corresponding type, with subtypes providing a modest sub-clustering. Astrocytes form a distinct cluster of networks, while Oli-Mal, OPC, and Mic cells form a separate cluster.

(c) Network archetype selection by cell type. Color corresponds to the mean weight of each archetype (interpreted as a probability) in constructing cell-specific networks for cells of each cell type. In addition to cell type-specific network archetypes, we also initialize 7 archetypes to uniformly random sparse networks (“Random”). Excitatory and inhibitory neurons share network archetypes between subclasses, but other cell types use distinct archetypes.

Figure 6: Cell-specific networks reveal cell-specific heterogeneity in adult mouse brain.

6 Conclusions

We have presented NOTMAD, a method to estimate context-specific BNs by learning to transform contextual data into BN parameters. Previously, the structured constraints of BNs prohibited graph generator functions; we show that these structured models can be efficiently modeled as interpolations between a small number of latent archetypes, and the non-convexity of the model domain can be enforced with regularization. As a result, NOTMAD enables simultaneous inference of sample-specific and archetypal BNs without any prior knowledge of BN structure, parameters, or sample relationships. NOTMAD scales to the extreme case of estimating sample-specific BNs, in which each context is observed for only a single statistical sample. BNs are a mainstay of analysis in many fields including biology, medicine, social sciences, and beyond; context-specific BNs enable new questions of heterogeneity of latent structure and provide axes of variation which can be used to describe community composition and influences at a new granularity.

Acknowledgements

B.L. was supported by the Alana Fellowship at MIT. C.E. and E.X. were supported by NIH NIGMS R01GM140467. B.A. was supported by NSF IIS-1956330, NIH R01GM140467, and the Robert H. Topel Faculty Research Fund at the University of Chicago Booth School of Business.

References

- [1] R. Agarwal, L. Melnick, N. Frosst, X. Zhang, B. Lengerich, R. Caruana, and G. E. Hinton. Neural additive models: Interpretable machine learning with neural nets. *arXiv preprint arXiv:2004.13912*, 2021.
- [2] I. I. Ageenko, K. A. Doherty, and A. P. Van Cleave. Personalized lifetime financial planning tool, June 24 2010. US Patent App. 12/316,967.
- [3] A. Ahmed and E. P. Xing. Recovering time-varying networks of dependencies in social and biological studies. *Proceedings of the National Academy of Sciences*, 106(29):11878–11883, 2009.
- [4] M. Al-Shedivat, A. Dubey, and E. Xing. Contextual explanation networks. *Journal of Machine Learning Research*, 21(194):1–44, 2020.
- [5] C. Boutilier, N. Friedman, M. Goldszmidt, and D. Koller. Context-specific independence in bayesian networks. *arXiv preprint arXiv:1302.3562*, 2013.
- [6] F. Buettner, K. N. Natarajan, F. P. Casale, V. Proserpio, A. Scialdone, F. J. Theis, S. A. Teichmann, J. C. Marioni, and O. Stegle. Computational analysis of cell-to-cell heterogeneity in single-cell rna-sequencing data reveals hidden subpopulations of cells. *Nature biotechnology*, 33(2):155, 2015.
- [7] K. L. Buschur, M. Chikina, and P. V. Benos. Causal network perturbations for instance-specific analysis of single cell and disease samples. *bioRxiv*, page 637710, 2019.
- [8] S. Chen, B. B. Lake, and K. Zhang. High-throughput sequencing of the transcriptome and chromatin accessibility in the same cell. *Nature biotechnology*, 37(12):1452–1457, 2019.
- [9] D. M. Chickering. Learning bayesian networks is np-complete. In *Learning from data*, pages 121–130. Springer, 1996.
- [10] 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.
- [11] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [12] A. J. Fisher, J. D. Medaglia, and B. F. Jeronimus. Lack of group-to-individual generalizability is a threat to human subjects research. *Proceedings of the National Academy of Sciences*, 115(27):E6106–E6115, 2018. ISSN 0027-8424. doi: 10.1073/pnas.1711978115. URL <https://www.pnas.org/content/115/27/E6106>.

- [13] S. A. Forbes, D. Beare, P. Gunasekaran, K. Leung, N. Bindal, H. Boutselakis, M. Ding, S. Bamford, C. Cole, and S. Ward. Cosmic: exploring the world’s knowledge of somatic mutations in human cancer. *Nucleic acids research*, 43(D1):D805–D811, 2014.
- [14] N. Friedman, M. Goldszmidt, et al. Discretizing continuous attributes while learning bayesian networks. In *ICML*, pages 157–165, 1996.
- [15] L. Gan, X. Yang, N. Narisetty, and F. Liang. Bayesian joint estimation of multiple graphical models. In *Advances in Neural Information Processing Systems*, pages 9802–9812, 2019.
- [16] J. Guinney, R. Dienstmann, X. Wang, A. De Reynies, A. Schlicker, C. Sonesson, L. Marisa, P. Roepman, G. Nyamundanda, P. Angelino, et al. The consensus molecular subtypes of colorectal cancer. *Nature medicine*, 21(11):1350–1356, 2015.
- [17] S. A. Hart. Precision education initiative: Moving toward personalized education. *Mind, Brain, and Education*, 10(4):209–211, 2016. doi: 10.1111/mbe.12109. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/mbe.12109>.
- [18] T. Hastie and R. Tibshirani. Varying-coefficient models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(4):757–779, 1993.
- [19] F. Jabbari and G. F. Cooper. An instance-specific algorithm for learning the structure of causal bayesian networks containing latent variables. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 433–441. SIAM, 2020.
- [20] F. Jabbari, S. Visweswaran, and G. F. Cooper. Instance-specific bayesian network structure learning. *Proceedings of machine learning research*, 72:169, 2018.
- [21] F. Jabbari, S. Visweswaran, and G. F. Cooper. An empirical investigation of instance-specific causal bayesian network learning. In *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 2582–2585. IEEE, 2019.
- [22] D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [23] M. L. Kuijjer, M. G. Tung, G. Yuan, J. Quackenbush, and K. Glass. Estimating sample-specific regulatory networks. *iScience*, 14:226–240, 2019.
- [24] F. Le Du, B. L. Eckhardt, B. Lim, J. K. Litton, S. Moulder, F. Meric-Bernstam, A. M. Gonzalez-Angulo, and N. T. Ueno. Is the future of personalized therapy in triple-negative breast cancer based on molecular subtype? *Oncotarget*, 6(15):12890, 2015.
- [25] K. Lee and X. Cao. Bayesian joint inference for multiple directed acyclic graphs. *arXiv preprint arXiv:2008.06190*, 2020.
- [26] B. J. Lengerich, B. Aragam, and E. P. Xing. Learning sample-specific models with low-rank personalized regression. In *Advances in Neural Information Processing Systems (In Press)*, 2019.
- [27] N. I. Lindeman, P. T. Cagle, M. B. Beasley, D. A. Chitale, S. Dacic, G. Giaccone, R. B. Jenkins, D. J. Kwiatkowski, J.-S. Saldivar, J. Squire, E. Thunnissen, and M. Ladanyi. Molecular testing guideline for selection of lung cancer patients for EGFR and ALK tyrosine kinase inhibitors: guideline from the College of American Pathologists, International Association for the Study of Lung Cancer, and Association for Molecular Pathology. *Journal of thoracic oncology : official publication of the International Association for the Study of Lung Cancer*, 8(7):823–859, jul 2013. ISSN 1556-1380. doi: 10.1097/JTO.0b013e318290868f. URL <https://pubmed.ncbi.nlm.nih.gov/23552377https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4159960/>.
- [28] T. P. Mourikis, L. Benedetti, E. Foxall, J. Perner, M. Cereda, J. Lagergren, M. Howell, C. Yau, R. Fitzgerald, P. Scaffidi, et al. Patient-specific detection of cancer genes reveals recurrently perturbed processes in esophageal adenocarcinoma. *bioRxiv*, page 321612, 2018.

- [29] K. Ng, J. Sun, J. Hu, and F. Wang. Personalized predictive modeling and risk factor identification using patient similarity. *AMIA Summits on Translational Science Proceedings*, 2015:132, 2015.
- [30] C. J. Oates, J. Q. Smith, S. Mukherjee, and J. Cussens. Exact estimation of multiple directed acyclic graphs. *Statistics and Computing*, 26(4):797–811, 2016.
- [31] A. P. Parikh, W. Wu, R. E. Curtis, and E. P. Xing. Treegl: reverse engineering tree-evolving gene networks underlying developing biological lineages. *Bioinformatics*, 27(13):i196–i204, 2011.
- [32] J. Pensar, H. Nyman, T. Koski, and J. Corander. Labeled directed acyclic graphs: a generalization of context-specific independence in directed graphical models. *Data mining and knowledge discovery*, 29(2): 503–533, 2015.
- [33] V. Relli, M. Trerotola, E. Guerra, and S. Alberti. Abandoning the Notion of Non-Small Cell Lung Cancer. *Trends in Molecular Medicine*, 25(7):585–594, 2019. ISSN 1471-4914. doi: <https://doi.org/10.1016/j.molmed.2019.04.012>. URL <http://www.sciencedirect.com/science/article/pii/S1471491419301017>.
- [34] 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*, 308(5721):523–529, 2005.
- [35] B. Saeed, S. Panigrahi, and C. Uhler. Causal structure discovery from distributions arising from mixtures of dags. In *International Conference on Machine Learning*, pages 8336–8345. PMLR, 2020.
- [36] T. Schaul and J. Schmidhuber. Metalearning. *Scholarpedia*, 5(6):4650, 2010.
- [37] L. Song, M. Kolar, and E. P. Xing. Keller: estimating time-varying interactions between genes. *Bioinformatics*, 25(12):i128–36, 2009. ISSN 1367-4811 (Electronic) 1367-4803 (Linking).
- [38] P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman. *Causation, prediction, and search*. MIT press, 2000.
- [39] N. Stransky, A. M. Egloff, A. D. Tward, A. D. Kostic, K. Cibulskis, A. Sivachenko, G. V. Kryukov, M. S. Lawrence, C. Sougnez, A. McKenna, E. Shefler, A. H. Ramos, P. Stojanov, S. L. Carter, D. Voet, M. L. Cortés, D. Auclair, M. F. Berger, G. Saksena, C. Guiducci, R. C. Onofrio, M. Parkin, M. Romkes, J. L. Weissfeld, R. R. Seethala, L. Wang, C. Rangel-Escareño, J. C. Fernandez-Lopez, A. Hidalgo-Miranda, J. Melendez-Zajgla, W. Winckler, K. Ardlie, S. B. Gabriel, M. Meyerson, E. S. Lander, G. Getz, T. R. Golub, L. A. Garraway, and J. R. Grandis. The mutational landscape of head and neck squamous cell carcinoma. *Science*, 333(6046):1157–1160, Aug. 2011.
- [40] E. V. Strobl. Improved causal discovery from longitudinal data using a mixture of dags. In *The 2019 ACM SIGKDD Workshop on Causal Discovery*, pages 100–133. PMLR, 2019.
- [41] B. Thiesson, C. Meek, D. M. Chickering, and D. Heckerman. Learning mixtures of dag models. *arXiv preprint arXiv:1301.7415*, 2013.
- [42] O. Ursu, J. T. Neal, E. Shea, P. I. Thakore, L. Jerby-Arnon, L. Nguyen, D. Dionne, C. Diaz, J. Bauman, M. Mosaad, et al. Massively parallel phenotyping of variant impact in cancer with perturb-seq reveals a shift in the spectrum of cell states induced by somatic mutations. *bioRxiv*, 2020.
- [43] R. K. Vinayak, W. Kong, G. Valiant, and S. Kakade. Maximum likelihood estimation for learning populations of parameters. In *International Conference on Machine Learning*, pages 6448–6457, 2019.
- [44] L. Wiechmann, M. Sampson, M. Stempel, L. M. Jacks, S. M. Patil, T. King, and M. Morrow. Presenting features of breast cancer differ by molecular subtype. *Annals of surgical oncology*, 16(10):2705–2710, 2009.

- [45] B. Zhang, C. Gaiteri, L.-G. Bodea, Z. Wang, J. McElwee, A. A. Podtelezhnikov, C. Zhang, T. Xie, L. Tran, R. Dobrin, E. Fluder, B. Clurman, S. Melquist, M. Narayanan, C. Suver, H. Shah, M. Mahajan, T. Gillis, J. Mysore, M. E. MacDonald, J. R. Lamb, D. A. Bennett, C. Molony, D. J. Stone, V. Gudnason, A. J. Myers, E. E. Schadt, H. Neumann, J. Zhu, and V. Emilsson. Integrated systems approach identifies genetic nodes and networks in late-onset alzheimer’s disease. *Cell*, 153(3):707—720, April 2013. ISSN 0092-8674. doi: 10.1016/j.cell.2013.03.030. URL <http://europepmc.org/articles/PMC3677161>.
- [46] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing. Dags with no tears: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, pages 9472–9483, 2018.