Electronic Journal of Statistics
Vol. 20 (2026) 1549–1566
ISSN: 1935-7524
https://doi.org/10.1214/26-EJS2524

Structure learning and causal effect estimation from unbalanced groups

Stefano Peluso[1] **and Antonietta Mira**[2]

[1]*Department of Statistical Sciences, Università Cattolica del Sacro Cuore and Università della Svizzera italiana, e-mail:* stefano.peluso@unicatt.it

[2]*Faculty of Economics, Università della Svizzera italiana and Università degli Studi dell'Insubria, e-mail:* antonietta.mira@usi.ch

Abstract: Causal implications can be undermined by control and treatment groups that are unbalanced by external confounders. Furthermore, the interventions on the treatment group can target unknown nodes of a Directed Acyclic Graph (DAG), with unknown alterations on the structure of dependencies. We propose a Bayesian methodology on a graph-driven multivariate Gaussian potential outcome model that identifies the unknown target nodes, quantifies the related causal effects and learns the network of dependencies pre- and post-interventions, accounting for observed confounders. For the purpose, we extend the DAG-Wishart prior to a Normal-DAG-Wishart prior in presence of covariates, whose conjugacy and marginal likelihood are derived. We first study the asymptotic properties of the propensity score Bayesian estimator, used to balance control and treatment groups from confounding effects. We then show the posterior ratio consistency of treated and untreated graphs, and the limiting distribution of the Average Treatment Effect estimator, under different asymptotic scenarios in terms of treated and control sample sizes. The theoretical results are validated on simulated data by an appropriately developed MCMC posterior sampler, and then implemented on Acute Myeloid Leukemia malignancies, to evaluate network dependencies, targets and causal effects of Histone Deacetylases inhibitor treatments.

MSC2020 subject classifications: Primary 62F15; secondary 62H22.
Keywords and phrases: Average treatment effect, causal inference, directed acyclic graphs, structure learning.

Received April 2025.

1. Introduction

Network models have been widely employed to understand dependencies between variables, a crucial problem in many scientific areas, especially in biology [18]. Typically, the network structure is inferred under the assumption that multivariate data have been generated by a stable system. More realistically, measurements can be heterogeneous, for instance due to exogenous interventions. This is the case of genomic medicine, where interactions between genes provide insights on the genesis and progression of diseases. In this setting, drug therapies capable of gene-inhibition can regulate and restore dependencies in the gene-network structure. Another example is in neurosciences, where the human brain is studied as a directional network [46]. In this context, many statistical issues remain unaddressed, with the study of dependencies among brain regions often based on simplistic inferential frameworks that involve various methods of crude thresholding and that oversight how activated target areas under stimuli can affect brain region dependencies [52].

Measurements produced after exogenous perturbations are called interventional data, opposed to observational samples. The effect of an intervention can have unknown consequences on the network; for instance, in the context of drug discovery, the effect of drug treatments at gene level can be uncertain or completely unknown [39]. The absence of clear treatment targets paves the way to the development of statistical models on data subject to interventions with unknown effects on the dependency network, that is with unknown interventions. Eaton and Murphy [15] develop a methodology in a categorical-data and Gaussian setting, based on the dynamic programming algorithm of Koivisto and Sood [28], for networks of controlled dimensionality. Zhang et al. [55] identify causal network structures subject to interventions that induce distribution shifts across datasets. Squires et al. [49] suggest a greedy search algorithm for causal network learning with unknown interventions, in the presence of multiple datasets, with one of the datasets bounded to be observational. Peters et al. [36], from prediction invariance considerations, propose a method to estimate causal effects that can be iteratively implemented with the purpose of network learning with uncertain interventions. In [22] a learning algorithm based on multiple tests is proposed that, building on the PC method of Spirtes et al. [47], recovers the causal structure of a network when multiple datasets are subject to interventions on unknown targets. The two control and treated graphs and the differential graph are estimated, respectively, in Yajima et al. [54] and Wang et al. [53], for balanced groups; see also Shojaie [45] for recent review with focus on undirected graphs. Finally, a Bayesian method for learning dependence structures and intervention targets from data subject to interventions on unknown variables is proposed by Castelletti and Peluso [10] in the simplifying setting of vanishing interventions, later extended by Castelletti and Peluso [11] to different asymptotic regimes.

All the outlined literature focuses on the identification of the effects that the treatment induces on the structure of dependencies, but, less often, on learning the target nodes of the intervention. Furthermore, even when targets are identified, the measurement of the direct intervention effect on the target is neglected, that the causal effect is not quantified. A notable exception is the recent Castelletti et al. [9], but limited to categorical data. Also, external confounders have been neglected in this context, leaving all current approaches exposed to treated and control data that are unbalanced, potentially altering any causal conclusions related to the interventions. Finally, networks in control and treated groups are bounded to have a specific relation, such as a treated graph bounded to be sparser than the control graph, or to respect the same topological order of the control graph.

To overcome all the stated limitations of the current literature, we will develop a new statistical model where we complement the potential outcome model [33, 42] with graph-driven dependences, to jointly estimate causal effects and targets of an intervention, and the pre- and post-intervention structures of dependencies, robustly to the presence of observed confounders. To develop the framework, we extend the DAG-Wishart prior [3, 6, 34] to a Normal-DAG-Wishart prior in presence of covariates, and we recover its conjugacy property and the closed-form marginal likelihood. We study the asymptotic properties of the propensity score Bayesian estimator, to balance the exposure of the two groups to the confounders. Then, we show that treated and untreated graphs can be consistently identified up to Markov equivalence classes, a property known as posterior ratio consistency. Furthermore, we recover the limiting distribution of the Average Treatment Effect estimator, under various regimes related to different asymptotic behaviors of sample sizes in treatment and control groups.

Simulated experiments and an appropriately built MCMC scheme validate the ability of the developed methodology in recovering the correct graphical structures of the intervened and control groups, propensity scores, causal effects, targets and additional model parameters, in comparison with recently proposed alternatives in the literature. We finally implement the methodology to learn networks, causal effects and targets of histone deacetylases (HDAC) inhibitors on acute myeloid leukemia (AML) malignancies.The model is outlined in Sections 2.1 (likelihood) and 2.2 (prior), whilst the theoretical properties are studied in Section 3. The sampling algorithm is applied on simulated (Section 4.1) and AML (Section 4.2) data. We conclude in Section 5 with a summary of the contributions and further directions of investigation. Proofs of propositions and details of the posterior sampler are provided in Supplemetary Material [35].

2. Model development

2.1. A potential outcome model with graph-driven dependences

Let the observations related to a statistical unit i be $(\boldsymbol{y}_i, \boldsymbol{x}_i, z_i)$, for $i = 1, \ldots, n$, where $\boldsymbol{y}_i = (y_{i1}, \ldots, y_{iq})^\top \in \mathcal{Y} \subseteq \mathbb{R}^q$ is a realization of a q-dimensional random vector $\boldsymbol{y}$ dependent on the vector of regressors $\boldsymbol{x}$, observed through the vector $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{ip})^\top \in \mathcal{X} \subseteq \mathbb{R}^p$, and with $z_i \in \{0, 1\}$ a label denoting absence or presence of an intervention, or a treatment, on the subject. We are interested in learning the dependence structure among the components of $\boldsymbol{y}$, accounting for the dependence on external potentially confounding features. An example is the one of a patient genetic profile, subject or not to some specific pharmacological treatment, showing different characteristics. We collect observations on all units in the $n \times q$ matrix $\boldsymbol{Y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)^\top$, the $n \times p$ matrix $\boldsymbol{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)^\top$ and in the n-dimensional binary vector $\boldsymbol{z} = (z_1, \ldots, z_n)$. The vector $\boldsymbol{z}$ partitions the n units into a treatment group $\mathcal{N}_1 = \{1 \leq i \leq n : \ z_i = 1\}$ of cardinality n_1 and a control group $\mathcal{N}_0 = \{1 \leq i \leq n : \ z_i = 0\}$ of cardinality $n_0 = n - n_1$. We will denote with $\boldsymbol{Y}_1$ the $n_1 \times q$ matrix of observations in the treatment group, and $\boldsymbol{Y}_0$ is the $n_0 \times q$ matrix of observations in the control group, with similar notation for $\boldsymbol{X}_1$ and $\boldsymbol{X}_0$.

Following the Neyman-Rubin potential outcome model [33, 42], we write $\boldsymbol{y}_i$ as a mixture of two components, $\boldsymbol{y}_i = z_i \boldsymbol{y}_i(1) + (1 - z_i)\boldsymbol{y}_i(0)$, where $\boldsymbol{y}_i(1)$ and $\boldsymbol{y}_i(0)$ are the potential outcomes for unit i, if respectively subject or not to the intervention: we only observe $\boldsymbol{y}_i(z_i)$, whilst $\boldsymbol{y}_i(1 - z_i)$ is the never observed counterfactual outcome. To allow the structure of dependences in $\boldsymbol{y}$, as well as the dependence between $\boldsymbol{y}$ and $\boldsymbol{x}$, to potentially change if under treatment or not, we impose, for $k = 0, 1$, that $\boldsymbol{y}_i(k) | \boldsymbol{x}, \boldsymbol{B}_k, \boldsymbol{\Omega}_k \sim N_q\left(\boldsymbol{x}^\top \boldsymbol{B}_k, \boldsymbol{\Omega}_k^{-1}\right)$, a q-variate Gaussian distribution, where $\boldsymbol{B}_k$ is a $p \times q$ matrix of regressor coefficients and $\boldsymbol{\Omega}_k$ is the precision $q \times q$ matrix, Markov to a group-specific Directed Acyclic Graph (DAG) $\mathcal{D}_k$. Then $\boldsymbol{y}(k)$ obeys the directed Markov property with respect to $\mathcal{D}_k$, i.e. the conditional independences among coordinates of $\boldsymbol{y}(k)$ are encoded by the parent-child relationships in $\mathcal{D}_k$ [30]. The probability of subject i being treated, given the set of subject covariates, is modelled as $z | \boldsymbol{x}, \beta \sim Ber(e(\boldsymbol{x}, \boldsymbol{\beta}))$, for some *propensity score* $e(\boldsymbol{x}, \boldsymbol{\beta}) := \Pr(z = 1 | \boldsymbol{x}, \boldsymbol{\beta})$ dependent on $\boldsymbol{x}$ and some parameter $\boldsymbol{\beta}$ that depends on the specific choice of e. Therefore, over all the n measurements, $\boldsymbol{Y}_k | \boldsymbol{X}, \boldsymbol{B}_k, \boldsymbol{\Omega}_k, \boldsymbol{z} \sim N_{n_k, q}\left(\boldsymbol{X}_k \boldsymbol{B}_k, \boldsymbol{I}_{n_k}, \boldsymbol{\Omega}_k^{-1}\right)$, for $k = 0, 1$, a matrix Gaussian distribution with

density proportional to $|\boldsymbol{\Omega}_k|^{n_k/2} \exp\left\{-\frac{1}{2}\text{tr}\left[(\boldsymbol{Y}_k - \boldsymbol{X}_k\boldsymbol{B}_k)\,\boldsymbol{\Omega}_k\,(\boldsymbol{Y}_k - \boldsymbol{X}_k\boldsymbol{B}_k)^\top\right]\right\}$, and $\boldsymbol{z}|\boldsymbol{X},\boldsymbol{\beta} \sim \prod_{i=1}^n e(\boldsymbol{x}_i,\boldsymbol{\beta})^{z_i}\,[1-e(\boldsymbol{x}_i,\boldsymbol{\beta})]^{1-z_i}$.

In the graph $\mathcal{D}_k$, a node i is a *parent* of a node j if the arrow $i \to j$ is in $\mathcal{D}_k$ or, in other terms, if $(i,j) \in E(\mathcal{D}_k)$, the edge set of $\mathcal{D}_k$. All parents of j in $\mathcal{D}_k$ form the *parent set* $\text{pa}_j(\mathcal{D}_k)$, of cardinality $|\text{pa}_j(\mathcal{D}_k)| =: \nu_j(\mathcal{D}_k)$, and the *family set* is $\text{fa}_j(\mathcal{D}) = \text{pa}_j(\mathcal{D}) \cup j$. In the sequel, it is assumed that there are no unobserved confounders other than $\boldsymbol{x}$ and the variables in the graphs. The graph $\mathcal{D}_k$ is equivalently represented by an adjacency binary matrix $\boldsymbol{A}_k$, having the (i,j) element equal to one if the edge $i \to j \in \mathcal{D}_k$ and zero otherwise. We measure the *structural effect* of the intervention by $\boldsymbol{A}_1 - \boldsymbol{A}_0$, to understand those relationships that are added or deleted following an intervention. See for instance Hauser and Bühlmann [19] for structural effects confined to post-intervention edge deletions. On the other hand, the *causal effect* of the intervention is measured by the Average Treatment Effect (ATE), see Imbens [24], defined as $\boldsymbol{ATE} := \mathbb{E}\,[\boldsymbol{y}(1) - \boldsymbol{y}(0)]$ that, in the model above, assume the form of $\mathbb{E}\,[\boldsymbol{x}^\top(\boldsymbol{B}_1 - \boldsymbol{B}_0)]$.

Note that two DAGs are called Markov equivalent if they encode the same set of conditional independencies. The set of all DAGs that are Markov equivalent to a given DAG $\mathcal{D}$ defines its Markov equivalence class [2], uniquely represented by a partially directed graph named the essential graph (EG).

2.2. Prior structure

The positive definite precision matrix $\boldsymbol{\Omega}_k$, $k = 0, 1$, has a unique *modified Choleski decomposition* $\boldsymbol{\Omega}_k = \boldsymbol{L}_k\boldsymbol{D}_k^{-1}\boldsymbol{L}_k^\top$, where $\boldsymbol{L}_k = \left(\boldsymbol{L}_{k,ij}\right)_{ij}$, for $i, j = 1, \ldots, q$, is a lower-triangular $q \times q$ matrix with $\boldsymbol{L}_{k,ii} = 1$, and $\boldsymbol{D}_k$ is a diagonal $q \times q$ matrix with positive entries on the main diagonal; see e.g. Pourahmadi [38]. There is a relationship between the Markov property of the distribution of $\boldsymbol{y}(k)$ and the structure of $\boldsymbol{L}_k$: $\boldsymbol{L}_{k,ij} = 0$ for $j < i$ if and only if $i \notin pa_j(\mathcal{D}_k)$, the parent set of node j in the graph $\mathcal{D}_k$. The parameter $(\boldsymbol{D}_k, \boldsymbol{L}_k)$ is known as the *Choleski parameter*: $(\boldsymbol{D}_k, \boldsymbol{L}_k) \in \mathcal{D}_+^q \times \mathcal{L}_\mathcal{D}$, where $\mathcal{D}_+^q$ is the set of $q \times q$ diagonal matrices with positive entries, and $\mathcal{L}_\mathcal{D}$ is the set of lower triangular matrices with unit main diagonal entries and with off-diagonal terms coherent with the parent ordering of $\mathcal{D}$. In absence of covariates and without distinctions between control and treatment groups, the *DAG-Wishart distribution* [3, 10] has been proposed as a prior distribution with nice theoretical properties [6, 34] for the Choleski parameter. In our setting, with a control and treatment groups exposed to potentially unbalanced covariates $\boldsymbol{x}$, we extend this prior to a *Normal-DAG-Wishart* prior on $(\boldsymbol{B}_k, \boldsymbol{D}_k, \boldsymbol{L}_k)$. The proposed prior will be the starting point for a covariate-adjusted learning of the conditional independence structures of $\mathcal{D}_1$ and $\mathcal{D}_2$, as developed below, along the lines of Consonni et al. [13].

Definition 2.1 (Normal-DAG-Wishart prior)**.** For $\boldsymbol{B} \in \mathbb{R}^{p\times q}$, $\boldsymbol{D} \in \mathcal{D}_+^q$, and $\boldsymbol{L} \in \mathcal{L}_\mathcal{D}$, we say that $(\boldsymbol{B}, \boldsymbol{D}, \boldsymbol{L}) \sim NW_\mathcal{D}(\underline{\boldsymbol{B}}, \boldsymbol{C}, \boldsymbol{U}, \boldsymbol{a}(\mathcal{D}))$, a Normal-DAG-Wishart distribution, with hyperparameters $\underline{\boldsymbol{B}} \in \mathbb{R}^{p\times q}$, $\boldsymbol{C} \in \mathbb{R}_+^{p\times p}$, $\boldsymbol{U} \in \mathbb{R}_+^{q\times q}$ and $\boldsymbol{a}(\mathcal{D}_k) = (a_1(\mathcal{D}_k), \ldots, a_q(\mathcal{D}_k)) \in \mathbb{R}^q$ if its density function is

$$\frac{\exp\left\{-\frac{1}{2}\text{tr}\left[\left(\boldsymbol{L}\boldsymbol{D}^{-1}\boldsymbol{L}^\top\right)\left(\boldsymbol{U} + (\boldsymbol{B} - \underline{\boldsymbol{B}})^\top\boldsymbol{C}(\boldsymbol{B} - \underline{\boldsymbol{B}})\right)\right]\right\} \prod_{i=1}^q \boldsymbol{D}_{ii}^{-\frac{a_i(\mathcal{D})+p}{2}}}{\mathcal{Z}_\mathcal{D}(\boldsymbol{C}, \boldsymbol{U}, \boldsymbol{a}(\mathcal{D}))}, \tag{1}$$

where, with $c_i(\mathcal{D}) := a_i(\mathcal{D}) - v_i(\mathcal{D})$, the normalizing constant is

$$\mathcal{Z}_{\mathcal{D}}(\boldsymbol{C}, \boldsymbol{U}, \boldsymbol{a}(\mathcal{D})) = \frac{(2\pi)^{pq/2}}{|\boldsymbol{C}|^{q/2}} \prod_{i=1}^{q} \Gamma\left(\frac{c_i(\mathcal{D})}{2} - 1\right) 2^{\frac{a_i(\mathcal{D})}{2}-1} \left(\sqrt{\pi}\right)^{v_i(\mathcal{D})} \frac{\left|\boldsymbol{U}_{pa_i(\mathcal{D})}\right|^{\frac{c_i(\mathcal{D})-3}{2}}}{\left|\boldsymbol{U}_{fa_i(\mathcal{D})}\right|^{\frac{c_i(\mathcal{D})-2}{2}}}.$$

Following the hierarchical representation of the DAG-Wishart distribution in Ben-David et al. [3], and the properties of the Normal-Wishart distribution, we note that $(\boldsymbol{B}, \boldsymbol{D}, \boldsymbol{L}) \sim NW_{\mathcal{D}}(\underline{\boldsymbol{B}}, \boldsymbol{C}, \boldsymbol{U}, \boldsymbol{a}(\mathcal{D}))$ can be alternatively expressed as

$$p(\boldsymbol{B}, \boldsymbol{D}, \boldsymbol{L}) = p(\boldsymbol{B}|\boldsymbol{D}, \boldsymbol{L}) \prod_{j=1}^{q} p\left(\boldsymbol{L}_{\mathrm{pa}_j(\mathcal{D}),j} \,|\, \boldsymbol{D}_{jj}\right) p\left(\boldsymbol{D}_{jj}\right)$$

in a hierarchical form, with

$$\begin{aligned}
\boldsymbol{B}|\boldsymbol{D}, \boldsymbol{L} &\sim N_{p,q}\left(\underline{\boldsymbol{B}}, \boldsymbol{C}^{-1}, (\boldsymbol{L}\boldsymbol{D}^{-1}\boldsymbol{L}^{\top})^{-1}\right), \\
\boldsymbol{L}_{\mathrm{pa}_j(\mathcal{D}),j} \,|\, \boldsymbol{D}_{jj} &\sim N_{v_j(\mathcal{D})}\left(-\boldsymbol{U}^{-1}_{\mathrm{pa}_j(\mathcal{D})}\boldsymbol{U}_{\mathrm{pa}_j(\mathcal{D}),j}, \boldsymbol{D}_{jj}\boldsymbol{U}^{-1}_{\mathrm{pa}_j(\mathcal{D})}\right), \\
\boldsymbol{D}_{jj} &\sim \text{I-Ga}\left(\frac{c_j(\mathcal{D})}{2} - 1, \frac{1}{2}\boldsymbol{U}_{jj|\mathrm{pa}_j(\mathcal{D})}\right),
\end{aligned}$$

where I-Ga(a, b) stands for an Inverse-Gamma distribution with shape $a > 1$ and rate $b > 0$ having expectation $b/(a-1)$, and where we have denoted, for a generic matrix $\boldsymbol{U}$ and sets A and B, $\boldsymbol{U}_{A,B} := (\boldsymbol{U})_{i\in A, j\in B}$, with $\boldsymbol{U}_A := \boldsymbol{U}_{A,A}$ when $A = B$, and $\boldsymbol{U}_{A|B} := \boldsymbol{U}_A - \boldsymbol{U}_{A,B}\boldsymbol{U}_B^{-1}\boldsymbol{U}_{B,A}$ is the Schur complement of $\boldsymbol{U}_B$ in $\boldsymbol{U}_{A\cup B}$. Hyperparameters $a_j(\mathcal{D})$ are specific to each DAG model, and it can be shown that the default choice (hereinafter adopted) $a_j(\mathcal{D}) = a + 2v_j(\mathcal{D})$ for some $a > 2$ guarantees compatibility among prior distributions for Markov equivalent DAGs; see Peluso and Consonni [34]. In particular, we set $a = 3$, the minimum integer value that guarantees a proper prior distribution, regardless of the specific $\mathcal{D}$. Also, a standard choice, hereinafter adopted, is $\boldsymbol{U} = g\boldsymbol{I}_q$, for some $g > 0$. Therefore, we fix $(\boldsymbol{B}_k, \boldsymbol{D}_k, \boldsymbol{L}_k) \sim NW_{\mathcal{D}_k}(\underline{\boldsymbol{B}}_k, \boldsymbol{C}_k, \boldsymbol{U}_k, \boldsymbol{a}(\mathcal{D}_k))$, for $k = 0, 1$, on the parameters associated to the control and treatment groups.

We impose $p(\mathcal{D}_0, \mathcal{D}_1) = p(\mathcal{D}_1|\mathcal{D}_0)p(\mathcal{D}_0)$ on pre- and post-intervention DAGs, where

$$p(\mathcal{D}_0) \propto \prod_{i<j} \text{Bern}\left((\boldsymbol{A}_0)_{ij} \big| \eta_0\right) = \eta_0^{|E(\mathcal{D}_0)|}(1-\eta_0)^{q(q-1)/2-|E(\mathcal{D}_0)|},$$

for some prior probability $\eta_0 \in (0, 1)$ of edge inclusion in $\mathcal{D}_0$, and where $q(q-1)/2$ is the maximum number of edges in $\mathcal{D}_0$. Following Castelletti et al. [8], such a prior only depends on the number of edges in the graph and can easily incorporate prior knowledge on the network sparsity. In fact, if we expect the intervention will destroy some of the parent-child relationships in $\mathcal{D}_0$ and increase sparsity, by keeping the rest of the network structure unaltered, as in gene-knockout experiments [21, 7], it can be proposed

$$p(\mathcal{D}_1|\mathcal{D}_0) \propto \prod_{(i,j)\in E(\mathcal{D}_0)} \text{Bern}\left((\boldsymbol{A}_1)_{ij} \big| \eta_1\right) = \eta_1^{|E(\mathcal{D}_1)|}(1-\eta_1)^{|E(\mathcal{D}_0)|-|E(\mathcal{D}_1)|},$$

for some $\eta_1 \in (0, 1)$, and with $|E(\mathcal{D}_0)|$ being the maximum number of edges in $\mathcal{D}_1$. On the other hand, if we believe that the intervention will spread its effects over the whole

network structure, with an increase in the density of connections, as in targeted stimuli of brain regions [32], or in diffuse perturbations of proteins' pathways [43], we can opt for $p(\mathcal{D}_1|\mathcal{D}_0) = p(\mathcal{D}_1) \propto \prod_{i<j} \text{Bern}\left((A_1)_{ij}|\eta_1\right)$, for some $\eta_1 > \eta_0$. This last one is the prior we will actually adopt below, since we do not want to impose an increased sparsity in the treated graph. Algorithmic consequences of independent prior choices on $(\mathcal{D}_0, \mathcal{D}_1)$ are discussed in Section 2 of Supplementary material, with the aim of avoiding computational inefficiencies in the exploration of unreliable parameter regions. We stress that, for ease of notation, in the graph priors the dependence of the parent ordering from the given DAG is not explicit, but we do not need to assume a known parent ordering of the true graphs; this assumption would be restrictive in the present context, since comparisons would be limited only to graphs respecting the ordering, and would be excluded between Markov equivalent graphs. We conclude the prior choices with $\boldsymbol{\beta} \sim N_p(\underline{\boldsymbol{\beta}}, \Sigma_{\boldsymbol{\beta}})$ for the propensity score parameter.

3. Theoretical properties: Propensity score, causal effects and network discovery

In the current section, we will study the theoretical properties of the proposed methodology: we start from studying the asymptotic behaviour of the propensity score Bayesian estimator, necessary for designing a balancing of the covariates among treated and untreated groups. Then we prove the conjugacy of our Normal-DAG-Wishart prior to a matrix Gaussian data with covariates. This result then suggests a Bayesian estimator of ATE, whose properties will be studied in different asymptotic regimes, related to different control and treatment sample size dynamics, to show how the correct causal effects and targets are identified. We also derive the closed-form expression for the marginal likelihood of a given network, to be useful below in Section 4 for the development of the MCMC algorithm. Furthermore, we will investigate the asymptotic behaviour of the posterior distribution on the network space, to prove the ability in recovering the true indirect structural effects of the intervention.

In the next proposition, we first focus on the propensity score, and we show that its Bayesian estimator is correctly centered and with a vanishing variance that can be appropriately estimated.

Proposition 1 (Asymptotic distribution of propensity score)**.** *Assume* $0 < e(\boldsymbol{x}_i, \boldsymbol{\beta}) < 1$ *and a finite* $e_{\boldsymbol{\beta}}(\boldsymbol{x}_i, \boldsymbol{\beta})$ *for all* $\boldsymbol{x}_i$ *and* $\boldsymbol{\beta}$, $i = 1, \dots, n$, *where* $e_{\boldsymbol{\beta}} = (e_{\beta_1}, \dots, e_{\beta_1})^\top$, *with* $e_{\beta_k} = \partial e / \partial \beta_k$, $k = 1, \dots, p$. *For a covariance profile* $\boldsymbol{x} \in \mathbb{R}^p$, *we have that*

$$\sqrt{n}\left(e(\boldsymbol{x}, \hat{\boldsymbol{\beta}}) - e(\boldsymbol{x}, \boldsymbol{\beta}^0)\right) \xrightarrow{d} N_1\left(0, e_{\boldsymbol{\beta}}(\boldsymbol{x}, \boldsymbol{\beta}^0)^\top \left(\mathbb{E}\boldsymbol{I}(\boldsymbol{\beta}^0)\right)^{-1} e_{\boldsymbol{\beta}}(\boldsymbol{x}, \boldsymbol{\beta}^0)\right), \tag{2}$$

where $\hat{\boldsymbol{\beta}}$ *is the posterior mean of* $\boldsymbol{\beta}^0$, *and*

$$\mathbb{E}\boldsymbol{I}(\boldsymbol{\beta}^0) = \operatorname*{plim}_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \frac{e_{\boldsymbol{\beta}}(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}}) e_{\boldsymbol{\beta}}(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}})^\top}{e(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}})(1 - e(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}}))}.$$

Proof. See Supplementary material. □

From the above result, we can therefore rely on a $\sqrt{n}-$consistent Bayesian estimator of the propensity score, whatever the covariance profile $\boldsymbol{x}$ of interest, and we can propose

$$\left(\sum_{i=1}^{n} \frac{e_{\boldsymbol{\beta}}(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}}) e_{\boldsymbol{\beta}}(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}})^\top}{e(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}})(1 - e(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}}))}\right)^{-1} e_{\beta_k}(\boldsymbol{x}, \hat{\boldsymbol{\beta}}) e_{\beta_l}(\boldsymbol{x}, \hat{\boldsymbol{\beta}})$$

as an estimate of propensity score variance. Interestingly, when we assume $e(\boldsymbol{x}, \boldsymbol{\beta}) = \Phi(\boldsymbol{x}^\top \boldsymbol{\beta})$, a Probit propensity score, we have $e_{\beta_k}(\boldsymbol{x}^\top \boldsymbol{\beta}) = \phi(\boldsymbol{x}^\top \boldsymbol{\beta})\boldsymbol{x}_k$, and the estimated asymptotic variance reduces to

$$\phi(\boldsymbol{x}^\top \hat{\boldsymbol{\beta}})^2 \boldsymbol{x}^\top \left(\sum_{i=1}^{n} \frac{\phi(\boldsymbol{x}_i^\top \hat{\boldsymbol{\beta}})^2}{\Phi(\boldsymbol{x}_i^\top \hat{\boldsymbol{\beta}})(1 - \Phi(\boldsymbol{x}_i^\top \hat{\boldsymbol{\beta}}))} \boldsymbol{x}_i \boldsymbol{x}_i^\top \right)^{-1} \boldsymbol{x},$$

and when, on the other hand, $e(\boldsymbol{x}, \boldsymbol{\beta}) = (1 + e^{-\boldsymbol{x}^\top \boldsymbol{\beta}})^{-1}$, a Logit propensity score, we have $e_{\beta_k}(\boldsymbol{x}^\top \boldsymbol{\beta}) = e(\boldsymbol{x}^\top \boldsymbol{\beta})(1 - e(\boldsymbol{x}^\top \boldsymbol{\beta}))\boldsymbol{x}_k$, and the estimated asymptotic variance simplifies to

$$e(\boldsymbol{x}, \hat{\boldsymbol{\beta}})^2 (1 - e(\boldsymbol{x}, \hat{\boldsymbol{\beta}}))^2 \boldsymbol{x}^\top \left(\sum_{i=1}^{n} e(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}})(1 - e(\boldsymbol{x}_i, \hat{\boldsymbol{\beta}})) \boldsymbol{x}_i \boldsymbol{x}_i^\top \right)^{-1} \boldsymbol{x}.$$

Not surprisingly, in the two examples of Probit and Logit, the asymptotic variance coincides with the one of the maximum likelihood propensity score estimator. Indeed, Proposition 1, through Theorem 4.1 of Ibragimov and Has' Minskii [23], also affirms the asymptotic equivalence between the maximum likelihood estimator and $\hat{\boldsymbol{\beta}}$, and therefore the asymptotic efficiency of the posterior mean $\hat{\boldsymbol{\beta}}$, but this result is valid for a large class of loss functions.

The above proposition is a theoretical support to a covariate balancing in treated and control groups. With estimated propensity scores at hand, we can design our study in a way that we can proceed *as if* we were in a completely randomized experiment with perfect balance between treated and control groups. For instance, in the application of Section 4.2 we perform propensity score matching [1], but alternatives are available; see for instances Imbens and Rubin [25]. Once an approprieate covariate balancing is reached, we turn our focus on the estimation of ATE and of the graphs; preliminary to this, in the next proposition we prove some useful properties of the proposed Normal-DAG-Wishart prior.

Proposition 2 (Conjugacy of $NW_{\mathcal{D}}$ and marginal likelihood)**.** *For* $\boldsymbol{B} \in \mathbb{R}^{p \times q}$, $\boldsymbol{D} \in \mathcal{D}_+^q$, *and* $\boldsymbol{L} \in \mathcal{L}_{\mathcal{D}}$, *let* $\boldsymbol{B}, \boldsymbol{D}, \boldsymbol{L} \,|\, \mathcal{D} \sim NW_{\mathcal{D}}(\underline{\boldsymbol{B}}, \boldsymbol{C}, \boldsymbol{U}, \boldsymbol{a}(\mathcal{D}))$, *and for* $\boldsymbol{Y} \in \mathbb{R}^{n \times q}$, $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, *let* $\boldsymbol{Y}|\boldsymbol{X}, \boldsymbol{B}, \boldsymbol{D}, \boldsymbol{L} \sim N_{n,q}\left(\boldsymbol{X}\boldsymbol{B}, \boldsymbol{I}_n, (\boldsymbol{L}\boldsymbol{D}\boldsymbol{L}^\top)^{-1}\right)$. *Then*

$$\boldsymbol{B}, \boldsymbol{D}, \boldsymbol{L} \,|\, \boldsymbol{Y}, \boldsymbol{X}, \mathcal{D} \sim NW_{\mathcal{D}}(\hat{\boldsymbol{B}}, \boldsymbol{C} + \boldsymbol{X}^\top \boldsymbol{X}, \hat{\boldsymbol{U}}, \boldsymbol{a}(\mathcal{D}) + n),$$

where $\hat{\boldsymbol{U}} := \boldsymbol{U} + \boldsymbol{Y}^\top \boldsymbol{Y} + \underline{\boldsymbol{B}}^\top \boldsymbol{C} \underline{\boldsymbol{B}} - \left(\boldsymbol{C}\underline{\boldsymbol{B}} + \boldsymbol{X}^\top \boldsymbol{Y}\right)^\top \hat{\boldsymbol{B}}$ *and where*

$$\hat{\boldsymbol{B}} := \left(\boldsymbol{C} + \boldsymbol{X}^\top \boldsymbol{X}\right)^{-1} \left(\boldsymbol{C}\underline{\boldsymbol{B}} + \boldsymbol{X}^\top \boldsymbol{Y}\right).$$

Also, we have $p(\boldsymbol{Y}|\boldsymbol{X}, \mathcal{D}) = (2\pi)^{-nq/2} \mathcal{Z}_{\mathcal{D}}(\boldsymbol{C} + \boldsymbol{X}^\top \boldsymbol{X}, \hat{\boldsymbol{U}}, \boldsymbol{a}(\mathcal{D}) + n) / \mathcal{Z}_{\mathcal{D}}(\boldsymbol{C}, \boldsymbol{U}, \boldsymbol{a}(\mathcal{D}))$.

Proof. See Supplementary material. □

From this proposition, the natural Bayesian estimator of the Average Treatment Effect $\boldsymbol{ATE}$ is $\widehat{\boldsymbol{ATE}} = \mathbf{1}_n^\top \boldsymbol{X} \left(\hat{\boldsymbol{B}}_1 - \hat{\boldsymbol{B}}_0\right)/n$, where, for $k = 0, 1$,

$$\hat{\boldsymbol{B}}_k = \left(\boldsymbol{C}_k + \boldsymbol{X}_k^\top \boldsymbol{X}_k\right)^{-1} \left(\boldsymbol{C}_k \underline{\boldsymbol{B}}_k + \boldsymbol{X}_k^\top \boldsymbol{Y}_k\right)$$

and $\mathbf{1}_n$ is a column vector of ones of dimension n. In the following result we show convergence in probability of $\widehat{\boldsymbol{ATE}}$ to the true value $\boldsymbol{ATE}^0$, and the asymptotic q-variate Gaussian distribution with vanishing covariance matrix. We also denote other true quantities with a zero-superscript, for instance $\boldsymbol{B}_k^0$ and $\boldsymbol{\Omega}_k^0$.

Proposition 3 (Asymptotic distribution of $\widehat{\boldsymbol{ATE}}$)**.** *For $k = 0, 1$, let $\boldsymbol{B}_k, \boldsymbol{D}_k, \boldsymbol{L}_k \mid \mathcal{D}_k \sim NW_{\mathcal{D}_k}(\underline{\boldsymbol{B}}_k, \boldsymbol{C}_k, \boldsymbol{U}_k, \boldsymbol{a}(\mathcal{D}_k))$, with $\boldsymbol{Y}_k|\boldsymbol{X}_k, \boldsymbol{B}_k, \boldsymbol{D}_k, \boldsymbol{L}_k \sim N_{n_k,q}\left(\boldsymbol{X}_k\boldsymbol{B}_k, \boldsymbol{I}_{n_k}, \left(\boldsymbol{L}_k\boldsymbol{D}_k\boldsymbol{L}_k^\top\right)^{-1}\right)$.*

If $\bar{\boldsymbol{X}} \xrightarrow{p} \boldsymbol{\mu} \in \mathbb{R}^p$ and $\operatorname{vec}(\boldsymbol{X}^\top\boldsymbol{X})/n \xrightarrow{p} \operatorname{vec}\boldsymbol{M}$, and $\bar{\boldsymbol{X}}_k \xrightarrow{p} \boldsymbol{\mu}_k$ and $\operatorname{vec}\left(\boldsymbol{X}_k^\top\boldsymbol{X}_k\right)/n \xrightarrow{p} \operatorname{vec}\boldsymbol{M}_k$ for each group, for some positive definite $p \times p$ matrices $\boldsymbol{M}$ and $\boldsymbol{M}_k$, we have

- *if $n_1/n \to \alpha \in (0, 1)$, $\sqrt{n}\left(\widehat{\boldsymbol{ATE}} - \boldsymbol{ATE}^0\right) \xrightarrow{d} N_q\left(\boldsymbol{0}, \boldsymbol{\Gamma}_0 + \boldsymbol{\Gamma}_1 + \frac{\alpha\boldsymbol{\Lambda}_0 + (1-\alpha)\boldsymbol{\Lambda}_1}{\alpha(1-\alpha)}\right)$,*
- *if $n_1/n \to 1$, $\sqrt{n_0}\left(\widehat{\boldsymbol{ATE}} - \boldsymbol{ATE}^0\right) \xrightarrow{d} N_q\left(\boldsymbol{0}, \boldsymbol{\Gamma}_0 + \boldsymbol{\Lambda}_0\right)$,*
- *if $n_1/n \to 0$, $\sqrt{n_1}\left(\widehat{\boldsymbol{ATE}} - \boldsymbol{ATE}^0\right) \xrightarrow{d} N_q\left(\boldsymbol{0}, \boldsymbol{\Gamma}_1 + \boldsymbol{\Lambda}_1\right)$,*

where $\boldsymbol{\Gamma}_k := \operatorname{diag}\left[\left(\boldsymbol{B}_k^0\right)^\top\left(\boldsymbol{M} - \boldsymbol{\mu}\boldsymbol{\mu}^\top\right)\boldsymbol{B}_k^0\right]$ and $\boldsymbol{\Lambda}_k := \boldsymbol{\mu}^\top\boldsymbol{M}_k^{-1}\boldsymbol{\mu}\left(\boldsymbol{\Omega}_k^0\right)^{-1}$.

Proof. See Supplementary material. □

Proposition 3 says that the proposed estimator of ATE is correctly centered and asymptotically q-variate Gaussian when correctly scaled. The convergence depends on the relative asymptotic behaviours of sample sizes in treatment and control groups. It is interesting to note that when the two sample sizes increases in n at the same rate, so that asymptotically there is a fixed size α of measurements in the treatment group and $1 - \alpha$ in the control group, then convergence is at rate $\sqrt{n}$. If on the other hand size unbalances between the two groups are so acute that only one will eventually dominate, convergence rate will deteriorate and be linked to the sample size of the dominated group. Once the causal effects are correctly identified, so are the targets of the interventions, being those nodes for which ATE is different from zero. We finally note that Proposition 3 can also be used to asymptotically test for specific null values of ATE, jointly on the q response variables or on a subset of them, using the consistent estimators suggested in the proposition to estimate all the terms in the asymptotic covariance matrices of $\widehat{\boldsymbol{ATE}}$. Marginally for single scalar response variables, the corresponding marginal asymptotically normal null distribution can be used, whilst for a joint test on more responses, a corresponding quadratic form would provide an asymptotically chi-square null distribution. We refer the reader to Section 4 of Supplementary material for more details on hypothesis test implementations.

In the next result we show the posterior ratio consistency [6] of the Normal-DAG-Wishart prior, to the correct graphical structures of treatment and control groups. Given two graphs $\mathcal{D}^0$ and $\mathcal{D}$, where $\mathcal{D}^0$ is the true one, this property guarantees that the ratio of the posterior distribution of $\mathcal{D}$ and $\mathcal{D}^0$ will converge in probability correctly to 0, for any $\mathcal{D}$ that does not entail the same conditional independence structure of $\mathcal{D}^0$.

Proposition 4 (Posterior ratio graph consistency)**.** *For $k = 0, 1$, let $\mathcal{D}_k^0$ and $\boldsymbol{\Omega}_k^0$ be respectively the true DAGs and precision matrices, and let $\boldsymbol{B}_k, \boldsymbol{D}_k, \boldsymbol{L}_k \mid \mathcal{D}_k \sim NW_{\mathcal{D}_k}(\underline{\boldsymbol{B}}_k, \boldsymbol{C}_k, \boldsymbol{U}_k, \boldsymbol{a}(\mathcal{D}_k))$ and $\boldsymbol{Y}_k|\boldsymbol{X}_k, \boldsymbol{B}_k, \boldsymbol{D}_k, \boldsymbol{L}_k \sim N_{n_k,q}\left(\boldsymbol{X}_k\boldsymbol{B}_k, \boldsymbol{I}_{n_k}, \left(\boldsymbol{L}_k\boldsymbol{D}_k\boldsymbol{L}_k^\top\right)^{-1}\right)$, where $\boldsymbol{U}_k$ and $\boldsymbol{\Omega}_k^0$ are non-singular positive-definite matrices with bounded eigenvalues and $a_j(\mathcal{D}_k) - 2\nu_j(\mathcal{D}_k) > 2$, for $j = 1, \ldots, q$. Assume vanishing prior edge inclusion probabilities $\eta_k = C_k n_k^{-\frac{1}{2+\nu_k}}$, for some $C_k, \nu_k > 0$. As $n_k \to \infty$ for $k = 0, 1$, we have*

$$\max_{(\mathcal{D}_0, \mathcal{D}_1) \notin [\mathcal{D}_0^0] \times [\mathcal{D}_1^0]} \frac{p(\mathcal{D}_0, \mathcal{D}_1|\boldsymbol{Y}, \boldsymbol{X})}{p(\mathcal{D}_0^0, \mathcal{D}_1^0|\boldsymbol{Y}, \boldsymbol{X})} \xrightarrow{p} 0,$$

where $[\mathcal{D}]$ *is the Markov equivalent class of a DAG* $\mathcal{D}$*, whilst for any* $\mathcal{D}_0 \in [\mathcal{D}_0^0]$ *and* $\mathcal{D}_1 \in [\mathcal{D}_1^0]$ *we have* $p(\mathcal{D}_0, \mathcal{D}_1|\boldsymbol{Y}, \boldsymbol{X})/p(\mathcal{D}_0^0, \mathcal{D}_1^0|\boldsymbol{Y}, \boldsymbol{X}) = p(\mathcal{D}_0, \mathcal{D}_1)/p(\mathcal{D}_0^0, \mathcal{D}_1^0)$..

Proof. See Supplementary material. □

Proposition 4 shows that the correct graphical structures will eventually be detected in both groups, when the comparison in made with DAGs outside $[\mathcal{D}_0^0] \times [\mathcal{D}_1^0]$, the equivalence class of the true DAGs. On the contrary, within the true equivalence class the posterior ratio coincides with the prior ratio and we are not able to discriminate between graphs with same dependences: in other words, when equivalent graphs are compared we have compatibility (Peluso and Consonni 34, or score equivalence), in that those graphs will have the same marginal likelihood and then a unit Bayes factor. Therefore identification of graphical structures is possible only up to Markov equivalence; indeed, our model can be rewritten as the Structural Equation Model $\boldsymbol{y}_i(k) - \boldsymbol{x}_i^\top \boldsymbol{B}_k = (\boldsymbol{I}_q - \boldsymbol{L}_k^\top)(\boldsymbol{y}_i(k) - \boldsymbol{x}_i^\top \boldsymbol{B}_k) + \boldsymbol{\epsilon}_i$, where $\boldsymbol{\epsilon}_i \sim N(0, \boldsymbol{D})$, a diagonal potentially heteroskedastic covariance matrix, and therefore specific identifiability conditions, such as homoskedasticity or heteroskedasticity known up to a constant [31, 4] would be needed to guarantee the identification of the single true graphical structure. Finally, we conjecture that a more restricting set of assumptions, more explicitely in line with those in [5], with adaptations required by possibly dependent graph priors, would extend Proposition 4 to the high-dimensional case where the number of nodes increases in n.

4. Numerical examples

4.1. Simulated experiments

In the first set of experiments, we assume known control and treatment graphs, and we evaluate the estimation of the propensity score and of ATE, at different levels of $\boldsymbol{\beta}$. Keeping fixed the generative process of $\boldsymbol{x}$, a different $\boldsymbol{\beta}$ corresponds to a more or less pronounced disparity between groups: for a propensity score increasing in $\boldsymbol{\beta}$, higher positive (lower negative) values $\boldsymbol{\beta}$ under-represent the control (treatment) group. Still, the distribution of $\boldsymbol{x}$ is the same in both groups, so that the preliminary covariance balancing is not necessary in this simulated setting. For $q = 3$, we fix $E(\mathcal{D}_0) = \{(2,1),(3,2)\}$ and $E(\mathcal{D}_1) = \{(2,1)\}$, so that in the intervened group the edge $3 \to 2$ is removed. Given these graphs, similarly to Hauser and Bühlmann [20] and Castelletti and Peluso [11], $\boldsymbol{D}_0 = \boldsymbol{D}_1 = 0.01 \cdot \boldsymbol{I}_q$, where $\boldsymbol{I}_q$ is the q-dimensional identity matrix, whilst non-zero elements of $\boldsymbol{L}_0$ and $\boldsymbol{L}_1$ are uniformly chosen in the interval $[-5, -0.5] \cup [0.5, 5]$. For $p = 1$, we generate $\boldsymbol{x}_i \sim N(1,1)$, for $i = 1, \ldots, n = 5000$, and the unit i is allocated to the treatment group with probability $e(\boldsymbol{x}_i, \boldsymbol{\beta}) = \Phi(\boldsymbol{x}_i^\top \boldsymbol{\beta})$, so that $z_i \sim Bern\left(\Phi(\boldsymbol{x}_i^\top \boldsymbol{\beta})\right)$. Finally, $\boldsymbol{y}_i \sim z_i N_q\left(\boldsymbol{x}_i^\top \boldsymbol{B}_1, \left(\boldsymbol{L}_1 \boldsymbol{D}_1^{-1} \boldsymbol{L}_1\right)^{-1}\right) + (1 - z_i) N_q\left(\boldsymbol{x}_i^\top \boldsymbol{B}_0, \left(\boldsymbol{L}_0 \boldsymbol{D}_0^{-1} \boldsymbol{L}_0\right)^{-1}\right)$, where $\boldsymbol{B}_0 = \boldsymbol{0}$ and $\boldsymbol{B}_1 = \boldsymbol{ATE} = (2,1,3)^\top$. We let $\boldsymbol{\beta}$ vary in $\{-0.9, -0.45, 0, 0.45, 0.9\}$ and for each value we repeat the whole procedure $M = 3000$ times. As a standard choice, we fix location and variance hyperparameter values to, respectively, zero and one: $\underline{\boldsymbol{B}_0} = \underline{\boldsymbol{B}_1} = \boldsymbol{0}_{p \times q}$, $\underline{\boldsymbol{\beta}} = \boldsymbol{0}_{p \times 1}$ and $\boldsymbol{C}_0 = \boldsymbol{C}_1 = \boldsymbol{\Sigma}_{\boldsymbol{\beta}} = \boldsymbol{I}_p$. In Figure 1 we report the results: the top panels show, for each value of $\boldsymbol{\beta}$ of the column, the histogram of $\hat{\boldsymbol{\beta}}$ recovered from the M repetitions, with highlighted as a vertical bar the true unknown $\boldsymbol{\beta}$. Similarly, the bottom panels show the M estimated $e(\boldsymbol{x}, \boldsymbol{\beta})$, for $\boldsymbol{x}$ in the range ∓ 2 standard deviations from its mean. We see that the estimated $\boldsymbol{\beta}$ is correctly centered, and the estimated propensity score functions are coherent to the depicted

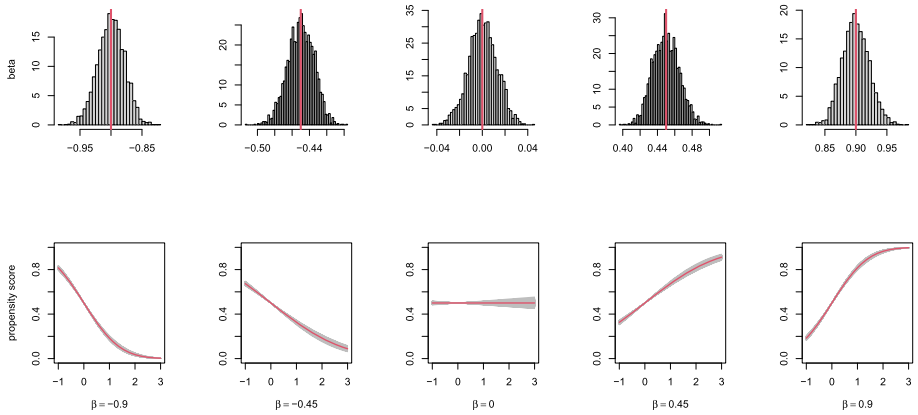

Fig 1. *Estimated slope β and propensity score, over $M = 3000$ simulated samples of size $n = 5000$, for $q = 3$ number of nodes and $p = 1$ covariate. Each column identifies a different value of true slope $\beta \in \{-0.9, -0.45, 0, 0.45, 0.9\}$. First row: distribution of $\hat{\beta}$, with true value as verticle line. Second row: estimated propensity score $e(x, \beta) = \Phi(x\beta)$, based on the average β over the M repetitions (red) and with all β within 2 standard deviations (grey).*

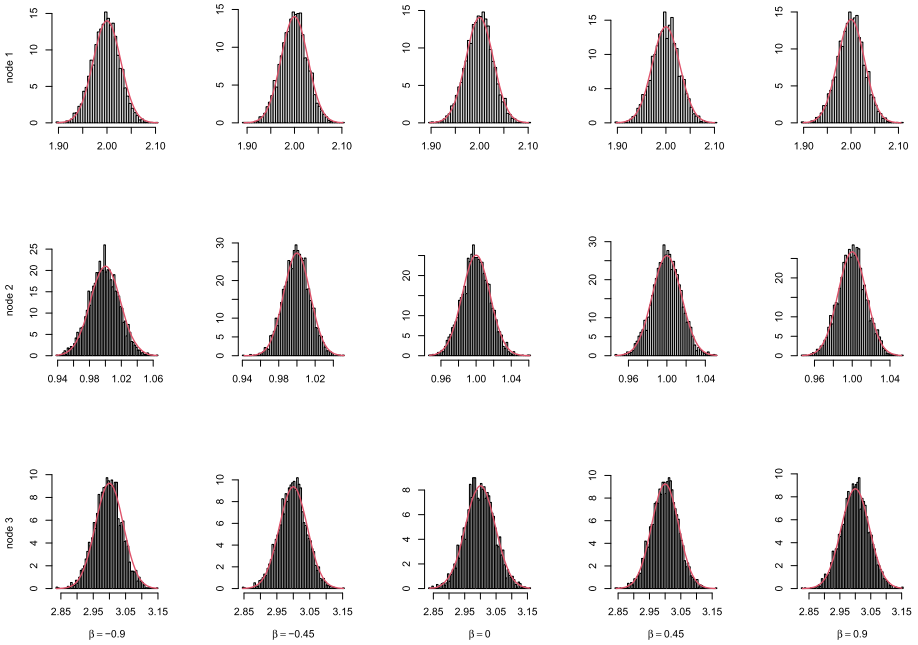

Fig 2. *Histogram of estimated ATE, over $M = 3000$ simulated samples of size $n = 5000$, for $q = 3$ number of nodes and $p = 1$ covariate, for each node (each row) and different values of $\beta \in \{-0.9, -0.45, 0, 0.45, 0.9\}$ (each column), with limiting distribution predicted by Proposition 3 (red curve).*

true unknown propensity score function; the average estimates are not reported, since they completely overlap the true values. Furthermore, in Figure 2 we show the histogram of $\widehat{\boldsymbol{ATE}}$ for node 1 (top), node 2 (middle) and node 3 (bottom), at different values of $\boldsymbol{\beta}$ at each column. The limiting Gaussian densities predicted by Proposition 3 are overimposed and perfectly match the simulation-based distributions.

In a second set of experiments, we relax the assumption of known graphs, now to be

learnt through the MCMC algorithm outlined in Supplementary Material. Furthermore, two q-dimensional $\mathcal{D}_0$ and $\mathcal{D}_1$ are randomly generated by the R function *pcalg::randomDAG*, fixing a probability of edge inclusion equal to $2/q$, as in Castelletti and Peluso [11], for $q \in \{5, 10, 20\}$, and correspondent number of covariates $p \in \{2, 3, 4\}$. The true $\boldsymbol{\beta}$ for the assignment to control or intervened groups is $(0.45, 0.6, 0.2, -0.5)$ whose first p elements are used in each simulated setting. The matrix $\boldsymbol{B}_0 = \mathbf{1}_{p\times q}$ and $\boldsymbol{B}_1$ is the same as $\boldsymbol{B}_0$, with the exception of the elements in its first two columns, randomly generated from a standard uniform, so that only the first two coordinates of $\boldsymbol{y}$ are targets of the intervention, with a non-null ATE corresponding to these coordinates. Everything else related to the generation of datasets and hyperparameter values are fixed as in the previous set of experiments with known graphs. The whole MCMC is run for $S = 5000$ iterations, after a burn-in period of equal length initialized to empty graphs. Finally, everything is repeated $M = 50$ times.

In Figure 3 we measure the performance in graph recovery in terms of Structural Hamming Distance (SHD, Kalisch and Bühlmann 26) between true and estimated EGs, for both groups separately. Our method is named BAYES in the figure. As the final graph estimate, we keep edges that are present in at least $S/2$ iterations, the so-called median probability model mentioned in Section 2 of Supplementary material. We also include the greedy equivalence search (GES) method of Chickering [12], a penalized maximum likelihood method that uses, following Foygel and Drton [17], three different penalization levels: the BIC (GES 1), the extended BIC with tuning coefficient $\gamma = 0.5$ (GES 2) and the extended BIC with tuning coefficient $\gamma = 1$ (GES 3). As further benchmarks, we also include the PC algorithm of [48], at the significance level 0.01 (PC 1), 0.05 (PC 2) and 0.10 (PC 3). In all scenarios for both groups, our Bayesian methodology consistently show the best SHDs, regardless of the tuning parameters adopted in the competing methods. Finally in Figure 2 (Supplementary Material) we report the 95% MCMC-based credible intervals of the estimated Average Treatment Effect, averaged over M simulated repetitions in each scenario (left panels), and the proportion of simulated datasets for which each node is estimated as causally affected by the intervention (right panels), that is those for which the ATE credible interval does not include the zero value: it is clearly identified in all settings which are the true targets.

4.2. Protein networks for acute myeloid leukemia

The acute myeloid leukemia (AML) is a hematological malignancy characterized by uncontrolled proliferation, differentiation arrest, and accumulation of immature myeloid progenitors in peripheral blood and bone marrow [50, 16]. Although clinical advances in AML have been made, especially in young patients, long-term disease-free survival remains poor, and the AML treatment has remained unchanged for the past four decades, still largely based on standard chemotherapy and the use of allogeneic hematopoietic stem cell transplantation [44]. While hypomethylating agents have already been approved in AML, the use of other epigenetic inhibitors, such as histone deacetylases (HDAC) inhibitors (HDACi), is under clinical development. HDACi such as Panobinostat, Vorinostat, and Tricostatin A have been shown to promote cell death, autophagy, apoptosis, or growth arrest in preclinical AML models, yet these inhibitors do not seem to be effective as monotherapies. So far the results with HDACi in clinical trials have been ambiguous, and we will implement our method to recover network dependency structures, causal effects and targets of HDAC-based therapies.

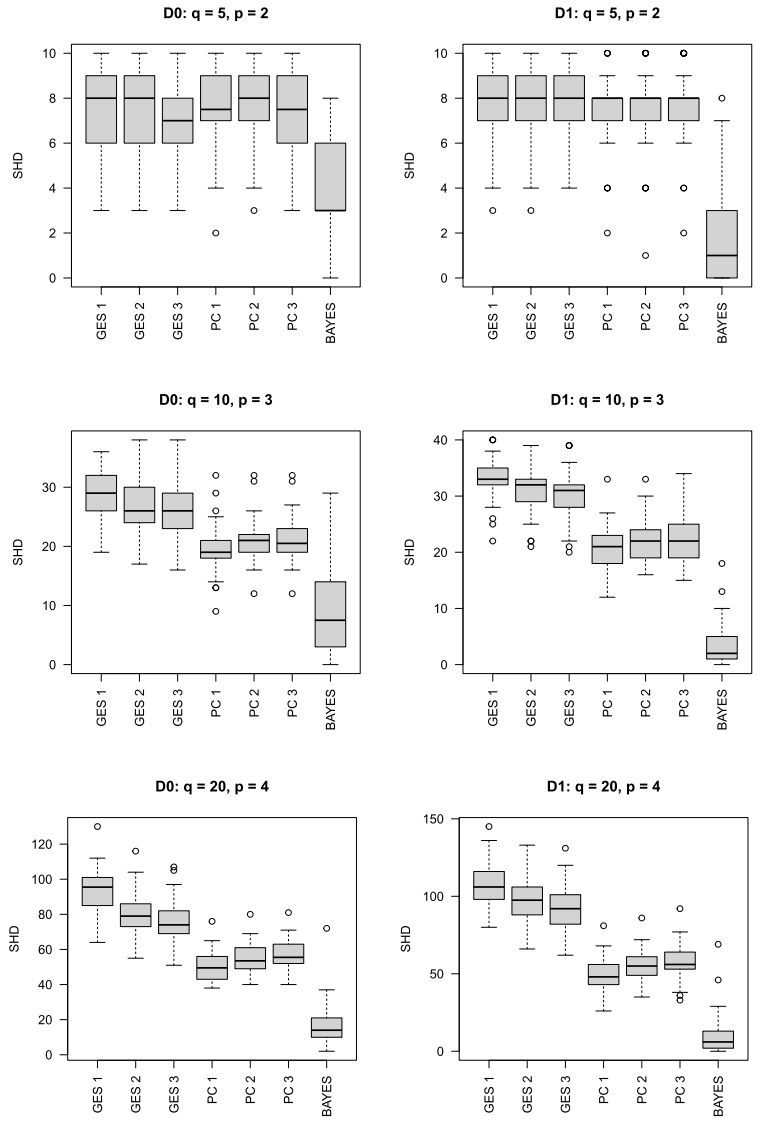

Fig 3. *Structural Hamming Distance (SHD) between estimated and true $\mathcal{I}$-Essential Graph, over $M = 50$ simulated repetitions of size $n = 5000$, for $(q, p) = (5, 2)$ (first row), $(q, p) = (10, 3)$ (second row) and $(q, p) = (20, 4)$ (third row). Methods under comparison, for control (first column) and treatment (second column) groups: the greedy equivalence search method of Chickering [12], using BIC (GES 1), extended BIC with $\gamma = 0.5$ (GES 2), extended BIC with $\gamma = 1$ (GES 3); the PC algorithm of Spirtes et al. [48], at significance level 0.01 (PC 1), 0.05 (PC 2) and 0.10 (PC 3); our method (BAYES).*

The dataset includes protein levels for 213 newly diagnosed AML patients, provided as a supplement to Kornblau et al. [29] and available from the MD Anderson Department of Bioinformatics and Computational Biology at `http://bioinformatics.mdanderson.org/Supplements/Kornblau-AML-RPPA/aml-rppa.xls`. Following Peterson et al. [37] and Castelletti et al. [8], we restrict the analysis to $n = 178$ subjects with AML subtype M0, M1, M2 or M4, according to the French-American-British (FAB) classification system, and to $q = 18$ proteins that are known to be involved in apoptosis and cell

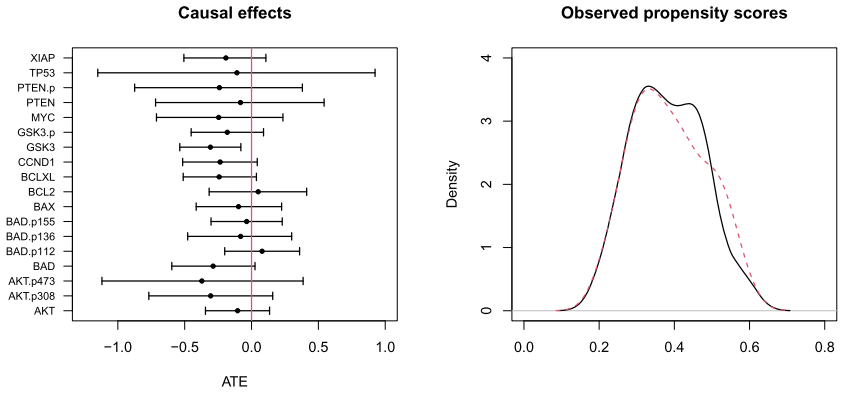

Fig 4. *Left panel: 95% MCMC-based credible intervals of the estimated Average Treatment Effect, for each gene in the AML data. Right panel: kernel density estimates of propensity scores for patients subject (red dashed line) or not (black solid line) to histone deacetylases therapies.*

cycle regulation [27]. The AML subtype is one of the regressors, together with age, sex and prior existence of a malignancy, then $p = 6$, with three binary variables coding the four categories of the FAB classification. By their nature, these features can't be modelled as Gaussian random variables, and therefore can't be included as additional nodes in the network of $\boldsymbol{y}$. The intervention vector is such that $z_i = 1$ if patient i receives HDAC inhibitors ($n_1 = 64$), and $z_i = 0$ otherwise ($n_0 = 114$). In Figure 3 (Supplementary Material) we show the probit propensity scores, based on the estimated $\hat{\boldsymbol{\beta}} = (-0.0058, -0.1512, -0.2281, -0.2429, 0.3593, 0.0248)$, as a function of age, for females (first row) and males (second row), in absence (first column) or presence (second column) of prior malignancy, for the four FAB categories M0, M1, M2 and M4. We then balance the estimated propensity scores in the two groups by matching each treated unit with its closest control unit, reducing the sample to $n_0 = n_1 = 64$ observations. Indeed, in the right panel of Figure 4, we report, separately for items subject or not to HDAC therapies, the densities of the estimated propensity scores after the matching: they clearly overlap, with no evidence of bias in the covariance profiles of treated and untreated subjects. As a further confirmation, not reported for brevity, on the average age, the proportion of males, the proportion of units with prior malignancy, and the AML subtype proportion, no test on the mean or proportion difference reveal any significant difference between the HDAC and no-HDAC groups, suggesting that groups are balanced on average, and that the significant average treatment effects genuinely stem from the treatment itself.

We run the MCMC sampler of Section 2 in Supplementary material for $M = 50000$ iterations, whose first $M/2$ discarded, initialized at the empty control and treatment graphs, and all the hyperprior specifications are fixed as in the simulated studies. In Figure 4 we report the 95% MCMC-based credible intervals of the estimated ATE, for each studied gene: we see a clear significant causal effect for GSK3, whilst BAD, CCND1 and BCLXL are close to significance. In Figure 5 we show the two estimated EGs, without ($\mathcal{D}_0$) and with ($\mathcal{D}_1$) HDAC treatments: the graphical structure is subject to clear changes: 33 directions are removed from $\mathcal{D}_0$ to $\mathcal{D}_1$, but other 29 directions are added, suggesting that the structural effects of HDAC therapies are not in line with those predicted by hard interventions. Indeed, the (interventional) EG estimated assuming hard unknown interventions [10] that destroy parental relationships

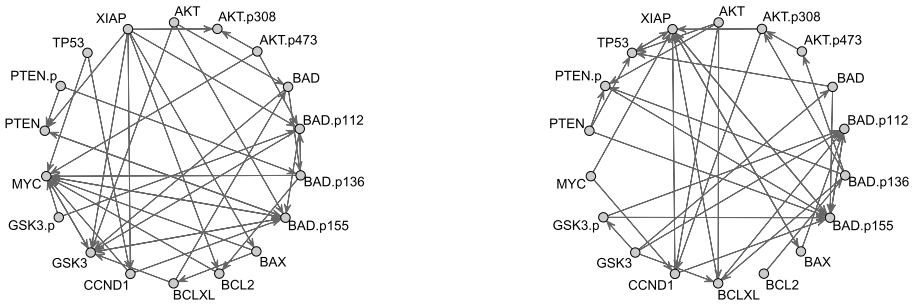

Fig 5. *Estimated Essential Graphs, without histone deacetylases treatment ($\mathcal{D}_0$, left) and with histone deacetylases treatment ($\mathcal{D}_1$, right).*

(not reported) is very close to the reported $\mathcal{D}_0$ and it is therefore not able to detect any change in the dependency structure stemming from HDAC therapies, further suggesting that the assumption of hard interventions would be restrictive in this context.

5. Conclusions and further developments

We have developed a Bayesian statistical methodology to learn pre- and post-intervention networks, causal effects and treatment targets, from data partly subject to external unknown interventions and to observed confounders. A multivariate potential outcome model is complemented with a newly proposed Normal-DAG-Wishart prior, whose theoretical properties are studied. We first derive the asymptotic properties of the propensity score Bayesian estimator, on the basis of which we balance the two treated and control groups. We then prove the ability of the model to correctly identify treated and untreated networks, and the asymptotic distribution of the Average Treatment Effect estimator, under different dynamics of the two groups. An implementation on Acute Myeloid Leukemia data sheds light on structural and causal effects of histone deacetylases inhibitors on a network of proteins known to be involved in apoptosis and cell cycle regulation.

Other design choices may avoid the subsampling implicit in the propensity matching. For instance, weighted estimators of the causal treatment effect can be studied, through an inverse probability weighting (IPW) ATE estimator $\boldsymbol{e}^{-1}\boldsymbol{X}_1\hat{\boldsymbol{B}}_1/n - (1-\boldsymbol{e})^{-1}\boldsymbol{X}_0\hat{\boldsymbol{B}}_0/n$, where $\boldsymbol{e}^{-1} = (e(\boldsymbol{x}_1,\hat{\boldsymbol{\beta}})^{-1},\dots,e(\boldsymbol{x}_n,\hat{\boldsymbol{\beta}})^{-1})^\top$, or a doubly robust Augmented IPW [40] alternative. Also, ATE estimators computed on propensity score-based stratifications can be analyzed along the lines of Proposition 3, for which we conjecture an asymptotic variance with explicit within- and between-strata variability terms. There may also be interest in alternative measures that deviate from ATE, as the Average Treatment effect on the Treated (ATT), see Imbens [24], defined as $\boldsymbol{ATT} := \mathbb{E}\left[\boldsymbol{y}(1)-\boldsymbol{y}(0)|\, z=1\right]$, that would be $\mathbb{E}\left[\boldsymbol{x}^\top(\boldsymbol{B}_1-\boldsymbol{B}_0)\,|z=1\right] = \mathbb{E}\left[\boldsymbol{x}^\top(\boldsymbol{B}_1-\boldsymbol{B}_0)\,e(\boldsymbol{x},\boldsymbol{\beta})\right]/\mathbb{E}\left[e(\boldsymbol{x},\boldsymbol{\beta})\right]$.

Furthermore, the assumed latent unconfoundedness is a limitation that can be violated for various reasons. An extension in this direction could be performed on both causal effect estimation and structure learning: for the former, the marginal sensitivity model of Tan [51] imposes a maximum threshold on the importance of a latent confounder, and could lead to the replacement of point estimated propensity scores and ATE with related sharp bounds [56, 14];

for the latter, the presence of hidden latent variables that are parents of the observed nodes would require to extend the currently proposed Normal-DAG-Wishart prior from a diagonal matrix $\boldsymbol{D}$ in the modified Cholesky decomposition of the precision matrix, to a non-diagonal matrix $\boldsymbol{D}$, since this extension would be equivalent to introduce non-null covariance terms in the errors of a Structural Equation Model [41]. The marginal likelihood of the new prior could then contribute to posterior ratios of graphs, whose asymptotic behaviour could reveal, along the lines of Proposition 4, the ability of consistently identifying networks in presence of missing confounders.

Funding

The authors acknowledge support from the Swiss National Science Foundation (SNF), Grant No. 200021_200557.

Supplementary Material

Structure learning and causal effect estimation from unbalanced groups: Supplementary Material
(doi: 10.1214/26-EJS2524SUPP; .pdf). It contains proofs of propositions, the algorithm for posterior sampling, additional details on the leukemia study, a second application on cardiac arrests in Ticino, and an additional implementation of the proposed methodology as hypothesis test on the multivariate average treatment effect.

References

[1] Abadie, A. and Imbens, G. W. (2006). Large sample properties of matching estimators for average treatment effects. *econometrica*, 74(1):235–267. MR2194325

[2] Andersson, S. A., Madigan, D., and Perlman, M. D. (1997). A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25(2):505–541. MR1439312

[3] Ben-David, E., Li, T., Massam, H., and Rajaratnam, B. (2015). High dimensional Bayesian inference for Gaussian directed acyclic graph models. *arXiv preprint arXiv: 1109.4371*.

[4] Berrevoets, J., Raymaekers, J., Van der Schaar, M., Verdonck, T., and Yao, R. (2025). Differentiable causal structure learning with identifiability by notime. In *International Conference on Artificial Intelligence and Statistics*, pages 3115–3123. PMLR.

[5] Cao, X., Khare, K., and Ghosh, M. (2019a). Posterior graph selection and estimation consistency for high-dimensional Bayesian DAG models. *The Annals of Statistics*, 47(1):319–348. MR3909935

[6] Cao, X., Khare, K., Ghosh, M., et al. (2019b). Posterior graph selection and estimation consistency for high-dimensional Bayesian DAG models. *The Annals of Statistics*, 47(1):319–348. MR3909935

[7] Castelletti, F. and Consonni, G. (2019). Objective Bayes model selection of Gaussian interventional essential graphs for the identification of signaling pathways. *The Annals of Applied Statistics*, 13(4):2289–2311. MR4037431

[8] Castelletti, F., Consonni, G., Della Vedova, M., and Peluso, S. (2018). Learning Markov equivalence classes of directed acyclic graphs: an objective Bayes approach. *Bayesian Analysis*, 13:1231–1256. MR3855370

[9] Castelletti, F., Consonni, G., and Della Vedova, M. L. (2023). Joint structure learning and causal effect estimation for categorical graphical models. *arXiv preprint arXiv:2306.16068*. MR4890873

[10] Castelletti, F. and Peluso, S. (2022). Network structure learning under uncertain interventions. *Journal of the American Statistical Association*, pages 1–12. MR4646630

[11] Castelletti, F. and Peluso, S. (2024). Bayesian learning of network structures from interventional experimental data. *Biometrika*, 111(1):195–214. MR4704565

[12] Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554. MR1991085

[13] Consonni, G., La Rocca, L., and Peluso, S. (2017). Objective Bayes covariate-adjusted sparse graphical model selection. *Scandinavian Journal of Statistics*, (3):741–764. MR3687971

[14] Dorn, J. and Guo, K. (2023). Sharp sensitivity analysis for inverse propensity weighting via quantile balancing. *Journal of the American Statistical Association*, 118(544):2645–2657. MR4681610

[15] Eaton, D. and Murphy, K. (2007). Bayesian structure learning using dynamic programming and MCMC. In *Proceedings of the Twenty-Third Annual Conference on Uncertainty in Artificial Intelligence (UAI-07)*, pages 101–108. AUAI Press.

[16] Estey, E. and Döhner, H. (2006). Acute myeloid leukaemia. *The Lancet*, 368(9550):1894–1907.

[17] Foygel, R. and Drton, M. (2010). Extended Bayesian information criteria for Gaussian graphical models. In *Advances in Neural Information Processing Systems 23*, pages 2020–2028.

[18] Friedman, N. (2004). Inferring cellular networks using probabilistic graphical models. *Science*, 303(5659):799–805.

[19] Hauser, A. and Bühlmann, P. (2012). Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13(79):2409–2464. MR2973606

[20] Hauser, A. and Bühlmann, P. (2012). Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464. MR2973606

[21] Hauser, A. and Bühlmann, P. (2015). Jointly interventional and observational data: estimation of interventional Markov equivalence classes of directed acyclic graphs. *Journal of the Royal Statistical Society. Series B (Methodology)*, 77(1):291–318. MR3299409

[22] He, Y. and Geng, Z. (2016). Causal network learning from multiple interventions of unknown manipulated targets. *arXiv preprint arXiv:1610.08611*.

[23] Ibragimov, I. A. and Has' Minskii, R. Z. (2013). *Statistical estimation: asymptotic theory*, volume 16. Springer Science & Business Media.

[24] Imbens, G. W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics*, 86(1):4–29.

[25] Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge university press. MR3309951

[26] Kalisch, M. and Bühlmann, P. (2007). Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8:613–36.

[27] Kanehisa, M., Goto, S., Sato, Y., Furumichi, M., and Tanabe, M. (2012). Kegg for integration and interpretation of large-scale molecular data sets. *Nucleic acids research*, 40(D1):D109–D114.

[28] Koivisto, M. and Sood, K. (2004). Exact Bayesian structure discovery in Bayesian networks. *Journal of Machine Learning Research*, 5:549–573. MR2247991

[29] Kornblau, S. M., Tibes, R., Qiu, Y. H., Chen, W., Kantarjian, H. M., Andreeff, M., Coombes, K. R., and Mills, G. B. (2009). Functional proteomic profiling of aml predicts response and survival. *Blood, The Journal of the American Society of Hematology*, 113(1):154–164.

[30] Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press. MR1419991

[31] Loh, P.-L. and Bühlmann, P. (2014). High-dimensional learning of linear causal networks via inverse covariance estimation. *Journal of Machine Learning Research*, 15(140):3065–3105. MR3277162

[32] Mohan, U. R., Watrous, A. J., Miller, J. F., Lega, B. C., Sperling, M. R., Worrell, G. A., Gross, R. E., Zaghloul, K. A., Jobst, B. C., Davis, K. A., et al. (2020). The effects of direct brain stimulation in humans depend on frequency, amplitude, and white-matter proximity. *Brain stimulation*, 13(5):1183–1195.

[33] Neyman, J. S. (1923). On the application of probability theory to agricultural experiments. essay on principles. section 9.(tlanslated and edited by dm dabrowska and tp speed, statistical science (1990), 5, 465-480). *Annals of Agricultural Sciences*, 10:1–51. MR1092986

[34] Peluso, S. and Consonni, G. (2020). Compatible priors for model selection of high-dimensional Gaussian DAGs. *Electronic Journal of Statistics*, 14(2):4110–4132. MR4170698

[35] Peluso, S. and Mira, A. (2026). Supplement to "Structure learning and causal effect estimation from unbalanced groups". DOI: 10.1214/26-EJS2524SUPP

[36] Peters, J., Bühlmann, P., and Meinshausen, N. (2016). Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 78(5):947–1012. MR3557186

[37] Peterson, C., Stingo, F. C., and Vannucci, M. (2015). Bayesian inference of multiple Gaussian graphical models. *Journal of the American Statistical Association*, 110(509):159–174. MR3338494

[38] Pourahmadi, M. (2007). Cholesky decompositions and estimation of a covariance matrix: orthogonality of variance–correlation parameters. *Biometrika*, 94(4):1006–1013. MR2376812

[39] Rawat, C., Kutum, R., Kukal, S., Srivastava, A., Dahiya, U. R., Kushwaha, S., Sharma, S., Dash, D., Saso, L., Srivastava, A. K., and Kukreti, R. (2020). Downregulation of peripheral PTGS2/COX-2 in response to valproate treatment in patients with epilepsy. *Scientific Reports*, 10(2546).

[40] Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866. MR1294730

[41] Rothenhäusler, D., Meinshausen, N., Bühlmann, P., and Peters, J. (2021). Anchor regression: Heterogeneous data meet causality. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(2):215–246. MR4250274

[42] Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688.

[43] Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D., and Nolan, G. (2005). Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308:523–529.

[44] San José-Enériz, E., Gimenez-Camino, N., Agirre, X., and Prosper, F. (2019). Hdac inhibitors in acute myeloid leukemia. *Cancers*, 11(11):1794.

[45] Shojaie, A. (2021). Differential network analysis: A statistical perspective. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(2):e1508. MR4218944

[46] Simpson, S. L., Bowman, F. D., and Laurienti, P. J. (2013). Analyzing complex functional brain networks: fusing statistics and network science to understand the brain. *Statistics surveys*, 7:1. MR3161730

[47] Spirtes, P., Glymour, C., and Scheines, R. (2000a). Causation, prediction and search (2nd edition). *Cambridge, MA: The MIT Press.*, pages 1–16. MR1227558

[48] Spirtes, P., Glymour, C. N., Scheines, R., and Heckerman, D. (2000b). *Causation, prediction, and search*. MIT press. MR1815675

[49] Squires, C., Wang, Y., and Uhler, C. (2020). Permutation-based causal structure learning with unknown intervention targets. In *Conference on Uncertainty in Artificial Intelligence*, pages 1039–1048. PMLR.

[50] Tallman, M. S., Gilliland, D. G., and Rowe, J. M. (2005). Drug therapy for acute myeloid leukemia. *Blood*, 106(4):1154–1163.

[51] Tan, Z. (2006). A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637. MR2279484

[52] Telesford, Q. K., Simpson, S. L., Burdette, J. H., Hayasaka, S., and Laurienti, P. J. (2011). The brain as a complex system: using network science as a tool for understanding the brain. *Brain connectivity*, 1(4):295–308. MR3161730

[53] Wang, Y., Squires, C., Belyaeva, A., and Uhler, C. (2018). Direct estimation of differences in causal graphs. *Advances in neural information processing systems*, 31.

[54] Yajima, M., Telesca, D., Ji, Y., and Müller, P. (2015). Detecting differential patterns of interaction in molecular pathways. *Biostatistics*, 16(2):240–251. MR3365426

[55] Zhang, K., Huang, B., Zhang, J., Glymour, C., and Schölkopf, B. (2017). Causal discovery from nonstationary/heterogeneous data: Skeleton estimation and orientation determination. In *IJCAI: Conference Proceedings*, volume 2017, page 1347. NIH Public Access.

[56] Zhao, Q., Small, D. S., and Bhattacharya, B. B. (2019). Sensitivity analysis for inverse probability weighting estimators via the percentile bootstrap. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 81(4):735–761. MR3997099