
An Analysis of the Adaptation Speed of Causal Models

Rémi Le Priol

Mila, Université de Montreal

Yoshua Bengio

Mila, Université de Montreal

Canada CIFAR AI Chair

Reza Babanezhad Harikandeh

SAIT AI Lab, Montreal

Simon Lacoste-Julien

Mila, Université de Montreal

Canada CIFAR AI Chair

SAIT AI Lab, Montreal

Abstract

Consider a collection of datasets generated by unknown interventions on an unknown structural causal model G . Recently, [Bengio et al. \(2020\)](#) conjectured that among all candidate models, G is the *fastest to adapt* from one dataset to another, along with promising experiments. Indeed, intuitively G has less mechanisms to adapt, but this justification is incomplete. Our contribution is a more thorough analysis of this hypothesis. We investigate the adaptation speed of cause-effect SCMs. Using convergence rates from stochastic optimization, we justify that a relevant proxy for adaptation speed is distance in parameter space after intervention. Applying this proxy to categorical and normal cause-effect models, we show two results. When the intervention is on the cause variable, the SCM with the correct causal direction is advantaged by a large factor. When the intervention is on the effect variable, we characterize the relative adaptation speed. Surprisingly, we find situations where the anticausal model is advantaged, falsifying the initial hypothesis.

1 INTRODUCTION

A learning agent interacting with its environment should be able to answer questions such as “what will happen to Y if I change X ”. Structural Causal Models




Figure 1: Two models of cause-effect data $X \rightarrow Y$.

(SCM) offer a formalism to answer this kind of questions ([Pearl, 2009](#); [Peters et al., 2017](#)). The simplest SCM is the model $X \rightarrow Y$ where X is the cause and Y the effect. Modifying X will modify Y but modifying Y will not alter X . In general, SCMs model the distribution of observations with a directed graph where edges represent *independent mechanisms* ([Janzing and Scholkopf, 2010](#)).

Modern machine learning methods can fail surprisingly when the test distribution differ from the training distribution ([Rosenfeld et al., 2018](#)). A recent line of work describes these distribution shifts as interventions in an underlying causal model ([Zhang et al., 2013](#); [Magliacane et al., 2018](#)). If this description is accurate, then an agent endowed with this hypothetical causal model could handle distribution shifts by updating the few mechanisms affected by the intervention. On contrary, an agent endowed with an incorrect model, would have to update many mechanisms. [Bengio et al. \(2020\)](#) infer that the causal agent will be the fastest to adapt to distribution shifts. Conversely, they use the speed of adaptation to unknown interventions as a criterion to learn the true causal model, showing promising empirical results on cause-effect models. Yet they lack a theoretical argument to connect interventions and fast adaptation. Thus we raise the question:

Do causal models adapt faster than non-causal models to distribution shifts induced by interventions?

Contributions. We theoretically and empirically answer this question for cause-effect SCMs with categor-

ical variables, and partially for multivariate normal distributions.

- For both settings, we use stochastic optimization convergence rates to show that the adaptation speed mostly depends on the distance in parameter space between the initialization (before intervention) and the optimum (after intervention).
- For categorical variables, we fully characterize this distance. We show that the causal model is faster by a large factor when the intervention is on the cause.
- When the intervention is on the effect, we surprisingly find settings where the anticausal model is systematically faster. As appealing as the fastest-to-adapt hypothesis may sound, it does not hold in every situations.

2 RELATED WORK

Causal relationships are asymmetric. These asymmetries are often visible in observations, so that one can identify which is cause and which is effect under relevant assumptions (Mooij et al., 2016). A common assumption is to constrain the set of functional dependencies between cause and effect. By contrast, in our work, we focus on two families of distributions which are notoriously unidentifiable from observational data: categorical and linear normal variables (Peters et al., 2017, Ch.4). With data coming from a generic directed acyclic graph (DAG), we can only hope to discover the Markov equivalence class of this DAG (Verma and Pearl, 1991). Many methods seek to achieve this goal, whether constraint-based such as the PC algorithm (Spirtes et al., 2000) or score-based methods using greedy search (Chickering, 2002) or more recently continuous optimization (Zheng et al., 2018; Lachapelle et al., 2020). However to discover the exact graph, we need access to interventional data.

Inferring causal links from interventions or experiments is the foundation of science. Inferring causal links from unknown interventions is a much harder and less principled problem. Tian and Pearl (2001) first studied this setting, proposing a constraint based method to infer the interventional equivalence class from a sequence of interventions. Then Eaton and Murphy (2007) proposed an exact Bayesian approach. More recently, Squires et al. (2019); Ke et al. (2019) proposed score based algorithms, improving in scalability and alleviating parametric assumptions. From a machine learning perspective, we are concerned with the predictive power that this structure will give us when faced with new data.

Distribution shifts are a common problem in machine learning, as well as in causal statistics (Zhang et al., 2013; Pearl and Bareinboim, 2014). Schölkopf et al. (2012) first brought up the idea of *invariance* to tackle this problem. Following up on this idea, Peters et al. (2016) designed an algorithm able to identify robust causal features from heterogeneous data. This work has set a fruitful line of research for robust machine learning (Heinze-Deml et al., 2018b,a; Rothenhäusler et al., 2019; Arjovsky et al., 2019). In a way, fast adaptation is the complementary idea of invariance: if most mechanisms are kept invariant, then only a few have to adapt. Schölkopf (2019) shed light on these approaches and the broader scope of causality research for machine learning.

3 BACKGROUND

In this section, we review the formalism of Bengio et al. (2020) on observations, interventions, models and adaptation.

Reference and Transfer Distributions. We assume perfect knowledge of a reference distribution \mathbf{p} over the pair (X, Y) sampled from an SCM $X \rightarrow Y$. This distribution is the object of interventions, which results in new *transfer* distributions \mathbf{p}^* . If the *intervention is on the cause*, X is sampled from a different marginal, then Y is sampled from the reference conditional

$$\mathbf{p}^*(x, y) = \mathbf{p}^*(x)\mathbf{p}(y|x). \quad (1)$$

If the *intervention is on the effect*, X is sampled from the reference marginal, then Y is sampled from another marginal independently of X

$$\mathbf{p}^*(x, y) = \mathbf{p}(x)\mathbf{p}^*(y). \quad (2)$$

For each transfer distribution, we observe a few samples.

Models. We parametrize two generative models of (X, Y) (Fig. 1):

$$\mathbf{p}_{\theta_{\rightarrow}}(x, y) = \mathbf{p}_{\theta_X}(x)\mathbf{p}_{\theta_{Y|X}}(y|x) \quad - \text{causal} \quad (3)$$

$$\mathbf{p}_{\theta_{\leftarrow}}(x, y) = \mathbf{p}_{\theta_Y}(y)\mathbf{p}_{\theta_{X|Y}}(x|y) \quad - \text{anticausal}. \quad (4)$$

For each model, we call mechanisms the marginal and conditional models. Each mechanisms has its own set of parameters, e.g. θ_X and $\theta_{Y|X}$. In the following we will use θ to denote interchangeably θ_{\rightarrow} and θ_{\leftarrow} .

Adaptation. Both models are initialized to fit perfectly the reference distribution $\mathbf{p}_{\theta_{\rightarrow}^{(0)}} = \mathbf{p}_{\theta_{\leftarrow}^{(0)}} = \mathbf{p}$. They observe fresh samples from \mathbf{p}^* one by one and update their parameters θ_{\rightarrow} and θ_{\leftarrow} to maximize the




Figure 2: **Intuition behind fast adaptation.** An intervention on X turns the reference distribution $\mathbf{p}^{(0)}$ into a transfer distribution \mathbf{p}^* . The causal model (blue) only has to adapt θ_X , whereas the anticausal model (red) has to adapt both its mechanisms. After adaptation, the causal model ends up the closest from the transfer in terms of KL, as visible in the abstract distribution space. Blue and red balls represent the *proximity prior* induced by taking a few steps of SGD from the reference in each parameter space. Convergence rate analysis reveals that they are spherical functions of the parameter distance, but they get mapped to non-trivial shapes in distribution space – ellipses in this sketch.

log-likelihood with a step of stochastic gradient (SGD). Thanks to the separate parameters, the causal model log-likelihood loss decomposes as

$$\begin{aligned} \mathcal{L}_{\text{causal}}(\theta_{\rightarrow}) &= \mathbb{E}_{(X,Y) \sim \mathbf{p}^*} [-\log \mathbf{p}_{\theta_{\rightarrow}}(X, Y)] \\ &= \mathbb{E}_{\mathbf{p}^*} [-\log \mathbf{p}_{\theta_X}(X)] + \mathbb{E}_{\mathbf{p}^*} [-\log \mathbf{p}_{\theta_{Y|X}}(Y|X)] \end{aligned} \quad (5)$$

When \mathbf{p}^* comes from an intervention, [Bengio et al. \(2020\)](#) observe that the causal model is often faster to adapt than the anticausal model. Intuitively, this is because the causal model has to adapt only the mechanism which was modified by the intervention. On the other hand, the anticausal model has to adapt both its mechanisms. In Figure 2, we compare these different scenarios and the concept of adaptation figuratively. While appealing, *this reasoning is not rigorous*, as sample complexity bounds of SGD typically do not depend on the number of parameters to update ([Bubeck et al., 2015](#), Th. 6.2 & 6.3). In the next section, we formalize and understand this phenomenon in the light of convergence rates of stochastic optimization methods.

Distribution Families. We study two of the simplest sub-families of the exponential family ([Wainwright and Jordan, 2008](#)): categorical and linear normal variables. Their negative log-likelihood is a convex function of their natural parameter. These families are interesting because the direction is not identifiable from observational data ([Peters et al., 2017](#), Ch.4) – e.g. $\mathbf{p}_{\theta_{\rightarrow}}$ and $\mathbf{p}_{\theta_{\leftarrow}}$ can model the same set of distributions – which makes them challenging for causal discovery.

AN OPTIMIZATION PERSPECTIVE One way to formalize adaptation speed is to characterize it via the convergence speed of the stochastic optimization procedure. An appealing aspect of stochastic optimization algorithms such as SGD (when only using fresh samples and running it on the true loss we care about) is that they come with convergence rate guarantees on the *population risk* in machine learning, thus giving us direct sample complexity results to obtain a specific generalization error. The convergence rate is an *upper bound* on the expected suboptimality after a given number of iterations. While these rates are about worst case performance and might also be loose, fortunately, for convex optimization, they tend to correspond well to actual empirical performance ([Nesterov, 2004](#)). We can thus use the convergence bounds as theoretical proxy for the convergence speed. In our experiments, we also verify empirically that the bounds correlate well with the observed convergence speed.

Here we provide a classical convergence rate on the expected suboptimality with Average Stochastic Gradient Descent (ASGD) under convexity and bounded gradient assumptions. We re-derive this rate in Appendix B.1 for completeness. This rate applies to log-likelihood maximization for categorical random variables (details in B.2). Since the target distribution is part of the model family, the log-likelihood suboptimality is equal to the KL-divergence – e.g. $\mathcal{L}(\theta) - \mathcal{L}(\theta^*) = D_{\text{KL}}(\mathbf{p}^* || \mathbf{p}_{\theta})$.

ASGD. Assume $\forall \theta, x, \|\nabla \log p_\theta(x)\| \leq B$. After T iterations of SGD on (5),

$$\theta^{(t+1)} = \theta^{(t)} + \gamma \nabla \log p_{\theta^{(t)}}(X_t, Y_t) \quad (6)$$

with learning rate $\gamma := \frac{c}{\sqrt{T}}$, starting from $\theta^{(0)}$, the average parameter's $\bar{\theta}^{(T)} = \frac{1}{T} \sum_{t=0}^{T-1} \theta^{(t)}$ suboptimality is upper bounded by

$$\mathbb{E}[D_{\text{KL}}(\mathbf{p}^* \| \mathbf{p}_{\bar{\theta}^{(T)}})] \leq \frac{c^{-1} \|\theta^{(0)} - \theta^*\|^2 + cB^2}{2\sqrt{T}} \quad (7)$$

where the expectation is taken over the sampling of $T-1$ training points X_t, Y_t and θ^* is the closest solution to $\theta^{(0)}$ in the solution set $\text{argmin}_\theta \mathcal{L}(\theta)$.

For categorical models, $B = 2$ (see B.2). Consequently, for a fixed T and with small enough c , the convergence upper bounds for causal and anticausal models differ mainly by $\delta := \|\theta^{(0)} - \theta^*\|^2$.

The bounded gradient assumption of (7) does not apply to the log-likelihood of normal variables. In Section 5.1, we provide an algorithm along with a convergence rate (22) that do apply to this case. Overall both bounds (7) and (22) carry the same message which can be summarized by:

The adaptation speed is dominated by the initial distance

$$\delta_{\text{causal}} = \left\| \theta_{\rightarrow}^{(0)} - \theta_{\rightarrow}^* \right\|^2 \quad (8)$$

$$\delta_{\text{anticausal}} = \left\| \theta_{\leftarrow}^{(0)} - \theta_{\leftarrow}^* \right\|^2. \quad (9)$$

Other optimization methods. Yang et al. (2016, Theorem 1) provides a unified convergence rate for stochastic heavy ball and Nesterov methods that is similar to (7), where the initial distance is the main difference between causal and anticausal models. Consequently, our theoretical analysis holds for a larger class of algorithms than ASGD. More generally, it applies to any stochastic optimization method whose sample complexity depends on parameter distance.

4 CATEGORICAL VARIABLES

In this section, both cause and effect come from categorical distribution. We provide theoretical bounds on δ_{causal} and $\delta_{\text{anticausal}}$. We consider different scenarios to generate reference and transfer data and explain the consequences of each scenario.

4.1 Definitions

Cause X and effect Y are now two categorical variables taking values in $\{1, \dots, K\}$. Categorical variables are




Figure 3: Parametrization of categorical models

an exponential family with mean parameters $\mathbf{p} \in \Delta_K$ the probability vector, and with natural parameter $\mathbf{s} \in \mathbb{R}^K$ – the logits or score parameters such that $p_z = \frac{e^{s_z}}{\sum_{z'} e^{s_{z'}}$. The causal model has parameters $\mathbf{s}_X := (s_x)_{x=1 \dots K}$ and $\mathbf{s}_{Y|X} := (s_{y|x})_{x,y=1 \dots K}$. We gather the causal parameters in the variable $\theta_{\rightarrow} = (\mathbf{s}_X, \mathbf{s}_{Y|X})$ and the anticausal parameters in $\theta_{\leftarrow} = (\mathbf{s}_Y, \mathbf{s}_{X|Y})$ (Fig. 3). The loss (5) becomes

$$\begin{aligned} \mathcal{L}_{\text{causal}}(\theta_{\rightarrow}) &= \mathbb{E}_{(X,Y) \sim \mathbf{p}^*} [-\log p_{\theta_{\rightarrow}}(X, Y)] \\ &= \mathbb{E}_{\mathbf{p}^*} \left[-s_X + \log \sum_x e^{s_x} - s_{Y|X} + \log \sum_y e^{s_{y|x}} \right]. \end{aligned} \quad (10)$$

Each mechanism's stochastic loss is the sum of a linear function and a softmax function. The softmax function is convex and 1-Lipschitz, so we can apply rate (7). To be self-contained, we include details in Appendix B.2.

4.2 Distance after Intervention

In this section, we prove that interventions on the cause advantage the causal model by a factor K , and we describe when interventions on the effect will advantage one model over another.

Intervention on cause X , $\mathbf{s}_X \leftarrow \mathbf{s}_X^*$. The causal conditional $\mathbf{s}_{Y|X}$ is left unchanged, but the effect marginal \mathbf{s}_Y is modified in a non-trivial way. Consequently the initial distances are

$$\delta_{\text{causal}} = \|\mathbf{s}_X - \mathbf{s}_X^*\|^2 \quad (11)$$

$$\delta_{\text{anticausal}} = \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 + \sum_y \|\mathbf{s}_{X|y} - \mathbf{s}_{X|y}^*\|^2. \quad (12)$$

The causal model has to update K parameters, whereas the anticausal model has to adapt $K^2 + K$ parameters. Therefore the causal model seems to be advantaged by a factor K . The following proposition – proved in Appendix C – shows that this is reflected by ℓ_2 distances.

Proposition 1. *When the intervention happens on the cause,*

$$\delta_{\text{anticausal}} \geq K \delta_{\text{causal}}. \quad (13)$$




Figure 4: **Illustration of Proposition 2.** c is on the line joining s_Y and $m := \frac{1}{K} \sum_x s_{Y|x}$. When s_Y^* is within the blue ball of radius R centered at c , $\Delta \leq 0$ and the causal model is advantaged, otherwise the anticausal model is advantaged (red area). This is a surprising counter-example to the adaptation-speed hypothesis.

Intervention on effect Y , $\forall x, s_{Y|x} \leftarrow s_Y^*$. Cause and effect become independent. The causal model is advantaged only if the intervention s_Y^* is close enough from the previous marginal, as formalized by the following proposition:

Proposition 2. *When the intervention happens on the effect*

$$\begin{aligned} \Delta &:= \delta_{causal} - \delta_{anticausal} \\ &= (K-1) \left(\|s_Y^* - c\|^2 - R^2 \right) \end{aligned} \quad (14)$$

where $R^2 \approx K \widehat{\text{Var}}_X[\log \sum_y e^{s_{y|x}}]$ and $c = \frac{(\sum_x s_{Y|x}) - s_Y}{K-1}$.

See Figure 4 for an illustration and Appendix C.3 for the exact formula of R and the proof. When the intervention s_Y^* is close enough to c , which depends on the reference, the causal model is advantaged. If s_Y^* is far from c or if R is small then the anticausal model is likely to be advantaged.

4.3 Simulating Reference Distributions

To evaluate the fast adaptation criterion, we are going to work on synthetic data, which raises the question : from which distribution should we sample $p = p_{\theta^{(0)}}$? We call this distribution *prior*. Following the independent mechanism assumption, the marginal on the cause p_X and the conditional of effect given cause $p_{Y|X}$ should not contain any information about each other.

Dense Prior. To sample causal mechanisms, a natural choice is

$$p_X \sim \text{Dir}(\mathbf{1}_K) \quad \text{and} \quad \forall x, p_{Y|x} \sim \text{Dir}(\mathbf{1}_K) \quad (15)$$

where Dir is the Dirichlet distribution and $\mathbf{1}_K$ is the all-one vector of dimension K . $\text{Dir}(\mathbf{1}_K)$ the uniform law over the simplex Δ_K . This prior leads to the K2 score from the Bayesian network literature (Cooper and Herskovits, 1991). We call this choice the *dense prior* by opposition to the sparse prior introduced next. This is the choice made in Bengio et al. (2020), as well as Chalupka et al. (2016). The latter work reports that distributions sampled from this prior exhibit some asymmetry between X and Y . In Appendix D.1, we complement their work, explaining how the effect marginal is likely to be closer from the uniform distribution than the cause marginal. This asymmetry means that *the causal direction is identifiable from observational data*.

Sparse Prior. To fix this issue, we study an alternative prior that is symmetric and ensures that both cause and effect marginals are sampled from a uniform prior over Δ_K . We sample the causal mechanisms as follows

$$p_X \sim \text{Dir}(\mathbf{1}_K) \quad \text{and} \quad \forall x, p_{Y|x} \sim \text{Dir}(\mathbf{1}_K / K). \quad (16)$$

The $\mathbf{1}_K / K$ parameter means that samples will be approximately sparse, hence the name. We show in Appendix D.2 that with this sampling scheme, the joint is sampled from a sparse Dirichlet over Δ_{K^2} : $p_{(X,Y)} \sim \text{Dir}(\mathbf{1}_{K^2} / K)$. This in turns means that we can switch the roles of X and Y in (16). The effect marginal has uniform density over the simplex. In general, *the causal direction is not identifiable from observational data*. In Bayesian Networks literature, this is known as the Bayesian Dirichlet equivalent uniform prior (Heckerman et al., 1995).

4.4 Categorical Variables Experiments

Goal. As discussed in Section 4.3, the prior over the joint distribution on (X, Y) is going to influence the behavior of ASGD. We are seeking answers to two questions:

1. Is the adaptation speed positively correlated with the initial distance, as suggested by the upper bound (7) on the convergence rate of ASGD?
2. Is there a clear difference in adaptation speed between causal and anticausal models?

Data. We consider categorical variables with $K = 20$. For each initialization method, we sample 100 different reference joint distributions. For each of these distributions, we sample an intervention by sampling a probability vector q *uniformly* from Δ_K . If the intervention is on the cause, we plug q instead of p_X . If the intervention is on the effect, we redefine $p_{Y|x} = q, \forall x$.

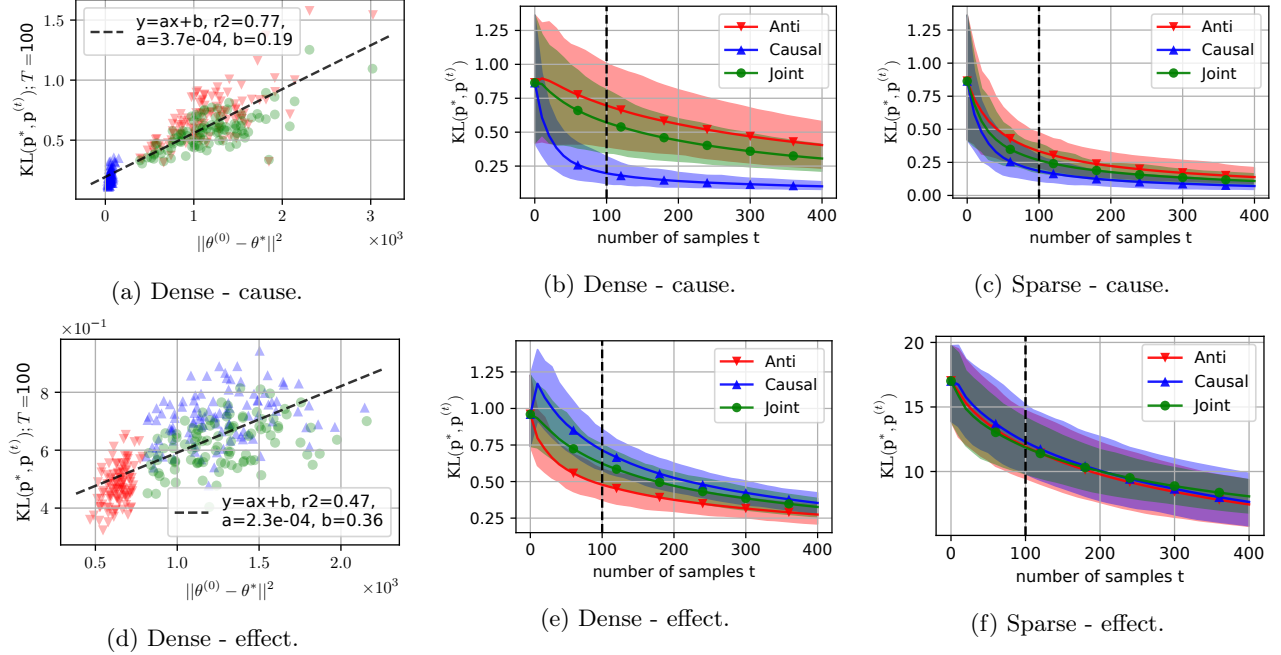


Figure 5: **Experimental results on categorical data.** Each plot is captioned with the prior and the intervention considered. **Scatter plots** are showing the positive correlation between the KL after 100 steps of SGD and the initial parameter distance. Each point represent one of 100 synthetic pairs $(p^{(0)}, p^*)$. **Training curves** show the average KL (solid line) and the (5,95) percentiles (shaded) over 100 runs. Remark how all models start from the same initial KL, but they converge at different speeds.

Models. We are comparing causal and anticausal models adaptation speed. We also report results for a model of the joint $p_{X,Y} = \text{softargmax}(s_{X,Y})$ as a reference model. We expect its results to be in between the performance of the causal and anticausal model as it expresses no prior over the direction. We optimize all models with Averaged SGD. In each iteration of SGD we get one fresh sample from the transfer distribution. For each model and each setting, we tune the (constant) learning rate so as to optimize the likelihood after seeing $\frac{K^2}{4} = 100$ samples, to explore the few samples regime. We present results in Figure 5

Dense prior. When the intervention is on the cause, the causal model is much closer from its optimum: in Fig. 5a the blue cluster is on the left of the scatter plots. This is well correlated with faster adaptation (Fig. 5b). On the contrary, *when the intervention is on the effect, the anticausal model starts closer from its optimum and it converges faster* (Fig. 5d, 5e). We can interpret this result in light of Proposition 2. In Appendix C.3, we explain why the radius R is small under the dense prior. As a result, s_Y^* is mostly sampled outside of the ball of radius R , consequently the anticausal model is advantaged. Overall, there is a wider gap between models in Fig. 5b than in Fig. 5e. Consequently, if we take a balanced average of a few interventions on the

cause and a few interventions on the effect, the causal model remains faster (details in Appendix C.4).

Sparse prior. When the intervention is on the cause, the causal model has a slight advantage (Fig. 5c). When the intervention is on the effect, no model has a set advantage (Fig. 5f), but the sparsity induces much higher KL values, as explained in Appendix D.3. This KL explosion drowns the signal coming from the cause intervention, calling for further algorithmic developments – such as inferring the intervention, as explored by [Ke et al. \(2019\)](#).

5 MULTIVARIATE NORMAL VARIABLES

In this section, we analyze the case of two multivariate normal variables with a linear relationship. Cause X and effect Y are sampled from the causal model

$$X \sim \mathcal{N}(\mu_X, \Sigma_X) \quad (17)$$

$$Y|X \sim \mathcal{N}(AX + a, \Sigma_{Y|X}) \quad (18)$$

with mean parameters $\mu_X, a \in \mathbb{R}^K$ and $\Sigma_X, A, \Sigma_{Y|X} \in \mathbb{R}^{K \times K}$. This parametrization is the most intuitive but it is unfortunately not appropriate to get convergence rates. We are going to introduce another parametriza-

tion along with an algorithm and a convergence rate (Sec. 5.1), before providing empirical results (Sec. 5.2).

5.1 Optimization Analysis

The negative log-likelihood of model (17) is notoriously non-convex. This is problematic for convergence results. For simplicity, we focus in this section on the simple marginal mechanism with mean parameters μ, Σ . We detail the full model in Appendix F. If we use the natural parameters $\eta = \Sigma^{-1}\mu$ and $\Lambda = \Sigma^{-1}$ (precision matrix), the negative log-likelihood is convex

$$\begin{aligned} \mathbb{E} [-\log \mathbf{p}_{(\eta, \Lambda)}(X)] \\ = \frac{1}{2} \left(\mathbb{E} [\text{Tr}(X X^\top \Lambda) - 2X^\top \eta] + \eta^\top \Lambda^{-1} \eta - \log |\Lambda| \right). \end{aligned} \quad (19)$$

This objective is composed of a pleasant stochastic linear term, and a difficult deterministic barrier objective which goes to infinity when $\Lambda \rightarrow 0$. This barrier is composed of a matrix inverse and a log determinant. The assumptions of Lipschitz or gradient-Lipschitz required to get SGD convergence do not hold for the barrier. While the empirical version of (19) has a close formed formula for its global minimum, quite surprisingly, gradient-based optimization of the normal likelihood is difficult to analyze. Convex optimization typically deals with non-smooth terms by introducing proximal operators (Parikh et al., 2014). However this barrier term is too complex to get an analytic formula for the proximal operator. We transform it into a more convenient form by introducing \mathbf{L} , the lower triangular Cholesky factor of the precision matrix $\Lambda = \mathbf{L}\mathbf{L}^\top$, and $\zeta = \mathbf{L}^{-1}\eta = \mathbf{L}^\top \mu$. Then (19) simplifies into

$$\begin{aligned} \mathbb{E} [-\log \mathbf{p}_{(\zeta, \mathbf{L})}(X)] \\ = \frac{1}{2} \mathbb{E} \left[\|\mathbf{L}^\top X - \zeta\|^2 \right] - \sum_i \log \mathbf{L}_{i,i}. \end{aligned} \quad (20)$$

We will refer to (ζ, \mathbf{L}) as *Cholesky parameters*. This objective is more suitable to gradient based optimization with a simple proximal operator, as detailed in the next section. We provide all details about the causal model in Appendix F.

Stochastic Proximal Gradient Algorithm We want to minimize the sum of a stochastic convex smooth function $f_X(\theta) := \frac{1}{2} \|\mathbf{L}^\top X - \zeta\|^2$ and convex non-smooth regularizer $g(\theta) = -\sum_i \log \mathbf{L}_{i,i}$. This is exactly the goal of the stochastic proximal gradient (Duchi et al., 2010) update

$$\theta_{t+1} = \underset{\theta}{\operatorname{argmin}} g(\theta) + \frac{1}{2\gamma_t} \|\theta_t - \gamma_t \nabla f_{X_t}(\theta_t) - \theta\|^2 \quad (21)$$

where γ_t is the step-size and X_t is randomly sampled. For objective (20), the proximal gradient update has a

closed form solution that amounts to updating all parameters with the stochastic gradient of the quadratic term, then updating the diagonal elements of \mathbf{L} with the mapping $x \mapsto \frac{1}{2}(x + \sqrt{x^2 + 4\gamma})$, thus ensuring that they remain strictly positive (details in Appendix E.2).

Convergence Rate. We assume that stochastic gradients are almost-surely B -Lipschitz. B is known as the smoothness constant. We show in Appendix E.1 that running the stochastic proximal gradient algorithm with step size $\gamma_t = \frac{\gamma}{3B\sqrt{T}}$ where $\gamma \leq 1$, for T iterations guarantees

$$\begin{aligned} \mathbb{E} [D_{\text{KL}}(\mathbf{p}^* \| \mathbf{p}_{\bar{\theta}(T)})] \\ \leq \frac{3B \|\theta^{(0)} - \theta^*\|^2}{\gamma\sqrt{T}} + \frac{D_{\text{KL}}(\mathbf{p}^* \| \mathbf{p}_{\theta^{(0)}})}{T}. \end{aligned} \quad (22)$$

Analysis. The term $KL(\mathbf{p}^* \| \mathbf{p}_{\theta^{(0)}})/T$ is equal for causal and anticausal models because we assume $\mathbf{p}_{\theta \rightarrow}^{(0)} = \mathbf{p}_{\theta \leftarrow}^{(0)}$. For normal variables, B depends only on the data and is a priori equal for both models (Appendix E.3). Similarly to (7), both models' rates differ mainly by $\delta = \|\theta^{(0)} - \theta^*\|^2$.

When the intervention is on the cause, we prove in Appendix F.5 that the anticausal model is farther away from its optimum in the natural parametrization

$$\delta_{\text{anticausal}}^{\text{natural}} \geq \delta_{\text{causal}}^{\text{natural}}. \quad (23)$$

Unfortunately, in the Cholesky parametrization (Fig. 6, 2nd column), or when the intervention is on the effect (Fig. 6, bottom row), we observe empirically that there is no such hard guarantee, although the causal distance tends to be smaller than the anticausal distance.

5.2 Experiments

Similarly to categorical variables, we need to decide on a prior over reference and transfer distributions. This choice is informed by two criteria. First the independent mechanism principle which states that we should sample θ_X independently of $\theta_{Y|X}$. Second we want θ_Y to have approximately the same distribution as θ_X – e.g. we want the distribution to be approximately symmetric so that we cannot identify the direction from observational data. These considerations lead us to a flavor of normal-Wishart prior (Geiger et al., 2002) described in Appendix G.

We sample 100 random joint distributions from this prior, and for each distribution we sample a random intervention on the cause, and a random intervention on the effect. We then run the stochastic proximal gradient on objective (20). We report results in Figure 6. Similarly to the categorical case, when the intervention

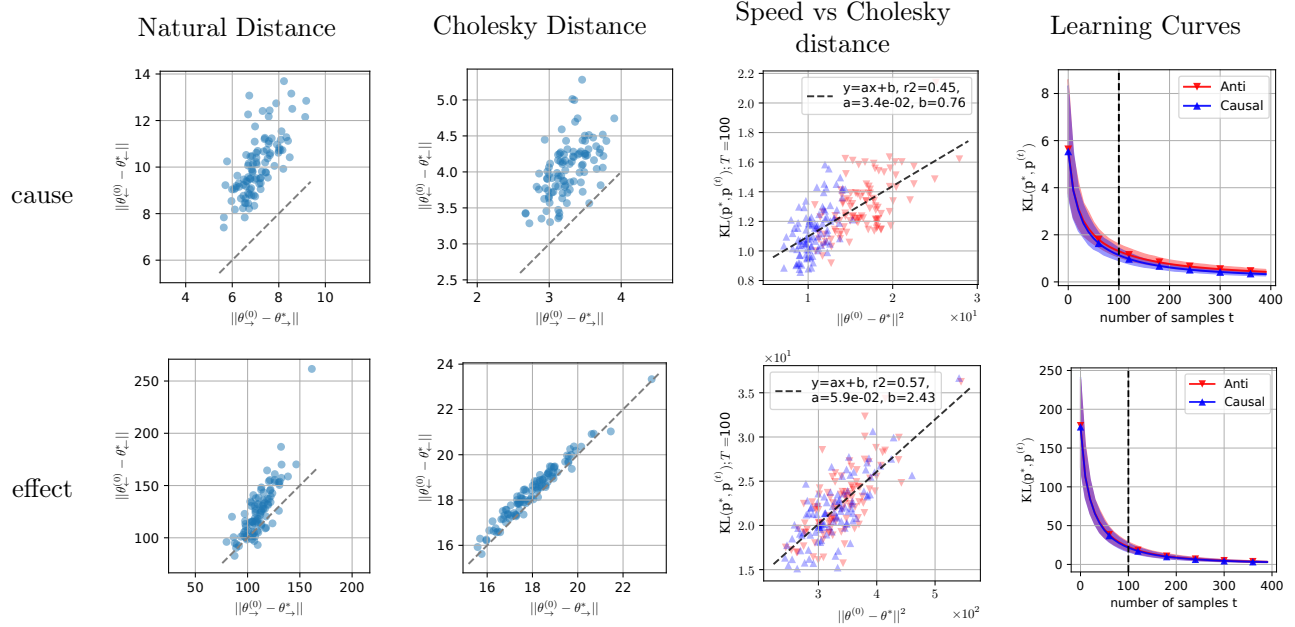


Figure 6: **Multivariate Normal Variables with dimension $K = 10$.** Row 1 and 2 correspond to interventions on cause and effect respectively. *Column 1 & 2:* scatter plot $\delta_{\text{anticausal}}$ vs δ_{causal} respectively in natural and Cholesky parametrization. The grey diagonal is the identity line. We observe a natural tendency for $\delta_{\text{anticausal}} > \delta_{\text{causal}}$ (points above the grey diagonal), but this is systematically true only for the natural distance when the intervention is on the cause. *Column 3 & 4:* same plot as in Figure 5. Once again we observe a correlation between initial distance and optimization speed. When the intervention is on the cause, the causal model is advantaged. When the intervention is on the effect, both curves overlap.

is on the cause, the causal model is advantaged by a slight margin (upper right figure). When the intervention is on the effect both models are learning at the same speed (bottom right figure).

Conclusion

We provided a first theoretical analysis of the adaptation speed in two-variables cause-effect SCMs under localized interventions for categorical and normal data. Convergence guarantees for stochastic optimization on the true population log-likelihood indicates that the adaptation speed is related to the distance between initial point and optimum in parameter space. We verified this correlation empirically. We proved analytically that this distance is lower for the causal model than for the anticausal model when the intervention is on the cause variable. This explains a surprising phenomenon: while both models start with the same suboptimality, one learns faster than the other. When the intervention is on the effect variable, we highlighted examples showing that either model can be advantaged. This observation challenges the intuition that the causal model should be the fastest to adapt, and it raises new questions for the approach of [Bengio et al. \(2020\)](#), such as: are there practical situations where the fastest-to-adapt

heuristic is useful? On a more theoretical note, is it possible to characterize the adaptation speed behavior for more general families of distributions?

Acknowledgments

We thank Sébastien Lachapelle and Tristan Deleu for insightful discussions and useful feedback that led to the main results of this paper. We thank Waïss Azizian for his help in debugging proofs. This work was partially supported by the Canada CIFAR AI Chair Program. Simon Lacoste-Julien is a CIFAR Associate Fellow and Yoshua Bengio is a Program Co-Director in the Learning in Machines & Brains program.

References

- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Yves F Atchadé, Gersende Fort, and Eric Moulines. On perturbed proximal gradient algorithms. *Journal of Machine Learning Research*, 18(1):310–342, 2017.
- Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Rosemary Ke, Sébastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher Pal. A meta-transfer

- objective for learning to disentangle causal mechanisms. In *ICLR*, 2020.
- S bastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends  in Machine Learning*, 2015.
- Krzysztof Chalupka, Frederick Eberhardt, and Pietro Perona. Estimating causal direction and confounding of two discrete variables. *arXiv preprint arXiv:1611.01504*, 2016.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.
- Gregory F Cooper and Edward Herskovits. A bayesian method for constructing bayesian belief networks from databases. In *Uncertainty Proceedings 1991*. Elsevier, 1991.
- John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *Journal of Machine Learning Research*, 10(Dec):2899–2934, 2009.
- John C Duchi, Shai Shalev-Shwartz, Yoram Singer, and Ambuj Tewari. Composite objective mirror descent. In *COLT*, pages 14–26, 2010.
- Daniel Eaton and Kevin Murphy. Exact bayesian structure learning from uncertain interventions. In *Artificial intelligence and statistics*, pages 107–114, 2007.
- Dan Geiger, David Heckerman, et al. Parameter priors for directed acyclic graphical models and the characterization of several probability distributions. *The Annals of Statistics*, 30(5):1412–1440, 2002.
- David Heckerman, Dan Geiger, and David M Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243, 1995.
- Christina Heinze-Deml, Marloes H Maathuis, and Nicolai Meinshausen. Causal structure learning. *Annual Review of Statistics and Its Application*, 2018a.
- Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 2018b.
- Dominik Janzing and Bernhard Scholkopf. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory*, 56(10), 2010.
- Nan Rosemary Ke, Olexa Bilaniuk, Anirudh Goyal, Stefan Bauer, Hugo Larochelle, Chris Pal, and Yoshua Bengio. Learning neural causal models from unknown interventions. *arXiv preprint arXiv:1910.01075*, 2019.
- S bastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien. Gradient-based neural DAG learning. In *ICLR*, 2020.
- Sara Magliacane, Thijs van Ommen, Tom Claassen, Stephan Bongers, Philip Versteeg, and Joris M Mooij. Domain adaptation by using causal inference to predict invariant conditional distributions. In *NeurIPS*, 2018.
- Joris M. Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Sch lkopf. Distinguishing cause from effect using observational data: Methods and benchmarks. *Journal of Machine Learning Research*, 2016.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2004.
- Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and Trends  in Optimization*, 1(3):127–239, 2014.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Judea Pearl and Elias Bareinboim. External validity: From do-calculus to transportability across populations. *Statistical Science*, 2014.
- Jonas Peters, Peter B hlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2016.
- Jonas Peters, Dominik Janzing, and Bernhard Sch lkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- Amir Rosenfeld, Richard Zemel, and John K Tsotsos. The elephant in the room. *arXiv preprint arXiv:1808.03305*, 2018.
- Dominik Rothenh usler, Peter B hlmann, and Nicolai Meinshausen. Causal dantzig: fast inference in linear structural equation models with hidden variables under additive interventions. *The Annals of Statistics*, 47(3):1688–1722, 2019.
- Bernhard Sch lkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris Mooij. On causal and anticausal learning. In *ICML*, 2012.
- Bernhard Sch lkopf. Causality for machine learning. *arXiv:1911.10500*, 2019.
- Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search*. MIT press, 2000.
- Chandler Squires, Yuhao Wang, and Caroline Uhler. Permutation-based causal structure learning with unknown intervention targets. *arXiv preprint arXiv:1910.09007*, 2019.
- Jin Tian and Judea Pearl. Causal discovery from changes. In *UAI*, 2001.

- Thomas Verma and Judea Pearl. *Equivalence and synthesis of causal models*. UCLA, Computer Science Department, 1991.
- Martin J Wainwright and Michael I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.
- Tianbao Yang, Qihang Lin, and Zhe Li. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. *arXiv preprint arXiv:1604.03257*, 2016.
- Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *ICML*, 2013.
- Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In *NeurIPS*, 2018.

A SOURCE CODE

Source code for all experiments is hosted at <https://github.com/remilepriol/causal-adaptation-speed>.

B CATEGORICAL OPTIMIZATION

In this section we prove a convergence rate of ASGD that applies to the categorical loss, and we show that the constants involved in this rate are the same for both causal and anticausal models.

B.1 Convergence of ASGD with Fixed Step-Size

Here we derive a classical convergence rate of Average SGD. This result is standard ; we include it to be self-contained. The objective is

$$\min_{\theta} F(\theta) = \mathbb{E}_i [f(\theta, i)] . \quad (24)$$

Theorem B.1. *If each $f_i(\theta) = f(\theta, i)$ has bounded gradient B , then after T steps of SGD with step-size $\gamma = \frac{c}{\sqrt{T}}$, starting from θ_0 , the expected sub-optimality verifies*

$$\mathbb{E} [F(\bar{\theta}_T) - F(\theta^*)] \leq \frac{1}{2c\sqrt{T}} \|\theta_0 - \theta^*\|^2 + \frac{cB^2}{2\sqrt{T}} \quad (25)$$

where $\bar{\theta}_T = \frac{1}{T} \sum_t \theta_t$.

Proof. First we relate the ℓ^2 distance to optimum at step $t + 1$ with the one at step t :

$$\begin{aligned} & \|\theta_{t+1} - \theta^*\|^2 \\ &= \|\theta_t - \theta^*\|^2 - 2\gamma \langle f'_i(\theta_t), \theta_t - \theta^* \rangle + \gamma^2 \|f'_i(\theta_t)\|^2 \\ &\leq \|\theta_t - \theta^*\|^2 - 2\gamma \langle f'_i(\theta_t), \theta_t - \theta^* \rangle + \gamma^2 B^2 . \end{aligned}$$

By convexity of f_i and rearranging the terms we get

$$\begin{aligned} 2\gamma(f_i(\theta_t) - f_i(\theta^*)) &\leq 2\gamma \langle f'_i(\theta_t), \theta_t - \theta^* \rangle \\ &\leq \|\theta_t - \theta^*\|^2 - \|\theta_{t+1} - \theta^*\|^2 + \gamma^2 B^2 . \end{aligned}$$

Now we take the expectation, sum up both sides for T iterations and divide by $2T\gamma$ to get

$$\begin{aligned} & \frac{1}{T} \sum_{i=1}^T \mathbb{E} [F(\theta_t) - F(\theta^*)] \\ &\leq \frac{1}{2\gamma T} (\mathbb{E} [\|\theta_0 - \theta^*\|^2] - \mathbb{E} [\|\theta_{T+1} - \theta^*\|^2]) + \frac{\gamma B^2}{2} \\ &\leq \frac{1}{2\gamma T} \|\theta_0 - \theta^*\|^2 + \frac{\gamma B^2}{2} \end{aligned}$$

Finally, we apply Jensen inequality to F in $\bar{\theta}_T = \frac{1}{T} \sum_t \theta_t$ to get the final result. \square

B.2 Categorical Loss Properties

We are now going to verify that assumptions of the rate (7) apply to the negative log-likelihood loss for the categorical distribution. This loss is standard and its properties are well-known, but we review them here to be self-contained.

Each mechanism has the same form of negative log-likelihood, with the same kind of stochastic gradients. The total loss is a sum over mechanisms, and the total stochastic gradient is a concatenation of each mechanisms stochastic gradient. To apply rate 7, we can either apply it separately on each mechanism, either apply to the

whole. Both path lead to the same result. In the end, we simply have to check that this loss is convex, and has bounded gradients for all z . The random functions coming from sampling Z are

$$f_z(\mathbf{s}) = -s_z + \log\left(\sum_{z'} e^{s_{z'}}\right) \quad (26)$$

This function is the softmax – or logsumexp – function minus a stochastic linear term.

Convexity We are going to show that it is convex but not strongly convex because it becomes flat for large score values. Its derivative is

$$\nabla f_z(\mathbf{s}) = -\mathbf{e}_z + \mathbf{p} . \quad (27)$$

where \mathbf{e}_z is the z -th canonical basis element and \mathbf{p} is the output of the softargmax function taken on \mathbf{s}_Z . The Hessian is the same for every z .

$$\nabla^2 f_z(\mathbf{s}) = \text{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\top . \quad (28)$$

We observe that for any vector \mathbf{v} ,

$$\begin{aligned} \mathbf{v}^\top \nabla^2 f_z(\mathbf{s}) \mathbf{v} &= \sum_z p_z v_z^2 - \left(\sum_z p_z v_z \right)^2 \\ &= \text{Var}_{Z \sim \mathbf{p}}[v_Z] \geq 0 \end{aligned} \quad (29)$$

which means that the logsumexp is convex. When s_0 tends toward positive infinity and the other components remain constant, \mathbf{p} tends toward a Dirac on the 0-th component. Then (29) is 0 for all \mathbf{v} , so the logsumexp is not strongly convex.

Bounded Gradients. The gradient norm is

$$\|\nabla f_z(\mathbf{s})\| = \|\mathbf{p} - \mathbf{e}_z\| \quad (30)$$

$$. \quad (31)$$

This norm is maximized for $\mathbf{p} = \mathbf{e}_{z'}, \forall z' \neq z$. The maximum is equal to $\sqrt{2}$. If there are d independent mechanisms (for d variables in the graph), then the total stochastic gradient which is a concatenation of all gradients has a norm bounded by $B = \sqrt{2}d$. In our case of cause-effect models, $d = 2$ and the gradients are bounded by $B = 2$, or in other words, all the f_z are 2-Lipschitz.

This bound is the same for causal and anticausal models. It depends on the part of space where \mathbf{p} is going to live. Assuming that it is going to live in most of the space for both directed models, both loss will have the same Lipschitz constants in practice.

Thanks to these properties, the sample complexity of $\mathbf{p}_{\theta \rightarrow}$ and $\mathbf{p}_{\theta \leftarrow}$ are bounded by (7). The difference in adaptation speed between causal and anticausal models is characterized by the distance in parameter space.

C CATEGORICAL ANALYSIS

In this section, we prove relationships between parameter distances induced by interventions between the causal and anticausal models. First we prove two useful lemmas. Then we establish that the causal model dominates the anticausal model by a factor K when the intervention is on the cause. Finally we show that no model has a set advantage when the intervention bears on the effect.

The logits or scores \mathbf{s} live in \mathbb{R}^K . They have one additional degree of freedom compared to the probability \mathbf{p} . More specifically, the softargmax is invariant by translations along the vector $\mathbf{1} = (1, \dots, 1)$. In other words, all scores $\{\mathbf{s} + \lambda \mathbf{1} \mid \forall \lambda \in \mathbb{R}\}$ are equivalent. Scores which move by following the gradient of this loss will remain in the same affine hyperplane orthogonal to $\mathbf{1}$. To ensure that the distances we measure are meaningful, we project all logits in the hyperplane such that $\sum_z s_z = 0$, by subtracting their mean.

Definition C.1 (Mean-zero score). *A score vector \mathbf{s} is mean-zero iff $\sum_z s_z = 0$.*

C.1 Switching Direction

In this section we are going to prove a few useful results relating cause and anticausal models. We know the causal parameters $X \rightarrow Y$, and we want to find the corresponding $X \leftarrow Y$ model, e.g. express $\mathbf{s}_Y, \mathbf{s}_{X|Y}$ as a function of $\mathbf{s}_X, \mathbf{s}_{Y|X}$. This will help us to find a relationship between δ_{causal} and $\delta_{\text{anticausal}}$. We first need to define a few useful variables

Definition C.2 (Average conditional score vectors). *For any x or y , define*

$$m(y) := \frac{1}{K} \sum_x s_{y|x}, \quad n(x) := \frac{1}{K} \sum_y s_{x|y}. \quad (32)$$

Definition C.3 (Conditional log-partition function).

$$A(x) = \log \sum_y e^{s_{y|x}} \quad (33)$$

With these variables, we can express the reverse conditional score from the causal parameters.

Lemma C.4 (Anticausal conditional score). *Let s_y, s_x be marginal scores, and $s_{y|x}, s_{x|y}$ be conditional scores. Then*

$$\boxed{s_{x|y} = s_x + (s_{y|x} - m(y)) - (A(x) - \alpha)}, \quad (34)$$

where $\alpha = \frac{1}{K} \sum_x A(x)$.

Proof. Let's apply Bayes rule to find the conditional probability mass function

$$\begin{aligned} \mathbf{p}(x|y) &\propto \mathbf{p}(y|x)\mathbf{p}(x) \\ &\propto \exp(s_{y|x} - A(x) + s_x) \end{aligned}$$

where $A(x)$ is the log-partition function of $\mathbf{p}(y|x)$. Taking the logarithm,

$$s_{x|y} = s_x + s_{y|x} - A(x) + C(y) \quad (35)$$

where $C(y)$ is a constant defined such that $\sum_x s_{x|y} = 0$ (see Definition C.1). We take the sum of (35) over x to find

$$\begin{aligned} &\sum_x s_x + \sum_x s_{y|x} - \sum_x A(x) + KC(y) \\ = &0 + Km(y) + K\alpha + KC(y) \end{aligned}$$

which simplifies into

$$C(y) = -m(y) - \alpha$$

We plug this in (35) to conclude the proof. \square

We conclude this section with an identity showing that conditional logits are equally close from their averages in both directions.

Lemma C.5. *For any x and y we have*

$$\boxed{s_{x|y} - n(x) = s_{y|x} - m(y)}. \quad (36)$$

where $n(x) := \frac{1}{K} \sum_y s_{x|y}$ and $m(y) := \frac{1}{K} \sum_x s_{y|x}$.

Proof. We apply Lemma C.4 with the roles of X and Y inverted to express $s_{y|x}$ as a function of the anti causal parameters

$$s_{y|x} = s_y + (s_{x|y} - n(x)) - (B(y) - \beta), \quad (37)$$

where $B(y) = \log \sum_x e^{s_{y|x}}$ and $\beta = \frac{1}{K} \sum_y B(y)$. We can add (34) and (37) to get rid of the conditional scores

$$s_x - n(x) - A(x) + \alpha = -(s_y - m(y) - B(y) + \beta),$$

for all x, y . The left hand side is constant in y whereas the right hand side is constant in x . Thus both sides are constants with respect to both x and y . In particular they are equal to their average

$$\begin{aligned} \forall x, s_x - n(x) - A(x) + \alpha &= \frac{1}{K} \sum_{x'} (s_{x'} - n(x')) \\ &\quad - \frac{1}{K} \sum_{x'} A(x') + \alpha \\ &= 0 - 0 - \alpha + \alpha \\ &= 0. \end{aligned}$$

We plug this equality into (34) to prove the lemma. \square

C.2 Intervention on Cause

In this section, we analyze the relationship between δ_{causal} and $\delta_{\text{anticausal}}$ after an intervention on the cause.

Proposition 1. *If an intervention happens on the cause X then we have*

$$\boxed{\delta_{\text{anticausal}} \geq K \delta_{\text{causal}}}, \quad (38)$$

where $\delta_{\text{causal}} = \|\mathbf{s}_X - \mathbf{s}_X^*\|^2$, and $\delta_{\text{anticausal}} = \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 + \sum_y \|\mathbf{s}_{X|y} - \mathbf{s}_{X|y}^*\|^2$

Proof. Given that $s_{y|x}^* = s_{y|x}$, $\mathbf{m}^* = \mathbf{m}$, $A^* = A$ and $\alpha^* = \alpha$, Lemma C.4 tells us that the anticausal conditional $\mathbf{s}_{X|Y}^*$ verifies

$$\begin{aligned} s_{x|y}^* - s_x^* &= (s_{y|x} - m(y)) - (A(x) - \alpha) = s_{x|y} - s_x \\ &\implies s_{x|y} - s_{x|y}^* = s_x - s_x^*. \end{aligned}$$

The distance between models before and after intervention are

$$\begin{aligned} \delta_{\text{causal}} &= \|\mathbf{s}_X - \mathbf{s}_X^*\|^2 \\ \delta_{\text{anticausal}} &= \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 + \sum_y \|\mathbf{s}_{X|y} - \mathbf{s}_{X|y}^*\|^2 \\ &\geq 0 + \sum_y \sum_x (s_x - s_x^*)^2 = K \|\mathbf{s}_X - \mathbf{s}_X^*\|^2. \end{aligned}$$

In conclusion,

$$\delta_{\text{anticausal}} \geq K \delta_{\text{causal}}. \quad (39)$$

\square


C.3 Intervention on Effect

The following proposition shows that when the intervention is on the effect, the causal model is advantaged only when the new effect marginal \mathbf{s}_Y^* is close enough from the previous marginal.

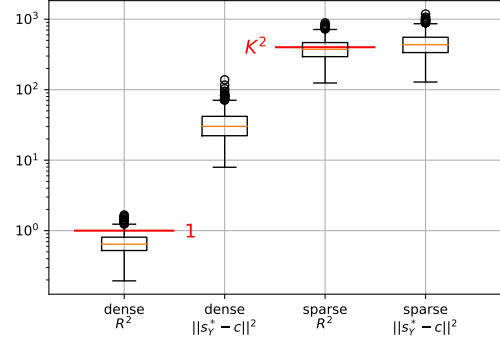
Proposition 2 *When an intervention happens on the effect*

$$\Delta := \delta_{\text{causal}} - \delta_{\text{anticausal}} \quad (40)$$

$$= (K - 1) \left(\|\mathbf{s}_Y^* - \mathbf{c}\|^2 - R^2 \right) \quad (41)$$



(a) \mathbf{m} is a convex combination of \mathbf{c} and \mathbf{s}_Y . The blue bubble is the sub-level set 0 of Δ . It is a circle of radius R centered at \mathbf{c} . Within this circle, $\delta_{\text{causal}} \leq \delta_{\text{anticausal}}$ the causal model is advantaged. Outside this circle, $\delta_{\text{causal}} \geq \delta_{\text{anticausal}}$ the anticausal model is advantaged.



(b) Box plots for the radius R^2 and deviations $\|\mathbf{s}_Y^* - \mathbf{c}\|^2$ for $K = 20$ with the dense prior (left) and the sparse prior (right). The y-axis is logarithmic. Red lines show analytical estimates for the expected radius.

Figure 7: Behavior of the hyper-sphere within which the causal model is advantaged, as presented in Proposition 2.

where the score vector \mathbf{c} and the scalar R are defined as

$$\mathbf{c} = \frac{K\mathbf{m} - \mathbf{s}_Y}{K - 1} \quad (42)$$

$$(K - 1)R^2 = K\|\mathbf{n} - \mathbf{s}_X\|^2 + (K - 1)\|\mathbf{c}\|^2 + \|\mathbf{s}_Y\|^2 - K\|\mathbf{m}\|^2 \quad (43)$$

with \mathbf{m} and \mathbf{n} as in Definition C.2.

We illustrate the relationship between \mathbf{m} , \mathbf{c} , \mathbf{s}_Y and R in Figure 7a.

Proof. First we expand the causal distance with a bias variance decomposition

$$\begin{aligned} \delta_{\text{causal}} &= \sum_x \|\mathbf{s}_{Y|x} - \mathbf{s}_Y^*\|^2 \\ &= \sum_{x,y} (s_{y|x} - m(y) + m(y) - s_y^*)^2 \\ &= \sum_{x,y} (s_{y|x} - m(y))^2 + K \sum_y (m(y) - s_y^*)^2 \\ &\quad + 2 \sum_y (m(y) - s_y^*) \underbrace{\sum_x (s_{y|x} - m(y))}_{=0} . \end{aligned} \quad (44)$$

Given that $\mathbf{s}_{X|y}^* = \mathbf{s}_X$, we can decompose the anticausal distance similarly

$$\begin{aligned} \delta_{\text{anticausal}} &= \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 + \sum_y \|\mathbf{s}_{X|y} - \mathbf{s}_{X|y}^*\|^2 \\ &= \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 + \sum_{x,y} (s_{x|y} - n(x) + n(x) - s_x)^2 \\ &= \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 + \sum_{x,y} (s_{x|y} - n(x))^2 \\ &\quad + K \sum_x (n(x) - s_x)^2 . \end{aligned} \quad (45)$$

Thanks to Lemma C.5, the variance of conditional score vectors (in blue) in (44) and (45) are equal

$$\sum_{x,y} (s_{x|y} - n(x))^2 = \sum_{x,y} (s_{y|x} - m(y))^2 .$$

What remains in the difference is the quadratic form

$$\begin{aligned} \Delta &= \delta_{\text{causal}} - \delta_{\text{anticausal}} \\ &= K \|\mathbf{m} - \mathbf{s}_Y^*\|^2 - K \|\mathbf{n} - \mathbf{s}_X\|^2 - \|\mathbf{s}_Y - \mathbf{s}_Y^*\|^2 , \end{aligned}$$

which we can expand to highlight the role of \mathbf{s}_Y^* as

$$\begin{aligned} \Delta &= (K-1) \|\mathbf{s}_Y^*\|^2 - 2 \langle \mathbf{s}_Y^*, K\mathbf{m} - \mathbf{s}_Y \rangle \\ &\quad + K \|\mathbf{m}\|^2 - \|\mathbf{s}_Y\|^2 - K \|\mathbf{n} - \mathbf{s}_X\|^2 \\ &= (K-1) \|\mathbf{s}_Y^* - \mathbf{c}\|^2 - (K-1) \|\mathbf{c}\|^2 \\ &\quad + K \|\mathbf{m}\|^2 - \|\mathbf{s}_Y\|^2 - K \|\mathbf{n} - \mathbf{s}_X\|^2 \end{aligned}$$

where \mathbf{c} appears as a non-convex interpolation of \mathbf{m} and \mathbf{s}_Y , $\mathbf{c} = \frac{K\mathbf{m} - \mathbf{s}_Y}{K-1}$. Define R^2 to conclude the proof. \square

Empirical estimates of the radius. We report values of R^2 and $\|\mathbf{s}_Y^* - \mathbf{c}\|^2$ observed for the dense and sparse priors in Figure 7b. For dense prior radii are much smaller than deviations, whereas for the sparse prior they have similar magnitude. This explains why the anticausal model systematically adapts faster when the intervention is on the effect and the prior is dense. We also observe that radii (and deviations) are much greater for the sparse prior than for the dense prior. In the following paragraph we provide some clues to explain this behaviour.

As illustrated by Figure 7a, \mathbf{m} is a convex combination of \mathbf{c} and \mathbf{s}_Y : $\mathbf{m} = \frac{(K-1)\mathbf{c} + \mathbf{s}_Y}{K}$ so by convexity of $\|\cdot\|^2$,

$$\begin{aligned} \frac{K-1}{K} \|\mathbf{c}\|^2 + \frac{1}{K} \|\mathbf{s}_Y\|^2 &\geq \|\mathbf{m}\|^2 \\ \implies (K-1) \|\mathbf{c}\|^2 + \|\mathbf{s}_Y\|^2 - K \|\mathbf{m}\|^2 &\geq 0 \\ \implies R^2 &\geq \|\mathbf{n} - \mathbf{s}_X\|^2 . \end{aligned}$$

As K grows larger, this inequality will get closer and closer to an equality. Indeed, \mathbf{m} will get closer and closer to \mathbf{c} and we will end up with

$$\frac{K-1}{K} \|\mathbf{c}\|^2 + \frac{1}{K} \|\mathbf{s}_Y\|^2 - \|\mathbf{m}\|^2 \ll \|\mathbf{n} - \mathbf{s}_X\|^2 \approx R^2 .$$

Before proceeding, let us prove a simple proposition that is a direct consequence of Lemma C.4.

Proposition 3. *The squared distance between marginal score and reverse average conditional is equal to the empirical variance of the conditional log-partition function*

$$\|\mathbf{s}_X - \mathbf{n}\|^2 = K \widehat{\text{Var}}_X[A(X)] . \quad (46)$$

Proof. Taking the average of Lemma C.4 over y yields

$$\begin{aligned} \frac{1}{K} \sum_y s_{x|y} &= s_x - (A(x) - \alpha) + \frac{1}{K} \sum_y (s_{y|x} - m(y)) \\ n(x) &= s_x - (A(x) - \alpha) , \end{aligned}$$

where we used the definition of $n(x)$ and the mean-zero scores. Reordering terms gives

$$s_x - n(x) = A(x) - \alpha \quad (47)$$

Recall that α is the average of $A(x)$ over x . Squaring this equation and summing over x concludes the proof. \square

Using this proposition we get that

$$\begin{aligned} R^2 &\approx \|\mathbf{s}_X - \mathbf{n}\|^2 \\ &= K \widehat{\text{Var}}_X[A(X)] \\ &= K \widehat{\text{Var}}_X[\log \sum_y e^{s_{y|X}}]. \end{aligned}$$

Conditional scores $s_{y|x}$ are taking much greater values with much higher variance under the sparse prior than under the dense prior. To be clear, we sample independently pseudo-scores $\tilde{s}_{y|x}$ from exp-gamma laws, and we subtract their mean to ensure that they sum to 0 (Definition C.1). $s_{y|x} = \tilde{s}_{y|x} - \frac{1}{K} \sum_{y'} \tilde{s}_{y'|x}$. This means that

$$\log \sum_y e^{s_{y|X}} = \log \sum_y e^{\tilde{s}_{y|X}} - \frac{1}{K} \sum_y \tilde{s}_{y|x}$$

If we make the approximation that the logsumexp term and the average term are independent then

$$\begin{aligned} &\widehat{\text{Var}}_X[\log \sum_y e^{s_{y|X}}] \\ &\approx \text{Var}_{s_{y|x}}[\log \sum_y e^{s_{y|x}}] \\ &\approx \text{Var}_{\tilde{s}_{y|x}}[\log \sum_y e^{\tilde{s}_{y|x}}] + \frac{1}{K} \text{Var}_{\tilde{s}_{y|x}}[\sum_y \tilde{s}_{y|x}] \\ &= \psi^{(1)}(K\lambda) + \frac{1}{K} \psi^{(1)}(\lambda) \end{aligned}$$

where this last step uses the formula for the variance of an exp-gamma variable twice. The variance of an exponential gamma with shape parameter λ is $\psi^{(1)}(\lambda)$ where $\psi^{(1)}$ is the trigamma function. The log-sum-exp of K independent exp-gamma with scale and shape parameters (λ, ζ) is another exp-gamma with scale and shape parameters $(K\lambda, \zeta)$. Finally we get the following approximation for the squared radius

$$R^2 \approx K\psi^{(1)}(K\lambda) + \psi^{(1)}(\lambda).$$

The dense prior uses a shape parameter $\lambda = 1$ while the sparse prior uses a shape parameter $\lambda = \frac{1}{K}$. We use two approximations of the trigamma function: $\psi^{(1)}(\frac{1}{K}) \approx K^2 + \pi^2/6$ and $\psi^{(1)}(K) \approx \frac{1}{K}$ when $K \geq 10$.

$$\begin{aligned} R_{\text{dense}}^2 &\approx K\psi^{(1)}(K) + \psi^{(1)}(1) \\ &\approx 1 + \pi^2/6 = O(1) \\ R_{\text{sparse}}^2 &\approx K\psi^{(1)}(1) + \psi^{(1)}(\frac{1}{K}) \\ &\approx K^2 + K\pi^2/6 + \pi^2/6 = O(K^2). \end{aligned}$$

In other words for dense prior the radius grows linearly with the dimension K . We report these estimates along with real data in Figure 7b.

Independent Special Case. if X is independent of Y in the reference distribution – e.g. $\forall x, y, \mathbf{p}(x, y) = \mathbf{p}(x)\mathbf{p}(y)$ – then $\forall x, y$,

$$\begin{aligned} s_x &= s_{x|y} = n(x) \\ s_y &= s_{y|x} = m(y) \end{aligned}$$

Plugging these equalities into Proposition 2 yields

$$\begin{aligned} \mathbf{c} &= \mathbf{m} = \mathbf{s}_Y \quad \text{and} \quad R = 0 \\ \implies \Delta &= (K-1)\|\mathbf{s}_Y^* - \mathbf{s}_Y\|^2 \geq 0 \end{aligned}$$

which means that the anticausal model is advantaged $\delta_{\text{causal}} \geq \delta_{\text{anticausal}}$. This is actually predictable from a simple parameter counting argument. When the reference distribution is made of independent distributions, the anticausal conditional mechanism is already optimal $s_x = s_{x|y}$. The anticausal model only has to adapt its marginal mechanism \mathbf{s}_Y of size K . On contrary, the causal model only has to adapt its conditional mechanism $s_{y|x} \neq s_y$ of size K^2 . Overall the causal model has to adapt K times more parameters than the anticausal model.

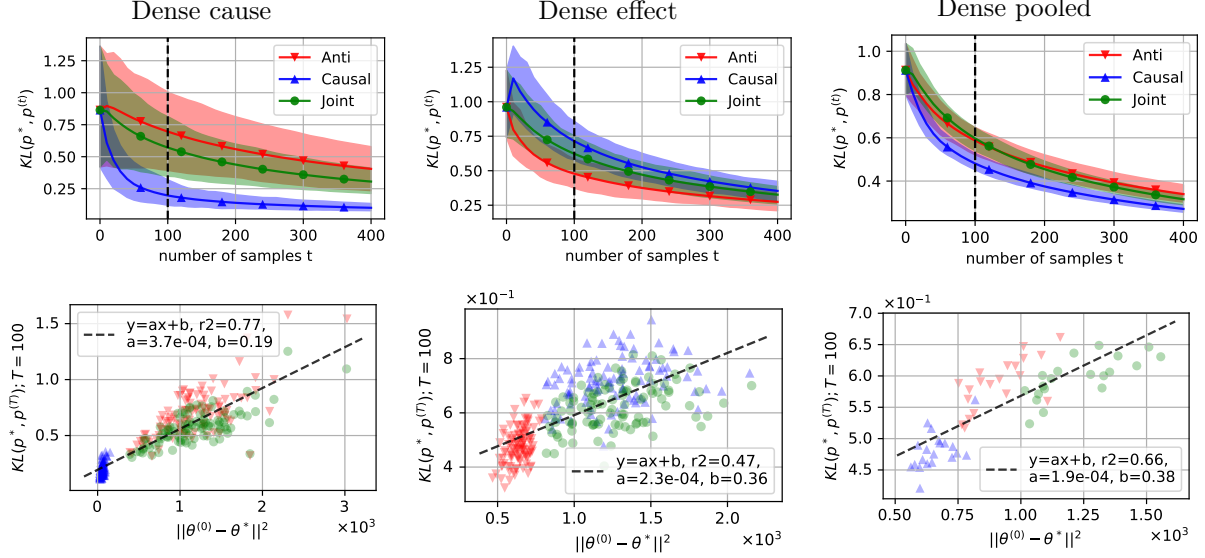


Figure 8: **Categorical dense prior with $K=20$.** Row 1: training curves. Solid lines are average KL over 100 runs. We tune hyper-parameters to minimize the average KL of each model at the black vertical dashed bar ($t=100$). Shaded areas are between (5,95) quantiles. Note that *all models start from the same initial KL, but they converge at different speeds*. Row 2: scatter plot of the KL at $t=100$ vs. initial distance. Note that the initial distance is well correlated with the KL after 100 steps of SGD. Columns: we report results for interventions on the cause on column 1, the effect on column 2, and an aggregation of both on column 3. We aggregate results by taking the average of 5 cause interventions and 5 effect interventions as one new trajectory. In total we have 20 such trajectories per model. We are reporting this result because the meta-learning criterion suggested by Bengio et al. (2020) is akin to the average adaptation speed over a small set of interventions.

C.4 Other Empirical Results for Cause and Effect Interventions

In this section, we present additional results for categorical variables. In Figure 8, compared to the main text, we add what happens with the dense prior when we average learning curves (pooled) from 5 interventions on the cause and 5 interventions on the effect : on average the causal model adapts the fastest. In Figure 9, compared to the main text we show what happens with the sparse prior, both in terms of distance (scatter plots) and in terms of pooled results. Because the intervention on the effect creates huge values of the KL, there is no set advantage for any of the models.

C.5 Single Mechanism Intervention

If only $\mathbf{s}_{Y|x_0}$ changes, for some x_0 , then from Lemma C.4 we get the following equality

$$\begin{aligned} \delta_{\text{anticausal}} = & \frac{K-1}{K} \delta_{\text{causal}} + \|\mathbf{s}_Y^* - \mathbf{s}_Y\|^2 \\ & + (K-1)(A^*(x_0) - A(x_0))^2. \end{aligned} \quad (48)$$

The causal and anticausal distances seem to be on the same scale, with a multiplicative factor $\frac{K-1}{K} \lesssim 1$ and a positive additive factor. This is interesting because the sparsity argument holds: the causal model needs to change K parameters whereas the anticausal model needs to change $K^2 + K$ parameters. That means we could expect an advantage by a factor K for the causal model, similarly to when the intervention is on the cause. However (48) tells another story: without further assumptions, it seems like both distances will have the same scale.

Experiments. For this kind of intervention to be detectable, we need to intervene on x_0 such that $\mathbf{p}(x_0)$ is quite large. To ensure this in our experiments, we pick $x_0 = \text{argmax}_x \mathbf{p}(x)$. We report results on dense and sparse priors in Figure 10. We observe no significant advantage for the causal model, in spite of the parameter counting prediction.

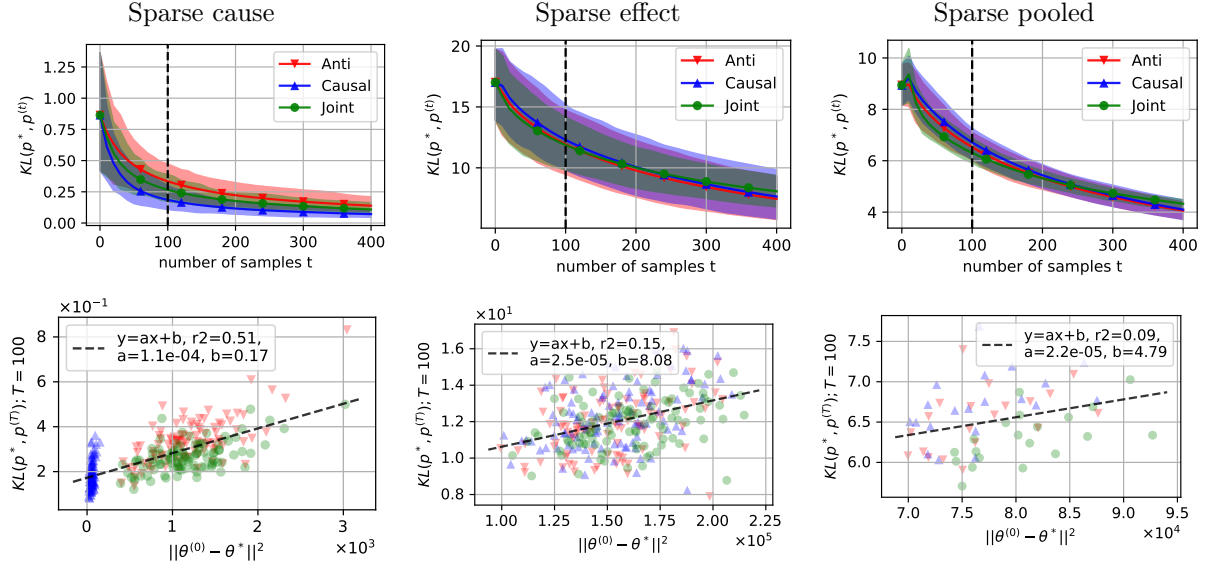


Figure 9: **Categorical sparse prior with $K=20$.** *Column 1:* intervention on the cause. The causal model starts closer from optimum and adapts slightly faster than others. *Column 2:* intervention on the effect. All models have the same initial distance and the same objective value. However the KL value is around 10. This is 10 times larger than when the intervention is on the cause. *Column 3:* we take the average of 5 effect and 5 cause interventions. The effect dominates this average because it is much larger. As a result there is no signal.

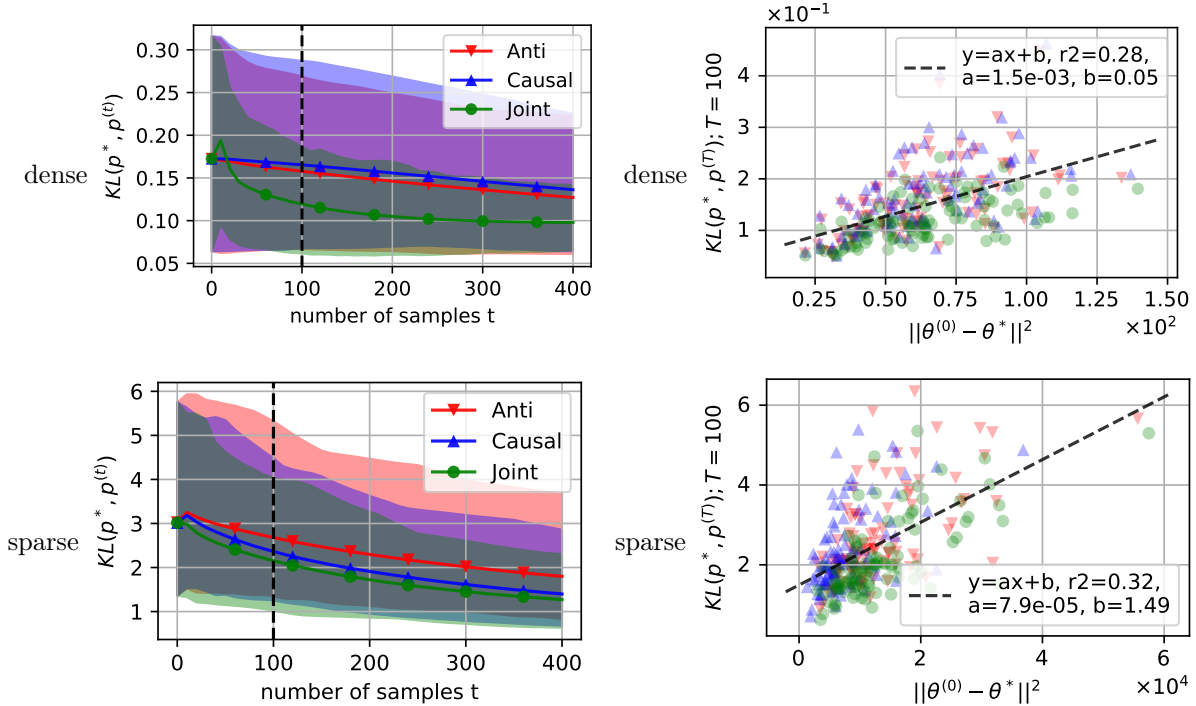


Figure 10: **Single mechanism intervention with $K=20$.** *Two first rows:* Dense prior. The only model to be slightly advantaged is the joint model. *Two last rows:* Sparse prior. This time there is a slight advantage for the causal model which performs comparably to the joint model. Overall the optimization is hard in both settings, since we are observing only K^2 samples for models with $O(K^2)$ parameters. The KL barely decreases.

K	2	3	4	5	6	7	8	9	10	11	12	13	14
error	.4	.2	.1	.07	.03	.01	.005	.002	.0007	.0003	7e-5	3e-5	4e-6

Table 1: Estimation of the Bayes error under the dense prior assumption for increasing categorical variables dimension K . We estimated these numbers by sampling one million joint distributions for each K . We report 1 significant figure.

D CATEGORICAL PRIORS

In this section we study the dense and sparse prior described in the main paper.

D.1 Causal Direction is Identifiable under the Dense Prior

Chalupka et al. (2016) study the prediction of causal direction from observational data under the dense prior assumption. The causal direction $X \rightarrow Y$ induces a certain prior over joint distributions $\pi(\mathbf{p} | \rightarrow)$. The anticausal direction $X \leftarrow Y$ induces another one $\pi(\mathbf{p} | \leftarrow)$. The Bayes classifier is predicting \rightarrow if

$$\log \pi(\rightarrow | \mathbf{p}) - \log \pi(\leftarrow | \mathbf{p}) > 0, \quad (49)$$

and \leftarrow otherwise. Under the dense prior assumption, this classifier makes an error of approximately 0.4 for $K = 2$, which decreases exponentially to 0.001 for $K = 10$. We reproduced their setting and report the error of the optimal classifier in Table 1 for varying K . In other words the dense prior induce very asymmetric distributions which makes the causal direction identifiable.

Is this Bayes classifier easy to estimate ? It turns out that under the dense prior, the criterion (49) can be simplified into the following criterion

$$D_{\text{KL}}(\mathbf{u} | \mathbf{p}_X) - D_{\text{KL}}(\mathbf{u} | \mathbf{p}_Y) > 0 \quad (50)$$

where $\mathbf{u} := \mathbf{1}/K$ is the uniform probability vector. The proof is left as an exercise to the reader. If (50) is positive, Bayes predicts that the cause is X , otherwise Y is the cause. In words, whichever variable has the most uniform marginal is the effect. This simple rule is optimal given the prior assumption (and if both directions are equally likely). We can understand it from a concentration of measure perspective. The effect marginal is written as a sum of quasi independent uniform variables

$$\mathbf{p}(y) = \sum_x \mathbf{p}(y|x) \mathbf{p}(x) \quad (51)$$

which ends up close from the uniform vector.

D.2 Joint Distribution with Sparse Prior

The following theorem shows how a Dirichlet prior over joint distributions $c_{x,y} = \mathbf{p}(x,y)$ is equal to independent Dirichlet priors over marginal $a_x = \mathbf{p}(x)$ and conditional $b_{y|x} = \mathbf{p}(y|x)$ probability mass functions. By applying this theorem, we find that the sparse prior is equivalent to $\text{Dir}(\frac{1}{K} \mathbf{1}_{K^2})$.

Theorem D.1 (Dirichlet and Factorization). *Let \mathbf{c} be a random square matrix of dimension K . Let's define \mathbf{a} as the random vector obtained by summing columns of \mathbf{c} , and \mathbf{b} as a copy of \mathbf{c} with rows normalized so that they sum to 1. $\forall (i, j) \in \{1, \dots, K\}^2$*

$$a_i = \sum_j c_{i,j} \quad (52)$$

$$b_{j|i} = \frac{c_{i,j}}{a_i}. \quad (53)$$

Let γ be a positive square matrix of parameters. The following equivalence holds

$$\mathbf{c} \sim \text{Dir}(\gamma) \iff \begin{cases} \mathbf{a} \sim \text{Dir}(\sum_i \gamma_i) \\ \mathbf{b}_{:,|i} \sim \text{Dir}(\gamma_{i,:}), \forall i \\ \mathbf{a} \perp \mathbf{b}_{:,|i} \perp \mathbf{b}_{:,|i'}, \forall i \neq i' \end{cases}. \quad (54)$$

Proof. First let's remark that the right side of (54) is entirely characterizing the joint distribution on (\mathbf{a}, \mathbf{b}) , and that the relationship between \mathbf{c} and (\mathbf{a}, \mathbf{b}) is a bijection with reverse $c_{i,j} = b_{j|i} a_i$. This means that the equivalence (54) is an equality between distributions. This means that we can prove the forward implication and the converse will hold automatically.

If $\mathbf{c} \sim \text{Dir}(\gamma)$, then there exist K^2 independent Gamma variables $\tilde{c}_{i,j} \sim \Gamma(\gamma_{i,j}, 1), \forall i, j$ such that

$$\mathbf{c} = \frac{\tilde{\mathbf{c}}}{S} \quad \text{where} \quad S = \sum_{i,j} \tilde{c}_{i,j}. \quad (55)$$

We know from properties of the Gamma distribution that \mathbf{c} is independent of S . Now let's define $\tilde{a}_i := \sum_j \tilde{c}_{i,j}$. This definition has three consequences. First \tilde{a}_i is a sum of independent gammas, so it is a gamma with parameters $(\alpha_i := \sum_j \gamma_{i,j}, 1)$. Second $S = \sum_i \tilde{a}_i$. Third $\mathbf{a} = \frac{\tilde{\mathbf{a}}}{S}$ is independent of S and is a Dirichlet with parameter vector $\boldsymbol{\alpha} = \sum_j \boldsymbol{\gamma}_{:,j}$.

That was for the marginal. Now for the conditional,

$$b_{j|i} = \frac{c_{i,j}}{a_i} = \frac{\tilde{c}_{i,j}/S}{\tilde{a}_i/S} = \frac{\tilde{c}_{i,j}}{\tilde{a}_i}. \quad (56)$$

Again from properties of the Gamma distribution, $\mathbf{b}_{:,|i} \perp \tilde{a}_i, \forall i$, and $\mathbf{b}_{:,|i} \sim \text{Dir}(\boldsymbol{\gamma}_{i,:})$. Each of the conditional $\mathbf{b}_{:,|i}$ is defined with independent gammas, so we also have the independence between conditionals. We verified all properties of the right side of (54), which concludes the proof. \square

D.3 Categorical Sparse Prior Explosion

In this section we explain why the KL takes large values with sparse prior and effect intervention.

On one hand the sparse prior samples probability vectors which are close from being Dirac. On the other hand the effect intervention creates an outer product between two samples drawn uniformly from the simplex $\text{Dir}(\mathbf{1})$.

For instance, for the uniform probability vector $\mathbf{u} = \frac{1}{K^2} \mathbf{1} \in \boldsymbol{\Delta}_{K^2}$ and an almost Dirac $\mathbf{p} = (1 - \varepsilon) \mathbf{e}_1 + \varepsilon \mathbf{u}$

$$D_{\text{KL}}(\mathbf{u}||\mathbf{p}) \in \Theta(\log(\frac{1}{\varepsilon}))$$

where ε is a small value. As we increase K , the sparse prior $\text{Dir}(\mathbf{1}_{K^2}/K)$ becomes more sparse. Conceptually, the value of ε decreases, and the value of $D_{\text{KL}}(\mathbf{u}||\mathbf{p})$ explodes. This is why we observe high KL values for sparse prior and effect intervention. Empirically, these values also increase with K .

E NORMAL OPTIMIZATION

In this Section we adapt the stochastic composite mirror-prox algorithm to our setting of unbounded multivariate normal optimization. First we describe the algorithm and prove a novel convergence rate that applies to our setting. Then we explicit the update formulas for the normal log-likelihood loss with Cholesky parameters. Finally we prove that worst case constants appearing in the rate are equal for both causal and anticausal models.

E.1 Stochastic Composite Mirror-Prox

We want to minimize the composite objective

$$F(\theta) = \mathbb{E}_i [f(\theta, i)] + g(\theta).$$

For simplicity we denote $f(\theta, i)$ by $f_i(\theta)$ and $f(\theta) = \mathbb{E}_i [f_i(\theta)]$. We assume that f_i is convex, ∇f_i is L -Lipschitz and g is a convex function. The stochastic mirror-prox algorithm update rule at time t is

$$\nu_t = \theta_{t+1} - \gamma_t f'_i(\theta_t) \quad (57)$$

$$\theta_{t+1} = \underset{\theta}{\text{argmin}} \left\{ g(\theta) + \frac{1}{\gamma_t} \mathcal{B}_h(\theta, \nu_t) \right\} \quad (58)$$

where f_i is sampled randomly. $B_h(x, y) = h(x) - h(y) - \langle h'(y), x - y \rangle$ denotes the Bregman divergence between x and y induced by the convex function h and we have $\|B_h(x, y)\| \geq \frac{\alpha}{2}\|x - y\|^2$. When we set $h(x) = 1/2\|x\|^2$, we recover something called the proximal stochastic gradient method (Duchi and Singer, 2009), also known as Perturbed proximal gradient algorithm (Atchadé et al., 2017). This last citation in particular has hypothesis very close to ours.

E.1.1 Convergence Rate

The following Theorem is a mild modification of the Theorem 8 in (Duchi et al., 2010). Our result is different in 2 ways. First, we remove the boundedness assumption for the Bregman divergence throughout the trajectory i.e. $\mathcal{B}_h(\theta^*, \theta_t) \leq D$ for all t . Second, we replace the f_i B -Lipschitz continuous assumption by ∇f_i B -Lipschitz continuous. We need this last modification for the result to hold on f_i quadratic.

First we prove the following lemma which is a modification of Lemma 1 in (Duchi et al., 2010).

Lemma E.1. *With f convex and B -smooth, g convex, and $\gamma \leq \frac{\alpha}{3B}$, at iteration t , if we sample i , we have:*

$$\begin{aligned} \gamma \{f_i(\theta_t) + g(\theta_{t+1}) - F(\theta^*)\} &\leq \mathcal{B}_h(\theta^*, \theta_t) - \mathcal{B}_h(\theta^*, \theta_{t+1}) \\ &\quad + \gamma \{f_i(\theta_t) - f_i(\theta_{t+1})\}. \end{aligned}$$

Proof. We have the following sequence of inequality

$$\begin{aligned} &\gamma \left(f_i(\theta_t) + g(\theta_{t+1}) - F(\theta^*) \right) \\ &\leq \gamma \langle \theta_t - \theta^*, f'_i(\theta_t) \rangle + \gamma \langle \theta_{t+1} - \theta^*, \partial g_t(\theta_t) \rangle \\ &\leq \mathcal{B}_h(\theta^*, \theta_t) - \mathcal{B}_h(\theta^*, \theta_{t+1}) \\ &\quad - \mathcal{B}_h(\theta_{t+1}, \theta_t) + \gamma \langle \theta_t - \theta_{t+1}, f'_i(\theta_t) \rangle \\ &\leq \mathcal{B}_h(\theta^*, \theta_t) - \mathcal{B}_h(\theta^*, \theta_{t+1}) \\ &\quad - \mathcal{B}_h(\theta_{t+1}, \theta_t) + \frac{B\gamma}{2} \|\theta_t - \theta_{t+1}\|^2 \\ &\quad + \gamma \left(f_i(\theta_t) - f_i(\theta_{t+1}) \right) \end{aligned}$$

where the first inequality comes from convexity of f_i and g , the second inequality comes from Eq.(6) of lemma 1 in (Duchi et al., 2010), and the third inequality comes from the smoothness of f_i . By $\gamma \leq \frac{\alpha}{3L}$, the term $-\mathcal{B}_h(\theta_{t+1}, \theta_t) + \frac{B\gamma}{2} \|\theta_t - \theta_{t+1}\|^2$ is negative : we can drop this term and get the required result. Note that ∂g is a subgradient of g . \square

Theorem E.2. *Given the above assumptions for f_i and g , after T iterations of the stochastic mirror prox algorithm with $\gamma = \frac{c}{\sqrt{T}}$, ($c \leq \frac{\alpha}{3B}$), we have*

$$\mathbb{E} [F(\bar{\theta}) - F(\theta^*)] \leq \frac{\mathcal{B}_h(\theta^*, \theta_0)}{c\sqrt{T}} + \frac{(F(\theta_0) - F(\theta^*))}{T}.$$

Proof. The proof is similar to the proof of the Theorem 8 in (Duchi et al., 2010) with some modifications. Take the expectation of lemma E.1 with respect to the samples $(i_u)_{u \leq t}$

$$\begin{aligned} &\mathbb{E} \left[\gamma \left(f(\theta_t) + g(\theta_{t+1}) - F(\theta^*) \right) \right] \\ &\leq \mathbb{E} \left[\mathcal{B}_h(\theta^*, \theta_t) - \mathcal{B}_h(\theta^*, \theta_{t+1}) + \gamma \left(f(\theta_t) - f(\theta_{t+1}) \right) \right] \end{aligned}$$

where θ_t and θ_{t+1} are random variable that depends on the samples. Sum up both side for T iterations:

$$\begin{aligned} &\gamma \sum_{t=0}^T \mathbb{E} [f(\theta_t) + g(\theta_{t+1}) - F(\theta^*)] \\ &\leq \mathcal{B}_h(\theta^*, \theta_0) - \mathbb{E} [\mathcal{B}_h(\theta^*, \theta_{T+1})] \\ &\quad + \gamma \left(f(\theta_0) - \mathbb{E} [f(\theta_{T+1})] \right) \end{aligned}$$

By adding $\gamma \mathbb{E} [g(\theta_0) - g(\theta_{t+1})]$ to both sides of the above inequality we get:

$$\begin{aligned} & \gamma \sum_{t=1}^T \mathbb{E} [F(\theta_t) - F(\theta^*)] \\ & \leq \mathcal{B}_h(\theta^*, \theta_0) - \mathbb{E} [\mathcal{B}_h(\theta^*, \theta_{T+1})] \\ & + \gamma \left(F(\theta_0) - \mathbb{E} [F(\theta_{T+1})] \right) \\ & \leq \mathcal{B}_h(\theta^*, \theta_0) + \gamma \left(F(\theta_0) - \mathbb{E} [F(\theta^*)] \right) \end{aligned}$$

where the last inequality is due to non-negativity of Bregman divergence and optimality of θ^* . Divide both sides by $\gamma T = c\sqrt{T}$ and use Jensen inequality on F to conclude the proof. \square

E.2 Normal Model Updates

The objective function at hand is:

$$\begin{aligned} F(L, \zeta) &= f(L, \zeta) + g(L) \\ f(L, \zeta) &= \frac{1}{2n} \sum_{i=1}^n \|L^T x_i - \zeta\|^2 \\ g(L) &= -\ln(|L|). \end{aligned}$$

Now the update rule for the ζ given that g is independent of ζ and we sample mini-batch B of size m :

$$\zeta_{t+1} = (1 - \gamma)\zeta_t + \gamma L_t^T \left(\frac{1}{m} \sum_{i \in B} x_i \right)$$

For the L the gradient update gives

$$L_{t+\frac{1}{2}} = L_t - \gamma \frac{1}{m} \sum_{i \in B} (x_i x_i^T L_t - x_i \zeta_t^T).$$

Since L is lower triangular, $g(L) = \ln(|L|) = \sum_{i=1}^d \log L_{i,i}$ and the proximal operator only applies to diagonal elements of L – e.g. when $i \neq j$ $[L_{t+1}]_{(i,j)} = [L_{t+\frac{1}{2}}]_{(i,j)}$. Otherwise we have to compute:

$$[L_{t+1}]_{(i,i)} = \operatorname{argmin}_{L_{ii}} \left\{ -\ln(L_{ii}) + \frac{1}{2\gamma} \|L_{ii} - [L_{t+\frac{1}{2}}]_{(i,i)}\|^2 \right\}.$$

Therefore the update rule for the $[L_{t+1}]_{(i,i)}$ is:

$$[L_{t+1}]_{(i,i)} = \frac{1}{2} \left\{ [L_{t+\frac{1}{2}}]_{(i,i)} + \sqrt{[L_{t+\frac{1}{2}}]_{(i,i)}^2 + 4\gamma} \right\}.$$

Remark how this proximal operator behaves as a smooth projection on the set of strictly positive numbers. If the diagonal is negative after the gradient update, it brings it to a small positive value. If it was already positive, it slightly increases its value.

E.3 Equality of Smoothness Constants

In this section, we show that the Lipschitz smoothness parameter B , which appears in the convergence rate (22), is the same for both causal and anticausal models. Similarly to the categorical case, we reason about marginals loss first because they have a simpler form.

The loss of a marginal mechanism is $f_x(L, \zeta) = \frac{1}{2} \sum_{i=1}^d (\zeta_i - L_i^T x)^2$ where x is a sample observation and the L_i are the columns of L . We need to show that its Hessian is upper-bounded $\|\nabla^2 f_x(L, \zeta)\| \leq B$. Thanks to the

objective f_x being quadratic, the Hessian is independent of the parameters (L, ζ) . It depends only on the data x . Since the data domain is *a priori* the same for causal and anticausal models – e.g. X and Y can live in the same range – the upper bound for the Hessian is the same. This holds true at least for marginal mechanisms, because their loss is written exactly like above.

For conditional mechanisms, this is a bit more complicated but the reasoning holds. The objective is similar, with extra parameters coming from the linear relationship between X and Y . The Cholesky parametrization is described in equation (68). The conditional model uses $\zeta_{Y|X} = MX + m$ where M is a matrix, m is a vector and X is a given sample. This means that the objective is still a quadratic and that 2nd order derivatives w.r.t. $L_{Y|X}$ are still independent of the parameters M and m . They depend only on the observed values of X and Y . We assume that these variables have the same domain *a priori*, therefore both models have similar worst case smoothness constants.

F NORMAL ANALYSIS

In this section we introduce three different parametrization of the multivariate normal cause-effect model. The mean parametrization is the most common and intuitive, but it yields a non-convex optimization problem. The natural parametrization yields a convex problem with convergence guarantees, but it has no closed update formulas for our optimization algorithm of choice. The Cholesky parametrization offers both a convex problem and simple updates.

Then we proceed to study how interventions induce distance in parameter space. We prove that in the natural parameter space, an intervention on the cause will create more distance in the anticausal model than in the causal model.

F.1 Mean Parameters

Cause X and effect Y are sampled from the causal model

$$\begin{aligned} X &\sim \mathcal{N}(\mu_X, \Sigma_X) \\ Y|X &\sim \mathcal{N}(AX + a, \Sigma_{Y|X}) \end{aligned}$$

with parameters $(\mu_X, \Sigma_X, A, a, \Sigma_{Y|X})$. All along, we will assume that all normal laws are non-degenerate – e.g. $\Sigma_X > 0, \Sigma_{Y|X} > 0$. We compute the marginal mean and covariance of Y as well as the covariance between X and Y

$$\begin{aligned} \mathbb{E}[Y] &= A\mu_X + a \\ \text{Cov}[Y] &= \Sigma_{Y|X} + A\Sigma_X A^\top \\ \text{Cov}[X, Y] &= \Sigma_X A^\top. \end{aligned}$$

From there we can derive the joint distribution as a function of the causal parameters

$$\begin{pmatrix} X \\ Y \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_X \\ A\mu_X + a \end{pmatrix}, \begin{pmatrix} \Sigma_X & \Sigma_X A^\top \\ A\Sigma_X & \Sigma_{Y|X} + A\Sigma_X A^\top \end{pmatrix}\right). \quad (59)$$

F.2 Natural Parameters

We want the negative log-likelihood objective to be convex, so we are going to use the natural parameters instead of the mean parameters

$$\mathcal{N}(\mu, \Sigma) = \mathcal{N}_{\text{nat}}(\eta = \Sigma^{-1}\mu, \Lambda = \Sigma^{-1}) \quad (60)$$

where we are using \mathcal{N}_{nat} to explicit that this is taking the natural parameters as arguments. Our causal model using natural parameters is:

$$\begin{aligned} X &\sim \mathcal{N}_{\text{nat}}(\eta_X, \Lambda_X) \\ Y|X &\sim \mathcal{N}_{\text{nat}}(BX + b, \Lambda_{Y|X}) \end{aligned} \quad (61)$$

where we get the natural parameters from the mean parameters with formulas

$$\begin{aligned}\Lambda_X &= \Sigma_X^{-1} \\ \Lambda_{Y|X} &= \Sigma_{Y|X}^{-1} \\ \eta_X &= \Lambda_X \mu_X \\ B &= \Lambda_{Y|X} A \\ b &= \Lambda_{Y|X} a\end{aligned}$$

F.2.1 Switching Direction

We want to get the natural parameters of the anticausal model as a function of the causal parameters. To do so we are going to express the natural parameters of the joint, and then we will simply have to swap rows and columns to invert the roles of X and Y . To get the joint precision matrix, we need to invert the joint covariance. We use the Schur complement and the blockwise matrix inversion formulas

$$\begin{aligned}M &= \begin{pmatrix} A & B \\ C & D \end{pmatrix} \\ M/D &:= D - CA^{-1}B \\ M^{-1} &= \begin{pmatrix} A^{-1} + A^{-1}B(M/D)^{-1}CA^{-1} & -A^{-1}B(M/D)^{-1} \\ (M/D)^{-1}CA^{-1} & (M/D)^{-1} \end{pmatrix}\end{aligned}$$

In our case, the Schur complement of Σ with respect to its lower right block Σ_Y is precisely

$$\begin{aligned}M/D &= \Sigma/\Sigma_Y \\ &= \Sigma_{Y|X} + A\Sigma_X A^\top - A\Sigma_X \Sigma_X^{-1} \Sigma_X A^\top \\ &= \Sigma_{Y|X}.\end{aligned}$$

By applying the formula and identifying the natural parameters, we get

$$\Lambda = \begin{pmatrix} \Lambda_X + B^\top \Lambda_{Y|X}^{-1} B & -B^\top \\ -B & \Lambda_{Y|X} \end{pmatrix}$$

To get the first natural parameter, all we have to do is to multiply the joint precision and the joint mean

$$\begin{aligned}\eta &= \Lambda \mu \\ &= \begin{pmatrix} \Lambda_X \mu_X + B^\top \Lambda_{Y|X}^{-1} B \mu_X - B^\top A \mu_X - B^\top a \\ -B \mu_X + \Lambda_{Y|X} A \mu_X + \Lambda_{Y|X} a \end{pmatrix} \\ &= \begin{pmatrix} \eta_X - B^\top \Lambda_{Y|X}^{-1} b \\ b \end{pmatrix}\end{aligned}$$

where $B^\top \Lambda_{Y|X}^{-1} B \mu_X - B^\top A \mu_X$, and $-B \mu_X + \Lambda_{Y|X} A \mu_X$ are zero and we express the other terms with natural parameters. Overall, the joint natural parameters are

$$\mathcal{N}_{\text{nat}} \left(\begin{pmatrix} \eta_X - B^\top \Lambda_{Y|X}^{-1} b \\ b \end{pmatrix}, \begin{pmatrix} \Lambda_X + B^\top \Lambda_{Y|X}^{-1} B & -B^\top \\ -B & \Lambda_{Y|X} \end{pmatrix} \right).$$

From there we can use a symmetry argument to switch from causal $X \rightarrow Y$ to anticausal $X \leftarrow Y$ model.

$$\begin{aligned}Y &\sim \mathcal{N}_{\text{nat}}(\eta_Y, \Lambda_Y) \\ X|Y &\sim \mathcal{N}_{\text{nat}}(CY + c, \Lambda_{X|Y})\end{aligned}$$

with the following formulas for the conditional mechanism

$$C = B^\top \quad (62)$$

$$c = \eta_X - B^\top \Lambda_{Y|X}^{-1} b \quad (63)$$

$$\Lambda_{X|Y} = \Lambda_X + B^\top \Lambda_{Y|X}^{-1} B \quad (64)$$

followed by these formulas for the marginal mechanisms

$$\Lambda_Y + C^\top \Lambda_{X|Y}^{-1} C = \Lambda_{Y|X} \quad (65)$$

$$\eta_Y - C^\top \Lambda_{X|Y}^{-1} c = b. \quad (66)$$

These formulas are going to be very useful to establish a relationship between the distance to optimum of the causal and anticausal models in Appendix F.5.

F.3 Cholesky Parameters

We call Cholesky parametrization of the normal law the parameters (\mathbf{L}, ζ) such that

$$\begin{aligned} \Lambda &= \mathbf{L} \mathbf{L}^\top \quad (\mathbf{L} \text{ is lower triangular}) \\ \zeta &= \mathbf{L}^{-1} \eta = \mathbf{L}^\top \mu \end{aligned}$$

We use $\mathcal{N}_{\text{cho}}(\zeta, \mathbf{L})$ to denote the normal law with Cholesky parameters ζ and \mathbf{L} . The full causal model (61) becomes

$$\begin{aligned} X &\sim \mathcal{N}_{\text{cho}}(\zeta_X, \mathbf{L}_X) \\ Y|X &\sim \mathcal{N}_{\text{cho}}(MX + m, \mathbf{L}_{Y|X}) \end{aligned} \quad (67)$$

where the 5 parameters are defined from the natural model by the equations

$$\begin{aligned} \mathbf{L}_X \mathbf{L}_X^\top &= \Lambda_X \\ \mathbf{L}_{Y|X} \mathbf{L}_{Y|X}^\top &= \Lambda_{Y|X} \\ \zeta_X &= \mathbf{L}_X^{-1} \eta_X \\ M &= \mathbf{L}_{Y|X}^{-1} B \\ m &= \mathbf{L}_{Y|X}^{-1} b \end{aligned} \quad (68)$$

There is no closed formula to express the Cholesky decomposition of a sum of matrix $A + B$ with the Cholesky decomposition of A and B . As a consequence, there is no simple formula to switch between causal and anticausal models with this parametrization.

F.3.1 Joint Cholesky

We derive a formula for the joint Cholesky for future reference. To get a closed form for the joint Cholesky parameters from the conditional parameters, we need to switch the positions of X and Y in the joint vector – e.g. we are using (Y, X) instead of (X, Y) . Indeed the Cholesky decomposition is very dependent on the orders of the rows and columns. That's also why we cannot simply switch the column orders in the joint representation.

$$\begin{pmatrix} Y \\ X \end{pmatrix} \sim \mathcal{N}_{\text{cho}} \left(\begin{pmatrix} m \\ \zeta_X \end{pmatrix}, \begin{pmatrix} \mathbf{L}_{Y|X} & 0 \\ -M^\top & \mathbf{L}_X \end{pmatrix} \right).$$

This joint representation is simply taking the conditional parameters and putting them in an array. It has the advantage that the distance is equal to the conditional distance. It hints towards the idea that for multivariate normal variables, knowing the right Cholesky decomposition is equivalent to knowing the right causal graph.

F.4 Kullback-Leibler Divergence

We express the KL divergence in all three parametrizations because we use them in the code.

$$\begin{aligned}
 2D_{\text{KL}}(\mathcal{N}_0||\mathcal{N}_1) &= (\mu_1 - \mu_0)^\top \Sigma_1^{-1} (\mu_1 - \mu_0) \\
 &\quad + \text{Tr}(\Sigma_1^{-1} \Sigma_0) - k - \log|\Sigma_1^{-1} \Sigma_0| \\
 &= \eta_1^\top \Lambda_1^{-1} \eta_1 - 2\eta_1^\top \Lambda_0^{-1} \eta_0 + \eta_0^\top \Lambda_0^{-1} \Lambda_1 \Lambda_0^{-1} \eta_0 \\
 &\quad + \text{Tr}(\Lambda_1 \Lambda_0^{-1}) - k - \log(|\Lambda_1 \Lambda_0^{-1}|) \\
 &= \|V^\top \zeta_0 - \zeta_1\|^2 + \|V\|_F^2 - k - 2\log|V|
 \end{aligned}$$

where $V := \mathbf{L}_0^{-1} \mathbf{L}_1$ is a lower triangular matrix which plays a special role.

F.5 Distance after Intervention

In this section we evaluate the effect on interventions on the cause and effect for both models. When the intervention happens on the cause, we replace μ_X, Σ_X by $\tilde{\mu}_X, \tilde{\Sigma}_X$, or equivalently we replace the natural parameters of the marginal on X . The natural causal distance is simply

$$\delta_{\text{causal}} = \|\eta_X - \tilde{\eta}_X\|^2 + \|\Lambda_X - \tilde{\Lambda}_X\|_F^2. \quad (69)$$

Unless indicated otherwise, we will consider the Frobenius distance between matrices. For the anticausal model, both marginal and conditional parameters need to change. Here similar to categorical case, we have

$$\delta_{\text{anticausal}} \geq \delta_{\text{causal}}.$$

However when the intervention happens on the effect, there is no clear formal relation between δ_{causal} and $\delta_{\text{anticausal}}$. More detail about the deriving the mathematical formula for δ_{causal} and $\delta_{\text{anticausal}}$ is presented in the following.

F.5.1 Intervention on Cause

When the intervention happens on the cause, the natural causal distance is

$$\delta_{\text{causal}} = \|\eta_X - \tilde{\eta}_X\|^2 + \|\Lambda_X - \tilde{\Lambda}_X\|_F^2. \quad (70)$$

How does this transformation affect the anticausal parameters? Both the marginal and the conditional have to adapt. The anticausal conditional is elegantly expressed with the causal natural parameters in (64), so we will start with the conditional

$$C - \tilde{C} = B^\top - \tilde{B}^\top = 0 \quad (71)$$

$$c - \tilde{c} = \eta_X - \tilde{\eta}_X \quad (72)$$

$$\Lambda_{X|Y} - \tilde{\Lambda}_{X|Y} = \Lambda_X - \tilde{\Lambda}_X. \quad (73)$$

In words, the linear transformation is invariant, the bias moves like the mean of X , and the conditional precision moves like the precision of X . This means that we can directly lower bound the anticausal distance with the causal distance

$$\begin{aligned}
 \delta_{\text{anticausal}} &= \|C - \tilde{C}\|^2 + \|c - \tilde{c}\|^2 + \|\Lambda_{X|Y} - \tilde{\Lambda}_{X|Y}\|^2 \\
 &\quad + \|\Lambda_Y - \tilde{\Lambda}_Y\|^2 + \|\eta_Y - \tilde{\eta}_Y\|^2 \\
 &= \delta_{\text{causal}} + \|\Lambda_Y - \tilde{\Lambda}_Y\|^2 + \|\eta_Y - \tilde{\eta}_Y\|^2 \\
 &> \delta_{\text{causal}}.
 \end{aligned} \quad (74)$$

We could get a stronger bound by bounding the anticausal marginal parameters, but any such bound would involve the value of the linearity B and make the result needlessly more complicated.

F.5.2 Intervention on Effect

We perform an intervention on Y such that the causal model become independent

$$\begin{aligned} X &\sim \mathcal{N}(\mu_X, \Sigma_X) \\ Y &\sim \mathcal{N}(\tilde{\mu}_Y, \tilde{\Sigma}_Y) . \end{aligned}$$

This independence means that the linear models have a slope 0

$$\tilde{A} = \tilde{B} = \tilde{C} = 0 . \quad (75)$$

The bias then has to account for the mean parameter

$$\tilde{a} = \tilde{\mu}_Y, \quad \tilde{b} = \tilde{\eta}_Y, \quad \tilde{c} = \eta_X \quad (76)$$

And the conditional precision have to match the marginal precision

$$\tilde{\Sigma}_{Y|X} = \tilde{\Sigma}_Y, \quad \tilde{\Lambda}_{Y|X} = \tilde{\Lambda}_Y, \quad \tilde{\Lambda}_{X|Y} = \Lambda_X \quad (77)$$

So the distances are written

$$\begin{aligned} \delta_{\text{causal}} &= \|B\|_F^2 + \|b - \tilde{\eta}_Y\|^2 + \|\Lambda_{Y|X} - \tilde{\Lambda}_Y\|_F^2 \\ \delta_{\text{anticausal}} &= \|C\|_F^2 + \|c - \eta_X\|^2 + \|\Lambda_{X|Y} - \Lambda_X\|_F^2 \\ &\quad + \|\eta_Y - \tilde{\eta}_Y\|^2 + \|\Lambda_Y - \tilde{\Lambda}_Y\|_F^2 \\ &= \|B\|_F^2 + \|B^\top \Lambda_{Y|X}^{-1} b\|^2 + \|B^\top \Lambda_{Y|X}^{-1} B\|_F^2 \\ &\quad + \|\eta_Y - \tilde{\eta}_Y\|^2 + \|\Lambda_Y - \tilde{\Lambda}_Y\|_F^2 \end{aligned}$$

We did not find any meaningful simplification of these formulas.

G NORMAL PRIOR

Exactly like in the categorical setting, the distributions we sample are going to impact the speed of adaptation and the distances we measure. Let K be the dimension of X and Y , and $n_0 = 2K + 2$ an arbitrary number of prior observations. We define a *pseudo-conjugate prior*

$$\Lambda_X \sim \mathcal{W}(n_0, \frac{I_K}{K}) \quad (78)$$

$$\eta_X | \Lambda_X \sim \mathcal{N}(0, \frac{\Lambda_X}{n_0}) \quad (79)$$

$$\Lambda_{Y|X} \sim \mathcal{W}(n_0, \frac{I_K}{K}) \quad (80)$$

$$b | \Lambda_{Y|X} \sim \mathcal{N}(0, \frac{\Lambda_{Y|X}}{n_0}) \quad (81)$$

$$B = \Lambda_{Y|X} A \text{ where } A \sim \mathcal{N}(0, \frac{I_{K^2}}{\sqrt{K}}) \quad (82)$$

where \mathcal{W} is the Wishart distribution with parameters: degrees of freedom and scale matrix. We picked these parameters such that η_Y, Λ_Y follows approximately the same law as η_X, Λ_X . Two important factors to get a symmetric relationship between X and Y are **10** and \sqrt{K} . First, we sample a larger conditional precision, so that their relationship is quite deterministic. Second we sample the linear layer such that it preserves the scale of X , so that X and Y have approximately the same variance. We also use appropriate covariance matrices to sample other parameters such that the prior is somewhat conjugate and gives proper variance formulas.

We sample interventions from the same distributions as the cause marginal in (82). For an intervention on the cause

$$\begin{aligned} \tilde{\Lambda}_X &\sim \mathcal{W}(n_0, \frac{I_K}{K}) \\ \tilde{\eta}_X | \tilde{\Lambda}_X &\sim \mathcal{N}(0, \frac{\tilde{\Lambda}_X}{n_0}) . \end{aligned}$$

For an intervention on the effect

$$\begin{aligned}\tilde{B} = 0\tilde{\Lambda}_{Y|X} = \tilde{\Lambda}_Y &\sim \mathcal{W}(n_0, \frac{I_K}{K}) \\ \tilde{b} = \tilde{\eta}_Y|\tilde{\Lambda}_Y &\sim \mathcal{N}(0, \frac{\tilde{\Lambda}_Y}{n_0}) .\end{aligned}$$