



 Latest updates: <https://dl.acm.org/doi/10.1145/3583780.3614774>

RESEARCH-ARTICLE

Adaptation Speed Analysis for Fairness-aware Causal Models

YUJIE LIN, Tianjin University, Tianjin, China

CHEN ZHAO, Baylor University, Waco, TX, United States

MINGLAI SHAO, Tianjin University, Tianjin, China

XUJIANG ZHAO, NEC Laboratories America, Inc., Princeton, NJ, United States

HAIFENG CHEN, NEC Laboratories America, Inc., Princeton, NJ, United States

Open Access Support provided by:

Baylor University

Tianjin University

NEC Laboratories America, Inc.



PDF Download
3583780.3614774.pdf
30 December 2025
Total Citations: 4
Total Downloads: 89

Published: 21 October 2023

[Citation in BibTeX format](#)

CIKM '23: The 32nd ACM International
Conference on Information and
Knowledge Management
October 21 - 25, 2023
Birmingham, United Kingdom

Conference Sponsors:

SIGIR
SIGWEB

Adaptation Speed Analysis for Fairness-aware Causal Models

Yujie Lin
School of New Media and
Communication, Tianjin University
Tianjin, China
linyujie_22@tju.edu.cn

Chen Zhao
Department of Computer Science,
Baylor University
Waco, Texas, USA
chen_zhao@baylor.edu

Minglai Shao*
School of New Media and
Communication, Tianjin University
Tianjin, China
shaoml@tju.edu.cn

Xujiang Zhao
NEC Lab
Princeton, New Jersey, USA
zhaouxj32@gmail.com

Haifeng Chen
NEC Lab
Princeton, New Jersey, USA
haifeng@nec-labs.com

ABSTRACT

For example, in machine translation tasks, to achieve bidirectional translation between two languages, the source corpus is often used as the target corpus, which involves the training of two models with opposite directions. The question of which one can adapt most quickly to a domain shift is of significant importance in many fields. Specifically, consider an original distribution p that changes due to an unknown intervention, resulting in a modified distribution p^* . In aligning p with p^* , several factors can affect the adaptation rate, including the causal dependencies between variables in p . In real-life scenarios, however, we have to consider the fairness of the training process, and it is particularly crucial to involve a sensitive variable (bias) present between a cause and an effect variable. To explore this scenario, we examine a simple structural causal model (SCM) with a cause-bias-effect structure, where variable A acts as a sensitive variable between cause (X) and effect (Y). The two models respectively exhibit consistent and contrary cause-effect directions in the cause-bias-effect SCM. After conducting unknown interventions on variables within the SCM, we can simulate some kinds of domain shifts for analysis. We then compare the adaptation speeds of two models across four shift scenarios. Additionally, we prove the connection between the adaptation speeds of the two models across all interventions.

CCS CONCEPTS

• **Computing methodologies** → *Causal reasoning and diagnostics.*

KEYWORDS

Adaptation speed, Fairness Learning, Causal Graph

ACM Reference Format:

Yujie Lin, Chen Zhao, Minglai Shao, Xujiang Zhao, and Haifeng Chen. 2023. Adaptation Speed Analysis for Fairness-aware Causal Models. In *Proceedings*

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '23, October 21–25, 2023, Birmingham, United Kingdom
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0124-5/23/10...\$15.00
<https://doi.org/10.1145/3583780.3614774>

of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23), October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3583780.3614774>

1 INTRODUCTION

The widespread adoption of machine learning systems, particularly in decision-critical domains like criminal sentencing and bank loans, has raised concerns regarding the fairness implications [7]. AI systems are increasingly employed in sensitive contexts where they make significant and life-altering decisions [18, 31]. A sensitive feature is defined as an attribute that contains protected information about individuals or groups within a dataset. This information may encompass characteristics such as race, gender, religion, or socioeconomic status, which are safeguarded by ethical considerations, legal regulations, or societal norms [3]. Machine learning models have the potential to inadvertently acquire discriminatory patterns if sensitive variables exhibit spurious correlations with the target variable or predictive outcomes [21, 25]. Consequently, this can lead to biased decisions or unfair predictions that disproportionately impact specific individuals or groups. Therefore, it is crucial to ensure that these decisions do not reflect discriminatory behavior towards particular groups or populations [19].

Causal models have been widely applied in machine learning to address issues related to model fairness. Structural Causal Models (SCM) [9] provide a means of explaining machine learning model predictions. Analyzing causal graphs and paths helps understand how the model's predictions for different groups are formed, thereby identifying and addressing potential unfair factors. A simple SCM is the model $X \rightarrow Y$ where X is the cause and Y is the effect, where it indicates that X determines Y. Modern machine learning methods encounter surprising failures when the test distribution differs from the training distribution, which is commonly known as *domain shift* [5]. Although many domain adaptation [4, 27, 32] and domain generalization [15, 20, 34] methods have been proposed in recent years to mitigate the problem of domain shift, in reality, we often need to achieve the best possible results and for the convenience of training, we still need to retrain the data in the new domain. In the process of relearning from the original distribution to the new distribution, the relative speed of adaptation between the causal model and the anti-causal model may differ. Considering this SCM ($X \rightarrow Y$), the previous work [13] analyzed the adaptation speed of two models trained in the direction of the causal dependency (i.e., from X to Y) and in the reverse direction of the causal




Figure 1: Two fairness-aware training models. The causal model has the same $X \rightarrow Y$ direction of the cause-bias-effect SCM, while the anti-causal model keeps the opposite direction.

dependency (i.e., from Y to X) concerning domain shift. While they have obtained some promising findings, the analysis of adaptation speed did not take into account fairness considerations (i.e., sensitive variables), which are crucial in real-life scenarios. There are currently theoretical and experimental gaps in the analysis of this aspect, which is precisely what our work aims to address.

To comprehensively analyze the aforementioned issue, it is necessary to consider the convergence speed of the two models under different domain shifts. To achieve this, we employed interventions on various variables to simulate different domain shifts. In our analysis, three variables need to be considered, namely: bias (or sensitive variable) A , cause X , and effect Y . We will perform interventions on one or more variables among these three. Specifically, we consider four scenarios, including interventions on bias and cause separately, interventions on both bias and cause, and interventions on the effect.

Having identified the four types of domain shift that we need to consider, now let us provide a more formal definition. The three variables (bias: A , cause: X , effect: Y) related to fairness can form a cause-bias-effect structural causal model ($A \rightarrow X$ and $A, X \rightarrow Y$). We consider two training models with sensitive variables: one that has the consistent cause-effect direction of the cause-bias-effect SCM, and the other that has the opposite direction (Fig. 1). Domain shift can be described as the transition from the initial joint distribution $p(A, X, Y)$ in the training set to the joint distribution $p^*(A, X, Y)$. During the alignment process of distribution p to p^* , we extract several samples from p^* in each round and utilize them to calculate the loss value. This approach aims to gradually converge the distribution towards p^* . Through experimentation and formula derivation, we were surprised to discover that interventions on bias causing domain shift only affect the absolute convergence speed of the two models, but do not impact the relative speed between the causal model and the anti-causal model. In other words, the faster adaptation to the new distribution by either the causal model or the anti-causal model is solely determined by interventions on other variables. We observed that in only one specific scenario of domain shift, the two models maintain a consistent relative speed relationship. When interventions are applied to the cause variable, the causal model consistently exhibits a speed advantage.

However, when interventions are applied to the effect variable, the convergence speed of the causal model is faster only under certain conditions. Otherwise, the anti-causal model has a faster adaptation speed. In conclusion, our **contributions** can be summarized as follows:

- Our work provides a significant contribution to the understanding of the adaptation speed of causal models in scenarios involving sensitive variables. Our work is the first of its kind and offers valuable information from the perspective of fairness.
- We considered domain shifts caused by four types of interventions on the data distribution. And we analyze the relationship between the adaptation speed of two different causal models and provide insights into this topic, using both synthetic and real data. The results can guide the development of more efficient algorithms for learning causal relationships and improving our understanding of complex systems.
- We conduct theoretical analyses for each of the four cases. The proof results are consistent with the conclusions drawn from the experiments, ensuring the rationality and correctness of the findings.

2 RELATED WORK

Domain shift, which is also known as dataset shift or covariate shift, has received considerable attention in recent years [17, 29, 33]. It is a common problem in machine learning where the training and testing data are drawn from different distributions. This difference in distribution can cause a significant drop in the performance of the learned model on the testing data. Domain shift can occur due to various reasons, such as differences in data collection procedures, environmental changes, or task-specific variations. To address domain shift, several approaches have been proposed, including re-weighting [30], importance sampling [6], and transfer learning [22]. These methods aim to either adjust the training data distribution to match the testing data distribution or to learn a mapping from the source domain to the target domain. Some deep learning-based methods, such as deep neural networks, have also shown promising results in domain adaptation tasks [16]. These methods learn a representation of the data that is invariant to the domain shift, enabling the model to generalize to new domains.

Causal fairness aims to prevent machine learning models from perpetuating or amplifying unfairness inherent in the underlying data-generating process. Numerous studies have proposed methods to achieve causal fairness, including the utilization of causal inference to identify and adjust for confounding variables that may introduce bias in predictions. For instance, previous work proposed a method based on counterfactual regression, allowing for the estimation of causal effects between protected attributes (e.g., gender or race) and the desired outcomes [12]. Other researchers have advocated for the use of causal graph structures to represent variable relationships and ensure the model adheres to specific causal consistency conditions [26]. Additionally, the application of causal models has been explored in reasoning about counterfactual fairness [10].

3 BACKGROUND

In this section, we first provide an overview of Structural Causal Models (SCM) and interventions that can cause domain shifts. We

then introduce the formalism used in the analysis, which includes observations, interventions, models, and adaptation, as presented in the analysis [13]. And the formalism was first proposed by the work [1].

3.1 Structural Causal Model (SCM)

Structural causal models (SCMs) are widely used in causal inference to model the causal relationships among variables. An SCM consists of a directed acyclic graph (DAG) and a set of structural equations that define the causal relationships among the variables in the graph [23, 24, 28]. The structural equation for an endogenous variable X_i can be expressed as follows:

$$X_i = f_i(\mathbf{Pa}_i, U_i) \quad (1)$$

where \mathbf{Pa}_i denotes the parent set of X_i in the graph, and U_i denotes the set of exogenous variables that directly affect X_i . The function f_i represents the causal relationship between the parent variables and X_i . SCMs are used to estimate causal effects and test causal hypotheses. By including sensitive variables in the graph and modeling their causal relationships with other variables, SCMs can adjust for sensitive and produce unbiased estimates of causal effects [8].

Interventions on SCMs involve changing the value of a variable to a specified value. This can be represented mathematically using the do-operator, denoted by $\mathbf{do}(X = x)$. The do-operator separates the effect of an intervention from the effect of other variables in the system. For example, if we want to investigate the effect of drug treatment on a disease outcome, we might use the do-operator to set the value of the treatment variable to "treated" and observe the effect on the outcome variable. In the following narrative, we will use \mathbf{p}^* to represent this modified distribution, such as

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

By controlling one or several variables in this way, we simulate domain shifts under different scenarios. And \mathbf{p}^* is the outcome after domain shifts.

3.2 Reference and Transfer Distributions

We obtain the initial reference distribution \mathbf{p} by sampling the triad (A, X, Y) from a structural causal model (SCM) constructed as follows: A is a bias, X is the cause, and Y is the effect. The SCM is defined by the following two equations: $A \rightarrow X$ and $A, X \rightarrow Y$. Now, we perturb the distribution \mathbf{p} by performing interventions on certain variables to obtain the following transfer distribution \mathbf{p}^* .

Table 1: Transfer distributions of different interventions.

Intervention	Transfer Distribution
Bias	$\mathbf{p}^*(a, x, y) = \mathbf{p}^*(a) \mathbf{p}(x a) \mathbf{p}(y a, x)$
Cause	$\mathbf{p}^*(a, x, y) = \mathbf{p}(a) \mathbf{p}^*(x) \mathbf{p}(y a, x)$
Bias and Cause	$\mathbf{p}^*(a, x, y) = \mathbf{p}^*(a) \mathbf{p}^*(x) \mathbf{p}(y a, x)$
Effect	$\mathbf{p}^*(a, x, y) = \mathbf{p}(a) \mathbf{p}(x a) \mathbf{p}^*(y)$

If the intervention is on the bias, we sample A from a different marginal distribution, while X and Y are sampled from the reference conditional distribution. If the intervention is on the cause, A is

Algorithm 1: Adaptation of Causal Model

Input: The initialized causal model $\mathbf{p}_{\theta_{\downarrow}^{(0)}} = \mathbf{p}$, the intervened variable v (e.g. a, x or y), training epochs T , learning rate α , num of samples in each epoch K .
Goal: Let initial parameter $\theta^{(0)}$ adapt to the parameter θ^* of the transfer distribution \mathbf{p}^* .
 $\mathbf{p}^*(a, x, y) = \mathbf{p}(a, x, y | \mathbf{do}(v = t))$ // t represents random interference on variable v
for $t=1$ to T **do**
 for $k=1$ to K **do**
 | Sample $\theta_k = (a_k, x_k, y_k) \sim \mathbf{p}^*(a, x, y)$
 $\theta_K = \{\theta_k\}_{k=0}^K$
 Calculate the loss $\mathcal{L}_{\text{causal}}$ in Eqn. 6 with θ_K
 $\theta_t = \theta_{t-1} - \alpha \nabla_{\theta_{t-1}} \mathcal{L}_{\text{causal}}$

sampled from the reference marginal distribution, X is sampled from a different marginal distribution independently of A , and Y is sampled from the reference conditional distribution. If the intervention is on both the bias and the cause, A is sampled from a different marginal distribution, X is sampled from a different marginal distribution independently of A , and Y is sampled from the reference conditional distribution. If the intervention is on the effect, A is sampled from the reference marginal distribution, X is sampled from the reference conditional distribution, and Y is sampled from a different marginal distribution independently of A and X . Thus, we obtain all the transfer joint distributions that arise from interventions on some of the variables (Table 1).

3.3 Fairness-aware Models for Training

The role of SCM mentioned earlier is to demonstrate the variables involved in the actual training of the model, as well as the causal dependencies between them. The models mentioned in this section (Fig. 1), which are distinct from the SCM and are used for training, are referred to as causal models and anti-causal models, respectively. The causal model and the anti-causal model are constructed with the variables (A, X, Y) . The causal model can be described as

$$\mathbf{p}_{\theta_{\downarrow}}(a, x, y) = \mathbf{p}_{\theta_A}(a) \mathbf{p}_{\theta_{X|A}}(x|a) \mathbf{p}_{\theta_{Y|A, X}}(y|a, x). \quad (3)$$

Meanwhile, the anti-causal model can be described as

$$\mathbf{p}_{\theta_{\leftarrow}}(a, x, y) = \mathbf{p}_{\theta_A}(a) \mathbf{p}_{\theta_{Y|A}}(y|a) \mathbf{p}_{\theta_{X|A, Y}}(x|a, y), \quad (4)$$

where the θ_{\downarrow} and θ_{\leftarrow} represent parameters of the two models respectively (e.g. θ_{\downarrow} includes θ_A , $\theta_{X|A}$, and $\theta_{Y|A, X}$).

3.4 Adaptation of Two Models

This section explains the most fundamental issue of this work. Assuming the initial distribution of two models is both \mathbf{p} , and due to certain factors, the distribution drifts to \mathbf{p}^* , the training process is to make the models approach the distribution \mathbf{p}^* (Algorithm 1). When the training is about to start, the training term is $T := 0$. The two models will be initialized to fit the same reference distribution \mathbf{p} like

$$\mathbf{p}_{\theta_{\downarrow}^{(0)}} = \mathbf{p}_{\theta_{\leftarrow}^{(0)}} = \mathbf{p}. \quad (5)$$

Then we get samples from the transfer distribution \mathbf{p}^* . Letting these samples join the training process, the log-likelihood will gradually




Figure 2: Parameters of the causal model in SGD. In which the length K of each dimension represents that the variables A , X , and Y can be divided into K categories.

increase in every step of stochastic gradient descent (SGD). The distribution p_θ adapts to p^* closest until we get the minimal log-likelihood loss. Taking the causal model for example, the loss is

$$\begin{aligned} \mathcal{L}_{\text{causal}}(\theta_{\rightarrow}) &= \mathbb{E}_{(A,X,Y) \sim p^*} [-\log p_{\theta_{\rightarrow}}(A, X, Y)] \\ &= \mathbb{E}_{p^*} [-\log p_{\theta_A}(A)] + \mathbb{E}_{p^*} [-\log p_{\theta_{A|X}}(A|X)] \\ &\quad + \mathbb{E}_{p^*} [-\log p_{\theta_{Y|A,X}}(Y|A, X)], \end{aligned} \quad (6)$$

where the log-likelihood suboptimality is equal to the KL-divergence

$$\mathcal{L}(\theta) - \mathcal{L}(\theta^*) = D_{\text{KL}}(p^* \| p_\theta). \quad (7)$$

4 PARAMETERS AND ANALYSIS

4.1 Relevant Inequality

Based on Average Stochastic Gradient Descent (ASGD) [2], the previous work [13] proves that the average parameter's $\bar{\theta}^{(T)} = \frac{1}{T} \sum_{t=0}^{T-1} \theta^{(t)}$ suboptimality is upper bounded by

$$\mathbb{E} [D_{\text{KL}}(p^* \| p_{\bar{\theta}^{(T)}})] \leq \frac{c^{-1} \|\theta^{(0)} - \theta^*\|^2 + cB^2}{2\sqrt{T}}, \quad (8)$$

where c is a small enough constant. This inequality indicates that the upper bound of the convergence is mainly determined by the distance $\delta := \|\theta^{(0)} - \theta^*\|^2$ between the reference distribution and the transfer distribution. Specifically, the initial distance of the causal model and the anti-causal model is respectively denoted by $\delta_{\text{causal}} = \|\theta_{\rightarrow}^{(0)} - \theta_{\rightarrow}^*\|^2$ and $\delta_{\text{anticausal}} = \|\theta_{\leftarrow}^{(0)} - \theta_{\leftarrow}^*\|^2$. And the two distances are the core basis of the subsequent discussion.

4.2 Parameters in Trainable Models

We assume that the bias A , the cause X , and the effect Y in the two models are multiclass variables with K classes. The natural parameter $\theta \in \mathbb{R}^K$ is generated from the distribution p by the inverse function of the softmax function

$$p_z = \frac{e^{s_z}}{\sum_{z'} e^{s_{z'}}}. \quad (9)$$

We set s as the trainable parameter in SGD. Taking causal model for example, the model has parameters $s_A := (s_a)_{a=1 \dots K}$, $s_{X|A} := (s_{x|a})_{a,x=1 \dots K}$ and $s_{Y|A,X} := (s_{y|a,x})_{a,x,y=1 \dots K}$. The parameters of

causal model can be represented as $\theta_{\rightarrow} = (s_A, s_{X|A}, s_{Y|A,X})$, while that's $\theta_{\leftarrow} = (s_A, s_{Y|A}, s_{X|A,Y})$ in anti-causal model. We describe the relationship between the parameter shapes in Fig. 2. Using the parameter s , the loss (6) can be written as

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

After deriving the new form of the loss, we can use inequality (8) to further compare the adaptation speeds of the two models. In other words, we only need to compare the initial distances δ_{causal} and $\delta_{\text{anticausal}}$ as described in Section 4.1.

4.3 Adaptation Speeds of Two Models

Based on the introduction in Section 4.1, the adaptation speed of two models depends on the distance between their initial distribution and the distribution after the change, which can be represented by the parameter "s" in Section 4.2. Furthermore, comparing the convergence speed of two models can be equivalently expressed as comparing the magnitudes of $\delta_{\text{causal}} = \|\theta_{\rightarrow}^{(0)} - \theta_{\rightarrow}^*\|^2$ and $\delta_{\text{anticausal}} = \|\theta_{\leftarrow}^{(0)} - \theta_{\leftarrow}^*\|^2$, the model with a smaller initial distance (δ) has a faster convergence speed.

Domain Shift by bias A , $s_A \leftarrow s_A^*$. In this scenario, both the causal and anti-causal models modify only the same sensitive marginal s_A . The initial distance between the two models can be expressed as

$$\delta_{\text{causal}} = \delta_{\text{anticausal}} = \|s_A - s_A^*\|^2, \quad (11)$$

which implies that the two models converge simultaneously.

Domain Shift by cause X , $\forall a, s_{X|a} \leftarrow s_{X|a}^*$. The conditional $s_{Y|A,X}$ remains unchanged while $s_{Y|A}$ and $s_{X|A,Y}$ of anti-causal model are modified. The initial distance can be written as

$$\delta_{\text{causal}} = \sum_a \|s_{X|a} - s_{X|a}^*\|^2, \quad (12)$$

$$\delta_{\text{anticausal}} = \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_a \sum_y \|s_{X|a,y} - s_{X|a,y}^*\|^2. \quad (13)$$

We can compare the initial distances of the two models based on the aforementioned distance and arrive at the following proposition.

PROPOSITION 1. *When the intervention is on the cause,*

$$\delta_{\text{anticausal}} \geq K \delta_{\text{causal}}, \quad (14)$$

where the specific proof process will be explained in Appendix B.2.

Domain Shift by both bias A and cause X , $s_A \leftarrow s_A^*$ and $\forall a, s_{X|a} \leftarrow s_{X|a}^*$. Compared to the case of intervening on cause X , the intervention on bias A introduces an additional equal distance to the initial distance of the two models. Hence, we can obtain a similar result as before and derive the following inequality:

$$\delta_{\text{anticausal}} \geq \delta_{\text{causal}}, \quad (15)$$

which will be simply explained in Appendix B.3. It can be observed that in all three scenarios, the initial distance of the causal model is consistently smaller than that of the anti-causal model, indicating that the causal model adapts to the domain more quickly. However,

if the domain shift is induced by interventions on the effect variable (Y), interestingly, different conclusions can be drawn.

Domain Shift by effect Y, $\forall a, x, s_{Y|a,x} \leftarrow s_Y^*$. The marginal s_A and the conditional $s_{X|a}$ remain unchanged under this intervention. On the other hand, the conditional $s_{Y|A}$ and $s_{Y|A,X}$ in the anti-causal model change with respect to effect Y. The initial distances for the two models can be expressed as:

$$\delta_{\text{causal}} = \|s_Y - s_{Y|A,X}^*\|^2, \quad (16)$$

$$\delta_{\text{anticausal}} = \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_a \sum_y \|s_{X|a,y} - s_{X|a,y}^*\|^2, \quad (17)$$

where the distance above does not maintain a constant relationship compared to the other three situations.

PROPOSITION 2. *When the intervention is on the effect, there will be two situations. If the following inequality:*

$$\|s_Y^* - c\|^2 < R^2 \quad (18)$$

is satisfied, $\delta_{\text{anticausal}} \geq \delta_{\text{causal}}$ (proved in Appendix B.4). Let us illustrate this inequality, where R^2 (see Appendix for specific formula) is a constant related with s_X , $s_{Y|A}$, $s_{X|A,Y}$ and $s_{Y|A,X}$, and $c = \frac{(\sum_x s_{Y|A,x}) - s_{Y|A}}{K-1}$.

The inequality (18) implies that the causal model has a comparative advantage only within a certain range, where the modified marginal s_Y is sufficiently close to c . However, if the distance goes beyond this range, the anti-causal model will converge faster.

5 EXPERIMENTS

5.1 Data and Settings

Synthetic Data. The first we need is to get the distributions $\mathbf{p} = \mathbf{p}_{\theta^{(0)}}$ which is called *prior*. Specifically, we get \mathbf{p}_A , $\mathbf{p}_{X|A}$ and $\mathbf{p}_{Y|A,X}$ from the Dirichlet distribution. The three distributions can be represented as:

$$\begin{aligned} \mathbf{p}_A &\sim \text{Dirichlet}(\mathbf{1}_K), \\ \forall a, \mathbf{p}_{X|a} &\sim \text{Dirichlet}(\mathbf{1}_K), \\ \forall a, x, \mathbf{p}_{Y|a,x} &\sim \text{Dirichlet}(\mathbf{1}_K), \end{aligned}$$

where $\mathbf{1}_K$ is the all-one vector of K-dimension. We can say that by using the aforementioned distributions, we ensure that these three distributions are mutually independent. Now we have thus obtained an initial joint distribution.

Besides employing synthetic data, we integrated two actual datasets to establish the initial distribution \mathbf{p} . This involved determining the frequency of each category for each variable. In the subsequent sections, we will elucidate the identifiable variables A, X, and Y present in each respective dataset. It is important to note that we did not directly train the causal model and anti-causal model using the adult and colored-MNIST datasets, nor did we compare their convergence speeds in this manner. Instead, all experiments in this article were conducted employing the two models described in Section 3.3, with the parameters mentioned in Section 4.2, and the convergence speeds were compared using the algorithm outlined in Section 3.4. To differentiate it from the strictly Dirichlet distribution in synthetic data, we aimed to acquire the initial distribution $\mathbf{p}(A, X, Y)$ from real datasets to showcase additional results.

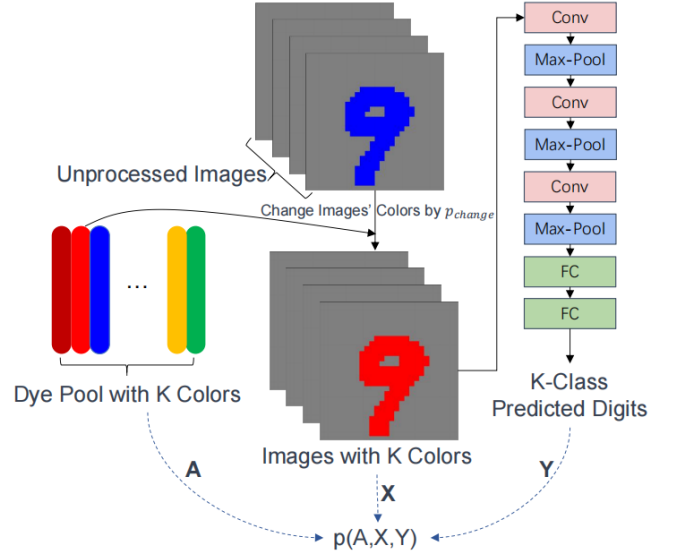


Figure 3: Initialize $\mathbf{p}(A,X,Y)$ with Colored-MNIST. The color pool consists of K RGB colors, and it can probabilistically alter the colors of the numbers in the image. The image is also segmented into K categories based on its color, and K prediction results are generated by CNN and linear layer.

Colored-MNIST. We use MNIST [14] to construct the initial distribution \mathbf{p} . The sensitive variable (A) is RGB color values, which consist of K kinds of colors represented by a three-dimensional vector. We obtain the initial $\mathbf{p}_{\text{color}}$ from the Dirichlet distribution. The cause (X) is the images' features of colored-MNIST, which have different colors and can be classified into K categories based on color. For each image, we randomly select a color from $\mathbf{p}_{\text{color}}$ to stain it. To avoid having the same distribution $\mathbf{p}_{\text{image}}$ as $\mathbf{p}_{\text{color}}$, we set the probability of changing the image's color as $\mathbf{p}_{\text{change}}$ (e.g., $\mathbf{p}_{\text{change}} = 0.5$). Then we need a model to input image features (X) and get predicted numbers (Y). Since we only need the initial distribution of Y obtained from X, the model used to obtain the predicted numbers is not crucial in this analysis. We process the images with a convolutional neural network (CNN) to predict the numbers, which represent the effect (Y) of the SCM. Finally, we obtain the initial distribution $\mathbf{p}(A, X, Y)$ by calculating the frequency of all categories (Fig.3).

Adult. The results from the MNIST dataset and synthetic data demonstrate the adaptation speed of multi-categorical variables. We use the Adult dataset [11] to demonstrate the adaptation speed of binary variables, where the bias (A) is gender, including male and female. The initial distribution $\mathbf{p}_{\text{gender}}$ is also obtained from the Dirichlet distribution. The cause (X) is the features of people in the dataset, which are divided into two categories based on gender. And for each individual, we select a gender from $\mathbf{p}_{\text{gender}}$. Similarly to the Colored-MNIST dataset, we change the initial gender of each feature to the selected gender with the probability $\mathbf{p}_{\text{change}}$. Using the modified features, we input them into a 2-layer MLP to obtain binary predicted labels. The binary labels are recognized as the effect (Y) in the SCM.

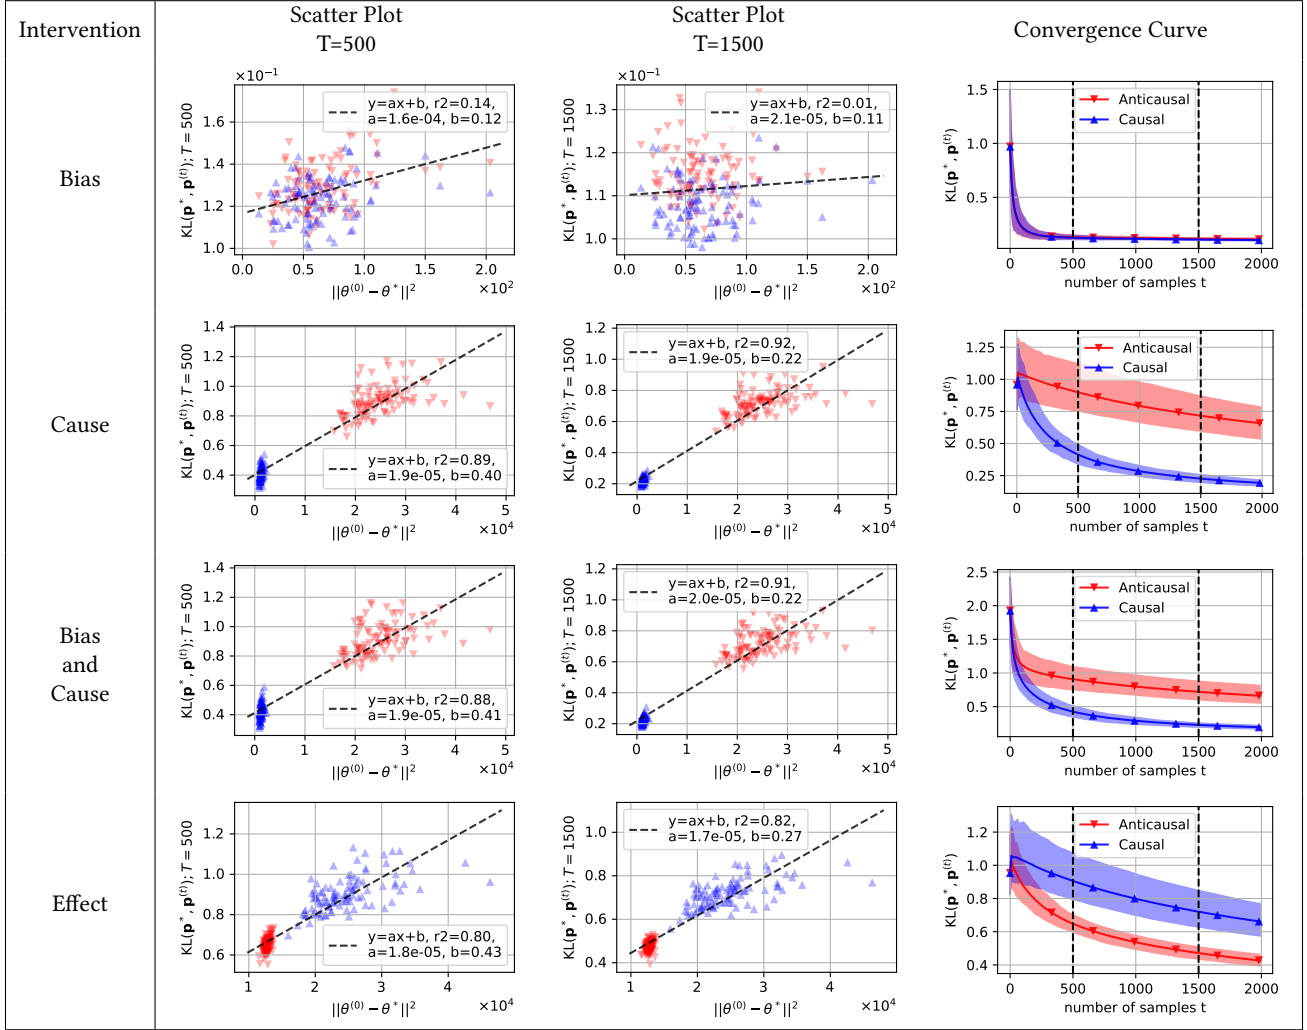


Table 2: Results on synthetic data. The scatter plots in the first column demonstrate the positive correlation between the KL-divergence after one-quarter of the training steps, while the second column shows the correlation after three-quarters of the training steps. Each point on the scatter plots represents a pair $(p^{(0)}, p^*)$ in the causal model (blue) or the anti-causal model (red). The parameters a , b , and r^2 in a scatter plot represent the slope, intercept, and the coefficient of determination, respectively, in the least squares method. The coefficient of determination measures the linear correlation between the independent and dependent variables. The curve in the third column shows the relative speeds of the two models. The shaded area indicates the 5th and 95th percentiles of the KL-divergence. The value of $K_{synthetic}$ is set to 20 in this experiment.

5.2 Results

This section demonstrates the adaptation speed of the causal and anti-causal models using the data from Section 5.1, as well as the positive relationship between the KL divergence and the initial parameter distance. And similar results were obtained in both multi-class and binary classification datasets.

Utilizing $K_{synthetic} = 20$, $K_{MNIST} = 10$, and $K_{Adult} = 2$, the outcomes for the three datasets are presented in Table 2, Table 3, and Table 4, respectively, to showcase the adaptation speed of the causal and anti-causal models. Additionally, scatter plots are employed to illustrate the positive correlation between the KL divergence

and the initial parameter distance throughout the training process. This positive correlation holds significant importance as it validates Inequality 8, which constitutes the foundation for deriving the relationship between the model’s adaptation speed. The results after one-quarter and three-quarters of the training steps are presented, while the curves depict the adaptation speed.

Domain Shift by bias A. When intervening on bias A, both the causal and anti-causal models undergo the same change concerning the bias, while holding other variables constant. Therefore, both models exhibit overlapping points on the scatter plots, resulting in the initial distance $\delta_{causal} = \delta_{anticausal}$ and coinciding curves in the third column. As a result, the convergence speed of both models is

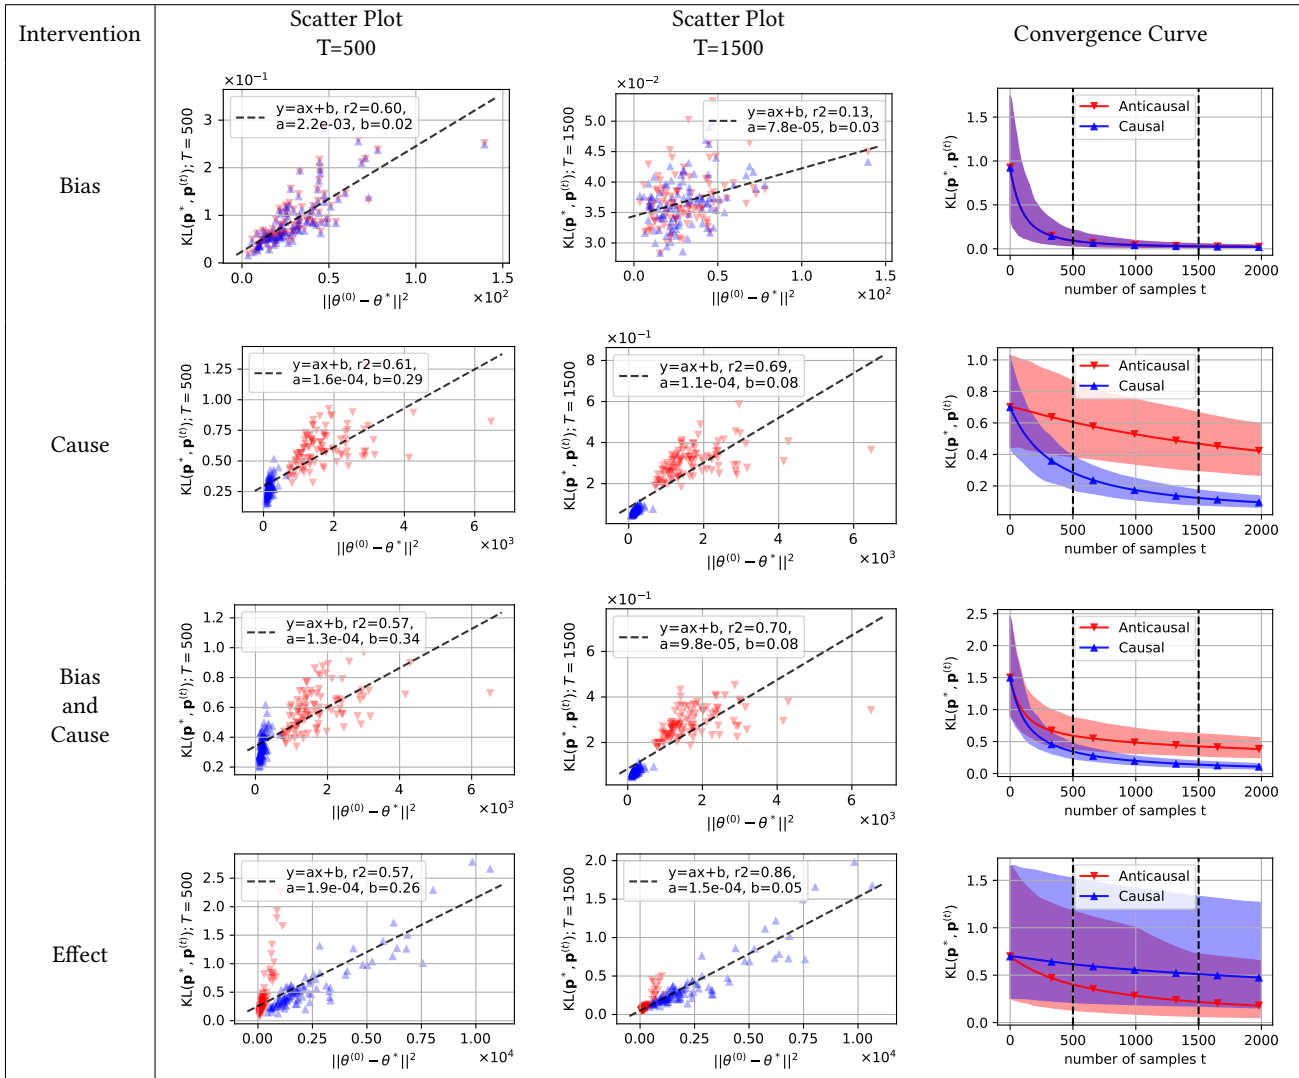


Table 3: Results on Color-MNIST data. The relative positional relationship between the two types of points in the scatter plots remains unchanged. Here we set $K_{MNIST} = 10$. The convergence speed of the two models is consistent with previous experiments on synthetic data.

the same. This is in line with the observation that interventions on bias do not change the causal structure of the model, and thus do not provide any new information that can be used to distinguish between the two models.

Domain Shift by cause X. In the case of an intervention on the cause X, the causal model has an advantage over the anti-causal model, as the points of the causal model cluster towards the bottom left as compared to the points of the anti-causal model. This relative positioning reflects the formula $\delta_{\text{causal}} < \delta_{\text{anticausal}}$, which is derived in our previous work. This advantage of the causal model is most pronounced in four intervention scenarios, where the causal model exhibits a significant advantage in terms of the curves.

Domain Shift by both bias A and cause X. When intervening on both the bias and the cause, the relative positions of points in the

causal and anti-causal models maintain their relative positions, with the anti-causal model being on the upper right of the causal model as a whole. However, the difference in adaptation speed between the two models shrinks, as the two curves have a short overlap in the early training steps. This result is somewhat counterintuitive, as one might expect that intervening on both the bias and the cause would provide more information that could be used to distinguish between the two models. However, our results suggest that this is not always the case, and that the advantage of the causal model may depend on the specific structure of the model.

Domain Shift by effect Y. When intervening on the effect Y, the relative positions of points in the causal and anti-causal models are reversed, with the points of the anti-causal model being

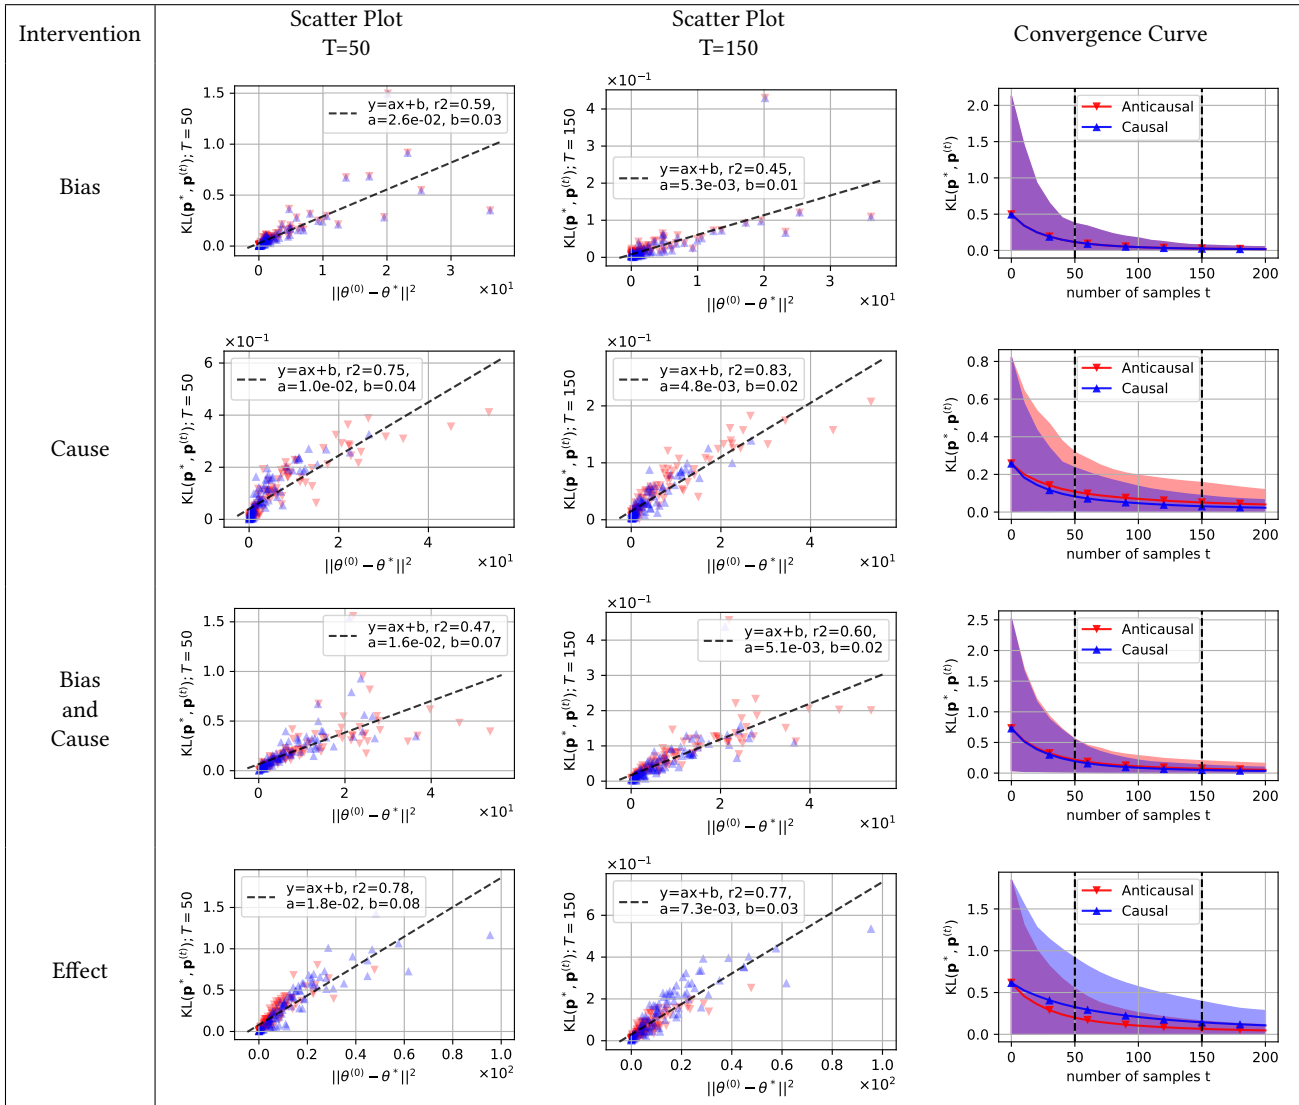


Table 4: Results on the Adult data. As an experiment using binary variables, we observed that both the causal and anti-causal models converge faster, similar to the results in the previous experiments. Additionally, the two adaptation curves are closer when $K_{Adult} = 2$. Furthermore, there is less separation between the points of the two models in the scatter plots.

concentrated in the lower-left corner. This advantage of the anti-causal model corresponds to the proposition that the anti-causal model is better suited to handle interventions on effect variables. Although the anti-causal model may not always have this advantage, it typically retains it in most cases. The curves also reflect this result, as the anti-causal model consistently converges faster when using the three datasets. Our results suggest that interventions on effect variables provide valuable information that can be used to distinguish between the causal and anti-causal models.

In addition to the analysis of multiple categorical variables, we also conducted a study on dichotomous variables using the Adult dataset. The results, as presented in Table 4, indicate that both models exhibit accelerated convergence rates due to the reduced

number of parameters in the training process. Except the bias intervention, the points of both models tend to converge and become less separated in the scatter plots in other scenarios. Interestingly, the relationship between the speed of the causal model and the anti-causal model remains consistent with the results from the previous two experiments. However, the two curves are significantly closer together in this study, suggesting that the number of categories of a variable may only affect the absolute adaptation, but not the relative speed of the causal model and anti-causal model. Based on these observations, we can conclude that dichotomous variables also have a significant impact on the performance of causal models. By reducing the complexity of the models, the convergence rate can be accelerated, and the overall accuracy of the models can be

improved. However, the relationship between the causal model and the anti-causal model remains unchanged, indicating that these two models are equally capable of capturing the underlying causal relationships in the data, irrespective of the type of variable used.

Returning to the example introduced in the abstract of this article, suppose that Japanese is a language developed from Chinese, in which Chinese is a causal factor in the development of Japanese. In the presence of the sensitive variable of polysemy, training a Japanese corpus using Chinese data (training in the causal direction of the variable) may not necessarily result in faster adaptation to domain shift compared to training a Chinese corpus using Japanese data (training in the anti-causal direction of the variable). This depends primarily on how domain shift occurs, i.e., which variable(s) undergo intervention.

6 CONCLUSION

This paper aims to explore spurious relationships in structural causal models (SCMs) that arise due to sensitive factors. We investigate the adaptation speed of both causal and anti-causal models in the presence of bias and build upon a theory that explains the relationship between the initial distance of parameters and the adaptation speed. Furthermore, it's a challenge to extend our analysis to models with more complex structural differences, such as those with varying numbers of variables and edges. To analyze the adaptation speed in such cases, an indicator to measure the difference between the models is necessary. However, finding an appropriate indicator can be challenging to analyze in the future.

ACKNOWLEDGMENTS

This work is supported by NSFC program (No. 62272338).

A PARAMETER

The score $s \in \mathbf{R}^K$ have one additional degree of freedom compared to the probability [13]. If we add or subtract any multiple of $\mathbf{l} = (1, 1, \dots, 1)$ to s , the softargmax will not change. That means all the scores $\{s + \lambda|\forall \lambda \in \mathbf{R}\}$ are equivalent. Thus we can subtract the score's mean and give the scores a good property: $\sum_z s_z = 0$.

B DISTANCE ANALYSIS

Due to space limitations, we are unable to provide complete proof, such as Lemma B.3 and the full details of the distance comparison.

B.1 Related Variables

It's useful to find the relation of $s_{y|a,x}$ and $s_{x|a,y}$, because δ_{causal} and $\delta_{anticausal}$ are represented by s .

Definition B.1 (Part-average conditional score vectors). For any a, x or y , define

$$\theta_1(a, y) := \frac{1}{K} \sum_x s_{y|a,x}, \quad \theta_2(a, x) := \frac{1}{K} \sum_y s_{x|a,y}. \quad (19)$$

Definition B.2 (Conditional log-partition function).

$$A_1(a, x) = \log \sum_y e^{s_{y|a,x}} \quad A_2(a) = \log \sum_x e^{s_{x|a}} \quad (20)$$

LEMMA B.3 (ANTICAUSAL CONDITIONAL SCORE). $s_{y|a,x}, s_{x|a,y}, s_{x|a}$ are all conditional scores in two models. Then we find the relation between $s_{y|a,x}$ and $s_{x|a,y}$.

$$s_{x|a,y} = s_{y|a,x} + s_{x|a} - A_1(a, x) - \theta_1(a, y) + \alpha(a), \quad (21)$$

where $\alpha = \frac{1}{K} \sum_x A_1(x)$.

B.2 Proof for Proposition 1

Given that $s_{y|a,x}^* = s_{y|a,x}$, $\theta_1^* = \theta_1$, $A_1^* = A_1$ and $\alpha^* = \alpha$, Lemma B.3 tells us that D-value between the anti-causal conditional $s_{X|A,Y}^*$ and the causal conditional $s_{X|A}^*$ remains unchanged.

$$\begin{aligned} s_{x|a,y}^* - s_{x|a}^* &= s_{y|a,x} - A_1(a, x) + \alpha_1(a) - \theta_1(a, y) \\ &= s_{x|a,y} - s_{x|a} \\ \implies s_{x|a,y} - s_{x|a}^* &= s_{x|a} - s_{x|a}^* \end{aligned}$$

The distance between models before and after intervention are

$$\begin{aligned} \delta_{causal} &= \sum_a \|s_{X|a} - s_{X|a}^*\|^2 = \sum_{a,x} (s_{x|a} - s_{x|a}^*)^2 \\ \delta_{anticausal} &= \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_{a,y} \|s_{X|a,y} - s_{X|a,y}^*\|^2 \\ &= \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_{a,x,y} (s_{x|a} - s_{x|a}^*)^2 \\ &\geq 0 + K \sum_{a,x} (s_{x|a} - s_{x|a}^*)^2 = K \delta_{causal}. \end{aligned}$$

B.3 Proof for Equation 15

$$\begin{aligned} \delta_{causal} &= \|s_A - s_A^*\|^2 + \sum_a \|s_{X|a} - s_{X|a}^*\|^2 \\ \delta_{anticausal} &= \|s_A - s_A^*\|^2 + \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_{a,y} \|s_{X|a,y} - s_{X|a,y}^*\|^2 \end{aligned}$$

Using the previous conclusion in Section B.2,

$$\sum_{a,y} \|s_{X|a,y} - s_{X|a,y}^*\|^2 = K \sum_a \|s_{X|a} - s_{X|a}^*\|^2 \geq \sum_a \|s_{X|a} - s_{X|a}^*\|^2,$$

we can get the relation: $\delta_{anticausal} \geq \delta_{causal}$.

B.4 Proof for Proposition 2

$$\begin{aligned} \delta_{causal} &= \sum_{a,x} \|s_{Y|x,a} - s_{Y|a}^*\|^2 \\ &= \sum_{a,x,y} (s_{y|x,a} - \theta_1(a, y))^2 + \sum_{a,x,y} (\theta_1(a, y) - s_{y|a}^*)^2 \\ \delta_{anticausal} &= \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_{a,y} \|s_{X|a,y} - s_{X|a,y}^*\|^2 \\ &= \sum_a \|s_{Y|a} - s_{Y|a}^*\|^2 + \sum_{a,x,y} (s_{x|a,y} - \theta_2(a, x))^2 + \sum_{a,x,y} (\theta_2(a, x) - s_{x|a})^2 \\ \delta_{causal} - \delta_{anticausal} &= (K-1) \|s_Y^* - \mathbf{c}\|^2 - (K-1) R^2, \end{aligned}$$

where $R^2 = \frac{(K-1) \|\mathbf{c}\|^2 - K \|\theta_1\|^2 + \|s_{Y|A}\|^2 + K \|\theta_2 - s_X\|^2}{K-1}$, $\mathbf{c} = \frac{K \theta_1 - s_{Y|A}}{K-1}$ and $\theta_1 = \frac{(K-1) \mathbf{c} + s_{Y|A}}{K}$. When $\|s_Y^* - \mathbf{c}\|^2 < R^2$, $\delta_{causal} \geq \delta_{anticausal}$.

REFERENCES

- [1] Y. Bengio, T. Deleu, N. Rahaman, N. R. Ke, S. Lachapelle, O. Bilaniuk, A. Goyal, and C. Pal. 2020. A Meta-Transfer Objective for Learning to Disentangle Causal Mechanisms. In *International Conference on Learning Representations*.
- [2] Léon Bottou. 2012. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*. Springer, 421–436.
- [3] Simon Caton and Christian Haas. 2020. Fairness in machine learning: A survey. *Comput. Surveys* (2020).
- [4] Chenhui Chu and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. *arXiv preprint arXiv:1806.00258* (2018).
- [5] Hady Elsahar and Matthias Gallé. 2019. To annotate or not? predicting performance drop under domain shift. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 2163–2173.
- [6] Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*. PMLR, 1180–1189.
- [7] Usman Gohar and Lu Cheng. 2023. A Survey on Intersectional Fairness in Machine Learning: Notions, Mitigation, and Challenges. arXiv:2305.06969 [cs.LG]
- [8] Miguel A Hernán and James M Robins. 2018. Causal inference. *International encyclopedia of statistical science* (2018), 1–10.
- [9] C. Hitchcock and J. Pearl. 2001. Causality: Models, Reasoning and Inference. *Philosophical Review* 110, 4 (2001), 639.
- [10] Niki Kilbertus, Giuseppe Parascandolo, Riedel Heiko, Bernhard Schölkopf, and Adrian Weller. 2017. Avoiding Discrimination through Causal Reasoning. *NIPS* (2017).
- [11] R. Kohavi. 1997. Scaling Up the Accuracy of Naive-Bayes Classifiers: a Decision-Tree Hybrid. *Proc of Kdd* (1997).
- [12] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. *NIPS* (2017).
- [13] Rémi Le Priol, Reza Babanezhad Harikandeh, Yoshua Bengio, and Simon Lacoste-Julien. 2020. An Analysis of the Adaptation Speed of Causal Models. *arXiv e-prints*, Article arXiv:2005.09136 (May 2020), arXiv:2005.09136 pages. arXiv:2005.09136 [stat.ML]
- [14] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324. <https://doi.org/10.1109/5.726791>
- [15] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. 2018. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [16] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. 2015. Learning transferable features with deep adaptation networks. In *International conference on machine learning*. PMLR, 97–105.
- [17] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. 2019. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2507–2516.
- [18] Trisha Mahoney, Kush Varshney, and Michael Hind. 2020. *AI fairness*. O'Reilly Media, Incorporated.
- [19] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [20] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. 2013. Domain generalization via invariant feature representation. In *International conference on machine learning*. PMLR, 10–18.
- [21] Luca Oneto and Silvia Chiappa. 2020. Fairness in machine learning. In *Recent trends in learning from data: Tutorials from the inns big data and deep learning conference (innsbdl2019)*. Springer, 155–196.
- [22] Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22, 10 (2010), 1345–1359.
- [23] Judea Pearl. 2009. *Causality*. Cambridge University Press.
- [24] Judea Pearl and Dana Mackenzie. 2018. *The book of why: the new science of cause and effect*. Basic books.
- [25] Dana Pessach and Erez Shmueli. 2022. A review on fairness in machine learning. *ACM Computing Surveys (CSUR)* 55, 3 (2022), 1–44.
- [26] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of Causal Inference*. MIT Press (2017).
- [27] Danielle Saunders. 2022. Domain adaptation and multi-domain adaptation for neural machine translation: A survey. *Journal of Artificial Intelligence Research* 75 (2022), 351–424.
- [28] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. 2000. *Causation, prediction, and search*. MIT press.
- [29] Karin Stacke, Gabriel Eilertsen, Jonas Unger, and Claes Lundström. 2020. Measuring domain shift for deep learning in histopathology. *IEEE journal of biomedical and health informatics* 25, 2 (2020), 325–336.
- [30] Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. 2007. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research* 8, 5 (2007).
- [31] Shari Trewin, Sara Basson, Michael Muller, Stacy Branham, Jutta Treviranus, Daniel Gruen, Daniel Hebert, Natalia Lyckowski, and Erich Manser. 2019. Considerations for AI fairness for people with disabilities. *AI Matters* 5, 3 (2019), 40–63.
- [32] Mei Wang and Weihong Deng. 2018. Deep visual domain adaptation: A survey. *Neurocomputing* 312 (2018), 135–153.
- [33] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. 2021. Adaptive risk minimization: Learning to adapt to domain shift. *Advances in Neural Information Processing Systems* 34 (2021), 23664–23678.
- [34] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. 2022. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).