
When is Transfer Learning Possible?

My Phan Kianté Brantley^{*1} Stephanie Milani^{*2} Soroush Mehri^{*3} Gokul Swamy^{*2} Geoffrey J. Gordon²

Abstract

We present a general framework for transfer learning that is flexible enough to capture transfer in supervised, reinforcement, and imitation learning. Our framework enables new insights into the fundamental question of *when* we can successfully transfer learned information across problems. We model the learner as interacting with a sequence of problem instances, or *environments*, each of which is generated from a common structural causal model (SCM) by choosing the SCM’s parameters from restricted sets. We derive a procedure that can propagate restrictions on SCM parameters through the SCM’s graph structure to other parameters that we are trying to learn. The propagated restrictions then enable more efficient learning (i.e., transfer). By analyzing the procedure, we are able to challenge widely-held beliefs about transfer learning. First, we show that having *sparse* changes across environments is neither necessary nor sufficient for transfer. Second, we show an example where the common heuristic of *freezing* a layer in a network causes poor transfer performance. We then use our procedure to select a more refined set of parameters to freeze, leading to successful transfer learning.

1. Introduction

Transfer learning is a well-studied problem with many variations such as covariate shift, representation learning, imitation learning, and finding invariant predictors. In this work, we present a framework that unifies multiple transfer learning problems, and can be used broadly for analyzing transfer learning problems. As a result, we challenge commonly-held assumptions about transfer, such as when sparse mechanism shift and layer freezing can be helpful.

^{*}Equal contribution ¹Cornell University ²Carnegie Mellon University ³Elementera AI. Correspondence to: My Phan <mail.myphan@gmail.com>.

Our framework also helps us design models and decide which components to transfer.

We model the transfer learning problem by viewing the *world* as a structural causal model (SCM, Pearl et al., 2000) with unknown constraints on its parameters. Learning about these constraints is what enables transfer. Within a world, the learner interacts with a sequence of *environments*. In each environment, the parameters of the SCM are chosen from the constrained domains, and the data is generated from the SCM with the chosen parameters. That is, each environment can be seen as a conditional intervention on a base SCM that represents the world. From each environment’s data, we learn its parameters, and thus gain information about the unknown constraints. This view can be used to describe covariate shift, representation learning, and transfer learning in supervised learning, reinforcement learning, and imitation learning (see Section 3).

Therefore the environments’ generation from the world follows the Independent Causal Mechanisms (ICM) Principle (Schölkopf et al., 2021), where the model is composed of independent modules. Within each module, we can have constraints on its parameters (called local constraints). We show how constraints on parameters across the model (called global constraints) can be computed from local constraints, allowing us to learn faster in new environments.

Contributions. In this work, we make the following contributions. First, we present a novel formalism for transfer learning that helps unify previously distinct areas of research, including covariate shift, representation, few-shot learning, and zero-shot learning, in which *environments* are generated from *worlds* by choosing local parameters of an SCM from constrained domains (Section 3). Second, we show how the constraints of global distributions can be computed from the constraints of local distributions, which enables us to analyze if transfer is feasible and derive a meta-algorithm to perform transfer (Section 4). Finally, we present case studies to illustrate how our work may be used to analyze and even challenge widely-held beliefs about transfer learning. We show that sparse mechanism shifts are neither necessary nor sufficient for transfer, and we show that freezing a layer of a network may either succeed or fail at transfer (Section 5).

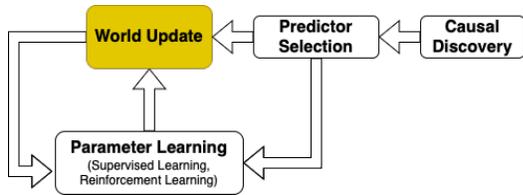


Figure 1. Relationships among problems in Section 3.1. We focus on transfer learning through the world update problem.

2. Motivation

Transfer learning means using knowledge from previous environments to speed up learning in a similar but different environment. For example, an agent can be trained to classify cows using photos where cows are often associated with pasture, and tested on photos where cows are often associated with sand (Beery et al., 2018). A robot may be trained to walk on Earth, then asked to perform on Mars where the gravity constant is different. Using our formalism, a world can consist of a set of possible associations between cows and the background, or a set of possible gravity constants, and environments correspond to specific values of the association or the gravity constant. Parameters learnt from previous environments can therefore be used as data points to learn about the constrained domain specified by the world. We call this the world update problem (Figure 1) and provide more details in Section 3.1. Formally, we define our **transfer learning** problem as the study of speedups we can get by *learning constraints* from previous environments to *reduce the hypothesis class* in the current environment — ruling out possible parameter values, so that we need less data or less computation for learning.

3. Problem Formalism

To that end, define a *world* as an SCM with a set of constrained domains for its parameters. We generate an *environment* from the world by choosing parameters from their domains. The learner starts out with some but not all information about the world: it might know something about the structure and/or the constraints. It then sees data from a sequence of environments. Parameters learnt in these environments become data points to estimate the domains, so that we can speed up learning in new environments.

For one example, the learner might know that some parameters of a model are constant in all environments but might not know their shared value. After seeing enough data in one environment, the learner can transfer these parameters to the next (parameter freezing). In this example, the transfer strategy is clear. It is not obvious in general how to learn unknown constraints for transfer; we discuss this question in App. A, and give a more complex example next.

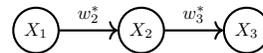


Figure 2. Causal graph for Example 3.1

Example 3.1. In the SCM of Figure 2, $X_3 = w_3^* X_2 \in \mathbb{R}^1$, $X_2 = w_2^* X_1 \in \mathbb{R}^2$, and $X_1 \in \mathbb{R}^3$ is drawn from a Gaussian distribution with mean w_1^* . In each environment, the parameters w_1^* , w_2^* , and w_3^* are selected from sets \mathcal{W}_1^* , \mathcal{W}_2^* , and \mathcal{W}_3^* . Here, \mathcal{W}_1^* is the set of all 2×1 matrices, $\mathcal{W}_2^* = \{[1, 2, 3; 4, 5, 6]\}$ is a singleton 2×3 matrix (same in all environments), and \mathcal{W}_3^* is \mathbb{R}^3 . Suppose the learner aims to predict X_3 from X_1 when X_2 is unobserved; it knows w_2^* is constant in all environments but doesn't know its value. Starting from a conditional linear model, $X_3 = w_{3|1} X_1$, the optimal (unknown) parameter value is $w_{3|1} = w_3^* w_2^*$ in each environment. Since w_2^* is the same in all environments, $w_{3|1}$ is always in the row space of w_2^* . Using examples of $w_{3|1}$ learnt from previous environments, the learner can discover this row space, then use it as the hypothesis class for $w_{3|1}$ in new environments to expedite learning.

We now present the formal definition of worlds and environments, extending Buesing et al. (2019). Upper-case letters denote random variables, lower-case letters their values, and bold-faced letters denote tuples.

Definition 3.1 (World). Let $\mathbf{X} = (X_1, \dots, X_n) \in \Omega_1 \times \dots \times \Omega_n$ be n variables where Ω_i denotes the domain of X_i . Let $\mathbf{O} \subseteq \mathbf{X}$ denote the observable variables and $\mathbf{Z} \subseteq \mathbf{X}$ denotes the unobservable ones. Let $\mathbf{U} = (U_1, \dots, U_n)$ where U_i is a vector of d_i independent uniform noise variables.¹ Let f_w be a parameterized function class and $\mathbf{P}_i \subseteq \{X_1, \dots, X_{i-1}\}$ be the set of possible causal parents of X_i , so that $X_i = f_{w_i^*}(\mathbf{P}_i, U_i)$, $1 \leq i \leq n$. (We call w_i^* an *atomic parameter*, to distinguish from aggregated parameters like $w_{3|1}$ in Ex. 3.1.) Let \mathcal{D}_i^* be the unconstrained domain of w_i^* and $\mathcal{W}_i^* \subseteq \mathcal{D}_i^*$ be the constrained domain of values for w_i^* . (When relevant, we call these *atomic domains*.) The *world* is the set of all possible SCMs obtained by selecting atomic parameters from their constrained domains: $w_i^* \in \mathcal{W}_i^*$, $\forall i, 1 \leq i \leq n$.

The learner sees a sequence of k environments $\{\mathcal{E}_e\}_{e=1}^k$.

Definition 3.2 (Environment). To make an *environment* \mathcal{E}_e , we divide the observable variables \mathbf{O} into ones the learner cannot intervene on (states) $\mathbf{S}^e \subseteq \mathbf{O}$ and ones the learner can intervene on (actions) $\mathbf{A}^e \subseteq \mathbf{O}$. We then select the parameter w_i^{*e} from \mathcal{W}_i^* for all i ($1 \leq i \leq n$).²

¹ U_i provides randomness for X_i . Using inverse CDF transform sampling, joint distributions can be expressed as SCMs (Buesing et al., 2019).

²This can be interpreted as conditional interventions, described in (Correa & Bareinboim, 2020).

In each environment, the learner learns about conditional distributions in order to make predictions. We leave the exact prediction task unspecified; instead we focus on the conditionals, and assume that discovering these will lead to task success. We represent conditional with parameter vectors:

Definition 3.3 (Description). Let \mathbb{I} and \mathbb{K} be tuples of indices, and write $\mathbf{X}_{\mathbb{I}}$ and $\mathbf{X}_{\mathbb{K}}$ for the corresponding tuples of variables. The learner represents the conditional distribution of $\mathbf{X}_{\mathbb{I}}$ given $\mathbf{X}_{\mathbb{K}}$, $P_{w_{\mathbb{I}|\mathbb{K}}}^{\mathbb{I}|\mathbb{K}}(\cdot)$, using a vector of parameters $w_{\mathbb{I}|\mathbb{K}}$.³ Many parameter vectors might be equivalent in a given environment \mathcal{E}_e , either because they yield identical conditional distributions, or because \mathcal{E}_e 's distribution over $\mathbf{X}_{\mathbb{K}}$ puts no weight on places where they differ (see Example 3.2 below). We write $P_w^{\mathbb{I}|\mathbb{K}} \simeq P_{w'}^{\mathbb{I}|\mathbb{K}}$, and we say that all of these vectors are *descriptions* of the same conditional distribution in \mathcal{E}_e . We use a bar to denote a set of equivalent descriptions: e.g., \bar{w} is the set of descriptions that are equivalent to w . If the conditional is for a single variable X_i given its parents, we call it an *atomic* description. For brevity, we sometimes write the conditional distribution as $\mathbf{X}_{\mathbb{I}}|\mathbf{X}_{\mathbb{K}}$ or $\mathbb{I}|\mathbb{K}$, leaving the parameters implicit.

After learning descriptions of some conditional distributions, the learner can update its estimates of the corresponding constrained domains. It can use these estimates to try to transfer to future environments. We discuss this process in more detail below (Sec. 4); there we address issues like how to reconcile descriptions learnt in different environments, and how to reason about the way different constrained domains relate to one another. To this end, we will give a toolbox of basic operations that can combine descriptions to build new ones step by step. These operations let us compute general domains from atomic ones: that is, $w_{\mathbb{I}|\mathbb{K}}$ lies in a domain $\mathcal{W}_{\mathbb{I}|\mathbb{K}} \subseteq \mathcal{D}_{\mathbb{I}|\mathbb{K}}$ that we can compute from the domains \mathcal{W}_i^* .

As mentioned above, whether two conditional distributions for $\mathbf{X}_{\mathbb{I}}|\mathbf{X}_{\mathbb{K}}$ are distinct can depend on the domain of $\mathbf{X}_{\mathbb{K}}$. The following example illustrates this:

Example 3.2 (Covariate Shift). We consider the SCM $X_2 = w_2^* X_1$ where $w_2^* = (1, 1)$. Let \mathcal{D}_1^* , the unrestricted domain of w_1^* , be $\{u, v\}$. If $w_1^* = u$, X_1^* is uniformly distributed on the line segment from $(0, 0)$ to $(1, 2)$, and if $w_1^* = v$, X_1^* is normally distributed on \mathcal{R}^2 . Therefore in Environment 1 when $w_1^* = u$, $\bar{w}_2^* = \{w \mid w(1, 2)^T = 3\}$ and in Environment 2 when $w_1^* = v$, $\bar{w}_2^* = \{(1, 1)\}$. Note that \bar{w}_2^* depends on w_1^* even when w_2^* is chosen independently from w_1^* . From Environment 1, the agent might learn $\hat{w}_2^* = (2, 0.5)$, which is a description of w_2^* and incorrectly transfers it to the next environment.

³For example, with discrete variables, $w_{\mathbb{I}|\mathbb{K}}$ could be the conditional probability table; or with Gaussian variables, $w_{\mathbb{I}|\mathbb{K}}$ could be the best linear predictor.

In Section 3.2 and Section 4.3, for simplicity, we assume that each w results in a unique distribution (Assumption 3.4); for example, this happens when the joint distribution is Gaussian. We discuss the case when the assumption does not hold in Section 4.4.

Assumption 3.4 (Unique Description). For any $w_1^*, \dots, w_n^* \in \mathcal{D}_1^*, \dots, \mathcal{D}_n^*$, for any two sets of disjoint indexes \mathbb{I}, \mathbb{K} , $\forall w', w \in \mathcal{D}_{\mathbb{I}|\mathbb{K}}$, if $P_w^{\mathbb{I}|\mathbb{K}}(\cdot) \simeq P_{w'}^{\mathbb{I}|\mathbb{K}}(\cdot)$ then $w = w'$.

3.1. Landscape around Transfer Learning

In this section, we differentiate transfer learning from related problems (Figure 1) and discuss how it relates to common concepts. Let $\mathbf{X}_{\mathbb{Q}}$ be some query variables that we wish to estimate: for example, the label in supervised learning or the expected return (Q function) in reinforcement learning. To help predict $\mathbf{X}_{\mathbb{Q}}$, the learner can solve several problems:

Causal discovery: Learn the graphical structure of the underlying SCM.

Predictor selection: Choose evidence variables $\mathbf{X}_{\mathbb{E}}$ to estimate $\mathbf{X}_{\mathbb{Q}}|\mathbf{X}_{\mathbb{E}}$.⁴

Parameter learning: Given a hypothesis set $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}$, learn $w_{\mathbb{Q}|\mathbb{E}}^e \in \widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}$.

World update: $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^0$ is the initial knowledge of the learner about $w_{\mathbb{Q}|\mathbb{E}}$ before any environment. Given estimated values of $w_{\mathbb{Q}|\mathbb{E}}^1, \dots, w_{\mathbb{Q}|\mathbb{E}}^e$ learnt from previous environments $\mathcal{E}_1, \dots, \mathcal{E}_e$, compute $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$ as an estimate of $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ to be used in parameter learning.

In this work, we focus on the world update problem, but the other problems are also important. The causal discovery problem is widely studied in the causal inference literature. Predictor selection is studied by works that focus on finding the invariant predictor (Zhang et al., 2020; Peters et al., 2015; Arjovsky et al., 2019). (Note that our framework is applicable even when there is no invariant predictor (see the Complex Colored MNIST example in Section 5.2.2).) Supervised learning and reinforcement learning methods focus on parameter learning.

Independent Causal Mechanism (ICM) principle and Sparse Mechanism Shift (SMS) hypothesis. The ICM principle (Schölkopf et al., 2021) states that the causal generative model consists of modules that do not inform or influence each other. The SMS hypothesis (Schölkopf et al., 2021) states that changes affect the causal factorization locally in-

⁴Evidence variables can be chosen to be either effects or causes of the query variables \mathbb{Q} . For example, medical tests (which measure effects) are often performed to diagnose an illness (the cause). We discuss in Example 5.3 how choosing the evidence variables can be a trade-off between complexity and achievable accuracy.

stead of affecting all factors simultaneously. Our framework follows the ICM principle, but does not assume the SMS hypothesis. In Section 5.1 we show that the SMS hypothesis is not enough to guarantee transfer.

Few-shot and Zero-shot Learning. Zero-shot learning happens when the constrained domain learnt from previous environments contains only one point (or a small neighborhood), and therefore can be used directly in the new environment. Few-shot learning happens when the learnt domain has low complexity, and therefore it takes a small number of samples to learn the parameter.

Covariate Shift. This is a special case of our setting where the model is $X \rightarrow Y$ and the conditional distribution $P(Y | X)$ is the same in all environments. In equations, $Y = f_{w_Y^*}(X, U_Y)$; $X = f_{w_X^*}(U_X)$, where w_Y^* is the same in all environments and w_X^* changes in different environments.

Representation Learning. Representation learning is a special case of our setting where X is the evidence and there exists a representation variable Z . Either $Z = f_{w_Z^*}(X, U_Z)$ and w_Z^* is the same in all environments or $X = f_{w_X^*}(Z, U_X)$ where w_X^* is the same in all environments.

Transfer in Supervised Learning. There is no action variable.

Time series. We have an SCM that consists of repeated copies of a base structure, one per time step; parameters and their domains are often shared across time steps. Examples include hidden Markov models and dynamic Bayes nets.

Transfer in Reinforcement Learning. There are action variables. There may be multiple time steps. The setting includes Markov decision processes (MDPs, Puterman, 1990) when there are no hidden variables. When there are hidden variables the setting includes partially-observable MDPs (POMDPs, Kaelbling et al., 1998).

Imitation Learning. Action variables change from non-intervenable in a previous environment to intervenable in the current environment.

3.2. Feasibility of Transfer Learning

In order for transfer to be feasible, our world updates must somehow reduce the difficulty of learning in new environments. Depending on the setting, we can measure this difficulty in different ways: e.g., estimates of the number of samples needed or of the computational resources required. For example, it is common to bound the required number of samples in terms of a complexity measure such as VC dimension in supervised learning (Vapnik et al., 1994) or Bellman-Eluder dimension in reinforcement learning (Jin et al., 2021). Transfer is feasible when our updated domain $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$ improves on the original one $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^0$ according to the appropriate complexity measure.

We formalize this idea below. For simplicity, we assume unique generators (Assumption 3.4), and we assume **realizability**, i.e., $\mathcal{W}_{\mathbb{Q}|\mathbb{E}} \subseteq \widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^0$. We believe these assumptions can be lifted, but we leave this direction to future work.

Definition 3.5 (Transfer feasibility of the world). Given a complexity measure $\text{comp} : 2^{\mathcal{D}_{\mathbb{Q}|\mathbb{E}}} \rightarrow \mathbb{R}$ we say that the world is transfer feasible if $\text{comp}(\mathcal{W}_{\mathbb{Q}|\mathbb{E}}) < \text{comp}(\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^0)$.

In a transfer-feasible world, our initial knowledge of the domains is meaningfully weaker than what we might eventually learn. This doesn't guarantee that learning will actually happen; for that, we need a transfer algorithm.

Definition 3.6 (Transfer feasibility of an algorithm). Let $\widehat{\mathcal{W}}_{\mathbb{I}|\mathbb{K}}^e \subseteq \mathcal{D}_{\mathbb{Q}|\mathbb{E}}$ be the estimate of $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ after environments $\mathcal{E}_1, \dots, \mathcal{E}_e$ output by algorithm A . Given a complexity measure $\text{comp} : 2^{\mathcal{D}_{\mathbb{Q}|\mathbb{E}}} \rightarrow \mathbb{R}$, if $\mathcal{W}_{\mathbb{Q}|\mathbb{E}} \subseteq \widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$ and $\text{comp}(\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e) < \text{comp}(\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^0)$ then A is called transfer feasible after e environments.

Zero-shot and few-shot learning are examples of transfer feasibility: they happen when $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$ is extremely simple, so that parameter learning is unnecessary (zero-shot) or takes only a small number of samples (few-shot).

4. Main Results

Following the ICM principle, the generator w_i^{*e} 's are chosen independently from their domains \mathcal{W}_i^* . In this section we show that the generator of interest, $w_{\mathbb{Q}|\mathbb{E}}^e$, is also chosen from a constrained domain $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ which can be computed from \mathcal{W}_i^* . We present a procedure to describe how to compute the constrained domain $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ from atomic constrained domains \mathcal{W}_i^* (Sections 4.1 to 4.3). Therefore, given some knowledge about the atomic constrained domains \mathcal{W}_i^* 's, we present a meta-algorithm to compute the constrained domain $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ (Section 4.4) for the world update problem. The meta-algorithm helps to design the architecture of the model and identify the component to transfer, which we demonstrate in Section 5.

4.1. Operations and Generators

In this section we define operations to compute a description of the distribution $\mathbb{Q}|\mathbb{E}$ from the atomic parameters w_i^* of the SCM. Using the operations, we define the computation procedure in Section 4.2.

First, we assume there exist operations to compute the product, sum-out and conditional \otimes , $\oplus_{\mathbb{J}}$ and $\odot_{\mathbb{J}}$ on the descriptions of the distributions. Using these operations, we show in Theorem 4.2 that we can compute a description of $\mathbb{Q}|\mathbb{E}$ from descriptions of X_i given its parents. Second, we show a description of X_i given its parents that only depends on w_i^* . Therefore, we can compute a description of $\mathbb{Q}|\mathbb{E}$ from

the atomic parameters w_i^* . In Section 4.3 we show how to use this computation to compute the constraint on this description of $\mathbb{Q}|\mathbb{E}$ from the atomic constraints \mathcal{W}_i^* .

First we define the description of the distribution of X_i given its parents that only depends on w_i^* . We call this the generator.

Definition 4.1 (Generator). Let Ω_i and Ω_i^p denote the domains of X_i and P_i (the parents of X_i) defined by the world. We assume that there exists a function $P_w^{i^*}(\cdot, \cdot) : \Omega_i \times \Omega_i^p \rightarrow \mathbb{R}$ parameterized by w such that for any $\mathbf{p} \in \Omega_i^p$, $P_w(\cdot, \mathbf{p}) : \Omega_i \rightarrow \mathbb{R}$ represents the distribution of $X_i = f_w(\mathbf{p}, U_i)$.⁵ Then $P_{w_i^*}^{i^*}(\cdot, \mathbf{p}) : \Omega_i \rightarrow \mathbb{R}$ represents the distribution of $X_i = f_{w_i^*}(\mathbf{p}, U_i)$ defined by the SCM. We call $P_{w_i^*}^{i^*}(\cdot, \cdot)$ the canonical (atomic) conditional distribution of X_i given its parents, and w_i^* its (atomic) generator.

Operation	Explanation
$\underbrace{w}_{\mathbb{I}, \mathbb{J}, \mathbb{K}} = \underbrace{u}_{\mathbb{I}, \mathbb{J}, \mathbb{K}} \otimes \underbrace{v}_{\mathbb{J}, \mathbb{K}}$	Computing the product $(\mathbb{I}, \mathbb{J}, \mathbb{K})$
$\underbrace{w}_{\mathbb{I}, \mathbb{K}} = \oplus_{\mathbb{J}} \underbrace{u}_{\mathbb{I}, \mathbb{J}, \mathbb{K}}$	Summing out \mathbb{J}
$\underbrace{w}_{\mathbb{I}, \mathbb{J}, \mathbb{K}} = \odot_{\mathbb{J}} \underbrace{u}_{\mathbb{I}, \mathbb{J}, \mathbb{K}}$	Conditioning on \mathbb{J}
$\underbrace{w}_{\mathbb{I}, \mathbb{K}} = \underbrace{u}_{\mathbb{I}, \mathbb{J}, \mathbb{K}} \odot \underbrace{v}_{\mathbb{J}, \mathbb{K}}$	Computing $\mathbb{I} \mathbb{K}$ from $\mathbb{I} \mathbb{J}, \mathbb{K}$ and $\mathbb{J} \mathbb{K}$

Table 1. Operations. $\underbrace{w}_{\mathbb{I}, \mathbb{K}}$ is equivalent to $(w; \mathbb{I}|\mathbb{K})$.

Let the expression $(u; \mathbb{I}|\mathbb{K})$, where $u \in \mathcal{D}_{\mathbb{I}|\mathbb{K}}$, denote that u is a description of $\mathbb{I}|\mathbb{K}$, and return u .

- Given $(u; \mathbb{I}|\mathbb{J}, \mathbb{K})$ where $u \in \mathcal{D}_{\mathbb{I}|\mathbb{J}, \mathbb{K}}$ and $(v; \mathbb{J}|\mathbb{K})$ where $v \in \mathcal{D}_{\mathbb{J}|\mathbb{K}}$, \otimes returns $(w; \mathbb{I}, \mathbb{J}|\mathbb{K})$ where $\mathbb{I}, \mathbb{J}|\mathbb{K}$ specifies the distribution \otimes computes. Assuming that u is a description of the distribution $\mathbb{I}|\mathbb{J}, \mathbb{K}$ and v is a description of the distribution $\mathbb{J}|\mathbb{K}$, \otimes returns a description w of $\mathbb{I}, \mathbb{J}|\mathbb{K}$. If u is the generator of $\mathbb{I}|\mathbb{J}, \mathbb{K}$ and v is the generator of $\mathbb{J}|\mathbb{K}$ then we call w the generator of $\mathbb{I}, \mathbb{J}|\mathbb{K}$.
- Given $(u; \mathbb{I}, \mathbb{J}|\mathbb{K})$ where $u \in \mathcal{D}_{\mathbb{I}, \mathbb{J}|\mathbb{K}}$, $\oplus_{\mathbb{J}}$ returns $(w; \mathbb{I}|\mathbb{K})$ where $\mathbb{I}|\mathbb{K}$ specifies the distribution $\oplus_{\mathbb{J}}$ computes. Assuming that u is a description of the distribution $\mathbb{I}, \mathbb{J}|\mathbb{K}$, $\oplus_{\mathbb{J}}$ returns a description of w of $\mathbb{I}|\mathbb{K}$. If u is the generator of $\mathbb{I}, \mathbb{J}|\mathbb{K}$ then we call w the generator of $\mathbb{I}|\mathbb{K}$.

Example: If $w_{\mathbb{I}, \mathbb{J}|\mathbb{K}}$ is a probability table then $\oplus_{\mathbb{J}}$ creates the probability table $w_{\mathbb{I}|\mathbb{K}}$ by summing over the

⁵For example, with a linear-Gaussian SCM, $P_w^{i^*}(\cdot, \mathbf{p})$ could be the Gaussian pdf and w could be the linear coefficient.

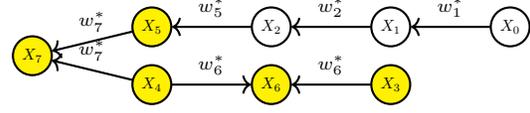


Figure 3. The causal graph and the scope of $w_{7|6,2}$ in yellow. The formula for $w_{7|6,2}$ only uses w_i^* of variables in the scope.

appropriate entries in $w_{\mathbb{I}, \mathbb{J}|\mathbb{K}}$.

- Given $(u; \mathbb{I}, \mathbb{J}|\mathbb{K})$ where $u \in \mathcal{D}_{\mathbb{I}, \mathbb{J}|\mathbb{K}}$, $\odot_{\mathbb{J}}$ returns $(w; \mathbb{I}|\mathbb{J}, \mathbb{K})$ where $\mathbb{I}|\mathbb{J}, \mathbb{K}$ specifies the distribution $\odot_{\mathbb{J}}$ computes. Assuming that u is a description of the distribution $\mathbb{I}, \mathbb{J}|\mathbb{K}$, $\odot_{\mathbb{J}}$ returns a description of w of $\mathbb{I}|\mathbb{J}, \mathbb{K}$. If u is the generator of $\mathbb{I}, \mathbb{J}|\mathbb{K}$ then we call w the generator of $\mathbb{I}|\mathbb{J}, \mathbb{K}$.

Example: Let $(X, Y) \sim Norm(\mu_{X,Y}, \Sigma_{X,Y})$ where

$$\mu_{X,Y} = \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \text{ and } \Sigma_{X,Y} = \begin{bmatrix} \Lambda_{XX} & \Lambda_{XY} \\ \Lambda_{YX} & \Lambda_{YY} \end{bmatrix}^{-1}.$$

Then $w_{X,Y} = (\mu_{X,Y}, \Sigma_{X,Y})$ and $w_{Y|X} = \odot_X w_{X,Y} = (\mu_Y - \Lambda_{YY}^{-1} \Lambda_{YX}(x - \mu_X), \Lambda_{YY}^{-1})$ (Bishop, 2006).

For convenience, we define the operation \odot from \otimes and \oplus :

$$\begin{aligned} (u; \mathbb{I}|\mathbb{J}, \mathbb{K}) \odot (v; \mathbb{J}|\mathbb{K}) &:= \oplus_{\mathbb{J}}((u; \mathbb{I}|\mathbb{J}, \mathbb{K}) \otimes (v; \mathbb{J}|\mathbb{K})) \quad (1) \\ &= (w; \mathbb{I}|\mathbb{K}). \quad (2) \end{aligned}$$

Given $(u; \mathbb{I}|\mathbb{J}, \mathbb{K})$ where $u \in \mathcal{D}_{\mathbb{I}|\mathbb{J}, \mathbb{K}}$ and $(v; \mathbb{J}|\mathbb{K})$ where $v \in \mathcal{D}_{\mathbb{J}|\mathbb{K}}$, \odot returns $(w; \mathbb{I}|\mathbb{K})$ where $\mathbb{I}|\mathbb{K}$ specifies the distribution \odot computes. Assuming that u is a description of the distribution $\mathbb{I}|\mathbb{J}, \mathbb{K}$ and v is a description of the distribution $\mathbb{J}|\mathbb{K}$, \odot returns a description w of $\mathbb{I}|\mathbb{K}$. If u is the generator of $\mathbb{I}|\mathbb{J}, \mathbb{K}$ and v is the generator of $\mathbb{J}|\mathbb{K}$ then we call w the generator of $\mathbb{I}|\mathbb{K}$.

We summarize the operations in Table 1. To simplify notations, we omit the indexes \mathbb{I}, \mathbb{J} and \mathbb{K} from the formula and write $u \otimes v$ instead of $(u; \mathbb{I}|\mathbb{J}, \mathbb{K}) \otimes (v; \mathbb{J}|\mathbb{K})$ whenever the context is clear. Similarly we write $\oplus_{\mathbb{J}} u$, $\odot_{\mathbb{J}} u$ and $u \odot v$ without the indexes. We reuse $w_{\mathbb{I}|\mathbb{K}}$ from Definition 3.3 to denote the generator of $\mathbb{I}|\mathbb{K}$ ⁶.

4.2. Scopes and Computation Trees

The result in this section does not depend on whether Assumption 3.4 holds. We present a theorem to compute the generator $w_{\mathbb{Q}|\mathbb{E}}$ from atomic generators w_i^* 's using the defined operations. Using the computation procedure, we can compute the constraint of this description of $\mathbb{Q}|\mathbb{E}$ from the constraints \mathcal{W}_i^* of the atomic parameters in Section 4.3.

⁶Note that w_i^* denotes the generator of X_i given its parents while w_i denotes the generator marginal of X_i .

Naively we can compute $w_{\mathbb{Q}|\mathbb{E}}$ from the joint distribution $w_{\mathbb{Q},\mathbb{E}}$, which could be calculated from the set of all w_i^* using variable elimination. However we derive a smaller set of variables w_i^* from which $w_{\mathbb{Q}|\mathbb{E}}$ can be computed, called the scope of $w_{\mathbb{Q}|\mathbb{E}}$ and denoted $\text{Scope}(\mathbb{Q}|\mathbb{E})$. We illustrate the scope in Figure 3 and Figure 4. Our method recursively computes $w_{\mathbb{Q}|\mathbb{E}}$ by building a computation tree top-down. It starts with $w_{\mathbb{Q}|\mathbb{E}}$ as the root, recursively chooses a node in the tree and decomposes it into its children using the operations until all of the leaf nodes are of the form w_i^* . Let $w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^* := \{w_i^* | i \in \text{Scope}(\mathbb{Q}|\mathbb{E})\}$.

Variable elimination is a special case of our method if we perform $w_{\mathbb{Q}|\mathbb{E}} = \circ_{\mathbb{E}} w_{\mathbb{Q},\mathbb{E}}$ in the first step and then compute the joint $w_{\mathbb{Q},\mathbb{E}}$ from the atomic generators using \oplus and \otimes in the remaining steps. We compare our method with variable elimination in Figure 4 and Appendix E.

Theorem 4.2 (Top-down computation tree, informal). *$w_{\mathbb{Q}|\mathbb{E}}$ can be calculated from w_i^* 's in the scope of $w_{\mathbb{Q}|\mathbb{E}}$ using \oplus, \otimes and \circ so that each w_i^* appear at most once in reverse topological order. We denote the resulting formula $g_{\mathbb{Q}|\mathbb{E}}(w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*)$, also called a computation tree of $w_{\mathbb{Q}|\mathbb{E}}$.*

For clarity, we show some examples of the computation trees in Figure 4. We also show other examples (including the computation tree for Q -learning) in App. E.

4.3. Computing Global Constraints from Local Constraints

Given the computation tree, we can compute the set $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ that $w_{\mathbb{Q}|\mathbb{E}}$ is chosen from. Let $\mathcal{W}_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^* = \{W_i^* | w_i^* \in w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*\}$ be the corresponding constrained domain of $w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*$. From Theorem 4.2, $w_{\mathbb{Q}|\mathbb{E}}$ is computed by selecting $w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*$ and computing $g_{\mathbb{Q}|\mathbb{E}}(w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*)$. Therefore $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ is the set of all possible $w_{\mathbb{Q}|\mathbb{E}}$ computed by selecting $w_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*$ from $\mathcal{W}_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*$ and then applying the computation tree.

We define a notation for the set of function values over a set of inputs: for any set \mathcal{S} and a function f such that \mathcal{S} is in the domain of f , $f(\mathcal{S}) := \{f(s) | s \in \mathcal{S}\}$. Another approach to compute $g_{\mathbb{Q}|\mathbb{E}}(\mathcal{W}_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*)$ in the case of supervised learning is discussed in Appendix E.4.

Corollary 4.3. *$\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ can be computed from $\mathcal{W}_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^* \subseteq \{W_1^*, \dots, W_n^*\}$: $\mathcal{W}_{\mathbb{Q}|\mathbb{E}} = g_{\mathbb{Q}|\mathbb{E}}(\mathcal{W}_{\text{Scope}(\mathbb{Q}|\mathbb{E})}^*)$. Therefore the transfer feasibility of the world can be determined from W_1^*, \dots, W_n^* by computing $\text{comp}(\mathcal{W}_{\mathbb{Q}|\mathbb{E}})$ ⁸.*

When the learner does not know the exact formula, we in-

⁷There can be multiple computation trees depending on the choice the learner made.

⁸The second statement of Corollary 4.3 assumes Assumption 3.4 holds while the first statement does not depend on Assumption 3.4.

roduce the following properties to analyze how the local constraints can propagate throughout the network. We illustrate in Sec. 5 how to use the properties to derive constraints on the quantity we would like to estimate.

Property 4.3.1 (Propagatable). Let $\mathcal{D}_1, \mathcal{D}_2$, and $\mathcal{D}_{1 \circ 2}$ be the unconstrained domain of w_1, w_2 , and $w_1 \circ w_2$ respectively where $\circ : \mathcal{D}_1 \circ \mathcal{D}_2 \rightarrow \mathcal{D}_{1 \circ 2}$ be an operation. We define how a constraint can propagate to the left and to the right. Let $\mathcal{F} \subseteq \mathcal{D}_2 \cup \mathcal{D}_{1 \circ 2}$. An operation \circ is called \mathcal{F} -left-propagatable if for any $w_2 \in \mathcal{F}, w_1 \circ w_2 \in \mathcal{F}$. Now, let $\mathcal{F} \subseteq \mathcal{D}_1 \cup \mathcal{D}_{1 \circ 2}$. An operation \circ is called \mathcal{F} -right-propagatable if for any $w_1 \in \mathcal{F}, w_1 \circ w_2 \in \mathcal{F}$.

4.4. Transfer Process

Now we are equipped to derive a process to learn $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$ from $\widehat{w}_{\mathbb{Q}|\mathbb{E}}^1, \dots, \widehat{w}_{\mathbb{Q}|\mathbb{E}}^e$ learnt from multiple environments $\mathcal{E}_1, \dots, \mathcal{E}_e$. Let $\widehat{\mathcal{W}}_i^*$'s for some $i, 1 \leq i \leq n$ be some local information the learner knows about \mathcal{W}_i^* (if the learner does not know anything about \mathcal{W}_i^* it can assume $\widehat{\mathcal{W}}_i^* = \mathcal{D}_i^*$). Let $\widehat{\text{Scope}}(\mathbb{Q}|\mathbb{E})$ be an estimation of $\text{Scope}(\mathbb{Q}|\mathbb{E})$ such that $\text{Scope}(\mathbb{Q}|\mathbb{E}) \subseteq \widehat{\text{Scope}}(\mathbb{Q}|\mathbb{E})$. Without any knowledge we can set $\widehat{\text{Scope}}(\mathbb{Q}|\mathbb{E})$ to be all variables. Let $\widehat{g}_{\mathbb{Q}|\mathbb{E}}$ be an approximation of $g_{\mathbb{Q}|\mathbb{E}}$ obtained by knowledge or by the propagation properties such that $g_{\mathbb{Q}|\mathbb{E}}(\mathcal{W}_{\text{Scope}(\mathbb{Q}|\mathbb{E})}) \subseteq \widehat{g}_{\mathbb{Q}|\mathbb{E}}(\mathcal{W}_{\widehat{\text{Scope}}(\mathbb{Q}|\mathbb{E})})$ for any input \mathcal{W} . For example, if all operations are \mathcal{F} -right propagatable and the left-most term is in \mathcal{F} , $\widehat{g}_{\mathbb{Q}|\mathbb{E}}$ can output \mathcal{F} . We call a set \mathcal{W} parameterized by a parameter θ if \mathcal{W} is determined by θ . For example, if \mathcal{W} is the set of all w such that $f_\theta(w) = 0$ for a known class of function f_θ , then \mathcal{W} is parameterized by θ . Algorithm 1 captures our procedure to compute $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$. In App. A, we describe an algorithm to learn the constraint from the approximations of the parameters in previous environments.

If Assumption 3.4 does not hold, Algorithm 1 needs to be modified. Consider the case in Example 3.2, when $w_2^* = (1, 1)$ but $\overline{w}_2^* = \{w : w^T(1, 2) = 3\}$, and w_2^* is the same in all environments. From Environment 1, the agent might learn $\widehat{w}_2^* = (2, 0.5)$, which is a description of w_2^* and incorrectly transfers it to the next environment. Instead, for all environments e , the agents need to learn the set of all possible descriptions in $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$, therefore in this case $\widehat{w}_{\mathbb{Q}|\mathbb{E}}^i$ does not denote just a point but a set of all possible, equivalent descriptions. In Example 3.2, it would be equivalent to learn \overline{w}_2^* from Environment 1. The learner then uses the set of all possible descriptions $\widehat{w}_{\mathbb{Q}|\mathbb{E}}^i$ learned from previous environment to learn the set Θ of θ such that $\forall \theta \in \Theta: \forall i = 1, \dots, e, \exists w^i \in \widehat{w}_{\mathbb{Q}|\mathbb{E}}^i$ such that $f_\theta(w^i) = 0$. Let $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e = \cup_{\theta \in \Theta} \{w : f_\theta(w) = 0\}$. Then the agent will learn the set of all possible descriptions $\widehat{w}_{\mathbb{Q}|\mathbb{E}}^{e+1}$ in $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$.

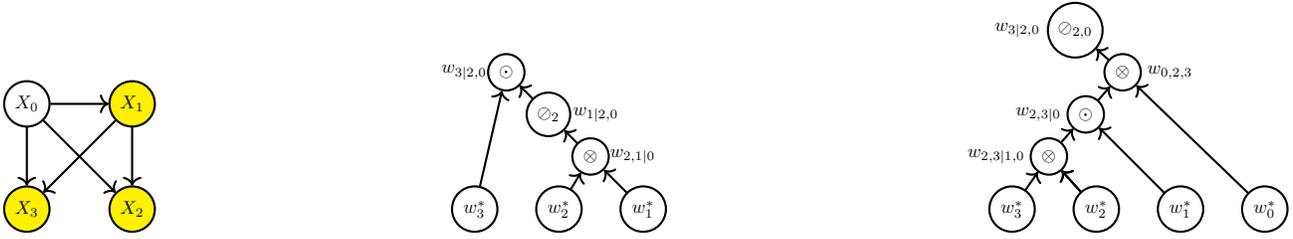


Figure 4. Causal graph and the computation trees for $w_{3|2,0}$. For clarity, we do not draw the uniform noise variables. **(Left)**: Causal graph and the scope of $w_{3|2,0}$ in yellow. **(Middle)** A computation tree of $w_{3|2,0}$ constructed from Thm. 4.2, which is equivalent to calculating $p(X_3|X_2, X_0) = \sum_{x_1} p(X_3|X_0, X_1 = x_1)p(X_1 = x_1|X_2, X_0)$; $p(X_1 = x_1|X_2, X_0) = \frac{p(X_1=x_1, X_2|X_0)}{\sum_{x_1'} p(X_1=x_1', X_2|X_0)}$ and $p(X_1 = x_1, X_2|X_0) = p(X_2|X_0, x_1)p(x_1|X_0)$ for discrete distributions. Variable elimination cannot construct this computation tree because conditioning (\odot) is performed to compute $p(X_1 = x_1|X_2, X_0)$ before sum-out and product (\odot). w_3^* goes through only 1 operation. **(Right)** A computation tree of $w_{3|2,0}$ constructed by naive variable elimination on the joint distribution $w_{0,2,3}$. Each leaf goes through at least 2 operations (at least one to construct the joint distribution and then the conditional \odot).

Algorithm 1 Meta-Algorithm for Transfer

Input: Estimated constrained domain $\widehat{\mathcal{W}}_1^*, \dots, \widehat{\mathcal{W}}_n^*$, parameters learnt from previous environments $\widehat{w}_{\mathbb{Q}|\mathbb{E}}^1, \dots, \widehat{w}_{\mathbb{Q}|\mathbb{E}}^e$; estimated $\widehat{\text{Scope}}(\mathbb{Q}|\mathbb{E})$; estimated computation tree $\widehat{g}_{\mathbb{Q}|\mathbb{E}}$ to compute $w_{\mathbb{Q}|\mathbb{E}}$ from w_i^* 's

Output: $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$, an estimation of $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$

Calculate $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e = \widehat{g}_{\mathbb{Q}|\mathbb{E}}(\mathcal{W}_{\widehat{\text{Scope}}(\mathbb{Q}|\mathbb{E})}^*)$ parameterized by unknown parameter θ ;

Learn θ using $\widehat{w}_{\mathbb{Q}|\mathbb{E}}^1, \dots, \widehat{w}_{\mathbb{Q}|\mathbb{E}}^e$ and compute $\widehat{\mathcal{W}}_{\mathbb{Q}|\mathbb{E}}^e$ from θ .

5. Case Studies

We now demonstrate how our framework can be used to analyze widely-held beliefs about transfer learning.

5.1. Sparse Mechanism Shift

The Sparse Mechanism Shift (SMS) Hypothesis states that changes are often sparse and local in a causal graph (Schölkopf et al., 2021; Lopez et al., 2023; Lachapelle et al., 2022; Bereket & Karalestos, 2024). We demonstrate that SMS does not always imply transfer feasibility and vice versa.

Example 5.1 (The Chain Problem). Assume X_0 is normally distributed on \mathbb{R}^2 . For $i = 1, \dots, 100$, $X_i = f_{w_i^*}(X_{i-1})$. X_1, \dots, X_{99} are unobserved. Consider the linear model where $f_{w_i^*}(X_{i-1}) = w_i^* X_{i-1}$. The learner observes $X = X_0$ and $Y = X_{100}$ and would like to predict Y from X .

The learner fits a linear model $Y = wX$ to predict Y from X , where w is any 2×2 matrix, making $\widehat{\mathcal{W}}_{100|0}^0$ the set of all 2×2 matrices. The complexity measure comp is the number of parameters that need to be learned, with $\text{comp}(\widehat{\mathcal{W}}_{100|0}^0) = 4$. App. F shows that if $w_i^* (\forall 2 \leq i \leq 99)$

is the same in all environments (sparse changes), transfer learning may not be feasible. However, if only w_1^* is the same in all environments (non-sparse changes), transfer may be feasible, reducing $\text{comp}(\widehat{\mathcal{W}}_{100|0}^1)$ — the dimension of the hypothesis class after environment 1 — to 2.

5.2. Freezing Layers of Neural Networks

Freezing a layer of a neural network learnt from previous environments is standard in transfer learning (Yosinski et al., 2014; Long et al., 2015; Kirichenko et al., 2023; Tajbakhsh et al., 2016; Guo et al., 2019). Our framework helps to explain when this practice works (Section 5.2.1). In the next few subsections, we show that this standard practice may be ill-advised: performance can depend strongly on exactly which parameters we freeze.

5.2.1. THE CHAIN PROBLEMS

We sketch an example where freezing layers of a NN is correct for transfer learning – see Appendix F for more details. We use the result in (Rolnick & Kording, 2020), where it was shown that if the Linear Region Assumption is satisfied, ReLU networks can be identified up to permutation and scaling.

Example 5.2 (Freezing layers for neural networks). Modify the SCM in Example 5.1 so that $f_{w_i^*}$ is a single ReLU layer with weights w_i^* . Assume that in all environments \mathcal{E}_e , the Linear Region Assumption (Rolnick & Kording, 2020) is satisfied and that X_0 's distribution satisfies that the description set of $w_{100|0}^e$ consists of only equivalent neural networks up to permutation and scaling. \odot is stacking layers: for $i > j > k$, $w_{i|k} = w_{i|j} \odot w_{j|k} = (w_{i|j}, w_{j|k})$. Therefore the generator $w_{100|0} = (w_{100}^*, \dots, w_1^*)$. For any i , if w_i^* is the same in all environments, $\mathcal{W}_{100|0}$ is the set of $(w_{100}^*, \dots, w_1^*)$ up to permutation and scaling with the same w_i^* in all environments. Thus, we can learn layer w_i^* from the first environment and freeze it in later environments.

We show in App. F.2 that even with slight modifications from the above example, θ no longer corresponds to an entire layer, and therefore freezing a layer results in poor performance. In the next section, we present a more complex example to explore this phenomenon.

5.2.2. THE COMPLEX COLORED MNIST PROBLEM

We now show when freezing an entire layer of the network leads to poor transfer, while freezing a subset of the layer leads to successful transfer. This example also demonstrates our approach where we treat transfer learning as a new problem separated from finding an invariant predictor which are often pursued in previous works (Zhang et al., 2020; Peters et al., 2015; Arjovsky et al., 2019). It also demonstrates a case of *representation learning*, where the distribution of the image X given its parents is assumed to be the same in all environments.

Example 5.3 (Complex Colored MNIST). This binary classification problem is based on MNIST: instances are colorized images of digits, and the colors are spuriously correlated with labels. In the original dataset (Arjovsky et al., 2020), binary labels are based on the digit with probability 0.25 and flipped otherwise. In 2 environments, the image is colored with a color that correlated with the binary label with probability 0.1 and 0.9 respectively. We modify the environments to make it more difficult. In \mathcal{E}_1 , the correlation between digit identity and label is 0.1, and the correlation between color and label is 0.2. In \mathcal{E}_2 , the correlation is 1 between digit and label, and 0.4 between color and label.

We illustrate the trade-off between achievable accuracy and sample complexity for choosing different predictors (Section 3.1) using the original Colored MNIST problem. Predicting the label using the digit in Environment 1 allows zero-shot learning in Environment 2 with a maximum accuracy of 0.75. Predicting the label using the color in Environment 2 requires samples from Environment 2 but enables a maximum theoretical accuracy of 0.9. Therefore, predictor choice affects the achievable accuracy and sample complexity in the new environment. In the Complex Colored MNIST problem, the label distribution given the digit changes across environments, making previous methods (Zhang et al., 2020; Peters et al., 2015; Arjovsky et al., 2019) that identify the invariant predictors for zero-shot learning ineffective.

Transfer Algorithm. Let $v(X)$ denote the flattened vector of X . We first derive the model for $\widehat{\mathcal{W}}_{Y|X}^1$. We assume that (1) the parents \mathbf{P}_X of the image X are discrete random variables (2) conditioning on $\mathbf{P}_X = p$, $v(X)$ is in the exponential family $p(x|p) = \frac{1}{s} h\left(\frac{1}{s}x\right) g(\lambda_p) \exp\left(\frac{1}{s}\lambda_p^T x\right)$ where the scale parameter s does not depend on p (Section 4 in (Bishop, 2006)); (3) X is d-separated from Y given \mathbf{P}_X , (4) the distribution $X|\mathbf{P}_X$ is unchanged in all environ-

ments.⁹

Via Theorem 4.2 we compute $P(Y|X)$ (Appendix F.3):

$$P(Y|X) = \frac{\sum_p \exp\left(\frac{1}{s}\lambda_p^T v(X) + \log g(\lambda_p) + \log P(p, Y)\right)}{\sum_y \sum_p \exp\left(\frac{1}{s}\lambda_p^T v(X) + \log g(\lambda_p) + \log P(p, y)\right)} \quad (3)$$

where \sum_p denotes the summation over all possible values p of the parents of X . The blue terms are affected by w_X^* while the black terms are not. Eq. 3 can be implemented as a 1-layer network, but in this implementation, freezing a layer is the same as freezing the entire model. This would lead to poor performance because the label in Environment 2 is almost the opposite of Environment 1. So, to allow more flexibility in freezing and give the competing model where the entire layer is frozen a better advantage, we derive a 2-layer implementation of Eq. 3. Rearranging the terms we get the following 2-layer network: the first layer implements

$$P(p|X) = \frac{\exp\left(\frac{1}{s}\lambda_p^T v(X) + \log g(\lambda_p) + \log P(p)\right)}{\sum_p \exp\left(\frac{1}{s}\lambda_p^T v(X) + \log g(\lambda_p) + \log P(p)\right)},$$

and the second layer implements

$$P(Y|X) = \sum_p P(p|X) \log P(Y|p).$$

To learn, we make two additional assumptions: (5) for any choice of the parameters in the constrained domain, the description set of the any network in the constrained domain is the set of all equivalent weights in which the linear coefficients differ only by permutation (the biases could differ by other means) and (6) the optima of the loss function are in the description set.

Given these assumptions, our method recommends transferring by freezing the linear coefficient in the first layer: from (6), minimizing the loss will result in a description. From (5), similar to the discussion in Section F.2, we can freeze the first layer’s linear coefficient to transfer. This reflects the fact that the linear coefficient represents $P(X|\mathbf{P}_X)$, which is the same in all environments. We compare to other transfer strategies that might seem reasonable without our derivations, such as freezing the entire first layer.

Transfer Feasibility of the Algorithm. Let the complexity measure comp be the number of parameters to be learned. Assuming that $n_p = 20$, without transfer learning, $\text{comp} = 47040 + 20 + 40$; with transfer learning, $\text{comp} = 20 + 40$. The details are provided in Appendix F.3.

Experiment. Fig. 5 compares transfer performance of several methods: learning the model in the new environment without transfer, freezing the first layer of the two-layer

⁹As an example, a detailed hypothetical model that satisfies our assumptions is given in Appendix F.3.

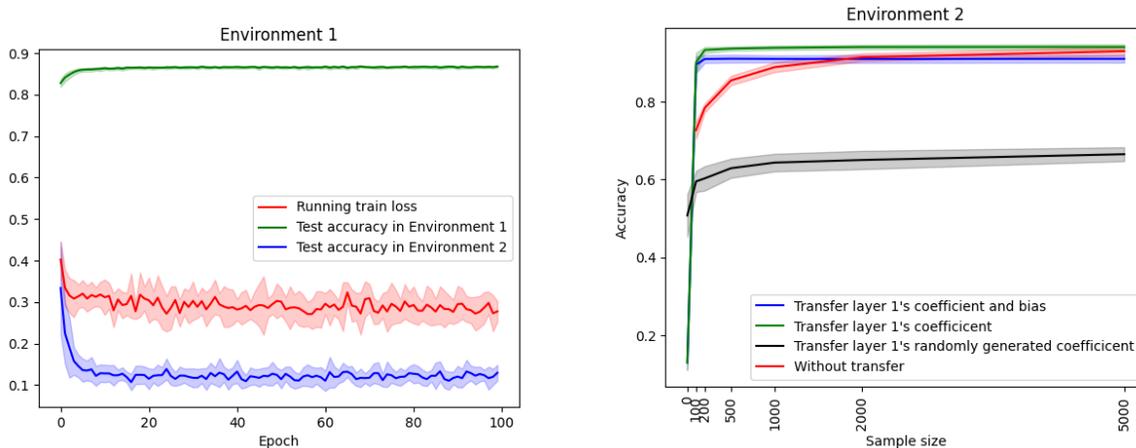


Figure 5. The Complex Colored MNIST problem, Ex. 5.3. **(Left)**: Zero-shot transfer performance of model trained in Env. 1. Accuracy in Env. 2 is poor because the labels in Env. 1 and Env. 2 are very different. **(Right)** Performance of various transfer methods as they train on Env. 2. Transferring the first layer’s linear coefficient results in the best performance.

model, and freezing just the linear coefficient of the first layer. App. F.3 gives implementation details using Ray Tune (Liaw et al., 2018). We show that the strategy recommended by our derivation (freezing just the linear coefficient of the first layer) performs best. We also compare to a random “transfer” strategy (freezing the first layer’s linear coefficient to a random value instead of a learned one) to provide a baseline. Note that when the sample size available in Env. 2 grows larger than in the plot (5000) it might be possible that the model without transfer might outperform the model that transfers layer 1’s coefficient due to limited samples in the first environment, failure of the model to find the true coefficient in the first environment or other factors.

5.3. Reinforcement Learning

We now illustrate how our framework applies to reinforcement learning with a simple example. Suppose the dynamics $P(s_{t+1}|s_t, a_t)$ remain constant across environments, and the expected reward is $E[R_t|s_t, a_t] = w^T \phi(s_t, a_t)$, where ϕ is constant across environments and w varies (Barreto et al., 2017). Suppose that states and actions are discrete, so that the dynamics are a table of size $n_s \times n_s n_a$ and ϕ is a table of size $d \times n_s n_a$, where n_s, n_a and d are the number of states, actions and dimension of ϕ . Observe that the Q function of any policy in any environment can be written as $Q^\pi(s, a) = E[R_t|s, a] + \gamma E[V(s_{t+1}|s, a)]$. Considered as a vector of length $n_s n_a$ indexed by (s, a) , the first term on the RHS is in the row space of the ϕ table, and the second term is in the row space of the dynamics table. So Q^π is in the span of these two row spaces.

Suppose there are $n_a = 4$ actions and $n_s = 5$ states, and ϕ has dimension $d = 2$. Then Q^π (of dimension 20) is in a subspace of dimension $n_s + d = 7$. Once we evaluate

enough policies in different environments, we can a basis for this subspace. In a new environment, we can constrain $Q(s, a)$ to be in this subspace, potentially saving both data and computation: e.g., we can estimate the environment’s w from observations (saving data compared to estimating the vector $E[R(s, a)]$), then run value iteration or Q iteration entirely in the subspace (saving data by not estimating the dynamics, and computation by working in the subspace). Barreto et al. (2017) compute an improved policy but not the optimal policy in the new environment. Brantley et al. (2021) compute the optimal policy, but do not address learning.

6. Conclusion

Our paper provides a framework for transfer learning that can guide how to perform transfer and analyze if common heuristics such as freezing a layer may or may not succeed. Our framework can be used to model a wide range of transfer scenarios including those in supervised learning, imitation learning and reinforcement learning. It can also help users design models that enable transfer or select components to transfer.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization, 2019. URL <https://arxiv.org/abs/1907.02893>.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization, 2020.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H., and Silver, D. Successor features for transfer in reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, pp. 4058–4068, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Beery, S., Van Horn, G., and Perona, P. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- Bereket, M. and Karaletsos, T. Modelling cellular perturbations with the sparse additive mechanism shift variational autoencoder. *Advances in Neural Information Processing Systems*, 36, 2024.
- Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. ISBN 0387310738.
- Brantley, K., Mehri, S., and Gordon, G. J. Successor feature sets: Generalizing successor representations across policies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11774–11781, 2021.
- Buesing, L., Weber, T., Zwols, Y., Heess, N., Racaniere, S., Guez, A., and Lespiau, J.-B. Woulda, coulda, shoulda: Counterfactually-guided policy search. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BJG0voC9YQ>.
- Chen, Y. and Bühlmann, P. Domain adaptation under structural causal models. *The Journal of Machine Learning Research*, 22(1):11856–11935, 2021.
- Correa, J. and Bareinboim, E. A calculus for stochastic interventions: causal effect identification and surrogate experiments. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(06):10093–10100, Apr. 2020. doi: 10.1609/aaai.v34i06.6567. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6567>.
- Dietterich, T. G. Hierarchical reinforcement learning with the MAXQ value function decomposition, 1999. URL <https://arxiv.org/abs/cs/9905014>.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks, 2016.
- Gulrajani, I. and Lopez-Paz, D. In search of lost domain generalization. *CoRR*, abs/2007.01434, 2020. URL <https://arxiv.org/abs/2007.01434>.
- Guo, Y., Shi, H., Kumar, A., Grauman, K., Rosing, T., and Feris, R. Spottune: Transfer learning through adaptive fine-tuning. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4800–4809, Los Alamitos, CA, USA, jun 2019. IEEE Computer Society. doi: 10.1109/CVPR.2019.00494. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2019.00494>.
- Huang, B., Feng, F., Lu, C., Magliacane, S., and Zhang, K. Adarl: What, where, and how to adapt in transfer reinforcement learning, 2021. URL <https://arxiv.org/abs/2107.02729>.
- Jin, C., Liu, Q., and Miryoosefi, S. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms, 2021.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- Kirichenko, P., Izmailov, P., and Wilson, A. G. Last layer re-training is sufficient for robustness to spurious correlations. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Zb6c8A-Fghk>.
- Lachapelle, S., Rodriguez, P., Sharma, Y., Everett, K. E., Le Priol, R., Lacoste, A., and Lacoste-Julien, S. Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ica. In *Conference on Causal Learning and Reasoning*, pp. 428–484. PMLR, 2022.
- Li, X., Grandvalet, Y., and Davoine, F. Explicit inductive bias for transfer learning with convolutional networks. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2825–2834. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/li18a.html>.
- Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J. E., and Stoica, I. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*, 2018.

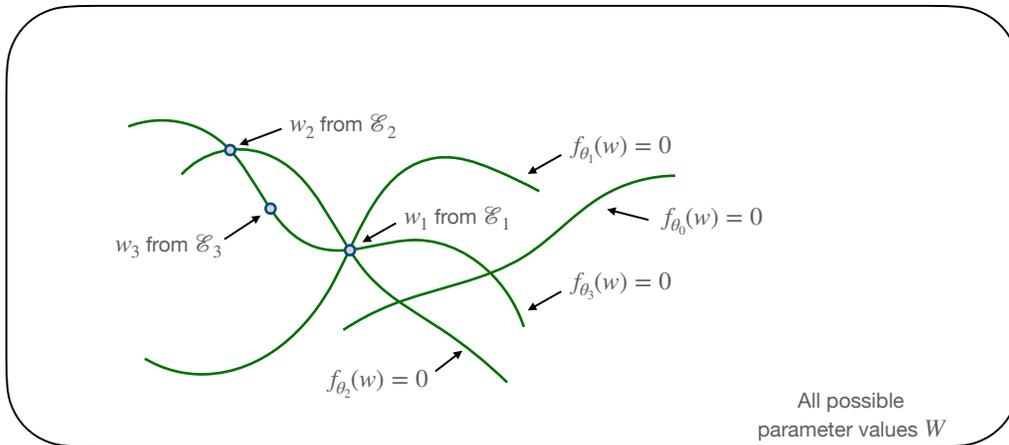


Figure 6. Constraint learning problem without noise. Each marker \circ represents a possible environment parameter (a value of w), and each teal curve represents a possible constraint (a value of θ).

Supplementary Material: When is Transfer Learning Possible?

In Section A we describe an algorithm to learn the constraint using w^e 's from previous environment. In Section B we discuss related works. In Section C we discuss future works. In Section D we discuss more details of the setting related to reinforcement learning. In Section E we provide details for Section 4 including the proof of Theorem 4.2 and more examples of computation trees including Q-Learning. In Section F we provide details for the case studies (Section 5) including the details of the Colored MNIST experiments. In Section G we discuss the limitation of our paper.

A. Constraint Learning

In this section we describe a simple method for learning constraints from examples. This is an important prerequisite for transfer learning: the tools we describe in the main paper can help us propagate constraints, so that we know how a constraint on one parameter of our model might imply a constraint on other parameters, but we have to combine this sort of propagation with a way to extract constraints from observed data.

There are many ways that one might hope to improve on the constraint learning method we describe here. The exact method of constraint learning is not central to the current paper; all that is important here is that some learning algorithm proposes constraints that fit into our tools for transfer. So, our intent is not to be state-of-the-art, but merely to demonstrate that this kind of learning is possible. We leave it to future work to refine and optimize our learning method.

We have two main goals in designing our constraint learning method. First, we'd like to be agnostic to how the environments are generated: for example, we don't want to assume that they are sampled i.i.d. from a fixed distribution over the set of allowed environments. Second, we'd like to maximize the benefits we see from transfer: we'd like to learn helpful constraints that narrow down the set of parameters that we have to learn from, but we'd like to avoid costly mistakes where incorrect constraints limit performance in new environments.

We'll build up the constraint learning problem in several steps, introducing new issues at each step. The first version of our learning problem is illustrated in Figure 6. We are trying to learn some parameter w that takes a different (unknown) value w_i in each new environment \mathcal{E}_i . We assume a family of possible constraints that might apply to w : $f_\theta(w) = 0$, where θ is an (unknown) parameter vector. If we want to enforce multiple constraints on w , we can have f return a vector; in this case θ contains the parameters of all constraints, concatenated together.

As we gain experience in a sequence of environments $\mathcal{E}_1, \mathcal{E}_2, \dots$, we can hope to learn a constraint that describes the feasible values of w_1, w_2, \dots . For example, after we see data from \mathcal{E}_1 , we can learn an estimate of w_1 , which lets us narrow down the possible values of θ . In the figure, θ_0 is inconsistent with w_1 , since w_1 does not lie on the curve $f_{\theta_0}(w) = 0$. On the other hand, θ_1, θ_2 , and θ_3 are all consistent with w_1 .

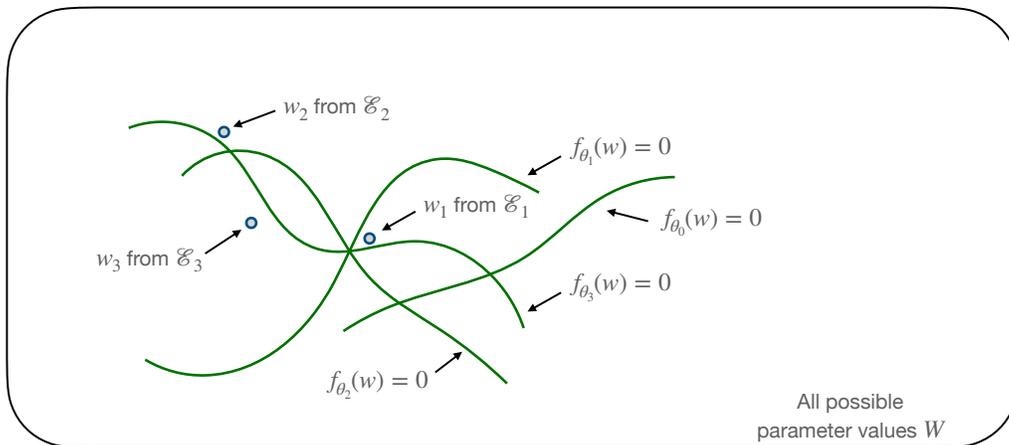


Figure 7. Constraint learning problem with noise.

As we see additional environments, we continue to narrow down the possible values of θ . If we next learn w_2 by seeing data from environment \mathcal{E}_2 , we can rule out θ_1 : the constraint $f_{\theta_1}(w) = 0$ passes through w_1 but not w_2 . If we then learn w_3 by visiting environment \mathcal{E}_3 , we rule out θ_2 , leaving $f_{\theta_3}(w) = 0$ as the only one of the illustrated constraints that is still consistent. If we assume that θ_3 is the correct constraint, then in \mathcal{E}_4 we can use this information to learn w_4 faster: the set $f_{\theta_3}(w) = 0$ has lower dimension than the full parameter space W , so we will need fewer samples to distinguish the value of w_4 that is the best fit within the constraint.

If we leave aside computational considerations, we can write a very simple algorithm that learns the correct constraint in our setup so far: we just keep track of which values of θ are still consistent after seeing w_1, w_2, \dots . When only a single value of θ remains consistent, we are done.

Of course, this algorithm will not work very well in practice, since we have ignored an important aspect of constraint learning: due to estimation error, we will not get a perfect estimate of w_i from any finite amount of data in \mathcal{E}_i . So, the picture will look more like Figure 7: the true constraint $f_{\theta_3}(w) = 0$ will not pass exactly through our estimates w_1, w_2, w_3 , but will only pass close to them.

To handle estimation errors like this, we propose a simple no-regret algorithm. We start from a set Θ of possible constraint parameters. Before seeing each environment \mathcal{E}_i , we propose a hypothesis $\theta_i \in \Theta$ as described below. Then we find out w_i , and we get a penalty

$$P_i(\theta_i) = \|f_{\theta_i}(w_i)\|^2$$

This penalty is zero if and only if w satisfies the constraint $f_{\theta}(w_i) = 0$. If we fail to satisfy the constraint, then $P_i(\theta)$ is strictly positive, with a larger value if we are farther from satisfying the constraint.

Our goal is to get a small total penalty $P(\theta) = \sum_i P_i(\theta)$. To accomplish this, we can choose θ_i using any no-regret algorithm — that is, any algorithm where the regret

$$\text{regret} = \max_{\theta \in \Theta} \left[\sum_i P_i(\theta_i) - P(\theta) \right]$$

grows sublinearly in T .

Depending on the sort of guarantees we want, we can pick our no-regret algorithm one of two ways. The first possibility is a learner like online gradient descent. Depending on the form of f_{θ} , OGD gives us different guarantees: if we have linear constraints, then $P_i(\theta)$ will be convex, and we can get convergence to the regret of the global optimum. On the other hand, for general f_{θ} , the functions $P_i(\theta)$ may not be convex, in which case we might get stuck in local optima. Either way, OGD

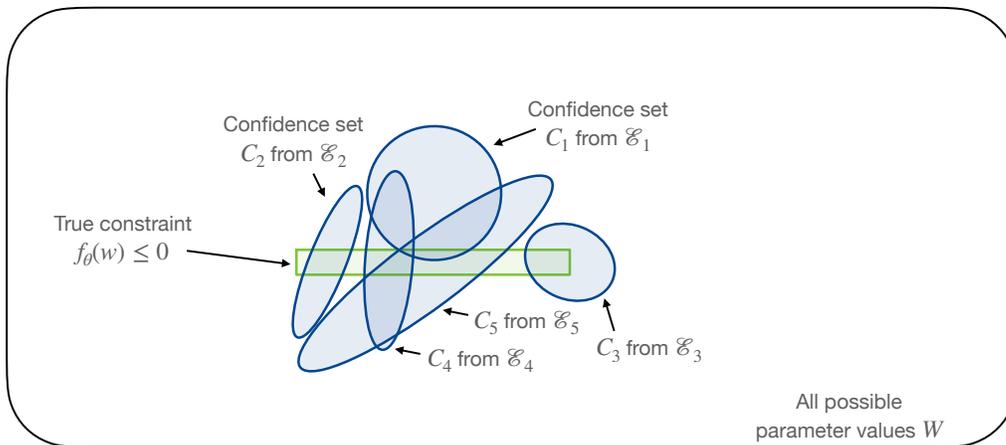


Figure 8. Constraint learning problem with inequality constraints and confidence sets.

is computationally cheap: the main cost is computing the gradient or subgradient of the penalty $P_i(\theta)$.

The second possibility is to choose a no-regret algorithm that has strong guarantees in the face of non-convex learning problems. One such approach is to tile the space of constraints Θ with an ϵ -net, and run a no-regret method such as multiplicative weights to pick among all of the points in the net. The size of the ϵ -net will be exponential in the dimension of Θ , so this approach will typically only be computationally feasible for low-dimensional constraint sets.¹⁰ On the other hand, we will get a much stronger guarantee: for MW, after T environments, we will propose at least one parameter θ_i whose penalty is at most

$$P(\theta_i) \leq P(\theta) + O(\sqrt{T} \ln |\Theta|)$$

where θ is the best (post-hoc) parameter value in Θ . Therefore, the average penalty $P(\theta_i)/T$ for the best proposal will approach the optimal value $P(\theta)/T$ as $T \rightarrow \infty$.

No matter how we choose our no-regret learner, as our estimate θ_i improves, we become more and more likely to see a benefit from transfer. But early in the learning process we may see *negative* transfer: if we enforce a poorly-estimated constraint, it can *reduce* our performance at estimating w_i in some environment \mathcal{E}_i . To guard against negative transfer, we can calculate both the best constrained w_i (the one that minimizes loss subject to $f_{\theta_i}(w_i) = 0$) as well as the best unconstrained w_i (the one that minimizes loss over the entire parameter space W). We can then choose between the constrained and unconstrained estimates: e.g., we can run an additional no-regret learner inside each new environment to predict w_i nearly as well as the better of the two estimates.

At this point we have a constraint learning method that can work well in practice. But, there are two additional wrinkles that can affect performance. So, for the final version of our constraint learning problem, we'll make two changes to our setup.

First, instead of just equality constraints of the form $f_{\theta}(w) = 0$, we'll allow inequality constraints like $f_{\theta}(w) \leq 0$. This makes it easier for us to represent a wider variety of constraint shapes. Second, we will let our learning algorithm give us more information: instead of just a point estimate w_i , we will let it return a confidence set C_i , with the intention that C_i covers the correct value of w_i with high probability. (We can still handle point estimates by treating them as singleton sets.)

With these changes, the picture looks like Figure 8. The biggest change here is that there is a more complicated relationship between the confidence sets and the constraint set. First, there will likely be parameter values in C_i that don't satisfy the constraint, due to limited data in environment \mathcal{E}_i . And second, there will likely be parameter values that satisfy the constraint but fall outside of C_i , since the constraint describes possible parameter values for *all* environments, not just \mathcal{E}_i .

To handle this situation, we can make some simple changes to our no-regret algorithm. First, the penalty for proposing θ_i

¹⁰If there is structure that we can take advantage of in the set of constraints Θ , we might be able to get the best of both worlds: computational efficiency and strong guarantees. We leave this sort of potential extension to future work.

becomes the penalty for the best value of w that falls within C_i . For an equality constraint, this is

$$P_i(\theta) = \min_{w \in C_i} \|f_\theta(w)\|^2$$

Second, if some coordinates of f_θ correspond to inequality constraints, we threshold so that negative values of f_θ are not penalized. If we write σ for the thresholding function, so that

$$\sigma_j(z) = \begin{cases} z_j & \text{if } z_j \geq 0 \text{ or coordinate } j \text{ is an equality constraint} \\ 0 & \text{otherwise} \end{cases}$$

then the penalty is

$$P_i(\theta) = \min_{w \in C_i} \|\sigma(f_\theta(w))\|^2$$

Finally, inequality constraints require a bit of additional care: different inequalities can be stronger or weaker, and therefore more or less desirable to learn. If we ignore this issue, we will tend to learn very weak constraints: in the limit, the trivial constraint that allows all of W will get zero penalty no matter what C_i is. (This issue is also present with some kinds of equality constraints, since we can equivalently incorporate σ into f_θ and treat every constraint as an equality; but essentially every family of inequality constraints needs to take account of this problem.)

Ideally we would directly measure the strength of transfer that each constraint provides: tighter constraints will get higher $P_i(\theta)$ but have the potential for stronger transfer. Unfortunately, this sort of measurement could be expensive; so, we leave it to future work to explore this avenue. Instead, here we suggest modifying $P_i(\theta)$ by adding an extra penalty that measures the weakness of the constraint that corresponds to θ : e.g., we could use a penalty proportional to the volume covered by the constraint. The scaling of such a penalty adds a hyperparameter; we hope that future work might show a way to eliminate the need for such a hyperparameter, or at least tune it at low cost.

To summarize, our final learning method is to run a no-regret algorithm over the set Θ of possible constraint parameters. For each environment \mathcal{E}_i , this outer-loop algorithm predicts a constraint θ_i . Given θ_i , we run two learners for w_i : one that enforces the predicted constraint, and one that is unconstrained. We switch between the two estimates of w_i using an inner-loop no-regret algorithm. We then use the unconstrained estimate of w_i to find a confidence set C_i , and pass this confidence set back to our outer-loop no-regret algorithm. To calculate this algorithm’s loss, we find the most favorable value of w within C_i : the value that minimizes the penalty $\|\sigma(f_{\theta_i}(w))\|^2$. Optionally, if different values of θ correspond to tighter or looser constraints, we add an additional penalty to the outer algorithm’s loss, in order to favor tighter constraints.

B. Related Works

There are a number of previous works that implicitly follow a special case of our setting and that approach our problem differently (e.g., with qualitatively different assumptions or goals). [Huang et al. \(2021\)](#) is a special case of our setting. This work assumes that a small set of variables is intervened on and therefore represents the changes in different environments. Therefore, they freeze the rest of the model learnt from previous environments, and only learn the changes in the new environment. [Li et al. \(2018\)](#); [Myung et al. \(2022\)](#) consider the case where there is limited data in the target domain but ample data in the source domain. Both papers assume we know what to transfer across environments, while our framework allows us to explicitly determine that from expert knowledge or the data.

MaxQ ([Dietterich, 1999](#)) approaches transfer differently: it is concerned with both *computational* and *statistical* benefits. Their setting for statistical benefits is reminiscent of ours: they propose a way to constrain the parameters of a policy in a new environment by freezing part of a “MAXQ hierarchy.” This same strategy also provides computational benefits, since decision-theoretic planning can be extremely expensive if we have to start from scratch. Our setup could address this benefit if we had a practical measure of the computational complexity of RL in different hypothesis classes, but such a measure is a subject for future research.

Deep Learning Heuristics for Transfer. Several heuristics for learning representations for domain adaptation have been proposed in the literature, including adversarial training ([Ganin et al., 2016](#)), self-supervised techniques ([Sohn et al., 2020](#)),

and contrastive learning (Shen et al., 2022). In contrast to these works, we focus on devising a framework via a causal perspective. However, it would be very interesting to apply our framework to these ideas.

Learning Invariant Features. Various works have explored the idea of learning “invariant” representations that use information available in all environments with the hope that this will lead to better transfer to downstream tasks (Arjovsky et al., 2020; Peters et al., 2015; Saengkyongam et al., 2021; Zhang et al., 2020; Magliacane et al., 2017). In our framework, taking advantage of an invariance is similar to choosing predictors \mathbb{E} such that $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ contains only one point, thus allowing zero-shot learning. As we discuss in Section 3.1 and Section 5.3, we consider the world update problem, which is different from choosing predictors.

Causal Transfer Learning. Our work builds on prior studies of transfer learning from a causal modeling perspective, without assuming causal faithfulness or sufficiency. Unlike Chen & Bühlmann (2021), who assume a linear SCM setting, our framework does not. Pearl & Bareinboim (2011) use selection variables S to model environmental differences, a special case of our framework. They allow any distribution selection, not constraining the variable given its parents. Our framework provides more details about when transfer learning is possible. Consider the causal graph among variables X, Y, Z, S with edges $Z \rightarrow X, Z \rightarrow Y, X \rightarrow Y$ and $S \rightarrow Y$. This example is a case of Figure 4(b) in Pearl & Bareinboim (2011) where it was determined that transfer learning of $P(Y|do(X))$ is not possible (Example 6 in Pearl & Bareinboim (2011)). Variable S here could be interpreted as determining the selection of the distribution of Y given its parents X and Z in different environments. However if $P(Y|X, Z)$ was selected from a constrained set, transfer learning of $P(Y|do(X))$ might still be feasible, which can be analyzed in an extension of our framework.

C. Future Works

We describe a more general approach to learn $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$.

We assume that the agent might have some expert knowledge about some $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$ even though \mathbb{I} or \mathbb{K} might not be observed. Thus, the agent wants to learn $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ from the set $\{\mathcal{W}_{\mathbb{I}|\mathbb{K}}\}$ for some \mathbb{I}, \mathbb{K} where it has some knowledge about $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$. In this paper we only consider a special case where the agent has knowledge about some atomic domains $\{\mathcal{W}_i^*\}$, and using Algorithm 1 we express $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ as a formula of \mathcal{W}_i^* . In the general case of learning $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ from the set $\{\mathcal{W}_{\mathbb{I}|\mathbb{K}}\}$, the problem can be divided into 2 steps:

1. *Choosing the terms.* Choosing the sets of (\mathbb{I}, \mathbb{K}) ’s to compute $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ from $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$ ’s.
2. *Calculation.* Compute $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$ from the set of chosen $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$ ’s.

While it might be intuitive to always choose \mathcal{W}_i^* ’s to compute $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$, we show an example where it is not the best way. Consider the graph $X \rightarrow Y \rightarrow Z$ where both $Y | X$ and $Z | Y$ are poorly conditioned and therefore estimating them yields higher errors while $Z | X$ is the identity. Therefore if the goal is to compute $Z | X$, it is better to compute it directly rather than from $Z | Y$ and $Y | X$.

Once the terms $w_{\mathbb{I}|\mathbb{K}}$ ’s are chosen, there are possible different approaches to calculate the target $w_{\mathbb{Q}|\mathbb{E}}$. One approach is to break both $w_{\mathbb{I}|\mathbb{K}}$ ’s and $w_{\mathbb{Q}|\mathbb{E}}$ into formulae of the atomic term w_i^* ’s using Theorem E.3 and then to reason about how the constraint on $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$ ’s implies constraint on the atomic \mathcal{W}_i^* , and then how the constraint on the atomic \mathcal{W}_i^* ’s implies constraint on $\mathcal{W}_{\mathbb{Q}|\mathbb{E}}$, converting the problem to the special case we already addressed in this paper. Another approach is to directly express $w_{\mathbb{Q}|\mathbb{E}}$ as a formula of the chosen term $w_{\mathbb{I}|\mathbb{K}}$ ’s.

We leave the extension to the future works.

Model Selection. If the agent has no expert knowledge to select the hypothesis class of function g_θ that describes $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$, or that the expert knowledge is not specific enough, selecting the hypothesis class of g_θ using the expert knowledge can result in a large over-estimate of $\mathcal{W}_{\mathbb{I}|\mathbb{K}}$. In this case the agent can employ model selection strategies to select the hypothesis class from a nested set of hypothesis class using the incoming data w^e . If the agent first assumes that w^e is in a large set but later notices that all incoming w^e ’s lie in a small subset, the agent can narrow down the hypothesis class. The expert knowledge, if exists, place an upper bound on the model selection strategy. For example, suppose that $w_{100|0} = w_{100}^* \cdots w_1^*$ is unchanged in all environments but the agent’s expert knowledge is only that w_1^* is unchanged in all environments, using which it infers that $w_{100|0}^e$ have the same row space in all environment. Then the agent can search through a nested hypothesis class for $\mathcal{W}_{100|0}$ with the largest class being the set of matrices with the same row space with unknown bases and the smallest class

being the set of an unknown single matrix. After the observing that all of the incoming data $w_{100|0}^1, \dots, w_{100|0}^e$ are the same matrix the agent can pick the smallest hypothesis class and learn the value of the unknown matrix. If the agent has no expert knowledge at all, it has to search through a nested hypothesis class with the largest class being the set of all matrices and the smallest class being the set of a single matrix.

We leave this extension to future works.

D. Details about Section 3

The SCM and the ordering λ generate a causal graph. Let T be the time horizon. Let λ be an ordering of the n variables \mathbf{X} . At time $t = 0$, each variable is initialized with an initial value \mathbf{X}^0 . At time t , at step i , $1 \leq i \leq n$, the SCM generates the variables $X_{\lambda(i)}^t$. Let $\mathbf{P}_i^{t-1,t} \subseteq \mathbf{X}^{t-1} \cup \mathbf{X}^t$ denote the set of parents of X_i^t . Then we have: $X_i^t = f_{w_i^{t*}}(\mathbf{P}_i^{t-1,t}, U_i^t)$ where $w_i^{t*} = w_i^*$.

For examples, consider the following SCM:

$$X_1 = f_{w_1^*}(U_1) \quad (4)$$

$$X_2 = f_{w_2^*}(X_1, U_2) \quad (5)$$

If the ordering $\lambda = (1, 2)$, then the parents of X_2^1 is $\mathbf{P}_2^{0,1} = \{X_1^0\}$ and $X_2^1 = f_{w_2^*}(X_1^0, U_2^1)$ since X_1^0 is generated before X_2^1 . If the ordering $\lambda = (2, 1)$, then the parents of X_2^1 is $\mathbf{P}_2^{0,1} = \{X_1^0\}$ and $X_2^1 = f_{w_2^*}(X_1^0, U_2^1)$ since X_1^0 is generated before and X_1^1 is generated after X_2^1 .

The SCM and the ordering λ induces a causal graph. The causal graph is defined such that there are directed edges from each node in the set of parents $\mathbf{P}_i^{t-1,t}$ of X_i^t to X_i^t (even when \mathcal{W}_i^* contains the value of w_i^* such that X_i^t does not depend on its parents). We provide an example of the graph for Q Learning in Figure 10b. The causal graph induces a partial order \preceq on the variables X_i^t 's: for any $A, B \in \{X_i^t, 1 \leq t \leq T, 1 \leq i \leq n\}$, $A \preceq B$ if there exists a directed path from A to B . To simplify the notations, in the main text we remove the superscript t by re-indexing the variables according to \preceq such that $X_i \preceq X_j$ if $i \leq j$.

E. Details about Section 4

E.1. Computation Trees

In this section we provide the proof of Theorem 4.2, which is restated by Theorem E.3.

d -separation between variables (denotes \perp_d) and the partial order \preceq is defined on the causal graph (Section D). For any 2 sets \mathbb{X} and \mathbb{Y} , $\mathbb{X} \preceq \mathbb{Y}$ if $X \preceq Y \forall X \in \mathbb{X}, Y \in \mathbb{Y}$. For any two variables $X \neq Y$, X is called an ancestor of Y and Y a descendant of X if there is a directed path from X to Y in the causal graph. Let the *ancestor set* $A(\mathbb{I})$ of \mathbb{I} be the set of all ancestors of variables in \mathbb{I} , excluding variables in \mathbb{I} . Let $A'(\mathbb{I}) := A(\mathbb{I}) \cup \mathbb{I}$. The *parent set* of \mathbb{I} is the set of all parents of variables in \mathbb{I} , excluding variables in \mathbb{I} .

A set \mathbb{J} is called a *sufficient ancestor set* of a set \mathbb{I} if $\mathbb{J} \subseteq A(\mathbb{I})$, $\mathbb{J} \preceq \mathbb{I}$ and $\mathbb{I} \perp_d A(\mathbb{J}) \mid \mathbb{J}$. The empty set is considered a sufficient ancestor set of any other set. Conditioning on a sufficient ancestor set \mathbb{J} of \mathbb{I} , we can compute \mathbb{I} from the atomic parameters without needing the ancestors of \mathbb{J} (Theorem E.3). Given \mathbb{I}, \mathbb{K} where $\mathbb{I} \cap \mathbb{K} = \emptyset$, a set \mathbb{J} is called a *sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$* if $\mathbb{J} \subseteq \mathbb{K}$ and \mathbb{J} is the sufficient ancestor set of $\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{J})$.

We show first that there is a maximal sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$.

Lemma E.1. *There exists a maximal sufficient ancestor set \mathbb{J}^* of $\mathbb{I} \mid \mathbb{K}$ such that \mathbb{J}^* is a sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$ and \mathbb{J}^* contains any other sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$.*

Proof. We will show that if \mathbb{J}_1 and \mathbb{J}_2 are sufficient ancestor sets of $\mathbb{I} \mid \mathbb{K}$ then $\mathbb{J} := \mathbb{J}_1 \cup \mathbb{J}_2$ is also a sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$. Since $\mathbb{J}_1 \subseteq \mathbb{K}, \mathbb{J}_2 \subseteq \mathbb{K}$, we have $\mathbb{J} \subseteq \mathbb{K}$.

By definition, $\mathbb{I} \cup \mathbb{K} \setminus \mathbb{J}_i$ is d -separated from $A(\mathbb{J}_i)$ given \mathbb{J}_i for $i = 1, 2$. Therefore $\mathbb{I} \cup \mathbb{K} \setminus \mathbb{J}$ is d -separated from $A(\mathbb{J}_i)$ given \mathbb{J} for $i = 1, 2$ by weak union by moving $\mathbb{J}_1 \setminus \mathbb{J}_2$ from $\mathbb{I} \cup \mathbb{K} \setminus \mathbb{J}_2$ to \mathbb{J}_2 (or $\mathbb{J}_2 \setminus \mathbb{J}_1$ from $\mathbb{I} \cup \mathbb{K} \setminus \mathbb{J}_1$ to \mathbb{J}_1). Therefore $\mathbb{I} \cup \mathbb{K} \setminus \mathbb{J}$ is d -separated from $A(\mathbb{J}_1) \cup A(\mathbb{J}_2)$ given \mathbb{J} by composition. Since $A(\mathbb{J}) = A(\mathbb{J}_1 \cup \mathbb{J}_2) \subseteq A(\mathbb{J}_1) \cup A(\mathbb{J}_2)$, $\mathbb{I} \cup \mathbb{K} \setminus \mathbb{J}$ is d -separated from $A(\mathbb{J})$ given \mathbb{J} by decomposition. By definition, \mathbb{J} is a sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$. \square

Therefore the scope of $\mathbb{I} \mid \mathbb{K}$ is defined by computing the scope of $\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{J}) \mid \mathbb{J}$ where $\mathbb{J} \subseteq \mathbb{K}$ is the maximal sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$.

Let \mathbb{K}' be an independent set of $\mathbb{I} \mid \mathbb{K}$ if $\mathbb{K}' \subseteq \mathbb{K}$ and $\mathbb{I} \perp_d \mathbb{K}' \mid \mathbb{K} \setminus \mathbb{K}'$. There exists a set \mathbb{K}' such that \mathbb{K}' is an independent set of $\mathbb{I} \mid \mathbb{K}$ and \mathbb{K}' contains any other independent set of $\mathbb{I} \mid \mathbb{K}$, called the *maximal independent set* of $\mathbb{I} \mid \mathbb{K}$. Such a maximal set exists because of the intersection property of d-separation: If $A \perp_d B \mid C \cup D$ and $A \perp_d C \mid B \cup D$ then $A \perp_d B \cup C \mid D$. Suppose there exist $\mathbb{K}_1 \subseteq \mathbb{K}$ and $\mathbb{K}_2 \subseteq \mathbb{K}$ such that $\mathbb{I} \perp_d \mathbb{K}_1 \mid \mathbb{K} \setminus \mathbb{K}_1$ and $\mathbb{I} \perp_d \mathbb{K}_2 \mid \mathbb{K} \setminus \mathbb{K}_2$. Then $\mathbb{I} \perp_d \mathbb{K}_2 \setminus \mathbb{K}_1 \mid \mathbb{K} \setminus (\mathbb{K}_1 \cup \mathbb{K}_2)$ by weak union.

Definition E.2 (Scope of $w_{\mathbb{I} \mid \mathbb{K}}$). Let \mathbb{I} and \mathbb{K} be 2 disjoint sets. Let $\mathbb{K}' \subseteq \mathbb{K}$ be the maximal independent set of $\mathbb{I} \mid \mathbb{K}$. Let $\mathbb{J} \subseteq \mathbb{K} \setminus \mathbb{K}'$ be the maximal sufficient ancestor set of $\mathbb{I} \mid \mathbb{K} \setminus \mathbb{K}'$. Let $\mathbb{I}' := \mathbb{I} \cup (\mathbb{K} \setminus \mathbb{K}') \setminus \mathbb{J}$. Then $\text{Scope}(\mathbb{I} \mid \mathbb{K}) := A'(\mathbb{I}') \setminus A'(\mathbb{J})$.

We show that $w_{\mathbb{Q} \mid \mathbb{E}}$ could be computed from variables in $\text{Scope}(\mathbb{Q} \mid \mathbb{E})$ using the operations \oplus , \otimes , and \circ , where w_s^{t*} 's appear in reverse topological order. Let $\mathbf{w}_{\text{Scope}(\mathbb{Q} \mid \mathbb{E})}^* := \{w_i^{t*} \mid (i, t) \in \text{Scope}(\mathbb{Q} \mid \mathbb{E})\}$ and $\mathcal{W}_{\text{Scope}(\mathbb{Q} \mid \mathbb{E})}^*$ be the corresponding constrained domain of $\mathbf{w}_{\text{Scope}(\mathbb{Q} \mid \mathbb{E})}^*$.

Theorem E.3 (Top-down computation tree). *Given $\mathbb{Q} \cap \mathbb{E} = \emptyset$, we initiate the formula with $w_{\mathbb{Q} \mid \mathbb{E}}$. For each term of the form $w_{\mathbb{I} \mid \mathbb{K}}$ in the formula, the agent repeatedly performing the following steps. First it removes the maximal independent set \mathbb{K}' of $\mathbb{I} \mid \mathbb{K}$ from \mathbb{K} . To make the notations simple, we re-use \mathbb{K} afterwards to denote the set $\mathbb{K} \setminus \mathbb{K}'$ after removing \mathbb{K}' . The agent then chooses one of the following computations.*

- Choose $\mathbb{J} \subseteq \mathbb{I}$, s.t. $\mathbb{J} \preceq \mathbb{I} \setminus \mathbb{J}$ and $\mathbb{J} \cup \mathbb{K} \setminus \mathbb{J}'$ is a sufficient ancestor set of $\mathbb{I} \setminus \mathbb{J}$ where \mathbb{J}' is the set of variables d-separated from $\mathbb{I} \setminus \mathbb{J}$ given $(\mathbb{J} \cup \mathbb{K}) \setminus \mathbb{J}'$. Then calculate the product:

$$w_{\mathbb{I} \mid \mathbb{K}} = w_{\mathbb{I} \setminus \mathbb{J} \mid (\mathbb{J} \cup \mathbb{K}) \setminus \mathbb{J}'} \otimes w_{\mathbb{J} \mid \mathbb{K}}. \quad (6)$$

- Choose $\mathbb{J} \subseteq \text{Scope}(\mathbb{I} \mid \mathbb{K})$ such that $\mathbb{I} \cap \mathbb{J} = \emptyset$, $\mathbb{K} \cap \mathbb{J} = \emptyset$ and sum out \mathbb{J} from the joint distribution:

$$w_{\mathbb{I} \mid \mathbb{K}} = \oplus_{\mathbb{J}} w_{\mathbb{I}, \mathbb{J} \mid \mathbb{K}}. \quad (7)$$

- Let $\mathbb{J} \subseteq \mathbb{K}$ be the maximal sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$. If $\mathbb{J} \neq \mathbb{K}$, the learner then compute $\mathbb{I} \mid \mathbb{K}$ from the joint $\mathbb{I}, \mathbb{K} \setminus \mathbb{J} \mid \mathbb{J}$ by conditioning on $\mathbb{K} \setminus \mathbb{J}$ in addition to \mathbb{J} :

$$w_{\mathbb{I} \mid \mathbb{K}} = \circ_{\mathbb{K} \setminus \mathbb{J}} w_{\mathbb{I}, \mathbb{K} \setminus \mathbb{J} \mid \mathbb{J}} \quad (8)$$

These options are executed until no such non-empty \mathbb{J} satisfying either condition exists.

All variables in the formula are $w_i^{t*} \in \text{Scope}(\mathbb{Q} \mid \mathbb{E})$ in reverse topological order and each w_i^{t*} appear at most once in the formula. We denote the resulting formula $g_{\mathbb{Q} \mid \mathbb{E}}(\mathbf{w}_{\text{Scope}(\mathbb{Q} \mid \mathbb{E})}^*)$, also called a computation tree¹¹ of $w_{\mathbb{Q} \mid \mathbb{E}}$.

Proof. We show that removing $\mathbb{K}' \subseteq \mathbb{K}$ from \mathbb{K} if $\mathbb{I} \perp_d \mathbb{K}' \mid \mathbb{K} \setminus \mathbb{K}'$ does not expand the scope, i.e. $\text{Scope}(\mathbb{I} \mid \mathbb{K} \setminus \mathbb{K}') \subseteq \text{Scope}(\mathbb{I} \mid \mathbb{K})$. Let $\mathbb{P} \subseteq \mathbb{K}$ and $\mathbb{P}' \subseteq \mathbb{K} \setminus \mathbb{K}'$ be the maximal sufficient ancestor set of $\mathbb{I} \mid \mathbb{K}$ and $\mathbb{I} \mid \mathbb{K} \setminus \mathbb{K}'$. Since $\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{P}) \perp_d A(\mathbb{P}) \mid \mathbb{P}$, we have $\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{K}' \setminus \mathbb{P}) \perp_d A(\mathbb{P}) \mid \mathbb{P}$ by decomposition. Therefore \mathbb{P} is a sufficient ancestor set of $\mathbb{I} \mid \mathbb{K} \setminus \mathbb{K}'$, and $\mathbb{P} \subseteq \mathbb{P}'$. Then we have $\underbrace{A'(\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{K}') \setminus \mathbb{P}') \setminus A'(\mathbb{P}')}_{\text{Scope}(\mathbb{I} \mid \mathbb{K} \setminus \mathbb{K}')} \subseteq \underbrace{A'(\mathbb{I} \cup \mathbb{K} \setminus \mathbb{P}) \setminus A'(\mathbb{P})}_{\text{Scope}(\mathbb{I} \mid \mathbb{K})}$ because $\mathbb{I} \cup ((\mathbb{K} \setminus \mathbb{K}') \setminus \mathbb{P}') \subseteq$

$\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{P})$ and $\mathbb{P}' \supseteq \mathbb{P}$.

We first show that if the agent terminates, then all the terms are in the form of w_i^{t*} in reverse topological order. Since the agent cannot execute the \circ operations, the term $w_{\mathbb{I} \mid \mathbb{K}}$ satisfies the condition that \mathbb{K} is a the maximal sufficient ancestor set of \mathbb{I} . Therefore $\text{Scope}(\mathbb{I} \mid \mathbb{K}) = A'(\mathbb{I}) \setminus A'(\mathbb{K})$. Since the agent cannot execute \oplus , $A'(\mathbb{I}) \setminus A'(\mathbb{K}) = \mathbb{I}$, and therefore the parent set of \mathbb{I} is a subset of $A'(\mathbb{K})$. Since \mathbb{K} is a sufficient ancestor set of \mathbb{I} and therefore $\mathbb{I} \perp_d A(\mathbb{K}) \mid \mathbb{K}$, the parent set of \mathbb{I} is a subset of \mathbb{K} . If \mathbb{I} has more than 1 element, the agent can choose \mathbb{J} to be the smallest element topologically to execute the \otimes operation because $\mathbb{J} \cup \mathbb{K}$ will contain the parent set of $\mathbb{I} \setminus \mathbb{J}$ and therefore $\mathbb{J} \cup \mathbb{K}$ will d-separated $\mathbb{I} \setminus \mathbb{J}$ from $A(\mathbb{J} \cup \mathbb{K})$. Therefore \mathbb{I} has only one element denoted (i, t) and \mathbb{K} contains all of its parents, and $w_{\mathbb{I} \mid \mathbb{K}}$ can be substituted by w_i^{t*} . Since the \otimes operation ensures that $\mathbb{J} \preceq \mathbb{I} \setminus \mathbb{J}$, w_i^{t*} comes before $w_j^{t'*}$ if $(j, t') \preceq (i, t)$, and therefore the terms are in reverse topological order.

¹¹There can be multiple computation trees depending on the choice the learner made.

We show now that for all operations \otimes , \oplus and \odot the scopes of the terms on the right-handed side are subset of the scope of the term on the left-hand side, and therefore all of the terms w_i^{t*} 's in the formula are in the scope of $w_{\mathbb{Q}|\mathbb{E}}$. We also show that in the \otimes operation the scopes of the 2 terms in the right-handed side are disjoint, and therefore each term w_i^{t*} appears at most once.

- First, consider the case when $w_{\mathbb{I}|\mathbb{K}} = w_{\mathbb{I}\setminus\mathbb{J}|\mathbb{J},\mathbb{K}} \otimes w_{\mathbb{J}|\mathbb{K}}$. Let $\mathbb{P} \subseteq \mathbb{K}$ denote the maximal sufficient ancestor set of $\mathbb{I} | \mathbb{K}$ and $\mathbb{P}' \subseteq \mathbb{K}$ denote the maximal sufficient ancestor set of $\mathbb{J} | \mathbb{K}$. Since $\mathbb{J} \cup (\mathbb{K} \setminus \mathbb{P}) \perp_d A(\mathbb{P}) | \mathbb{P}$ by decomposition, \mathbb{P} is a sufficient ancestor set of $\mathbb{J} | \mathbb{K}$. Therefore $\mathbb{P} \subseteq \mathbb{P}'$. We have $\underbrace{A'(\mathbb{J} \cup (\mathbb{K} \setminus \mathbb{P}')) \setminus A'(\mathbb{P}')}_{\text{Scope}(\mathbb{J}|\mathbb{K})} \subseteq \underbrace{A'(\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{P})) \setminus A'(\mathbb{P})}_{\text{Scope}(\mathbb{I}|\mathbb{K})}$.

We will now show that $\underbrace{A'(\mathbb{I}\setminus\mathbb{J}) \setminus A'(\mathbb{J} \cup \mathbb{K}\setminus\mathbb{J}')}_{\text{Scope}(\mathbb{I}\setminus\mathbb{J}|\mathbb{J} \cup \mathbb{K}\setminus\mathbb{J}'')} \subseteq A'(\mathbb{I}\setminus\mathbb{J}) \setminus A'(\mathbb{J} \cup \mathbb{K})$. We will show that for any mutually disjoint

sets $\mathbb{X}, \mathbb{Y}, \mathbb{Z}$, if $\mathbb{X} \subseteq A(\mathbb{Y})$ and $\mathbb{Y} \perp_d \mathbb{Z} | \mathbb{X}$ then $A'(\mathbb{Y}) \setminus A'(\mathbb{X}) \subseteq A'(\mathbb{Y}) \setminus A'(\mathbb{X} \cup \mathbb{Z})$. To do so, we show that any node $N \in A'(\mathbb{Y}) \setminus A'(\mathbb{X})$ is not in $A'(\mathbb{X} \cup \mathbb{Z})$. Since $A'(\mathbb{X}) \subseteq A'(\mathbb{Y})$ (because $\mathbb{X} \subseteq A(\mathbb{Y})$) and $N \in A'(\mathbb{Y}) \setminus A'(\mathbb{X})$, we have $N \notin A'(\mathbb{X})$. Suppose N is in $A'(\mathbb{X} \cup \mathbb{Z})$. Then there must exist $Z' \in \mathbb{Z}$ such that $N \in A'(Z')$ and there is a directed path from N to Z' not going through \mathbb{X} . Since $N \in A'(\mathbb{Y}) \setminus A'(\mathbb{X})$, there is a directed path from N to \mathbb{Y} not going through \mathbb{X} . Therefore conditioning on \mathbb{X} , there is a Bayes' ball path from \mathbb{Y} to Z' through N . This is a contradiction, therefore $N \notin A'(\mathbb{X} \cup \mathbb{Z})$. We have $\underbrace{A'(\mathbb{I}\setminus\mathbb{J}) \setminus A'(\mathbb{J} \cup \mathbb{K})}_{\text{A superset of Scope}(\mathbb{I}\setminus\mathbb{J}|\mathbb{J} \cup \mathbb{K}\setminus\mathbb{J}'')} \subseteq \underbrace{A'(\mathbb{I} \cup (\mathbb{K} \setminus \mathbb{P})) \setminus A'(\mathbb{P})}_{\text{Scope}(\mathbb{I}|\mathbb{K})}$ because $\mathbb{I}\setminus\mathbb{J} \subseteq \mathbb{I} \cup \mathbb{K}\setminus\mathbb{P}$ and

$\mathbb{J} \cup \mathbb{K} \supseteq \mathbb{P}$, and therefore $\text{Scope}(\mathbb{I}\setminus\mathbb{J} | \mathbb{J} \cup \mathbb{K}\setminus\mathbb{J}'') \subseteq \text{Scope}(\mathbb{I} | \mathbb{K})$.

We also have $\text{Scope}(\mathbb{I}\setminus\mathbb{J} | \mathbb{J} \cup \mathbb{K}\setminus\mathbb{J}'') \cap \text{Scope}(\mathbb{J} | \mathbb{K}) = \emptyset$ because $A'(\mathbb{J} \cup \mathbb{K})$, a superset of the second scope, is removed from the first scope. Therefore all leaves appear at most once in the formula.

- Second, consider the case when $w_{\mathbb{I}|\mathbb{K}} = \oplus_{\mathbb{J}} w_{\mathbb{I},\mathbb{J}|\mathbb{K}}$. Let \mathbb{P} denote the maximal sufficient ancestor set of $\mathbb{I} | \mathbb{K}$. We will first show that \mathbb{P} is a sufficient ancestor set of $\mathbb{I}, \mathbb{J} | \mathbb{K}$. Let $\mathbb{I}' := \mathbb{I} \cup (\mathbb{K} \setminus \mathbb{P})$. By the definition of sufficient ancestor set, \mathbb{I}' is d-separated from $A(\mathbb{P})$ given \mathbb{P} . Since $\mathbb{J} \subseteq A'(\mathbb{I}') \setminus A'(\mathbb{P})$, there are directed paths from \mathbb{J} to \mathbb{I}' not through \mathbb{P} . Suppose there are Bayes ball's paths from $A(\mathbb{P})$ to \mathbb{J} given \mathbb{P} . Then there are Bayes ball's paths from $A(\mathbb{P})$ to \mathbb{I} given \mathbb{P} by combining the Bayes ball's paths from $A(\mathbb{P})$ to \mathbb{J} given \mathbb{P} and the directed path from \mathbb{J} to \mathbb{I} not through \mathbb{P} . Therefore $A(\mathbb{P})$ is not d-separated from \mathbb{I} given \mathbb{P} , a contradiction. Therefore there is no Bayes' ball path from $A(\mathbb{P})$ to \mathbb{J} given \mathbb{P} , so \mathbb{J} is d-separated from $A(\mathbb{P})$ given \mathbb{P} . Therefore $\mathbb{I}' \cup \mathbb{J} = \mathbb{I} \cup \mathbb{J} \cup (\mathbb{K} \setminus \mathbb{P})$ is also d-separated from $A(\mathbb{P})$ given \mathbb{P} , which implies \mathbb{P} is a sufficient ancestor set of $\mathbb{I}, \mathbb{J} | \mathbb{K}$.

Let \mathbb{P}'' be the maximal sufficient ancestor set of $\mathbb{I}, \mathbb{J} | \mathbb{K}$, then we have $\mathbb{P} \subseteq \mathbb{P}''$. Therefore $\underbrace{A'(\mathbb{I} \cup \mathbb{J} \cup (\mathbb{K} \setminus \mathbb{P}''))}_{\text{Scope}(\mathbb{I},\mathbb{J}|\mathbb{K})} \subseteq \underbrace{A'(\mathbb{I}') \setminus A'(\mathbb{P})}_{\text{Scope}(\mathbb{I}|\mathbb{K})}$. Since $\mathbb{J} \subseteq A'(\mathbb{I}')$, $A'(\mathbb{I}' \cup \mathbb{J}) = A'(\mathbb{I}')$. Therefore $A'(\mathbb{I} \cup \mathbb{J} \cup (\mathbb{K} \setminus \mathbb{P}'')) \subseteq A'(\mathbb{I}')$, and we have $\underbrace{A'(\mathbb{I} \cup \mathbb{J} \cup (\mathbb{K} \setminus \mathbb{P}'')) \setminus A'(\mathbb{P}'')}_{\text{Scope}(\mathbb{I},\mathbb{J}|\mathbb{K})} \subseteq \underbrace{A'(\mathbb{I}') \setminus A'(\mathbb{P})}_{\text{Scope}(\mathbb{I}|\mathbb{K})}$.

- Finally, when $w_{\mathbb{I}|\mathbb{K}} = \odot_{\mathbb{K}\setminus\mathbb{J}} w_{\mathbb{I},\mathbb{K}\setminus\mathbb{J}|\mathbb{J}}$, the scope of the term on the right-handed side is the scope of the term on the left-handed side by definition.

Therefore all of the terms w_i^{t*} 's in the formula are in the scope of $w_{\mathbb{Q}|\mathbb{E}}$, and each term appear at most once. \square

E.2. Comparisons with Variable Elimination

There are several differences: (1) Variable elimination computes the tree from the bottom-up by starting with the leaves w_i^* 's and choosing several nodes to be combine to create their parent; (2) Variable elimination does not explicitly describe the intermediate nodes; (3) Variable elimination does not describe a smaller scope and (4) In variable elimination, each leaf goes through at least 2 operations: at least 1 to build the joint distribution and the conditional operation \odot while in our method it is possible for a leaf to go through only one operation (w_3^* in Figure 4), which could make it easier to analyze if w_3^* is the variable of interest.

E.3. Details about Examples in Section 4.1

We provide another example of causal graph, scope and computation tree in Figure 9, Example E.1 and Example E.2.

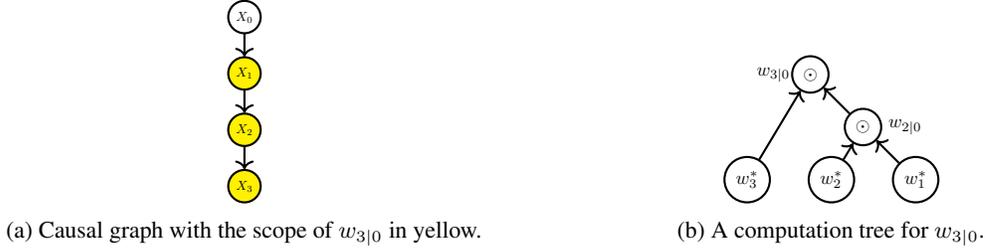


Figure 9. Causal graphs and their computation trees. For clarity, we do not draw the uniform noise variables.

Example E.1 (Q Learning). Let A_t, S_t, R_t , and G_t denote the action, state, reward, and return at time t . The $Q(s, a)$ function is the expected return at time t given the state at $t = 0$ parameterized by $w_{G_t|S_0, A_0}$. The SCM for Q-Learning is the following:

$$A = f_{w_A^*}(U_A) \quad (9)$$

$$S = f_{w_S^*}(A, S, U_S) \quad (10)$$

$$R = f_{w_R^*}(A, S, U_R) \quad (11)$$

$$G = f_{w_G^*}(R, G) \quad (12)$$

and the ordering to generate the variables $\lambda = (S, A, R, G)$.

A_t, S_t, R_t and G_t are interpreted as the action, state, reward and return at time t . The $Q(s, a)$ function is the expected return at time G_t given S_0 and A_0 . The causal graph is shown in Figure 10a and the below computation tree for $w_{G_2|A_0, S_0}$ is shown in Figure 10b. For any variable X_i , since $w_i^{t*} = w_i^*$ for any t , we can compress all nodes w_i^{t*} into a single node w_i^* .

$$w_{G_2|S_0, A_0} = w_{G_2|G_1, R_1, A_1, S_1} \odot \left[w_{G_1, R_1, A_1, S_1|G_0, R_0, A_0, S_0} \odot w_{G_0, R_0|S_0, A_0} \right] \quad (13)$$

$$= (w_G^* \odot (w_R^* \odot (w_S^* \otimes w_A^*))) \odot \left[(w_G^* \otimes (w_R^* \otimes (w_S^* \otimes w_A^*))) \odot (w_G^* \otimes w_R^*) \right] \quad (14)$$

Example E.2 (Figure 3). $\text{Scope}(7 | 6, 3) = \{3, 4, 5, 6, 7\}$. We show below a computation tree with variables in the scope:

$$w_{7|6,2} = \odot_6 w_{7,6|2} \quad (15)$$

$$= \odot_6 (w_{7,6|5,4,3,2} \odot w_{5,4,3|2}) \quad (16)$$

$$= \odot_6 ((w_7^* \otimes w_6^*) \odot (w_5^* \otimes w_4^* \otimes w_3^*)) \quad (17)$$

E.4. Details about Section 4.3

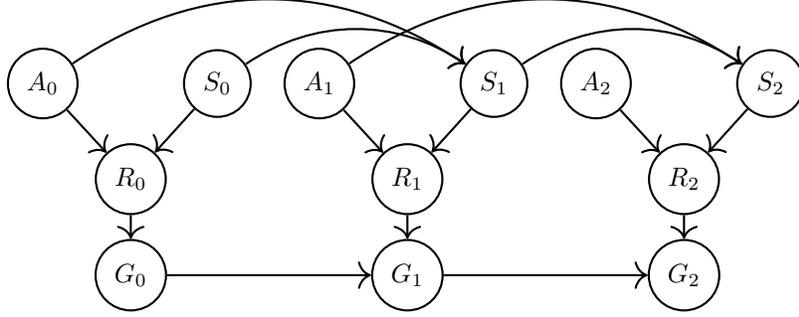
If each atomic parameter w_i^* appear at most once after the substitution of w_i^{t*} by w_i^* as in the case of supervised learning, the formula will be simple, and constrained domain can be computed iteratively: if $w_{Y|X} = (w_3^* \circ w_2^*) \circ w_1^*$ where \circ denote some operations, then we can first compute the constrained domain \mathcal{W}_{32} of $w_3^* \circ w_2^*$, and then compute $\mathcal{W}_{Y|X} = \mathcal{W}_{32} \circ \mathcal{W}_1^*$ where $\mathcal{W}_{32} \circ \mathcal{W}_1^*$ is define to be the set $\{u \circ v \mid u \in \mathcal{W}_{32}, v \in \mathcal{W}_1^*\}$ since \mathcal{W}_{32} are independently chosen according to the Independent Causal Mechanism. If $w_{Y|X} = (w_1^* \circ w_2^*) \circ w_1^*$ where w_1^* appears twice, we cannot compute $\mathcal{W}_{Y|X} = \mathcal{W}_{12} \circ \mathcal{W}_1^*$ because \mathcal{W}_{12} and \mathcal{W}_1 are not chosen independently. Instead we need to define $\mathcal{W}_{Y|X} = \{u \circ v \circ u \mid u \in \mathcal{W}_1^*, v \in \mathcal{W}_2^*\}$.

F. Details about Section 5

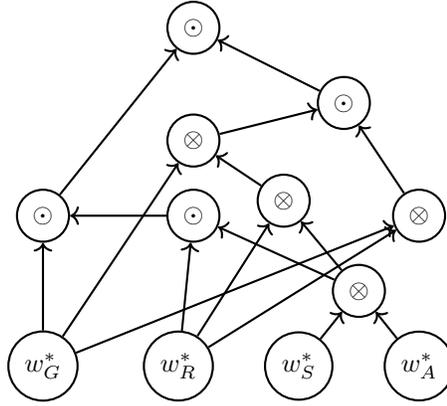
F.1. Details about Example 5.1

Sparse mechanism shifts do not imply transfer feasibility. Consider the first scenario where $\forall 2 \leq i \leq 99, w_i^*$ is the same matrix satisfying $\Pi_{i=2}^{99} w_i^*$ has rank 2 in all environments, while w_1^* and w_{100}^* are re-selected from all possible 2×2 matrices in each new environment. Since only w_1^* and w_{100}^* change in different environments, the changes are sparse. We show below that the constraints on w_2^*, \dots, w_{99}^* do not result in any constraint on $w_{100|0}^e$ by showing the following property of \odot .

Lemma F.1. *Let $\mathcal{M}^{m \times n}$ denote the set of all matrices with dimension $m \times n$. Let w be a 2×2 matrix of rank 2. Then: $\{AwB \mid A \in \mathcal{M}^{2 \times 2}, B \in \mathcal{M}^{2 \times 2}\} = \mathcal{M}^{2 \times 2}$.*



(a) The variables created by the SCM of Example E.1 described in Eq. 9 and the ordering $\lambda = (S, A, R, G)$. G_t is the return at time t . For clarity, we do not draw the uniform noise variables.



(b) A computation graph of $w_{G_2|A_0,S_0}$ for Example E.1. For all variables X , we compress all nodes w_X^{t*} into a single node w_X^* .

Figure 10. The causal graph and a computation tree for Q-Learning in Example E.1.

Since $w_{99}^* \cdots w_2^*$ has rank 2, $\{Aw_{99}^* \cdots w_2^*B \mid A \in \mathcal{M}^{2 \times 2}, B \in \mathcal{M}^{2 \times 2}\} = \mathcal{M}^{2 \times 2}$. The constraint that w_{99}^*, \dots, w_2^* is the same matrix in all environments does not enforce any constraint on $w_{100|0}^e$.

Transfer feasibility does not imply sparse mechanism shifts. Consider the second scenario in Example 5.1 where in each environment, $\forall 2 \leq i \leq 100$, w_i^* is selected from the set of all 2×2 matrices while w_1^* is the same matrix of rank 1. Since w_1^* is the right-most term in the computation graph, we analyze the left-propagatable property for \odot . Matrix multiplication is \mathcal{F} -left-propagatable where \mathcal{F} is the set of all matrices with rows in the row space of the right term w_1^* . Therefore, $w_{100|0}^e$ is in the set of all matrices with rows in the row space of w_1^* . Since w_1^* has rank 1, θ in Algorithm 1 is the 1×2 basis vector of the row space, which can be learnt from $w_{100|1}^1$ of Environment 1.

Transfer feasibility of the algorithm. Without transfer learning, $\text{comp} = 4$ since the matrix $w_{100|1}$ is 2×2 . With our transfer learning procedure, the learner needs to learn a 2×1 vector to compute $w_{100|0}$ after knowing the 1×2 basis θ of its row space, therefore $\text{comp} = 2$.

Representation Learning. The distribution of $X_1 \mid X_0$ where X_0 is the predictor could be consider the representation. In the second scenario, the representation is the same across the environments, and transfer feasibility is analyzed by the left-propagatable property of the operation since the representation is the right-most term.

Proof of Lemma F.1. Given $w \in \mathcal{M}^{2 \times 2}$ of rank 2, we will show that for any $C \in \mathcal{M}^{2 \times 2}$, there exist $A, B \in \mathcal{M}^{2 \times 2}$ such that:

$$AwB = C. \quad (18)$$

Let $w = U\Sigma V^T$ and $C = U'\Sigma'V'^T$ be the SVD decompositions of w and C . Since $\text{rank}(C) \leq \text{rank}(w)$, there exists a

diagonal matrix D such that $\Sigma' = \Sigma D$. Let $A = U'U^T$ and $B = VDV'^T$. Then:

$$AwB = (U'U^T)(U\Sigma V^T)(VDV'^T) \quad (19)$$

$$= U'\Sigma DV'^T \quad (20)$$

$$= U'\Sigma'V'^T \quad (21)$$

$$= C \quad (22)$$

□

F.2. Details about Example 5.2

We derive a computation graph:

$$w_{100|0} = w_{100}^* \odot w_{99}^* \odot \cdots \odot w_1^* \quad (23)$$

where the computations are performed from left to right.

For neural network. \odot is stacking layer: for $i > j > k$, the generator $w_{i|k} = w_{i|j} \odot w_{j|k} = (w_{i|j}, w_{j|k})$. Therefore $w_{100|0} = (w_{100}^*, \dots, w_1^*)$

(Rolnick & Kording, 2020) show that under the Linear Regions Assumption, except for a measure-zero set of networks, ReLU networks can be recovered up to permutations and rescaling if the domain of the input is the entire multidimensional real space. The description set \bar{w} can be obtained from w by replacing $w_k^{(\cdot, i)}, w_{k+1}^{(\cdot, i)}$ by $cw_k^{(\cdot, i)}, \frac{1}{c}w_{k+1}^{(\cdot, i)}$ and replacing w_k, w_{k+1} by a permutation σ of the rows of w_k and the columns of w_{k+1} . Under the Linear Region Assumption, in all environments, from Eq. 23, the description set $\bar{w}_{100|0} = \{(w_{100}^*, \dots, w_1^*), \sigma, c \mid \sigma \in \Sigma_{100|0}, c \in \mathcal{C}_{100|0}\}$ where $\Sigma_{100|0}$ and $\mathcal{C}_{100|0}$ are sets of permutations and scaling applying to $(w_{100}^*, \dots, w_1^*)$ to generate the descriptions.

For any $i, 1 \leq i \leq 100$, if w_i^* is the same in all environments, then for any description $u = (u_1^*, \dots, u_{100}^*)$ of $w_{100|0}$, u_i^* can be obtained from w_i^* by permutation or rescaling, and therefore is layer i of a description of $w_{100|0}^k$ in another environment \mathcal{E}_k . Therefore we can learn layer u_i^* of a description from the first environment and freeze it in later environments. We do not use the exact algorithm described at the end of Section 4.4 (when Assumption 3.4 does not hold) because this discussion is shorter, but in both the algorithm in Section 4.4 and this discussion, we still need to learn or infer all possible descriptions from previous environments in order to transfer correctly.

We now show below that even with slight modifications from the above example, θ no longer corresponds to an entire layer, and therefore freezing a layer learnt from previous environments lead to sub-optimal performance.

Example F.1. Consider the following SCM: $X_0, X_1, X_2 \in \mathbb{R}^2$ and $X_i = f_{w_i^*}(X_{i-1}) + \epsilon_i, i \in \{1, 2\}$ where ϵ_i 's are independent noise satisfying $E[\epsilon_i] = 0$ and $f_{w_i^*}$ is a 1-layer ReLU neural network where w_i^* denotes the weight. In all

environments, $w_2^* = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ is the same and w_1^* can be selected from all 2×2 matrices. Suppose the learner wants to transfer knowledge of $w_{2|0}$, which is the weight of a 2-layer ReLU network, from Environment 1 to Environment 2. Consider the following cases:

1. In environment \mathcal{E}_1 , $w_1^{*1} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ and $X_0 \in \mathbb{R}^2$ is a multivariate Gaussian on \mathcal{R}^2 .
2. In environment \mathcal{E}_1 , $w_1^{*1} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $X_0 \in \mathbb{R}^2$ is a multivariate Gaussian on the line $\{[1, 1]^T c \mid c \in \mathcal{R}\}$.

In Case 1, in Environment 1, the following 2-layer weights are both descriptions of $w_{2|0} = (w_1^*, w_2^*)$: $\left(\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right)$

and $\left(\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 6 \end{bmatrix} \right)$. In Case 2, in Environment 1, the following 2-layer weights are both descriptions of $w_{2|0} =$

(w_2^*, w_1^*) on the domain of X_0 : $\left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right)$ and $\left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0.3 & 0.7 \\ 0.5 & 0.5 \end{bmatrix} \right)$. Therefore in Case 1 and Case 2 the

agent cannot freeze w_2^* learnt from Environment 1 because it may freeze an incorrect value such as

1	0
0	6

 in Case 1 and

0.3	0.7
0.5	0.5

 in Case 2. This is because in Case 1 and Case 2, the constrained domain that can be learnt from Environment 1 is no longer the set of multi-layer networks with a common layer w_2^* . Instead, the agent needs to learn the set of all possible descriptions from Environment 1.

F.3. Details about Section 5.2.2

In this section we provide details for the case study of the Colored MNIST problem in Section 5.2.2.

F.3.1. DERIVATION OF THE ARCHITECTURE

We apply Theorem E.3 to derive the following architecture:

$$P(Y | X) = \frac{P(X, Y)}{\sum_y P(y, X)} \tag{24}$$

$$= \frac{\sum_p P(X, p, Y)}{\sum_y \sum_p P(X, p, y)} \tag{25}$$

$$= \frac{\sum_p P(X | p)P(p, Y)}{\sum_y \sum_p P(X | p)P(p, y)} \tag{26}$$

$$= \frac{\sum_p \exp\left(\frac{1}{s}\lambda_p^T v(X) + \log g(\lambda_p) + \log P(p, Y)\right)}{\sum_y \sum_p \exp\left(\frac{1}{s}\lambda_p^T v(X) + \log g(\lambda_p) + \log P(p, y)\right)} \tag{27}$$

F.3.2. DATA GENERATION

Dataset and the pre-processing code to generate multiple environments of colored images from the original MNIST dataset from Gulrajani & Lopez-Paz (2020)¹² were used with image size $s = 28$ and number of color channels $k = 2$. We add one channel to make $k = 3$. Environment 1 and 2 has 60000 and 10000 samples respectively. We divide Environments 1 to train, val and test set with ratio 0.8, 0.1, 0.1. In Environment 2, we use 5000 samples for training and validation and 5000 samples for testing. Among the 5000 samples for training and validation, we gradually take the first 100, 200, 500, 1000, 2000 and 5000 samples as the number of samples available to the agent. For each available sample size 100, 200, 500, 1000, 2000 and 5000, we divide the train and val sets with ratio 7 : 3.

F.3.3. TRANSFER FEASIBILITY OF THE ALGORITHM

To learn $\widehat{\mathcal{W}}_{Y|X}^1$, we need to learn the linear coefficient θ of $v(x)$. There are n_p coefficients $\theta \in \mathcal{R}^{n \cdot n \cdot k}$ for each value of p . There are $n_y \times n_p$ biases $\text{bias}_{y,p} \in \mathcal{R}$ for each value of y and p . Since we do not know n_p , we use $n_p = 20$ as an upper bound. The number of cells for θ is $n_p \cdot n \cdot n \cdot k = 20 \cdot 28 \cdot 28 \cdot 3 = 47040$. The number of cells for the bias of the first layer is $n_p = 20$. The number of cells of the weight of the second layer is $n_p \cdot n_y = 20 \cdot 2 = 40$. Let the complexity measure comp be the dimension of the parameters needed to learn. Then, without transfer learning, $\text{comp} = 47040 + 20 + 40$; with transfer learning, $\text{comp} = 20 + 40$.

F.3.4. IMPLEMENTATION

In all of our training procedures, we use stochastic gradient descent and the negative log likelihood loss. We repeat for 10 times the following procedure and plot the average and the standard deviation values in the plot. First we create Environments 1 and 2 as described above. We then train a model on the train set of Environment 1 and plot its accuracy the test set of Environment 1 and Environment 2 in Figure 5. Next we perform transfer from Environment 1 to Environment 2 in Figure 5. Each line is an average over 10 runs with the shaded region denoted the standard deviation. We plot the accuracy on the test set of Env. 2.

The weight in the second layer is $P(y | p)$ and therefore needs to satisfy $\sum_y P(y | p) = 1$. We satisfy this requirement by placing a softmax function on the weight.

¹²URL: <https://github.com/facebookresearch/DomainBed/blob/main/domainbed/datasets.py>

The linear coefficients of the first layer and second are initialized uniformly from $(0, 1)$. The bias of the first layer is initialized uniformly from $(-\sqrt{k}, \sqrt{k})$ where k is the number of incoming features.

F.3.5. HYPERPARAMETERS TUNING

We use Ray Tune to tune the hyperparameters. We sample the parameters from the configuration range in the table below, run the training process until the number of epochs is 100 or the standard deviation of the validation loss of the last 10 epoch is at most 0.01. We select the hyperparameter with the lowest validation loss, and retrain the model with the selected hyperparameter on the combined train and validation set for 100 epochs. We use the retrain model to predict the test set and to transfer. Due to limited computation resource, we use less resource to tune our method by using a smaller hyperparameter range (and therefore a smaller number of sampled hyperparameters) and use more resource to tune competing methods by using a larger range and more samples.

	Train in Env. 1	Train in Env.3 (No Transfer)	Transfer Coefficient	Transfer Layer, Transfer Random Coefficient
Learning rate	loguniform(0.01, 5)			
Momentum	0.9			
Weight decay	0.0001	loguniform(0.00001, 0.01)	0.0001	loguniform(0.00001, 0.01)
n_p	20	20 or 50 uniformly		
Batch size	512			
Number of sampled hyperparameters	8	100	8	50

F.3.6. PLOT LEGEND

Without transfer. The model is train from scratch in Env. 2 the available samples 100, 200, 500, 1000, 2000, 5000 successively.

Transfer coefficient. We freeze the first layer’s coefficient learnt from Env. 1 and retrain the first layer’s bias and the second layer in Env. 2 with the available samples 0, 100, 200, 500, 1000, 2000, 5000 successively.

Transfer layer. We freeze the first layer’s coefficient and bias learnt from Env. 1 and retrain the second layer in Env. 2 with the available samples 0, 100, 200, 500, 1000, 2000, 5000 successively.

Transfer random coefficient. We generate the first layer’s coefficient randomly, freeze it and retrain the first layer’s bias and the second layer in Env. 2 with the available samples 0, 100, 200, 500, 1000, 2000, 5000 successively.

F.3.7. COMPUTE

The experiments that produce the figures in this paper were performed on a personal laptop.

F.3.8. HYPOTHETICAL MODEL SATISFYING THE ASSUMPTIONS

We present a hypothetical model that satisfies the 4 assumption stated in Section 5.2.2. Note that this model is only for illustration and to show that the 4 assumptions could be reasonably satisfied. We did not use this model in our analysis or experiment.

Example F.2 (A hypothetical model for Colored MNIST). We present a hypothetical generative model that satisfied all of our assumptions (Figure 11). Let $X \in \mathcal{R}^{s \times s \times k}$ denote the image, $D \in \{0, \dots, 9\}$ denote the digit, $Y \in \{0, 1\}$ denote the binary digit label, $C \in \{0, 1\}$ denote the color, $M \in \mathcal{R}^{s \times s \times n}$ denote the digit matrix and $N \in \mathcal{R}^{k \times k}$ denote the color matrix. The binary digit Y is constructed by flipping D with probability 0.25. The color C is constructed by flipping Y with probability 0.1, 0.2, and 0.9 respectively in environments 1, 2, and 3. Each image x is a $s \times s \times k$ tensor where s is the image size and k is the number of color channels. For each digit d , there is a digit matrix $m_d \in \mathcal{R}^{s \times s \times k}$ which draws the digit in red while other color channels of m_d are 0. For each color c , there is an associated color matrix $n_c \in \mathcal{R}^{k \times k}$. The image x is constructed by first computing $m_d \circ n_c$ where \circ denotes multiplication where each $s \times k$ slice of m_d is multiplied with n_c . Then independent standard Gaussian noise is added to each cell.

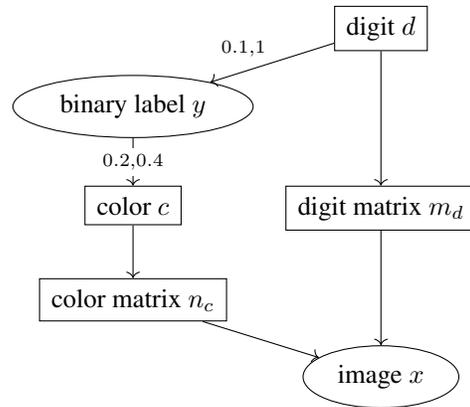


Figure 11. A hypothetical generative model of Complex Colored MNIST dataset. Ellipse variables are observable while square variables are not observable.

G. Limitations

Our work assumes realizability, a common assumption which does not always hold in practice. Nevertheless, experiments on the Colored MNIST example show that our method still works. Relaxation of the assumption is a topic for future works.

Our model does not include the case in reinforcement learning where the reward R_t at time t is a function of both the previous state S_{t-1} and the current state S_t .