
Model Transferability With Responsive Decision Subjects

Yang Liu¹ Yatong Chen¹ Zeyu Tang² Kun Zhang^{2,3}

Abstract

This paper studies model transferability when human decision subjects respond to a deployed machine learning model. In our setting, an agent or a user corresponds to a sample (X, Y) drawn from a distribution \mathcal{D} and will face a model h and its classification result $h(X)$. Agents can modify X to adapt to h , which will incur a distribution shift on (X, Y) . Therefore, when training h , the learner will need to consider the subsequently “induced” distribution when the output model is deployed. Our formulation is motivated by applications where the deployed machine learning models interact with human agents, and will ultimately face *responsive* and *interactive* data distributions. We formalize the discussions of the transferability of a model by studying how the model trained on the available source distribution (data) would translate to the performance on the induced domain. We provide both upper bounds for the performance gap due to the induced domain shift, as well as lower bound for the trade-offs that a classifier has to suffer on either the source training distribution or the induced target distribution. We provide further instantiated analysis for two popular domain adaptation settings with *covariate shift* and *target shift*.

1. Introduction

Decision makers are increasingly required to be transparent on their decision making to offer the “right to explanation” (Goodman & Flaxman, 2017; Selbst & Powles, 2018; Usun et al., 2019)¹. Being transparent also invites potential adaptations from the population, leading to potential shifts. We are motivated by settings where the deployed machine

¹Department of Computer Science and Engineering, University of California, Santa Cruz, CA, USA. ²Department of Philosophy, Carnegie Mellon University, USA. ³Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, United Arab Emirates. Correspondence to: Yang Liu <yangliu@ucsc.edu>.

Preprint.

¹See Appendix A.1 (supplemental material) for more detailed discussions.

learning models interact with human agents, which will ultimately face data distributions that reflect how human agents respond to the models. For instance, when a model is used to decide loan applications, candidates may adapt their features based on the model specification in order to maximize their chances of approval; thus the loan decision classifier observes a data distribution caused by its own deployment (e.g., see Figure 1 for a demonstration). Similar observations can be articulated for application in insurance sector (i.e. developing policy s.t. customers’ behaviors might adapt to lower premium (Haghtalab et al., 2020)), education sector (i.e. developing courses when students are less incentivized to cheat (Kleinberg & Raghavan, 2020)) and so on.

FEATURE	WEIGHT		ORIGINAL VALUE	→	ADAPTED VALUE
Income	2		\$ 6,000	→	\$ 6,000
Education Level	3		College	→	College
Debt	-10		\$40,000	→	\$20,000
Savings	5		\$20,000	→	\$0

Figure 1: An example of an agent who originally has both savings and debt, observes that the classifier penalizes debt (weight -10) more than it rewards savings (weight +5), and concludes that their most efficient adaptation is to use their savings to pay down their debt.

This paper investigates model transferability when the underlying distribution shift is induced by the model being deployed. What we would like is to have some guarantee on the *transferability* of a classifier — that is, how training on the available source distribution \mathcal{D}_S translates to performance on the induced domain $\mathcal{D}(h)$, which depends on the model h being deployed. A key concept in our setting is the *induced risk*, defined as the error a model incurs on the distribution induced by itself:

$$\text{Induced Risk} : \text{Err}_{\mathcal{D}(h)}(h) := \mathbb{P}_{\mathcal{D}(h)}(h(X) \neq Y) \quad (1)$$

Most relevant to the above formulation is the strategic classification literature (Hardt et al., 2016a; Chen et al., 2020a). In this literature, agents are modeled as rational utility maximizers and game theoretical solutions were proposed to characterize the induced risk. However, our results are motivated by the following challenges in more general scenarios:

- **Modeling assumptions being restrictive** In many practical situations, it is often hard to faithfully characterize agents’ utilities. Furthermore, agents might not be fully rational when they response. All the uncertainties can lead to a far more complicated distribution change in (X, Y) , as compared to often-made assumptions that agents only change X but not Y (Chen et al., 2020a).
- **Lack of access to response data** Machine learning practitioners may only have access to data from the source distribution during training, and although they anticipate changes in the population due to human agents’ responses, they cannot observe this new distribution until the model is actually deployed.
- **Retraining being costly** Even when samples from the induced data distribution are available, retraining the model from scratch may be impractical due to computational constraints.

The above observations motivate us to understand the transferability of a model trained on the source data to the domain induced by the deployment of itself. We study several fundamental questions:

- **Source risk \Rightarrow Induced risk** For a given model h , how different is $\text{Err}_{\mathcal{D}(h)}(h)$, the error on the distribution induced by h , from $\text{Err}_{\mathcal{D}_S}(h) := \mathbb{P}_{\mathcal{D}_S}(h(X) \neq Y)$, the error on the source distribution?
- **Induced risk \Rightarrow Minimum induced risk** How much higher is $\text{Err}_{\mathcal{D}(h)}(h)$, the error on the induced distribution, than $\min_{h'} \text{Err}_{\mathcal{D}(h')}(h')$, the minimum achievable induced error?
- **Induced risk of source optimal \Rightarrow Minimum induced risk** Of particular interest, and as a special case of the above, how does $\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*)$, the induced error of the optimal model trained on the source distribution $h_S^* := \min_h \text{Err}_{\mathcal{D}_S}(h)$, compare to $\min_h \text{Err}_{\mathcal{D}(h)}(h)$?
- **Lower bound for learning tradeoffs** What is the minimum error a model must incur on either the source distribution $\text{Err}_{\mathcal{D}_S}(h)$ or its induced distribution $\text{Err}_{\mathcal{D}(h)}(h)$?

For the first three questions, we prove upper bounds on the additional error incurred when a model trained on a source distribution is transferred over to its induced domain. We also provide lower bounds for the trade-offs a classifier has to suffer on either the source training distribution or the induced target distribution. We then show how to specialize our results to two popular domain adaptation settings: *covariate shift* and *target shift* All omitted proofs can be found in the Appendix (supplementary materials).

1.1. Related works

Most relevant to us are three topics: strategic classification (Hardt et al., 2016a; Chen et al., 2020a; Dekel et al., 2010;

Dong et al., 2018; Chen et al., 2020b; Miller et al., 2020; Kleinberg & Raghavan, 2020), a recently proposed notion of *performative prediction* (Perdomo et al., 2020; Mendler-Dünner et al., 2020), and domain adaptation (Jiang, 2008; Ben-David et al., 2010; Sugiyama et al., 2008; Zhang et al., 2019; Kang et al., 2019; Zhang et al., 2020). Hardt et al. (Hardt et al., 2016a) pioneered the formalization of strategic behavior in classification based on a sequential two-player game between agents and classifiers. Most of the existing literature focuses on finding the optimal classifier by assuming fully rational agents (and by characterizing the equilibrium response). In contrast, we do not make these assumptions and primarily study the transferability when only having knowledge of source data.

Our result was inspired by the transferability results in domain adaptations (Ben-David et al., 2010; Crammer et al., 2008; David et al., 2010). Later works examined specific domain adaptation models, such as covariate shift (Shimodaira, 2000; Zadrozny, 2004; Gretton et al., 2009; Sugiyama et al., 2008; Zhang et al., 2013a;b) and target/label shift (Lipton et al., 2018; Azizzadenesheli et al., 2019). A commonly established solution is to perform reweighted training on the source data, and robust and efficient solutions have been developed to estimate the weights accurately (Sugiyama et al., 2008; Zhang et al., 2013a;b; Lipton et al., 2018; Guo et al., 2020).

Our work, at the first sight, looks similar to several other area of studies. For instance, the notion of observing an “induced distribution” resembles similarity to the adversarial machine learning literature (Lowd & Meek, 2005; Huang et al., 2011; Vorobeychik & Kantarcioglu, 2018). One of the major differences between us and adversarial machine learning is the true label Y stays the same for the attacked feature while in our paper, both X and Y might change in the adapted distribution $\mathcal{D}(h)$. In Appendix A.2, we provide detailed comparisons with some areas in domain adaptations, including domain generalization, adversarial attack and test-time adaptation.

2. Formulation

Suppose we are learning a parametric model $h \in \mathcal{H}$ for a binary classification problem. Its training data set $S := \{x_i, y_i\}_{i=1}^N$ is drawn from a *source* distribution \mathcal{D}_S , where $x_i \in \mathbb{R}^d$ and $y_i \in \{-1, +1\}$. However, h will then be deployed in a setting where the samples come from a *test* or *target* distribution \mathcal{D}_T that can differ substantially from \mathcal{D}_S . Therefore instead of minimizing the prediction error on the source distribution $\text{Err}_{\mathcal{D}_S}(h) := \mathbb{P}_{\mathcal{D}_S}(h(X) \neq Y)$, the goal is to find h^* that minimizes $\text{Err}_{\mathcal{D}_T}(h) := \mathbb{P}_{\mathcal{D}_T}(h(X) \neq Y)$. This is often referred to as the *domain adaptation problem*, where typically, the transition from \mathcal{D}_S to \mathcal{D}_T is assumed to be independent of the model h being deployed.

We consider a setting in which the distribution shift depends on h , or is thought of as being *induced* by h . We will use $\mathcal{D}(h)$ to denote the *induced domain* by h :

$$\mathcal{D}_S \rightarrow \text{encounters model } h \rightarrow \mathcal{D}(h)$$

Strictly speaking, the induced distribution is a function of both \mathcal{D}_S and h and should be better denoted by $\mathcal{D}_S(h)$. To ease the notation, we will stick with $\mathcal{D}(h)$, but we shall keep in mind of its dependency of \mathcal{D}_S .

The challenge in the above setting is that when training h , the learner needs to carry the thoughts that $\mathcal{D}(h)$ should be the distribution it will be evaluated on and that the training cares about. Formally, we define the *induced risk* of a classifier h as the 0-1 error on the distribution h induces:

$$\text{Induced risk : } \text{Err}_{\mathcal{D}(h)}(h) := \mathbb{P}_{\mathcal{D}(h)}(h(X) \neq Y) \quad (2)$$

Denote by $h_T^* := \arg \min_{h \in \mathcal{H}} \text{Err}_{\mathcal{D}(h)}(h)$ the classifier with minimum induced risk. More generally, when the loss may not be the 0-1 loss, we define the *induced ℓ -risk* as

$$\text{Induced } \ell\text{-risk : } \text{Err}_{\ell, \mathcal{D}(h)}(h) := \arg \min_{h \in \mathcal{H}} \mathbb{E}_{z \sim \mathcal{D}(h)}[\ell(h; z)]$$

The induced risks will be the primary quantities that we are interested in minimizing. The following additional notation will also be helpful:

- Distributions of Y on a distribution \mathcal{D} : $\mathcal{D}_Y := \mathbb{P}_{\mathcal{D}}(Y = y)^2$, and in particular $\mathcal{D}_Y(h) := \mathbb{P}_{\mathcal{D}(h)}(Y = y)$, $\mathcal{D}_{Y|S} := \mathbb{P}_{\mathcal{D}_S}(Y = y)$.
- Distribution of h on a distribution \mathcal{D} : $\mathcal{D}_h := \mathbb{P}_{\mathcal{D}}(h(X) = y)$, and in particular $\mathcal{D}_h(h) := \mathbb{P}_{\mathcal{D}(h)}(h(X) = y)$, $\mathcal{D}_{h|S} := \mathbb{P}_{\mathcal{D}_S}(h(X) = y)$.
- Marginal distribution of X for a distribution \mathcal{D} : $\mathcal{D}_X := \mathbb{P}_{\mathcal{D}}(X = x)$, and in particular $\mathcal{D}_X(h) := \mathbb{P}_{\mathcal{D}(h)}(X = x)$, $\mathcal{D}_{X|S} := \mathbb{P}_{\mathcal{D}_S}(X = x)^3$.
- Total variation distance defined between \mathcal{D} and \mathcal{D}' (Ali & Silvey, 1966): $d_{\text{TV}}(\mathcal{D}, \mathcal{D}') := \sup_{\mathcal{O}} |\mathbb{P}_{\mathcal{D}}(\mathcal{O}) - \mathbb{P}_{\mathcal{D}'}(\mathcal{O})|$.

2.1. Examples of Distribution Shifts Induced by Model Deployment

We provide two exemplary models to demonstrate the use cases for the distribution shift models described in our paper.

Strategic Classification An example of distribution shift is the setting where decision subjects perform *strategic response* to a decision rule. It is well-known that when human agents are subject to a decision rule, they will adapt their

²The “:=” defines the RHS as the probability measure function for the LHS.

³For continuous X , the probability measure shall be read as the density function.

feature so as to get a favorable prediction outcome. In the literature of strategic classification, we say the human agents perform strategic adaptation (Hardt et al., 2016a).

It is natural to assume that the feature distribution before and after the human agents’ best response satisfies *covariate shift*: namely the feature distribution $\mathbb{P}(X)$ will change, but $\mathbb{P}(Y|X)$, the mapping between Y and X , remain unchanged. We use Figure 2 (Up) as a demonstrating of how distribution might shift for strategic response setting. In Section 4.3, we will use the strategic classification setup to verify our obtained results.




Figure 2: Example causal graph annotated to demonstrate covariate shift (**Up**) / target shift (**Down**) as a result of the deployment of h . Grey nodes indicate observable variables and transparent nodes are not observed at the training stage. Red arrow emphasises h induces changes of certain variables.

Replicator Dynamics Replicator dynamics is a commonly used model to study the evolution of an adopted “strategy” in evolutionary game theory (Tuyls et al., 2006; Friedman & Sinervo, 2016; Taylor & Jonker, 1978; Raab & Liu, 2021). The core notion of it is the growth or decline of the population of each strategy depends on its “fitness”. Consider the label $Y = \{-1, +1\}$ as the strategy, and the following behavioral response model to capture the induced target shift:

$$\frac{\mathbb{P}_{\mathcal{D}(h)}(Y = +1)}{\mathbb{P}_{\mathcal{D}_S}(Y = +1)} = \frac{\text{Fitness}(Y = +1)}{\mathbb{E}[\text{Fitness}(Y)]}$$

In short, the change of the $Y = +1$ population depends on how predicting $Y = +1$ “fits” a certain utility function. For instance, the “fitness” can take the form of the prediction accuracy of h for class +1. With assuming $\mathbb{P}(X|Y)$ stays unchanged, this instantiates one example of a specific induced *target shift*. We will specify the condition for target shift in Section 5. We use Figure 2 (Down) as a demonstrating of how distribution might shift for the replicator dynamic setting. In Section 5.3, we will use a detailed replicator dynamics model to further instantiate our results.

3. Transferability of Learning to Induced Domains

In this section, we first provide upper bounds for the transfer error of a classifier h (that is, the difference between $\text{Err}_{\mathcal{D}(h)}(h)$ and $\text{Err}_{\mathcal{D}_S}(h)$), as well as between $\text{Err}_{\mathcal{D}(h)}(h)$

and $\text{Err}_{\mathcal{D}(h_T^*)}(h_T^*)$. We then provide lower bounds for $\max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\}$; that is, the minimum error a model h must incur on either the source distribution \mathcal{D}_S or the induced distribution $\mathcal{D}(h)$.

3.1. Upper Bound

We first investigate upper bounds for the transfer errors. We begin by showing generic upper bounds, and further strengthen the bound for specific domain adaptation settings in Section 4 and 5. We begin with answering a central question in domain adaptation:

How does a model h trained on its training dataset fare on the induced distribution $\mathcal{D}(h)$?

To that end, define the minimum and maximum combined error of two distributions \mathcal{D} and \mathcal{D}' as:

$$\begin{aligned}\lambda_{\mathcal{D} \rightarrow \mathcal{D}'} &:= \min_{h' \in \mathcal{H}} \text{Err}_{\mathcal{D}'}(h') + \text{Err}_{\mathcal{D}}(h') \\ \Lambda_{\mathcal{D} \rightarrow \mathcal{D}'} &:= \max_{h' \in \mathcal{H}} \text{Err}_{\mathcal{D}'}(h') + \text{Err}_{\mathcal{D}}(h')\end{aligned}$$

and the \mathcal{H} -divergence (Ben-David et al., 2010) as

$$\begin{aligned}d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}, \mathcal{D}') \\ = 2 \sup_{h, h' \in \mathcal{H}} |\mathbb{P}_{\mathcal{D}}(h(X) \neq h'(X)) - \mathbb{P}_{\mathcal{D}'}(h(X) \neq h'(X))|\end{aligned}$$

The \mathcal{H} -divergence is celebrated measure proposed in the domain adaptation literature (Ben-David et al., 2010) which will be useful for bounding the difference in errors of two classifiers. Repeating classical arguments from (Ben-David et al., 2010), we can easily prove the following:

Theorem 3.1 (Source risk \Rightarrow Induced risk). *The difference between $\text{Err}_{\mathcal{D}(h)}(h)$ and $\text{Err}_{\mathcal{D}_S}(h)$ is upper bounded by: $\text{Err}_{\mathcal{D}(h)}(h) \leq \text{Err}_{\mathcal{D}_S}(h) + \lambda_{\mathcal{D}_S \rightarrow \mathcal{D}(h)} + \frac{1}{2}d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}_S, \mathcal{D}(h))$.*

The transferability of a model h between $\text{Err}_{\mathcal{D}(h)}(h)$ and $\text{Err}_{\mathcal{D}_S}(h)$ looks precisely the same as in the classical domain adaptation setting. The above practice informs us that the classical transferability bounds under domain adaptation still hold when the adaptation is induced by the model too. Nonetheless, an arguably more interesting quantity in our setting to understand is the difference between the induced error of a given model h and the error induced by a globally optimal model:

$$\text{Err}_{\mathcal{D}(h)}(h) - \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) \quad (3)$$

The proof is slightly more involved, and the bound differs from the one in Theorem 3.1:

Theorem 3.2 (Induced risk \Rightarrow Minimum induced risk). *The difference between $\text{Err}_{\mathcal{D}(h)}(h)$ and $\text{Err}_{\mathcal{D}(h_T^*)}(h_T^*)$ is upper bounded by:*

$$\begin{aligned}\text{Err}_{\mathcal{D}(h)}(h) - \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) \\ \leq \frac{\lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)} + \Lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)}}{2} + \frac{1}{2} \cdot d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)).\end{aligned}$$

The above theorem informs us that the induced transfer error is bounded by the ‘‘average’’ achievable error on both distributions $\mathcal{D}(h)$ and $\mathcal{D}(h_T^*)$, as well as the $\mathcal{H} \times \mathcal{H}$ divergence between the two distributions. Reflecting on the difference between the bounds of Theorem 3.1 and Theorem 3.2, we see that the primary change is replacing the minimum achievable error λ with the average of λ and Λ .

3.2. Lower Bound

Now we provide a lower bound on the induced transfer error. We particularly want to show that at least one of the two errors $\text{Err}_{\mathcal{D}_S}(h)$, $\text{Err}_{\mathcal{D}(h)}(h)$ must be lower-bounded by a certain quantity.

Theorem 3.3 (Lower bound for learning tradeoffs). *Any model h must incur the following error on either the source or induced distribution:*

$$\max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \geq \frac{d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))}{2}.$$

The proof leverages the triangle inequality of d_{TV} . This bound is dependent on h ; however, by the data processing inequality of d_{TV} (and f -divergence functions in general) (Liese & Vajda, 2006), we have

$$d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) \leq d_{\text{TV}}(\mathcal{D}_{X|S}, \mathcal{D}_X(h))$$

Applying this to Theorem 3.3 gives the following model-independent bound:

Corollary 3.4. *For any model h ,*

$$\max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \geq \frac{d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{X|S}, \mathcal{D}_X(h))}{2}.$$

4. Covariate Shift

In this section, we focus on a particular domain adaptation setting known as *covariate shift*, in which the distribution of features changes, but the distribution of labels conditioned on features does not:

$$\mathbb{P}_{\mathcal{D}(h)}(Y = y|X = x) = \mathbb{P}_{\mathcal{D}_S}(Y = y|X = x) \quad (4)$$

$$\mathbb{P}_{\mathcal{D}(h)}(X = x) \neq \mathbb{P}_{\mathcal{D}_S}(X = x) \quad (5)$$

Thus with covariate shift, we have

$$\begin{aligned}\mathbb{P}_{\mathcal{D}(h)}(X = x, Y = y) \\ = \mathbb{P}_{\mathcal{D}(h)}(Y = y|X = x) \cdot \mathbb{P}_{\mathcal{D}(h)}(X = x) \\ = \mathbb{P}_{\mathcal{D}_S}(Y = y|X = x) \cdot \mathbb{P}_{\mathcal{D}(h)}(X = x)\end{aligned}$$

Let $\omega_x(h) := \frac{\mathbb{P}_{\mathcal{D}(h)}(X=x)}{\mathbb{P}_{\mathcal{D}_S}(X=x)}$ be the *importance weight* at x , which characterizes the amount of adaptation induced by h at instance x . Then for any loss function ℓ we have

Proposition 4.1 (Expected Loss on the New Distribution).

$$\mathbb{E}_{\mathcal{D}(h)}[\ell(h; X, Y)] = \mathbb{E}_{\mathcal{D}_S}[\omega_x(h) \cdot \ell(h; x, y)].$$

The above derivation was not new and offered the basis for performing importance reweighting when learning under covariate shift (Sugiyama et al., 2008). The particular form informs us that $\omega_x(h)$ controls the generation of $\mathcal{D}(h)$ and encodes its dependency of both \mathcal{D}_S and h , and is critical for deriving our results below.

4.1. Upper Bound

We now derive an upper bound for transferability under covariate shift. We will focus particularly on the optimal model trained on the source data \mathcal{D}_S , which we denote as $h_S^* := \arg \min_{h \in \mathcal{H}} \text{Err}_S(h)$. Recall that the classifier with minimum induced risk is denoted as $h_T^* := \arg \min_{h \in \mathcal{H}} \text{Err}_{\mathcal{D}(h)}(h)$. We can upper bound the difference between h_S^* and h_T^* as follows:

Theorem 4.2 (Suboptimality of h_S^*). *Let X be distributed according to \mathcal{D}_S . We have:*

$$\begin{aligned} & \text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) - \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) \\ & \leq \sqrt{\text{Err}_{\mathcal{D}_S}(h_T^*)} \cdot \left(\sqrt{\text{Var}(\omega_X(h_S^*))} + \sqrt{\text{Var}(\omega_X(h_T^*))} \right). \end{aligned}$$

This result can be interpreted as follows: h_T^* incurs an irreducible amount of error on the source dataset, represented by $\sqrt{\text{Err}_{\mathcal{D}_S}(h_T^*)}$. Moreover, the difference in error between h_S^* and h_T^* is at its maximum when the two classifiers induce adaptations in “opposite” directions; this is represented by the sum of the standard deviations of their importance weights, $\sqrt{\text{Var}(\omega_X(h_S^*))} + \sqrt{\text{Var}(\omega_X(h_T^*))}$.

4.2. Lower Bound

Recall from Theorem 3.3, for the general setting, it is unclear whether the lower bound is strictly positive or not. In this section, we provide further understanding for when the lower bound $\frac{d_{TV}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{TV}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))}{2}$ is indeed positive under covariate shift.

We show under several assumptions, our previously provided lower bound in Theorem 3.3 is indeed strictly positive in the covariate shift setting. Details of the required conditions are specified in the Appendix, but the intuitions of the conditions are:

- Increased $\omega_x(h)$ value points are more likely to have positive labels. (Assumption A.3)
- Increased $\omega_x(h)$ value points are more likely to be classified as positive instances. (Assumption A.4)
- For a classifier h , within all $h(X) = +1$ or $h(X) = -1$, a higher $\mathbb{P}_{\mathcal{D}}(Y = +1|X = x)$ associates with a higher $\omega_x(h)$. (Assumption A.5)

Theorem 4.3. *With Assumption A.3 - A.5, the following*

lower bound is strictly positive for covariate shift:

$$\begin{aligned} & \max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \\ & \geq \frac{d_{TV}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{TV}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))}{2} > 0. \end{aligned}$$

4.3. Example Using Strategic Classification

As introduced in Section 2.1, we consider a setting caused by *strategic response* in which agents are classified by and adapt to a binary threshold classifier.

Consider a setup where each agent is associated with a d dimensional continuous feature $x \in \mathbb{R}^d$ and a binary true qualification $y(x) \in \{-1, +1\}$, where $y(x)$ is a function of the feature vector x . Consistent with the literature in strategic classification (Hardt et al., 2016a), a simple case where after seeing the threshold binary decision rule $h(x) = 2 \cdot \mathbb{1}[x \geq \tau_h] - 1$, the agents will *best response* to it by maximizing the following utility function:

$$u(x, x') = h(x') - h(x) - c(x, x')$$

where $c(x, x')$ is the *cost function* for decision subjects to modify their feature from x to x' . Assume all agents are rational utility maximizers: they will only *attempt* to change their features when the benefit of manipulation is greater than the cost (i.e. when $c(x, x') \leq 2$) and agent will not change their feature if they are already accepted (i.e. $h(x) = +1$). For a given threshold τ_h and manipulation budget B , the theoretical best response of an agent with original feature x is:

$$\Delta(x) = \arg \max_{x'} u(x, x') \text{ s.t. } c(x, x') \leq B \quad (6)$$

We show that under some further characterizations of the agents’ responsive behaviors (see Assumption A.6 - A.9 in Appendix A.9), we can specify the bound in Theorem 4.2 for the strategic response setting as follows:

Proposition 4.4 (Upper bound for the Strategic Response Setting). *Under assumption Assumption A.6 - A.9, we can bound the differences between $\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*)$ and $\text{Err}_{\mathcal{D}(h_T^*)}(h_T^*)$ by*

$$\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) - \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) \leq \sqrt{\frac{2B}{3} \text{Err}_{\mathcal{D}_S}(h_T^*)}.$$

To interpret this result, we can see that the upper bound for strategic response depends on the manipulation budget B , and the error the ideal classifier made on the source distribution $\text{Err}_{\mathcal{D}_S}(h_T^*)$. This aligns with our intuition that the smaller manipulation budget is, the less agents will change their features, thus leading to a tighter upper bound on the difference between $\text{Err}_{h_S^*}(h_S^*)$ and $\text{Err}_{h_T^*}(h_T^*)$. This bound also allows us to bound this quantity even without the knowledge of the mapping between $\mathcal{D}(h)$ and h , since we can directly compute $\text{Err}_{\mathcal{D}_S}(h_T^*)$ from the source distribution and an estimated optimal classifier h_T^* .

5. Target Shift

We consider another popular domain adaptation setting known as *target shift*, in which the distribution of labels changes, but not the distribution of features conditioned on the label:

$$\mathbb{P}_{\mathcal{D}(h)}(X = x|Y = y) = \mathbb{P}_{\mathcal{D}_S}(X = x|Y = y) \quad (7)$$

$$\mathbb{P}_{\mathcal{D}(h)}(Y = y) \neq \mathbb{P}_{\mathcal{D}_S}(Y = y) \quad (8)$$

In the case of binary classification, let $\omega(h) := \mathbb{P}_{\mathcal{D}(h)}(Y = +1)$, and $\mathbb{P}_{\mathcal{D}(h)}(Y = -1) = 1 - \omega(h)$. Again, $\omega(h)$ encodes the induced adaptation from \mathcal{D}_S and h . Then we have for any proper loss function ℓ :

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}(h)}[\ell(h; X, Y)] \\ &= \omega(h) \cdot \mathbb{E}_{\mathcal{D}(h)}[\ell(h; X, Y)|Y = +1] \\ & \quad + (1 - \omega(h)) \cdot \mathbb{E}_{\mathcal{D}(h)}[\ell(h; X, Y)|Y = -1] \\ &= \omega(h) \cdot \mathbb{E}_{\mathcal{D}_S}[\ell(h; X, Y)|Y = +1] \\ & \quad + (1 - \omega(h)) \cdot \mathbb{E}_{\mathcal{D}_S}[\ell(h; X, Y)|Y = -1] \end{aligned}$$

We will adopt the following shorthands:

$$\text{Err}_+(h) := \mathbb{E}_{\mathcal{D}_S}[\ell(h; X, Y)|Y = +1]$$

$$\text{Err}_-(h) := \mathbb{E}_{\mathcal{D}_S}[\ell(h; X, Y)|Y = -1]$$

Note that $\text{Err}_+(h), \text{Err}_-(h)$ are both defined on the conditional source distribution, which is invariant under the target shift assumption.

5.1. Upper bound

We again upper bound the transferability of h_S^* under target shift. Denote by \mathcal{D}_+ the positive label distribution on \mathcal{D}_S ($\mathbb{P}_{\mathcal{D}_S}(X = x|Y = +1)$) and \mathcal{D}_- the negative label distribution on \mathcal{D}_S ($\mathbb{P}_{\mathcal{D}_S}(X = x|Y = -1)$). Let $p := \mathbb{P}_{\mathcal{D}_S}(Y = +1)$.

Theorem 5.1. *Under target shift, the difference between $\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*)$ and $\text{Err}_{\mathcal{D}(h_T^*)}(h_T^*)$ bounds as:*

$$\begin{aligned} & \text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) - \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) \leq |\omega(h_S^*) - \omega(h_T^*)| \\ & + (1 + p) \cdot (d_{TV}(\mathcal{D}_+(h_S^*), \mathcal{D}_+(h_T^*)) + d_{TV}(\mathcal{D}_-(h_S^*), \mathcal{D}_-(h_T^*))) \end{aligned}$$

The above upper bound consists of two components. The first quantity captures the difference between the two induced distributions $\mathcal{D}(h_S^*)$ and $\mathcal{D}(h_T^*)$. The second quantity characterizes the difference between the two classifiers h_S^*, h_T^* on the source distribution.

5.2. Lower Bound

Now we discuss lower bounds. Denote by $\text{TPR}_S(h)$ and $\text{FPR}_S(h)$ the true positive and false positive rates of h on the source distribution \mathcal{D}_S . We prove the following:

Theorem 5.2. *Under target shift, any model h must incur the following error on either the \mathcal{D}_S or $\mathcal{D}(h)$:*

$$\begin{aligned} & \max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \\ & \geq \frac{|p - \omega(h)| \cdot (1 - |\text{TPR}_S(h) - \text{FPR}_S(h)|)}{2}. \end{aligned}$$

The proof extends the bound of Theorem 3.3 by further explicating each of $d_{TV}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h))$, $d_{TV}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))$ under the assumption of target shift. Since $|\text{TPR}_S(h) - \text{FPR}_S(h)| < 0$ unless we have a trivial classifier that has either $\text{TPR}_S(h) = 1, \text{FPR}_S(h) = 0$ or $\text{TPR}_S(h) = 0, \text{FPR}_S(h) = 1$, the lower bound is strictly positive. Taking a closer look, the lower bound is determined linearly by how much the label distribution shifts: $p - \omega(h)$. The difference is further determined by the performance of h on the source distribution through $1 - |\text{TPR}_S(h) - \text{FPR}_S(h)|$.

5.3. Example Using Replicator Dynamics

Let us instantiate the discussion using a specific fitness function for the replicator dynamics model (Section 2.1), which is the prediction accuracy of h for class +1:

$$[\text{Fitness of } Y = +1] := \mathbb{P}_{\mathcal{D}_S}(h(X) = +1|Y = +1) \quad (9)$$

Then we have $\mathbb{E}[\text{Fitness of } Y] = \text{Err}_{\mathcal{D}_S}(h)$, and

$$\frac{\omega(h)}{\mathbb{P}_{\mathcal{D}_S}(Y = +1)} = \frac{\Pr_{\mathcal{D}_S}(h(X) = +1|Y = +1)}{\text{Err}_{\mathcal{D}_S}(h)}$$

Plugging the result back to our Theorem 5.1 we have

Proposition 5.3. *Under the replicator dynamics model in Eqn. (9), $|\omega(h_S^*) - \omega(h_T^*)|$ further bounds as:*

$$\begin{aligned} & |\omega(h_S^*) - \omega(h_T^*)| \leq \mathbb{P}_{\mathcal{D}_S}(Y = +1) \\ & \cdot \frac{|\text{Err}_{\mathcal{D}_S}(h_S^*) - \text{Err}_{\mathcal{D}_S}(h_T^*)| \cdot |\text{TPR}_S(h_S^*) - \text{TPR}_S(h_T^*)|}{\text{Err}_{\mathcal{D}_S}(h_S^*) \cdot \text{Err}_{\mathcal{D}_S}(h_T^*)}. \end{aligned}$$

That is, the difference between $\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*)$ and $\text{Err}_{\mathcal{D}(h_T^*)}(h_T^*)$ is further dependent on the difference between the two classifiers' performances on the source data \mathcal{D}_S . This offers an opportunity to evaluate the possible error transferability using the source data only.

Concluding Remarks We presented a sequence of model transferability results for settings where agents will respond to a deployed model. The response leads to an induced distribution that the learner would not know before deploying the model. Our results cover for both a general response setting and for specific ones (covariate shift and target shift). Unawareness of the potential distribution shift might lead to unintended consequence when training a machine learning model. One goal of this paper is to raise awareness of this issue for a safe deployment of machine learning methods in high-stake societal applications. Our contributions are mostly theoretical. A future direction is to collect real human experiment data to support the findings.

Acknowledgment This work is partially supported by the National Science Foundation (NSF) under grants IIS-2007951, IIS-2143895, IIS-2040800 (FAI program in collaboration with Amazon), and CCF-2023495.

References

- Ali, S. M. and Silvey, S. D. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society: Series B (Methodological)*, 28(1):131–142, 1966.
- Azizzadenesheli, K., Liu, A., Yang, F., and Anandkumar, A. Regularized learning for domain adaptation under label shifts. *arXiv preprint arXiv:1903.09734*, 2019.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. A theory of learning from different domains. *Machine Learning*, 79:151–175, 2010.
- Board of Governors of the Federal Reserve System (US). *Report to the congress on credit scoring and its effects on the availability and affordability of credit*. Board of Governors of the Federal Reserve System, 2007.
- Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., and Mukhopadhyay, D. Adversarial attacks and defences: A survey, 2018.
- Chen, Y., Liu, Y., and Podimata, C. Learning strategy-aware linear classifiers, 2020a.
- Chen, Y., Wang, J., and Liu, Y. Strategic recourse in linear classification. *arXiv preprint arXiv:2011.00355*, 2020b.
- Crammer, K., Kearns, M., and Wortman, J. Learning from multiple sources. *Journal of Machine Learning Research*, 9(8), 2008.
- David, S. B., Lu, T., Luu, T., and Pál, D. Impossibility theorems for domain adaptation. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 129–136. JMLR Workshop and Conference Proceedings, 2010.
- Dekel, O., Fischer, F., and Procaccia, A. D. Incentive compatible regression learning. *J. Comput. Syst. Sci.*, 76(8): 759–777, December 2010.
- Dong, J., Roth, A., Schutzman, Z., Waggoner, B., and Wu, Z. S. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC ’18, pp. 55–70, New York, NY, USA, 2018. Association for Computing Machinery.
- Friedman, D. and Sinervo, B. *Evolutionary games in natural, social, and virtual worlds*. Oxford University Press, 2016.
- Goodman, B. and Flaxman, S. European union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine*, 38(3):50–57, Oct 2017.
- Gretton, A., Smola, A., Huang, J., Schmittfull, M., Borgwardt, K., and Schölkopf, B. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4): 5, 2009.
- Guo, J., Gong, M., Liu, T., Zhang, K., and Tao, D. Ltf: A label transformation framework for correcting label shift. In *International Conference on Machine Learning*, pp. 3843–3853. PMLR, 2020.
- Haghtalab, N., Immorlica, N., Lucier, B., and Wang, J. Z. Maximizing welfare with incentive-aware evaluation mechanisms. In Bessiere, C. (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pp. 160–166. International Joint Conferences on Artificial Intelligence Organization, 2020. doi: 10.24963/ijcai.2020/23. URL <https://doi.org/10.24963/ijcai.2020/23>.
- Hardt, M., Megiddo, N., Papadimitriou, C., and Wootters, M. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pp. 111–122, New York, NY, USA, 2016a. Association for Computing Machinery.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pp. 3315–3323, 2016b.
- Huang, L., Joseph, A. D., Nelson, B., Rubinstein, B. I., and Tygar, J. D. Adversarial machine learning. In *ACM Workshop on Security and Artificial Intelligence*, pp. 43–58, 2011.
- Jiang, J. A literature survey on domain adaptation of statistical classifiers. URL: <http://sifaka.cs.uiuc.edu/jiang4/domainadaptation/survey>, 3:1–12, 2008.
- Kang, G., Jiang, L., Yang, Y., and Hauptmann, A. G. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4893–4902, 2019.
- Kleinberg, J. and Raghavan, M. How do classifiers induce agents to invest effort strategically? *ACM Transactions on Economics and Computation (TEAC)*, 8(4):1–23, 2020.
- Li, D., Yang, Y., Song, Y.-Z., and Hospedales, T. M. Learning to generalize: Meta-learning for domain generalization, 2017.
- Liese, F. and Vajda, I. On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412, 2006.

- Lipton, Z., Wang, Y.-X., and Smola, A. Detecting and correcting for label shift with black box predictors. In *International conference on machine learning*, pp. 3122–3130. PMLR, 2018.
- Liu, Y. and Liu, M. An online learning approach to improving the quality of crowd-sourcing. *ACM SIGMETRICS Performance Evaluation Review*, 43(1):217–230, 2015.
- Lowd, D. and Meek, C. Adversarial learning. In *ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, pp. 641–647, 2005.
- Mendler-Dünner, C., Perdomo, J., Zrnic, T., and Hardt, M. Stochastic optimization for performative prediction. In *Advances in Neural Information Processing Systems*, pp. 4929–4939. Curran Associates, Inc., 2020.
- Miller, J., Milli, S., and Hardt, M. Strategic classification is causal modeling in disguise. In *International Conference on Machine Learning*, pp. 6917–6926. PMLR, 2020.
- Muandet, K., Balduzzi, D., and Schölkopf, B. Domain generalization via invariant feature representation, 2013.
- Nado, Z., Padhy, S., Sculley, D., D’Amour, A., Lakshminarayanan, B., and Snoek, J. Evaluating prediction-time batch normalization for robustness under covariate shift, 2021.
- Papernot, N., McDaniel, P., and Goodfellow, I. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples, 2016.
- Perdomo, J., Zrnic, T., Mendler-Dünner, C., and Hardt, M. Performative prediction. In *International Conference on Machine Learning*, pp. 7599–7609. PMLR, 2020.
- Raab, R. and Liu, Y. Unintended selection: Persistent qualification rate disparities and interventions. *Advances in Neural Information Processing Systems*, 34, 2021.
- Selbst, A. and Powles, J. “meaningful information” and the right to explanation. In Friedler, S. A. and Wilson, C. (eds.), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pp. 48–48. PMLR, 23–24 Feb 2018.
- Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.
- Song, C., He, K., Wang, L., and Hopcroft, J. E. Improving the generalization of adversarial training with domain adaptation, 2019.
- Sugiyama, M., Suzuki, T., Nakajima, S., Kashima, H., von Bünau, P., and Kawanabe, M. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008.
- Taylor, P. D. and Jonker, L. B. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1): 145–156, 1978. ISSN 0025-5564.
- Tuyls, K., Hoen, P. J., and Vanschoenwinkel, B. An evolutionary dynamical analysis of multi-agent learning in iterated games. *Autonomous Agents and Multi-Agent Systems*, 12(1):115–153, 2006.
- Ustun, B., Spangher, A., and Liu, Y. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 10–19, 2019.
- Varsavsky, T., Orbes-Arteaga, M., Sudre, C. H., Graham, M. S., Nachev, P., and Cardoso, M. J. Test-time unsupervised domain adaptation, 2020.
- Vorobeychik, Y. and Kantarcioglu, M. *Adversarial Machine Learning*. Morgan & Claypool Publishers, 2018.
- Wang, D., Shelhamer, E., Liu, S., Olshausen, B., and Darrell, T. Tent: Fully test-time adaptation by entropy minimization, 2021a.
- Wang, J., Lan, C., Liu, C., Ouyang, Y., Qin, T., Lu, W., Chen, Y., Zeng, W., and Yu, P. S. Generalizing to unseen domains: A survey on domain generalization, 2021b.
- Zadrozny, B. Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 114, 2004.
- Zhang, K., Schölkopf, B., Muandet, K., and Wang, Z. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pp. 819–827. PMLR, 2013a.
- Zhang, K., Zheng, V., Wang, Q., Kwok, J., Yang, Q., and Marsic, I. Covariate shift in hilbert space: A solution via surrogate kernels. In *International Conference on Machine Learning*, pp. 388–395. PMLR, 2013b.
- Zhang, K., Gong, M., Stojanov, P., Huang, B., LIU, Q., and Glymour, C. Domain adaptation as a problem of inference on graphical models. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 4965–4976. Curran Associates, Inc., 2020.
- Zhang, Y., Liu, T., Long, M., and Jordan, M. Bridging theory and algorithm for domain adaptation. In *International Conference on Machine Learning*, pp. 7404–7413. PMLR, 2019.

A. Appendix

We arrange the appendix as follows:

- Appendix A.1 provides some real life scenarios where transparent models are useful or required.
- Appendix A.2 provides comparisons of our setting and other sub-areas in domain adaptation.
- Appendix A.3 provides proof for Theorem 3.1.
- Appendix A.4 provides proof for Theorem 3.2.
- Appendix A.5 provides proof of Theorem 3.3.
- Appendix A.6 provides proof for Proposition 4.1.
- Appendix A.7 provides proof for Theorem 4.2.
- Appendix A.8 provides proof for Theorem 4.3.
- Appendix A.9 provides omitted assumptions and proof for Section 4.3.
- Appendix A.10 provides proof for Theorem 5.1.
- Appendix A.11 provides proof for Theorem 5.2.
- Appendix A.12 provides proof for Proposition 5.3.
- Appendix B provides missing experimental details.

A.1. Example Usages of Transparent Models

As we mentioned in Section 1, there is an increasing requirement of making the decision rule to be transparent due to its potential consequences impacts to individual decision subject. Here we provide the following reasons for using transparent models:

- Government regulation may require the model to be transparent, especially in public services;
- In some cases, companies may want to disclose their models so users will have explanations and are incentivized to better use the provided services.
- Regardless of whether models are published voluntarily, model parameters can often be inferred via well-known query “attacks”.

In addition, we name some concrete examples of some real-life applications:

- Consider the *Medicaid health insurance program* in the United States, which serves low-income people. There is an obligation to provide transparency/disclose the rules (model to automate the decisions) that decide whether individuals qualify for the program — in fact, most public services have “terms” that are usually set in stone and explained in the documentation. Agents can observe the rules and will adapt their profiles to be qualified if needed. For instance, an agent can decide to provide additional documentation they need to guarantee approval. For more applications along these lines, please refer to this report⁴.
- Credit score companies directly publish their criteria for assessing credit risk scores. In loan application settings, companies actually have the incentive to release criteria to incentivize agents to meet their qualifications and use their services. Furthermore, making decision models transparent will gain the trust of users.
- It is also known that it is possible to steal model parameters, if agents have incentives to do so⁵. For instance, spammers frequently infer detection mechanisms by sending different email variants; they then adjust their spam content accordingly.

⁴<https://datasociety.net/library/poverty-lawgorithms/>

⁵<https://www.wired.com/2016/09/how-to-steal-an-ai/>

A.2. Comparison of our setting and Some Areas in Domain Adaptation

We compare our setting (We address it as IDA, representing “induced domain adaptation”) with the following areas:

- Adversarial attack (Chakraborty et al., 2018; Papernot et al., 2016; Song et al., 2019): in adversarial attack, the true label Y stays the same for the attacked feature, while in IDA, we allow the true label to change as well. One can think of adversarial attack as a specific form of IDA where the induced distribution has a specific target, that is to maximize the classifier’s error by only perturbing/modifying. Our transferability bound does, however, provide insights for how standard training results transfer to the attack setting.
- Domain generalization (Wang et al., 2021b; Li et al., 2017; Muandet et al., 2013): the goal of domain generalization is to learn a more general model that can be generalized to any unseen distribution; On the contrary, our focus is to understand how the performance of a classifier trained on the source distribution degrades when evaluated on the induced distribution (which depends on how the population of decision subjects responds); this degradation depends on the classifier itself.
- Test-time adaptation (Varsavsky et al., 2020; Wang et al., 2021a; Nado et al., 2021): the issue of test-time adaptation falls into the classical domain adaptation setting where the adaptation is independent of the model being deployed. Applying this technique to solve our problem requires accessing data (either unsupervised or supervised) drawn from $\mathcal{D}_S(h)$ for each h being evaluated during different training epochs.

A.3. Proof of Theorem 3.1

Proof. We first establish two lemmas that will be helpful for bounding the errors of a pair of classifiers. Both are standard results from the domain adaption literature (Ben-David et al., 2010).

Lemma A.1. For any hypotheses $h, h' \in \mathcal{H}$ and distributions $\mathcal{D}, \mathcal{D}'$,

$$|\text{Err}_{\mathcal{D}}(h, h') - \text{Err}_{\mathcal{D}'}(h, h')| \leq \frac{d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}, \mathcal{D}')}{2}.$$

Proof. Define the-cross prediction disagreement between two classifiers h, h' on a distribution \mathcal{D} as $\text{Err}_{\mathcal{D}}(h, h') := \mathbb{P}_{\mathcal{D}}(h(X) \neq h'(X))$. By the definition of the \mathcal{H} -divergence,

$$\begin{aligned} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}, \mathcal{D}') &= 2 \sup_{h, h' \in \mathcal{H}} |\mathbb{P}_{\mathcal{D}}(h(X) \neq h'(X)) - \mathbb{P}_{\mathcal{D}'}(h(X) \neq h'(X))| \\ &= 2 \sup_{h, h' \in \mathcal{H}} |\text{Err}_{\mathcal{D}}(h, h') - \text{Err}_{\mathcal{D}'}(h, h')| \\ &\geq 2 |\text{Err}_{\mathcal{D}}(h, h') - \text{Err}_{\mathcal{D}'}(h, h')|. \end{aligned}$$

□

Another helpful lemma for us is the well-known fact that the 0-1 error obeys the triangle inequality (see, e.g., (Cramer et al., 2008)):

Lemma A.2. For any distribution \mathcal{D} over instances and any labeling functions f_1, f_2 , and f_3 , we have $\text{Err}_{\mathcal{D}}(f_1, f_2) \leq \text{Err}_{\mathcal{D}}(f_1, f_3) + \text{Err}_{\mathcal{D}}(f_2, f_3)$.

Denote by \bar{h}^* the ideal joint hypothesis, which minimizes the combined error:

$$\bar{h}^* := \arg \min_{h' \in \mathcal{H}} \text{Err}_{\mathcal{D}(h)}(h') + \text{Err}_{\mathcal{D}_S}(h')$$

We have:

$$\begin{aligned} \text{Err}_{\mathcal{D}(h)}(h) &\leq \text{Err}_{\mathcal{D}(h)}(\bar{h}^*) + \text{Err}_{\mathcal{D}(h)}(h, \bar{h}^*) && \text{(Lemma A.2)} \\ &\leq \text{Err}_{\mathcal{D}(h)}(\bar{h}^*) + \text{Err}_{\mathcal{D}_S}(h, \bar{h}^*) + |\text{Err}_{\mathcal{D}(h)}(h, \bar{h}^*) - \text{Err}_{\mathcal{D}_S}(h, \bar{h}^*)| \\ &\leq \text{Err}_{\mathcal{D}(h)}(\bar{h}^*) + \text{Err}_{\mathcal{D}_S}(h) + \text{Err}_{\mathcal{D}_S}(\bar{h}^*) + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}_S, \mathcal{D}(h)) && \text{(Lemma A.1)} \\ &= \text{Err}_{\mathcal{D}_S}(h) + \lambda_{\mathcal{D}_S \rightarrow \mathcal{D}(h)} + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}_S, \mathcal{D}(h)). && \text{(Definition of } \bar{h}^*) \end{aligned}$$

□

A.4. Proof of Theorem 3.2

Proof. Invoking Theorem 3.1, and replacing h with h_T^* and S with $\mathcal{D}(h_T^*)$, we have

$$\text{Err}_{\mathcal{D}(h)}(h_T^*) \leq \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) + \lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)} + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) \quad (10)$$

Now observe that

$$\begin{aligned} \text{Err}_{\mathcal{D}(h)}(h) &\leq \text{Err}_{\mathcal{D}(h)}(h_T^*) + \text{Err}_{\mathcal{D}(h)}(h, h_T^*) \\ &\leq \text{Err}_{\mathcal{D}(h)}(h_T^*) + \text{Err}_{\mathcal{D}(h_T^*)}(h, h_T^*) + \left| \text{Err}_{\mathcal{D}(h)}(h, h_T^*) - \text{Err}_{\mathcal{D}(h_T^*)}(h, h_T^*) \right| \\ &\leq \text{Err}_{\mathcal{D}(h)}(h_T^*) + \text{Err}_{\mathcal{D}(h_T^*)}(h, h_T^*) + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) && \text{(by Lemma A.1)} \\ &\leq \text{Err}_{\mathcal{D}(h)}(h_T^*) + \text{Err}_{\mathcal{D}(h_T^*)}(h) + \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) && \text{(by Lemma A.2)} \\ &\leq \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) + \lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)} + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) && \text{(by equation 10)} \\ &\quad + \text{Err}_{\mathcal{D}(h_T^*)}(h) + \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) + \frac{1}{2} d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) \end{aligned}$$

Adding $\text{Err}_{\mathcal{D}(h)}(h)$ to both sides and rearranging terms yields

$$\begin{aligned} 2\text{Err}_{\mathcal{D}(h)}(h) - 2\text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) &\leq \text{Err}_{\mathcal{D}(h)}(h) + \text{Err}_{\mathcal{D}(h_T^*)}(h) + \lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)} + d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) \\ &\leq \Lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)} + \lambda_{\mathcal{D}(h) \rightarrow \mathcal{D}(h_T^*)} + d_{\mathcal{H} \times \mathcal{H}}(\mathcal{D}(h_T^*), \mathcal{D}(h)) \end{aligned}$$

Dividing both sides by 2 completes the proof. □

A.5. Proof of Theorem 3.3

Proof. Using the triangle inequality of d_{TV} , we have

$$d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) \leq d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_{h|S}) + d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) + d_{\text{TV}}(\mathcal{D}_h(h), \mathcal{D}_Y(h)) \quad (11)$$

and by the definition of d_{TV} , the divergence term $d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h))$ becomes

$$\begin{aligned} d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_{h|S}) &= |\mathbb{P}_{\mathcal{D}_S}(Y = +1) - \mathbb{P}_{\mathcal{D}_S}(h(X) = +1)| \\ &= \left| \frac{\mathbb{E}_{\mathcal{D}_S}[Y] + 1}{2} - \frac{\mathbb{E}_{\mathcal{D}_S}[h(X)] + 1}{2} \right| \\ &= \left| \frac{\mathbb{E}_{\mathcal{D}_S}[Y]}{2} - \frac{\mathbb{E}_{\mathcal{D}_S}[h(X)]}{2} \right| \\ &\leq \frac{1}{2} \cdot \mathbb{E}_{\mathcal{D}_S}[|Y - h(X)|] \\ &= \text{Err}_{\mathcal{D}_S}(h) \end{aligned}$$

Similarly, we have

$$d_{\text{TV}}(\mathcal{D}_h(h), \mathcal{D}_Y(h)) \leq \text{Err}_{\mathcal{D}(h)}(h)$$

As a result, we have

$$\begin{aligned} \text{Err}_{\mathcal{D}_S}(h) + \text{Err}_{\mathcal{D}(h)}(h) &\geq d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_{h|S}) + d_{\text{TV}}(\mathcal{D}_h(h), \mathcal{D}_Y(h)) \\ &\geq d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) && \text{(by equation 11)} \end{aligned}$$

which implies

$$\max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \geq \frac{d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))}{2}.$$

□

A.6. Proof of Proposition 4.1

Proof.

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{D}(h)}[\ell(h; X, Y)] \\
 &= \int \mathbb{P}_{\mathcal{D}(h)}(X = x, Y = y) \ell(h; x, y) \, dx dy \\
 &= \int \mathbb{P}_{\mathcal{D}_S}(Y = y | X = x) \cdot \mathbb{P}_{\mathcal{D}(h)}(X = x) \ell(h; x, y) \, dx dy \\
 &= \int \mathbb{P}_{\mathcal{D}_S}(Y = y | X = x) \cdot \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot \frac{\mathbb{P}_{\mathcal{D}(h)}(X = x)}{\mathbb{P}_{\mathcal{D}_S}(X = x)} \cdot \ell(h; x, y) \, dx dy \\
 &= \int \mathbb{P}_{\mathcal{D}_S}(Y = y | X = x) \cdot \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot \omega_x(h) \cdot \ell(h; x, y) \, dx dy \\
 &= \mathbb{E}_{\mathcal{D}_S}[\omega_x(h) \cdot \ell(h; x, y)]
 \end{aligned}$$

□

A.7. Proof of Theorem 4.2

Proof. We start from the error induced by h_S^* . Let the average importance weight induced by h_S^* be $\bar{\omega}(h_S^*) = \mathbb{E}_{\mathcal{D}_S}[\omega_x(h_S^*)]$; we add and subtract this from the error:

$$\begin{aligned}
 \text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) &= \mathbb{E}_{\mathcal{D}_S}[\omega_x(h_S^*) \cdot \mathbf{1}(h_S^*(x) \neq y)] \\
 &= \mathbb{E}_{\mathcal{D}_S}[\bar{\omega}(h_S^*) \cdot \mathbf{1}(h_S^*(x) \neq y)] + \mathbb{E}_{\mathcal{D}_S}[(\omega_x(h_S^*) - \bar{\omega}(h_S^*)) \cdot \mathbf{1}(h_S^*(x) \neq y)]
 \end{aligned}$$

In fact, $\bar{\omega}(h_S^*) = 1$, since

$$\begin{aligned}
 \bar{\omega}(h_S^*) &= \mathbb{E}_{\mathcal{D}_S}[\omega_x(h_S^*)] = \int \omega_x(h_S^*) \mathbb{P}_{\mathcal{D}_S}(X = x) dx \\
 &= \int \frac{\mathbb{P}_{\mathcal{D}(h)}(X = x)}{\mathbb{P}_{\mathcal{D}_S}(X = x)} \mathbb{P}_{\mathcal{D}_S}(X = x) dx = \int \mathbb{P}_{\mathcal{D}(h)}(X = x) dx = 1
 \end{aligned}$$

Now consider any other classifier h . We have

$$\begin{aligned}
 & \text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) \\
 &= \mathbb{E}_{\mathcal{D}_S}[\mathbf{1}(h_S^*(x) \neq y)] + \mathbb{E}_{\mathcal{D}_S}[(\omega_x(h_S^*) - \bar{\omega}(h_S^*)) \cdot \mathbf{1}(h_S^*(x) \neq y)] \\
 &\leq \mathbb{E}_{\mathcal{D}_S}[\mathbf{1}(h(x) \neq y)] + \mathbb{E}_{\mathcal{D}_S}[(\omega_x(h_S^*) - \bar{\omega}(h_S^*)) \cdot \mathbf{1}(h_S^*(x) \neq y)] && \text{(by optimality of } h_S^* \text{ on } \mathcal{D}_S) \\
 &= \mathbb{E}_{\mathcal{D}_S}[\bar{\omega}(h) \cdot \mathbf{1}(h(x) \neq y)] + \mathbb{E}_{\mathcal{D}_S}[(\omega_x(h_S^*) - \bar{\omega}(h_S^*)) \cdot \mathbf{1}(h_S^*(x) \neq y)] && \text{(multiply by } \bar{\omega}(h_S^*) = 1) \\
 &= \mathbb{E}_{\mathcal{D}_S}[\omega_x(h) \cdot \mathbf{1}(h(x) \neq y)] + \mathbb{E}_{\mathcal{D}_S}[(\bar{\omega}(h) - \omega_x(h)) \cdot \mathbf{1}(h(x) \neq y)] && \text{(add and subtract } \bar{\omega}(h_S^*)) \\
 &\quad + \mathbb{E}_{\mathcal{D}_S}[(\omega_x(h_S^*) - \bar{\omega}(h_S^*)) \cdot \mathbf{1}(h_S^*(x) \neq y)] \\
 &= \text{Err}_{\mathcal{D}(h)}(h) + \text{Cov}(\omega_x(h_S^*), \mathbf{1}(h_S^*(x) \neq y)) - \text{Cov}(\omega_x(h), \mathbf{1}(h(x) \neq y))
 \end{aligned}$$

Moving the error terms to one side, we have

$$\begin{aligned}
 & \text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) - \text{Err}_{\mathcal{D}(h)}(h) \\
 &\leq \text{Cov}(\omega_x(h_S^*), \mathbf{1}(h_S^*(x) \neq y)) - \text{Cov}(\omega_x(h), \mathbf{1}(h(x) \neq y)) \\
 &\leq \sqrt{\text{Var}(\omega_x(h_S^*)) \cdot \text{Var}(\mathbf{1}(h_S^*(x) \neq y))} && (|\text{Cov}(X, Y)| \leq \sqrt{\text{Var}(X) \cdot \text{Var}(Y)}) \\
 &\quad + \sqrt{\text{Var}(\omega_x(h)) \cdot \text{Var}(\mathbf{1}(h(x) \neq y))} \\
 &= \sqrt{\text{Var}(\omega_x(h_S^*)) \cdot \text{Err}_S(h_S^*)(1 - \text{Err}_S(h_S^*))} + \sqrt{\text{Var}(\omega_x(h)) \cdot \text{Err}_{\mathcal{D}_S}(h)(1 - \text{Err}_{\mathcal{D}_S}(h))} \\
 &\leq \sqrt{\text{Var}(\omega_x(h_S^*)) \cdot \text{Err}_S(h_S^*)} + \sqrt{\text{Var}(\omega_x(h)) \cdot \text{Err}_{\mathcal{D}_S}(h)} && (1 - \text{Err}_{\mathcal{D}_S}(h) \leq 1) \\
 &\leq \sqrt{\text{Err}_{\mathcal{D}_S}(h)} \cdot \left(\sqrt{\text{Var}(\omega_x(h_S^*))} + \sqrt{\text{Var}(\omega_x(h))} \right)
 \end{aligned}$$

Since this holds for any h , it certainly holds for $h = h_T^*$.

□

A.8. Omitted Assumptions and Proof of Theorem 4.3

Denote $X_+(h) = \{x : \omega_x(h) \geq 1\}$ and $X_-(h) = \{x : \omega_x(h) < 1\}$. First we observe that

$$\begin{aligned} & \int_{X_+(h)} \mathbb{P}_{\mathcal{D}_S}(X = x)(1 - \omega_x(h))dx \\ & + \int_{X_-(h)} \mathbb{P}_{\mathcal{D}_S}(X = x)(1 - \omega_x(h))dx = 0 \end{aligned}$$

This is simply because of $\int_x \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot \omega_x(h)dx = \int_x \mathbb{P}_{\mathcal{D}(h)}(X = x)dx = 1$.

Now we provide detailed specifications of the assumptions for proving Theorem 4.3:

Assumption A.3 (increased $\omega_x(h)$ value points are more likely to have $Y = +1$).

$$\left| \int_{X_+(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1, X = x)(1 - \omega_x(h))dx \right| \geq \left| \int_{X_-(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1, X = x)(1 - \omega_x(h))dx \right|$$

Assumption A.4 (increased $\omega_x(h)$ value points are more likely to be classified as +1).

$$\left| \int_{X_+(h)} \mathbb{P}_{\mathcal{D}_S}(h(x) = +1, X = x)(1 - \omega_x(h))dx \right| \geq \left| \int_{X_-(h)} \mathbb{P}_{\mathcal{D}_S}(h(x) = +1, X = x)(1 - \omega_x(h))dx \right|$$

Assumption A.5. $\mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) - \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x)$ and $\omega_x(h)$ is positively correlated:

$$\text{Cov}(\mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) - \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x), \omega_x(h)) > 0$$

The above assumption states that for a deterministic classifier h , within all $h(X) = +1$ or $h(X) = -1$, a higher $\mathbb{P}_{\mathcal{D}}(Y = +1|X = x)$ associates with a higher $\omega_x(h)$. With the help of Assumption A.3 - Assumption A.5, we proceed to proof for Theorem 4.3:

Proof. Notice that in the setting of binary classification, we can write the total variation distance between $\mathcal{D}_{Y|S}$ and $\mathcal{D}_Y(h)$ as the difference between the probability of $Y = +1$ and the probability of $Y = -1$:

$$\begin{aligned} & d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) \\ & = |\mathbb{P}_{\mathcal{D}_S}(Y = +1) - \mathbb{P}_{\mathcal{D}(h)}(Y = +1)| \\ & = \left| \int \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x)\mathbb{P}_{\mathcal{D}_S}(X = x)dx - \int \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x)\mathbb{P}_{\mathcal{D}_S}(X = x)\omega_x(h)dx \right| \\ & = \left| \int \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x)\mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h))dx \right| \end{aligned} \tag{12}$$

Similarly we have

$$d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) = \left| \int \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x)\mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h))dx \right| \tag{13}$$

We can further expand the total variation distance between $\mathcal{D}_{Y|S}$ and $\mathcal{D}_Y(h)$ as follows:

$$\begin{aligned}
 & d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) \\
 &= \left| \int \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h)) dx \right| \\
 &= \underbrace{\left| \int_{X_+(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h)) dx \right|}_{\leq 0} \\
 &\quad + \underbrace{\left| \int_{X_-(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h)) dx \right|}_{> 0} \\
 &= - \int_{X_+(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h)) dx \\
 &\quad - \int_{X_-(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (1 - \omega_x(h)) dx && \text{(by Assumption A.3)} \\
 &= \int_{X_+(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx \\
 &\quad + \int_{X_-(h)} \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx && \text{(by equation 12)} \\
 &= \int \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx
 \end{aligned}$$

Similarly, by assumption A.4 and equation equation 13, we have

$$d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) = \int \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx$$

Thus we can bound the difference between $d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h))$ and $d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))$ as follows:

$$\begin{aligned}
 & d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) \\
 &= \int \mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx \\
 &\quad - \int \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x) \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx \\
 &= \int [\mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) - \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x)] \mathbb{P}_{\mathcal{D}_S}(X = x) \cdot (\omega_x(h) - 1) dx \\
 &= \mathbb{E}_{\mathcal{D}_S}[(\mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) - \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x)) (\omega_x(h) - 1)] && \text{(by Assumption A.5)} \\
 &> \mathbb{E}_{\mathcal{D}_S}[\mathbb{P}_{\mathcal{D}_S}(Y = +1|X = x) - \mathbb{P}_{\mathcal{D}_S}(h(x) = +1|X = x)] \mathbb{E}_{\mathcal{D}_S}[\omega_x(h) - 1] \\
 &= 0
 \end{aligned}$$

Combining the above with Theorem 3.3, we have

$$\max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \geq \frac{d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))}{2} > 0$$

□

A.9. Omitted details for Section 4.3

To make the problem tractable and meaningful, we make the following assumptions:

Assumption A.6. (Agent’s Initial Feature Distribution) Agents’ initial features are uniformly distributed between $[0, 1] \in \mathbb{R}^1$:

$$\Pr_{\mathcal{D}_S}(x) = \begin{cases} 1, & \text{if } x \in [0, 1] \\ 0, & \text{otherwise} \end{cases}$$

Assumption A.7. (Agent’s Cost Function) The cost of changing from x to x' is proportional to the distance between them: $c(x, x') = \|x - x'\|$.

Under Assumption A.7, only agents whose features are in between $[\tau_h - B, \tau_h)$ will *attempt* to change their feature. We also assume that feature updates are *probabilistic*, such that agents with features closer to the decision boundary τ_h have a greater *chance* of updating their feature and each updated feature x' is sampled from a uniform distribution depending on τ_h , B , and x (see Assumption A.8 and Assumption A.9):

Assumption A.8 (Agent’s Success Manipulation Probability). For agents who *attempt* to update their features, the probability of a successful feature update is

$$\Pr(X' \neq X) = 1 - \frac{|x - \tau_h|}{B} \quad (14)$$

Intuitively this assumption means that the closer the agent’s original feature x is to the decision boundary τ_h , the more likely they can successfully change their feature to cross the decision boundary.

Assumption A.9 (Newly Adapted Feature’s Distribution). An agent’s updated feature x' , given original feature x , manipulation budget B , and classification boundary τ_h , is sampled as

$$X' \sim \text{Unif}(\tau_h, \tau_h + B - x) \quad (15)$$

This assumption aims to capture the fact that even though agent targets to change their feature to the decision boundary τ_h (i.e. the least cost action to get a favorable prediction outcome), they might end up reaching to a feature that is beyond the decision boundary.

With Assumption A.7 - Assumption A.9, we can further specify the important weight $w_x(h)$ for the strategic response setting:

Lemma A.10. Recall the definition for the covariate shift important weight coefficient $\omega_x(h) := \frac{\mathbb{P}_{D(h)}(X=x)}{\mathbb{P}_{D_S}(X=x)}$, for our strategic response setting, we have,

$$w_x(h) = \begin{cases} 1, & x \in [0, \tau_h - B) \\ \frac{\tau_h - x}{B}, & x \in [\tau_h - B, \tau_h) \\ \frac{1}{B}(-x + \tau_h + 2B), & x \in [\tau_h, \tau_h + B) \\ 1, & x \in [\tau_h + B, 1] \end{cases} \quad (16)$$

Proof for Lemma A.10:

Proof. We discuss the induced distribution $\mathcal{D}(h)$ by cases:

- For the features distributed between $[0, \tau_h - B]$: since we assume the agents are rational, under assumption A.7, agents with feature that is smaller than $[0, \tau_h - B]$ will not perform any kinds of adaptations, and no other agents will adapt their features to this range of features either, so the distribution between $[0, \tau_h - B]$ will remain the same as before.
- For the target distribution between $[\tau_h - B, \tau_h]$ can be directly calculated from assumption A.8.
- For distribution between $[\tau_h, \tau_h + B]$, consider a particular feature $x^* \in [\tau_h, \tau_h + B]$, under Assumption A.9, we know

its new distribution becomes:

$$\begin{aligned}\Pr_{\mathcal{D}(h)}(x = x^*) &= 1 + \int_{x^*-B}^{\tau_h} \frac{1 - \frac{\tau_h - z}{B}}{B - \tau_h + z} dz \\ &= 1 + \int_{x^*-B}^{\tau_h} \frac{1}{B} dz \\ &= \frac{1}{B}(-x^* + \tau_h + 2B)\end{aligned}$$

- For the target distribution between $[\tau_h + B, 1]$: under assumption A.7 and A.9, we know that no agents will change their feature to this feature region. So the distribution between $[\tau_h + B, 1]$ remains the same as the source distribution.

Recall the definition for the covariate shift important weight coefficient $\omega_x(h) := \frac{\mathbb{P}_{\mathcal{D}(h)}(X=x)}{\mathbb{P}_{\mathcal{D}_S}(X=x)}$, the distribution of $\omega_x(h)$ after agents' strategic responding becomes:

$$\omega_x(h) = \begin{cases} 1, & x \in [0, \tau_h - B] \text{ and } x \in [\tau_h + B, 1] \\ \frac{\tau_h - x}{B}, & x \in [\tau_h - B, \tau_h] \\ \frac{1}{B}(-x + \tau_h + 2B), & x \in [\tau_h, \tau_h + B] \\ 0, & \text{otherwise} \end{cases} \quad (17)$$

□

Proof for Proposition 4.4:

Proof. According to Lemma A.10, we can compute the variance of $w_x(h)$ as $\text{Var}(w_x(h)) = \mathbb{E}(w_x(h)^2) - \mathbb{E}(w_x(h))^2 = \frac{2}{3}B$. Then by plugging it to the general bound for Theorem 4.2 gives us the desirable result. □

A.10. Proof of Theorem 5.1

Proof. Defining $p := \mathbb{P}_{\mathcal{D}_S}(Y = +1)$, we have

$$\begin{aligned}\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) &= \omega(h_S^*) \cdot \text{Err}_+(h_S^*) + (1 - \omega(h_S^*)) \cdot \text{Err}_-(h_S^*) \quad (\text{by definitions of } \omega(h_S^*), \text{Err}_+(h_S^*), \text{ and } \text{Err}_-(h_S^*)) \\ &= \underbrace{p \cdot \text{Err}_+(h_S^*) + (1 - p) \cdot \text{Err}_-(h_S^*)}_{(I)} + (\omega(h_S^*) - p)[\text{Err}_+(h_S^*) - \text{Err}_-(h_S^*)] \quad (18)\end{aligned}$$

We can expand (I) as follows:

$$\begin{aligned}& p \cdot \text{Err}_+(h_S^*) + (1 - p) \cdot \text{Err}_-(h_S^*) \\ & \leq p \cdot \text{Err}_+(h_T^*) + (1 - p) \cdot \text{Err}_-(h_T^*) \quad (\text{by optimality of } h_S^* \text{ on } \mathcal{D}_S) \\ & = \omega(h_T^*) \cdot \text{Err}_+(h_T^*) + (1 - \omega(h_T^*)) \cdot \text{Err}_-(h_T^*) + (p - \omega(h_T^*)) \cdot [\text{Err}_+(h_T^*) - \text{Err}_-(h_T^*)] \\ & = \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) + (p - \omega(h_T^*)) \cdot [\text{Err}_+(h_T^*) - \text{Err}_-(h_T^*)].\end{aligned}$$

Plugging this back into equation 18, we have

$$\text{Err}_{\mathcal{D}(h_S^*)}(h_S^*) - \text{Err}_{\mathcal{D}(h_T^*)}(h_T^*) \leq (\omega(h_S^*) - p)[\text{Err}_+(h_S^*) - \text{Err}_-(h_S^*)] + (p - \omega(h_T^*)) \cdot [\text{Err}_+(h_T^*) - \text{Err}_-(h_T^*)]$$

Notice that

$$\begin{aligned}0.5(\text{Err}_+(h) - \text{Err}_-(h)) &= 0.5 \cdot 1 - 0.5 \cdot \mathbb{P}(h(X) = +1 | Y = +1) - 0.5 \cdot \mathbb{P}(h(X) = +1 | Y = -1) \\ &= 0.5 - \mathbb{P}_{\mathcal{D}_u}(h(X) = +1)\end{aligned}$$

where \mathcal{D}_u is a distribution with uniform prior. Then

$$\begin{aligned}(\omega(h_S^*) - p)[\text{Err}_+(h_S^*) - \text{Err}_-(h_S^*)] &= 2(\omega(h_S^*) - p) \cdot (0.5 - \mathbb{P}_{\mathcal{D}_u}(h(X) = +1)) \\ (p - \omega(h_T^*))[\text{Err}_+(h_T^*) - \text{Err}_-(h_T^*)] &= 2(p - \omega(h_T^*)) \cdot (0.5 - \mathbb{P}_{\mathcal{D}_u}(h(X) = +1))\end{aligned}$$

Adding together these two equations yields

$$\begin{aligned}
 & (\omega(h_S^*) - p)[\text{Err}_+(h_S^*) - \text{Err}_-(h_S^*)] + (p - \omega(h_T^*)) \cdot [\text{Err}_+(h_T^*) - \text{Err}_-(h_T^*)] \\
 &= 2(\omega(h_S^*) - p) \cdot (0.5 - \mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1)) + 2(p - \omega(h_T^*)) \cdot (0.5 - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)) \\
 &= (\omega(h_S^*) - \omega(h_T^*)) - 2(\omega(h_S^*)\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \omega(h_T^*)\mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)) \\
 &\quad + 2p \cdot (\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)) \\
 &\leq |\omega(h_S^*) - \omega(h_T^*)| \cdot (1 + 2|\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)|) \\
 &\quad + 2p \cdot |\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)|
 \end{aligned} \tag{19}$$

Meanwhile,

$$\begin{aligned}
 & |\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)| \\
 &\leq 0.5 \cdot |\mathbb{P}_{\mathcal{D}|Y=+1}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}|Y=+1}(h_T^*(X) = +1)| \\
 &\quad + 0.5 \cdot |\mathbb{P}_{\mathcal{D}|Y=-1}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}|Y=-1}(h_T^*(X) = +1)| \\
 &= 0.5 (d_{\text{TV}}(\mathcal{D}_+(h_S^*), \mathcal{D}_+(h_T^*)) + d_{\text{TV}}(\mathcal{D}_-(h_S^*), \mathcal{D}_-(h_T^*)))
 \end{aligned} \tag{20}$$

Combining equation 19 and equation 20 gives

$$\begin{aligned}
 & |\omega(h_S^*) - \omega(h_T^*)| \cdot (1 + 2 \cdot |\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)|) \\
 &\quad + 2p \cdot |\mathbb{P}_{\mathcal{D}_u}(h_S^*(X) = +1) - \mathbb{P}_{\mathcal{D}_u}(h_T^*(X) = +1)| \\
 &\leq |\omega(h_S^*) - \omega(h_T^*)| \cdot (1 + d_{\text{TV}}(\mathcal{D}_+(h_S^*), \mathcal{D}_+(h_T^*)) + d_{\text{TV}}(\mathcal{D}_-(h_S^*), \mathcal{D}_-(h_T^*))) \\
 &\quad + p \cdot (d_{\text{TV}}(\mathcal{D}_+(h_S^*), \mathcal{D}_+(h_T^*)) + d_{\text{TV}}(\mathcal{D}_-(h_S^*), \mathcal{D}_-(h_T^*))) \\
 &\leq |\omega(h_S^*) - \omega(h_T^*)| + (1 + p) \cdot (d_{\text{TV}}(\mathcal{D}_+(h_S^*), \mathcal{D}_+(h_T^*)) + d_{\text{TV}}(\mathcal{D}_-(h_S^*), \mathcal{D}_-(h_T^*))) .
 \end{aligned}$$

□

A.11. Proof of Theorem 5.2

We will make use of the following fact:

Lemma A.11. *Under label shift, $\text{TPR}_S(h) = \text{TPR}_h(h)$ and $\text{FPR}_S(h) = \text{FPR}_h(h)$.*

Proof. We have

$$\begin{aligned}
 \text{TPR}_h(h) &= \mathbb{P}_{\mathcal{D}(h)}(h(X) = +1 | Y = +1) \\
 &= \int \mathbb{P}_{\mathcal{D}(h)}(h(X) = +1, X = x | Y = +1) dx \\
 &= \int \mathbb{P}_{\mathcal{D}(h)}(h(X) = +1 | X = x, Y = +1) \mathbb{P}_{\mathcal{D}(h)}(X = x | Y = +1) dx \\
 &= \int \mathbf{1}(h(x) = +1) \mathbb{P}_{\mathcal{D}(h)}(X = x | Y = +1) dx \\
 &= \int \mathbf{1}(h(x) = +1) \mathbb{P}_{\mathcal{D}_S}(X = x | Y = +1) dx && \text{(by definition of label shift)} \\
 &= \int \mathbb{P}_{\mathcal{D}_S}(h(X) = +1 | X = x, Y = +1) \mathbb{P}_{\mathcal{D}_S}(X = x | Y = +1) dx \\
 &= \text{TPR}_S(h)
 \end{aligned}$$

The argument for $\text{TPR}_h(h) = \text{TPR}_S(h)$ is analogous. □

Now we proceed to prove the theorem.

Proof of Theorem 5.2. In section 3.2 we showed a general lower bound on the maximum of $\text{Err}_{\mathcal{D}_S}(h)$ and $\text{Err}_{\mathcal{D}(h)}(h)$:

$$\max\{\text{Err}_{\mathcal{D}_S}(h), \text{Err}_{\mathcal{D}(h)}(h)\} \geq \frac{d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h))}{2}$$

In the case of label shift, and by the definitions of p and $\omega(h)$,

$$d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) = |\mathbb{P}_{\mathcal{D}_S}(Y = +1) - \mathbb{P}_{\mathcal{D}(h)}(Y = +1)| = |p - \omega(h)| \quad (21)$$

In addition, we have

$$\mathcal{D}_{h|S} = \mathbb{P}_S(h(X) = +1) = p \cdot \text{TPR}_S(h) + (1 - p) \cdot \text{FPR}_S(h) \quad (22)$$

Similarly

$$\begin{aligned} \mathcal{D}_h(h) &= \mathbb{P}_{\mathcal{D}(h)}(h(X) = +1) \\ &= \omega(h) \cdot \text{TPR}_h(h) + (1 - \omega(h)) \cdot \text{FPR}_h(h) \\ &= \omega(h) \cdot \text{TPR}_S(h) + (1 - \omega(h)) \cdot \text{FPR}_S(h) \end{aligned} \quad (\text{by Lemma A.11}) \quad (23)$$

Therefore

$$\begin{aligned} d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) &= |\mathbb{P}_{\mathcal{D}_S}(h(X) = +1) - \mathbb{P}_{\mathcal{D}(h)}(h(X) = +1)| \\ &= |(p - \omega(h)) \cdot \text{TPR}_S(h) + (\omega(h) - p) \cdot \text{FPR}_S(h)| \quad (\text{By equation 23 and equation 22}) \\ &= |p - \omega(h)| \cdot |\text{TPR}_S(h) - \text{FPR}_S(h)| \quad (24) \end{aligned}$$

which yields:

$$d_{\text{TV}}(\mathcal{D}_{Y|S}, \mathcal{D}_Y(h)) - d_{\text{TV}}(\mathcal{D}_{h|S}, \mathcal{D}_h(h)) = |p - \omega(h)|(1 - |\text{TPR}_S(h) - \text{FPR}_S(h)|) \quad (\text{By equation 21 and equation 24})$$

completing the proof. \square

A.12. Proof of Proposition 5.3

Proof.

$$\begin{aligned} &|\omega(h_S^*) - \omega(h_T^*)| \cdot \frac{1}{\mathbb{P}_{\mathcal{D}_S}(Y = +1)} \\ &= \frac{|(1 - \text{Err}_{\mathcal{D}_S}(h_S^*))\text{TPR}_S(h_S^*) - (1 - \text{Err}_{\mathcal{D}_S}(h_T^*))\text{TPR}_S(h_T^*)|}{(1 - \text{Err}_{\mathcal{D}_S}(h_S^*)) \cdot (1 - \text{Err}_{\mathcal{D}_S}(h_T^*))} \\ &\leq \frac{|\text{Err}_{\mathcal{D}_S}(h_S^*) - \text{Err}_{\mathcal{D}_S}(h_T^*)| \cdot |\text{TPR}_S(h_S^*) - \text{TPR}_S(h_T^*)|}{(1 - \text{Err}_{\mathcal{D}_S}(h_S^*)) \cdot (1 - \text{Err}_{\mathcal{D}_S}(h_T^*))} \end{aligned} \quad (25)$$

The inequality above is due to Lemma 7 of (Liu & Liu, 2015). \square

B. Missing Experimental Details

B.1. Synthetic Experiments Using DAG

Covariate Shift We specify the causal DAG for covariate shift setting in the following way:

$$\begin{aligned} X_1 &\sim \text{Unif}(-1, 1) \\ X_2 &\sim 1.2X_1 + \mathcal{N}(0, \sigma_2^2) \\ X_3 &\sim -X_1^2 + \mathcal{N}(0, \sigma_3^2) \\ Y &:= 2\text{sign}(X_2 > 0) - 1 \end{aligned}$$

where σ_2^2 and σ_3^2 are parameters of our choices.

Adaptation function We assume the new distribution of feature X'_1 will be generated in the following way:

$$X'_1 = \Delta(X) = X_1 + c \cdot (h(X) - 1)$$

where $c \in \mathbb{R}^1 > 0$ is the parameter controlling how much the prediction $h(X)$ affect the generating of X'_1 , namely the magnitude of distribution shift. Intuitively, this adaptation function means that if a feature x is predicted to be positive ($h(x) = +1$), then decision subjects are more likely to adapt to that feature in the induced distribution; Otherwise, decision subjects are more likely to be moving away from x since they know it will lead to a negative prediction.

Target Shift We specify the causal DAG for target shift setting in the following way:

$$\begin{aligned} (Y + 1)/2 &\sim \text{Bernoulli}(\alpha) \\ X_1|Y = y &\sim \mathcal{N}_{[0,1]}(\mu_y, \sigma^2) \\ X_2 &= -0.8X_1 + \mathcal{N}(0, \sigma_2^2) \\ X_3 &= 0.2Y + \mathcal{N}(0, \sigma_3^2) \end{aligned}$$

where $\mathcal{N}_{[0,1]}$ represents a truncated Gaussian distribution taken value between 0 and 1. $\alpha, \mu_y, \sigma^2, \sigma_2^2$ and σ_3^2 are parameters of our choices.

Adaptation function We assume the new distribution of the qualification Y' will be updated in the following way:

$$\mathbb{P}(Y' = +1|h(X) = h, Y = y) = c_{hy}, \text{ where } \{h, y\} \in \{-1, +1\}$$

where $0 \leq c_{hy} \in \mathbb{R}^1 \leq 1$ represents the likelihood for a person with original qualification $Y = y$ and get predicted as $h(X) = h$ to be qualified in the next step ($Y' = +1$).

B.2. Synthetic Experiments Using Real-world Data

On the preprocessed FICO credit score data set (Board of Governors of the Federal Reserve System (US), 2007; Hardt et al., 2016b), we convert the cumulative distribution function (CDF) of TransRisk score among demographic groups (denoted as A , including Black, Asian, Hispanic, and White) into group-dependent densities of the credit score. We then generate a balanced sample where each group has equal representation, with credit scores (denoted as Q) initialized by sampling from the corresponding group-dependent density. The value of attributes for each data point is then updated under a specified dynamics (as detailed below) to model the real-world scenario of repeated resource allocation (with decision denoted as D). Since we are considering the dynamic setting, we further specify the data generating process in the following way (from time step $T = t$ to $T = t + 1$):

$$\begin{aligned} X_{t,1} &\sim 1.5Q_t + U[-\epsilon_1, \epsilon_1] \\ X_{t,2} &\sim 0.8A_t + U[-\epsilon_2, \epsilon_2] \\ X_{t,3} &\sim A_t + \mathcal{N}(0, \sigma^2) \\ Y_t &\sim \text{Bernoulli}(q_t) \text{ for a given value of } Q_t = q_t \\ D_t &= f_t(A_t, X_{t,1}, X_{t,2}, X_{t,3}) \\ Q_{t+1} &= \{Q_t \cdot [1 + \alpha_D(D_t) + \alpha_Y(Y_t)]\}_{(0,1]} \\ A_{t+1} &= A_t \text{ (fixed population)} \end{aligned}$$

where $\{\cdot\}_{(0,1]}$ represents truncated value between the interval $(0, 1]$, $f_t(\cdot)$ represents the decision policy from input features, and $\epsilon_1, \epsilon_2, \sigma$ are parameters of our choices.

Within the same time step, i.e., for variables that share the subscript t , Q_t and A_t are root causes for all other variables ($X_{t,1}, X_{t,2}, X_{t,3}, D_t, Y_t$). For different time steps, e.g., from $T = t$ to $T = t + 1$, the new distribution at $T = t + 1$ is induced by the deployment of the decision policy D_t . Such impact is modeled by a multiplicative update in Q_{t+1} from Q_t with parameters (or functions) $\alpha_D(\cdot)$ and $\alpha_Y(\cdot)$ that depend on D_t and Y_t , respectively.