
On the Value of Target Data in Transfer Learning

Steve Hanneke

Toyota Technological Institute at Chicago
steve.hanneke@gmail.com

Samory Kpotufe

Columbia University, Statistics
skk2175@columbia.edu

Abstract

We aim to understand the value of additional labeled or unlabeled target data in transfer learning, for any given amount of source data; this is motivated by practical questions around minimizing sampling costs, whereby, target data is usually harder or costlier to acquire than source data, but can yield better accuracy.

To this aim, we establish the first minimax-rates in terms of both source and target sample sizes, and show that performance limits are captured by new notions of discrepancy between source and target, which we refer to as *transfer exponents*.

Interestingly, we find that attaining minimax performance is akin to ignoring one of the source or target samples, provided distributional parameters were known a priori. Moreover, we show that practical decisions – w.r.t. minimizing sampling costs – can be made in a minimax-optimal way *without* knowledge or estimation of distributional parameters nor of the discrepancy between source and target.

1 Introduction

The practice of transfer-learning often involves acquiring some amount of target data, and involves various practical decisions as to how to best combine source and target data; however much of the theoretical literature on transfer only addresses the setting where no target labeled data is available.

We aim to understand the value of target labels, that is, given n_P labeled data from some source distribution P , and n_Q labeled target labels from a target Q , what is the best Q error achievable by any classifier in terms of *both* n_Q and n_P , and which classifiers achieve such optimal transfer. In this first analysis, we mostly restrict ourselves to a setting, similar to the traditional *covariate-shift* assumption, where the best classifier – from a fixed VC class \mathcal{H} – is the same under P and Q .

We establish the first minimax-rates, for bounded-VC classes, in terms of both source and target sample sizes n_P and n_Q , and show that performance limits are captured by new notions of discrepancy between source and target, which we refer to as *transfer exponents*.

The first notion of transfer-exponent, called ρ , is defined in terms of discrepancies in excess risk, and is most refined. Already here, our analysis reveals a surprising fact: the best possible rate (matching upper and lower-bounds) in terms of ρ and both sample sizes n_P, n_Q is - up to constants - achievable by an oracle which simply ignores the least informative of the source or target datasets. In other words, if \hat{h}_P and \hat{h}_Q denote the ERM on data from P , resp. from Q , one of the two achieves the optimal Q rate over any classifier having access to both P and Q datasets. However, which of \hat{h}_P or \hat{h}_Q is optimal is not easily decided without prior knowledge: for instance, cross-validating on a holdout target-sample would naively result in a rate of $n_Q^{-1/2}$ which can be far from optimal given large n_P . Interestingly, we show that the optimal (n_P, n_Q) -rate is achieved by a generic approach, akin to so-called *hypothesis-transfer* [1, 2], which optimizes Q -error under the constraint of low P -error, and does so without knowledge of distributional parameters such as ρ .

We then consider a related notion of *marginal* transfer-exponent, called γ , defined w.r.t. marginals P_X, Q_X . This is motivated by the fact that practical decisions in transfer often involve the use of

cheaper unlabeled data (i.e., data drawn from P_X, Q_X). We will show that, when practical decisions are driven by observed changes in marginals P_X, Q_X , the marginal notion γ is then most suited to capture performance as it does not require knowledge (or observations) of label distribution $Q_{Y|X}$.

In particular, the marginal exponent γ helps capture performance limits in the following scenarios of current practical interest:

- **Minimizing sampling cost.** Given different costs of labeled source and target data, and a desired target excess error at most ϵ , how to use unlabeled data to decide on an optimal sampling scheme that minimizes labeling costs while achieving target error at most ϵ . (Section 6)
- **Choice of transfer.** Given two sources P_1 and P_2 , each at some unknown distance from Q , given unlabeled data and some or no labeled data from Q , how to decide which of P_1, P_2 transfers best to the target Q . (Appendix A.2)
- **Reweighting.** Given some amount of unlabeled data from Q , and some or no labeled Q data, how to optimally re-weight (out of a fixed set of schemes) the source P data towards best target performance. While differently motivated, this problem is related to the last one. (Appendix A.1)

Although optimal decisions in the above scenarios depend tightly on unknown distributional parameters such as different label noise in source and target data, and on unknown *distance* from source to target (as captured by γ), we show that such practical decisions can be made, near optimally, with no knowledge of distributional parameters, and perhaps surprisingly, without ever estimating γ . Furthermore, the unlabeled sampling complexity can be shown to remain low. Finally, the procedures described in this work remain of a theoretical nature, but yield new insights into how various practical decisions in transfer can be made near-optimally in a data-driven fashion.

Related Work. Much of the theoretical literature on transfer can be subdivided into a few main lines of work. As mentioned above, the main distinction with the present work is in that they mostly focus on situations with no labeled target data, and consider distinct notions of discrepancy between P and Q . We contrast these various notions with the transfer-exponents ρ and γ in Section 3.1.

A first direction considers refinements of total-variation that quantify changes in error over classifiers in a fixed class \mathcal{H} . The most common such measures are the so-called $d_{\mathcal{A}}$ -divergence [3, 4, 5] and the \mathcal{Y} -discrepancy [6, 7, 8]. In this line of work, the rates of transfer, largely expressed in terms of n_P alone, take the form $o_p(1) + C \cdot \text{divergence}(P, Q)$. In other words, transfer down to 0 error seems impossible whenever these divergences are non-negligible; we will carefully argue that such intuition can be overly pessimistic.

Another prominent line of work, which has led to many practical procedures, considers so-called density ratios f_Q/f_P (importance weights) as a way to capture the similarity between P and Q [9, 10]. A related line of work considers information-theoretic measures such as KL-divergence or Renyi divergence [11, 12] but has received relatively less attention. Similar to these notions, the transfer-exponents ρ and γ are *asymmetric* measures of distance, attesting to the fact that it could be easier to transfer from some P to Q than the other way around. However, a significant downside to these notions is that they do not account for the specific structure of a hypothesis class \mathcal{H} as is the case with the aforementioned divergences. As a result, they can be sensitive to issues such as minor differences of support in P and Q , which may be irrelevant when learning with certain classes \mathcal{H} .

On the algorithmic side, many approaches assign importance weights to source data from P so as to minimize some prescribed *metric* between P and Q [13, 14]; as we will argue, *metrics*, being symmetric, can be inadequate as a measure of discrepancy given the inherent asymmetry in transfer.

The importance of unlabeled data in transfer-learning, given the cost of target labels, has always been recognized, with various approaches developed over the years [15, 16], including more recent research efforts into so-called *semisupervised* or *active* transfer, where, given unlabeled target data, the goal is to request as few target labels as possible to improve classification over using source data alone [17, 18, 19, 20, 21].

More recently, [22, 23, 24] consider nonparametric transfer settings (unbounded VC) allowing for changes in conditional distributions. Also recent, but more closely related, [25] proposed a nonparametric measure of discrepancy which successfully captures the interaction between labeled source and target under nonparametric conditions and 0-1 loss; these notions however ignore the additional structure afforded by transfer in the context of a fixed hypothesis class.

2 Setup and Definitions

We consider a classification setting where the input $X \in \mathcal{X}$, some measurable space, and the output $Y \in \{0, 1\}$. We let $\mathcal{H} \subset 2^{\mathcal{X}}$ denote a fixed hypothesis class over \mathcal{X} , denote $d_{\mathcal{H}}$ the VC dimension [26], and the goal is to return a classifier $h \in \mathcal{H}$ with low error $R_Q(h) \doteq \mathbb{E}_Q[h(X) \neq Y]$ under some joint distribution Q on X, Y . The learner has access to two independent labeled samples $S_P \sim P^{n_P}$ and $S_Q \sim Q^{n_Q}$, i.e., drawn from *source* distributions P and target Q , of respective sizes n_P, n_Q .

Our aim is to bound the excess error, under Q , of any \hat{h} learned from both samples, in terms of n_P, n_Q , and (suitable) notions of discrepancy between P and Q . We will let $P_X, Q_X, P_{Y|X}, Q_{Y|X}$ denote the corresponding marginal and conditional distributions under P and Q .

Definition 1. For $D \in \{Q, P\}$, denote $\mathcal{E}_D(h) \doteq R_D(h) - \inf_{h' \in \mathcal{H}} R_D(h')$, the **excess error** of h .

Distributional Conditions. We consider various traditional assumptions in classification and transfer. The first one is a so-called *Bernstein Class Condition* on noise [27, 28, 29, 30, 31].

(NC). Let $h_P^* \doteq \operatorname{argmin}_{h \in \mathcal{H}} R_P(h)$ and $h_Q^* \doteq \operatorname{argmin}_{h \in \mathcal{H}} R_Q(h)$ exist. $\exists \beta_P, \beta_Q \in [0, 1], c_P, c_Q > 0$ s.t.

$$P_X(h \neq h_P^*) \leq c_P \cdot \mathcal{E}_P^{\beta_P}(h), \quad \text{and} \quad Q_X(h \neq h_Q^*) \leq c_Q \cdot \mathcal{E}_Q^{\beta_Q}(h). \quad (1)$$

For instance, the usual Tsybakov noise condition, say on P , corresponds to the case where h_P^* is the Bayes classifier, with corresponding regression function $\eta_P(x) \doteq \mathbb{E}[Y|x]$ satisfying $P_X(|\eta_P(X) - 1/2| \leq \tau) \leq C\tau^{\beta_P/(1-\beta_P)}$. Classification is easiest w.r.t. P (or Q) when β_P (resp. β_Q) is largest. We will see that this is also the case in Transfer.

The next assumption is stronger, but can be viewed as a relaxed version of the usual *Covariate-Shift* assumption which states that $P_{Y|X} = Q_{Y|X}$.

(RCS). Let h_P^*, h_Q^* as defined above. We have $\mathcal{E}_Q(h_P^*) = \mathcal{E}_Q(h_Q^*) = 0$. We then define $h^* \doteq h_P^*$.

Note that the above allows $P_{Y|X} \neq Q_{Y|X}$. However, it is not strictly weaker than *Covariate-Shift*, since the latter allows $h_P^* \neq h_Q^*$ provided the Bayes $\notin \mathcal{H}$. The assumption is useful as it serves to isolate the sources of hardness in transfer beyond just shifts in h^* . We will in fact see later that it is easily removed, but at the additive (necessary) cost of $\mathcal{E}_Q(h_P^*)$.

3 Transfer-Exponents from P to Q .

We consider various notions of *discrepancy* between P and Q , which will be shown to tightly capture the complexity of transfer P to Q .

Definition 2. We call $\rho > 0$ a **transfer-exponent** from P to Q , w.r.t. \mathcal{H} , if there exists C_ρ such that

$$\forall h \in \mathcal{H}, \quad C_\rho \cdot \mathcal{E}_P(h) \geq \mathcal{E}_Q^\rho(h). \quad (2)$$

We are interested in the smallest such ρ with small C_ρ . We generally would think of ρ as at least 1, although there are situations – which we refer to as *super-transfer*, to be discussed, where we have $\rho < 1$; in such situations, data from P can yield faster \mathcal{E}_Q rates than data from Q .

While the transfer-exponent will be seen to tightly capture the two-samples minimax rates of transfer, and can be *adapted to*, practical learning situations call for *marginal* versions that can capture the rates achievable when one has access to unlabeled Q data.

Definition 3. We call $\gamma > 0$ a **marginal transfer-exponent** from P to Q if $\exists C_\gamma$ such that

$$\forall h \in \mathcal{H}, \quad C_\gamma \cdot P_X(h \neq h_P^*) \geq Q_X^\gamma(h \neq h_P^*). \quad (3)$$

The following simple proposition relates γ to ρ .

Proposition 1 (From γ to ρ). *Suppose Assumptions (NC) and (RCS) hold, and that P has marginal transfer-exponent (γ, C_γ) w.r.t. Q . Then P has transfer-exponent $\rho \leq \gamma/\beta_P$, where $C_\rho = C_\gamma^{\gamma/\beta_P}$.*


Proof. $\forall h \in \mathcal{H}$, we have $\mathcal{E}_Q(h) \leq Q_X(h \neq h_P^*) \leq C_\gamma \cdot P_X(h \neq h_P^*)^{1/\gamma} \leq C_\gamma \cdot \mathcal{E}_P(h)^{\beta_P/\gamma}$. \square

3.1 Examples and Relation to other notions of discrepancy.

In this section, we consider various examples that highlight interesting aspects of ρ and γ , and their relations to other notions of distance $P \rightarrow Q$ considered in the literature. Though our results cover noisy cases, in all these examples we assume no noise for simplicity, and therefore $\gamma = \rho$.


Example 1. (Non-overlapping supports) This first example emphasizes the fact that, unlike in much of previous analyses of transfer, the exponents γ, ρ do not require that Q_X and P_X have overlapping support. This is a welcome property shared also by the d_A and \mathcal{Y} discrepancy.

In the example shown on the right, \mathcal{H} is the class of homogeneous linear separators, while P_X and Q_X are uniform on the surface of the spheres depicted (e.g., corresponding to different scalings of the data). We then have that $\gamma = \rho = 1$ with $C_\gamma = 1$, while notions such as *density-ratios*, KL-divergences, or the recent nonparameteric notion of [25], are ill-defined or diverge to ∞ .



Example 2. (Large d_A, d_Y) Let \mathcal{H} be the class of one-sided thresholds on the line, but now we let $P_X \doteq \mathcal{U}[0, 2]$ and $Q_X \doteq \mathcal{U}[0, 1]$. Let h^* be thresholded at $1/2$. We then see that for all h_t thresholded at $t \in [0, 1]$, $2P_X(h_t \neq h^*) = \frac{1}{2}Q_X(h_t \neq h^*)$, where for $t > 1$, $P_X(h_t \neq h^*) = \frac{1}{2}(t - 1/2) \geq \frac{1}{2}Q_X(h_t \neq h^*) = \frac{1}{4}$. Thus, the marginal transfer exponent $\gamma = 1$ with $C_\gamma = 2$, so we have fast transfer at the same rate $1/n_P$ as if we were sampling from Q (Theorem 3).

On the other hand, recall that the d_A -divergence takes the form $d_A(P, Q) \doteq \sup_{h \in \mathcal{H}} |P_X(h \neq h^*) - Q_X(h \neq h^*)|$, while the \mathcal{Y} -discrepancy takes the form $d_Y(P, Q) \doteq \sup_{h \in \mathcal{H}} |\mathcal{E}_P(h) - \mathcal{E}_Q(h)|$. The two coincide whenever there is no noise in Y .




Now, take h_t as the threshold at $t = 1/2$, and $d_A = d_Y = \frac{1}{4}$ which would wrongly imply that transfer is not feasible at a rate faster than $\frac{1}{4}$; we can in fact make this situation worse, i.e., let $d_A = d_Y \rightarrow \frac{1}{2}$ by letting h^* correspond to a threshold close to 0. A first issue is that these divergences get large in large disagreement regions; this is somewhat mitigated by *localization*, as discussed in Example 4.

Example 3. (Minimum γ, ρ , and the inherent **asymmetry** of transfer) Suppose \mathcal{H} is the class of one-sided thresholds on the line, $h^* = h_P^* = h_Q^*$ is a threshold at 0. The marginal Q_X has uniform density f_Q (on an interval containing 0), while, for some $\gamma \geq 1$, P_X has density $f_P(t) \propto t^{\gamma-1}$ on $t > 0$ (and uniform on the rest of the support of Q , not shown). Consider any h_t at threshold $t > 0$, we have $P_X(h_t \neq h^*) = \int_0^t f_P \propto t^\gamma$, while $Q_X(h_t \neq h^*) \propto t$. Notice that for any fixed $\epsilon > 0$,

$$\lim_{t>0, t \rightarrow 0} \frac{Q_X(h_t \neq h^*)^{\gamma-\epsilon}}{P_X(h_t \neq h^*)} = \lim_{t>0, t \rightarrow 0} C \frac{t^{\gamma-\epsilon}}{t^\gamma} = \infty.$$

We therefore see that γ is the smallest possible marginal transfer-exponent (similarly, $\rho = \gamma$ is the smallest possible transfer exponent). Interestingly, now consider transferring instead from Q to P : we would have $\gamma(Q \rightarrow P) = 1 \leq \gamma \doteq \gamma(P \rightarrow Q)$, i.e., it could be easier to transfer from Q to P than from P to Q , which is not captured by symmetric notions of distance (d_A , Wassertein, etc ...).



Finally note that the above example can be extended to more general hypothesis classes as it simply plays on how fast f_P decreases w.r.t. f_Q in regions of space.

Example 4. (*Super-transfer* and localization). We continue on the above Example 2. Now let $0 < \gamma < 1$, and let $f_P(t) \propto |t|^{\gamma-1}$ on $[-1, 1] \setminus \{0\}$, with $Q_X \doteq \mathcal{U}[-1, 1]$, h^* at 0. As before, γ is a transfer-exponent $P \rightarrow Q$, and following from Theorem 3, we attain transfer rates of $\mathcal{E}_Q \lesssim n_P^{-1/\gamma}$, faster than the rates of n_Q^{-1} attainable with data from Q . We call these situations *super-transfer*, i.e., ones where the source data gets us faster to h^* ; here P concentrates more mass close to h^* , while more generally, such situations can also be constructed by letting $P_{Y|X}$ be less noisy than $Q_{Y|X}$ data, for instance corresponding to controlled lab data as source, vs noisy real-world data.

Now consider the following ϵ -localization fix to the $d_A = d_Y$ divergences over h 's with small P error (assuming we only observe data from P): $d_Y^* \doteq \sup_{h \in \mathcal{H}: \mathcal{E}_P(h) \leq \epsilon} |\mathcal{E}_P(h) - \mathcal{E}_Q(h)|$. This is no longer worst-case over all h 's, yet it is still not a complete fix. To see why, consider that, given n_P data from P , the best P -excess risk attainable is n_P^{-1} so we might set $\epsilon \propto n_P^{-1}$. Now the subclass $\{h \in \mathcal{H} : \mathcal{E}_P(h) \leq \epsilon\}$ corresponds to thresholds $t \in [\pm n_P^{-1/\gamma}]$, since $\mathcal{E}_P(h_t) = P([0, t]) \propto |t|^\gamma$.

We therefore have $d_{\mathcal{Y}}^* \propto \left| n_P^{-1} - n_P^{-1/\gamma} \right| \propto n_P^{-1}$, wrongly suggesting a transfer rate $\mathcal{E}_Q \lesssim n_P^{-1}$, while the super-transfer rate $n_P^{-1/\gamma}$ is achievable as discussed above. The problem is that, even after localization, $d_{\mathcal{Y}}^*$ treats errors under P and Q symmetrically.

4 Lower-Bounds

Definition 4 ((NC) Class). Let $\mathcal{F}_{(\text{NC})}(\rho, \beta_P, \beta_Q, C)$ denote the class of pairs of distributions (P, Q) with transfer-exponent ρ , $C_\rho \leq C$, satisfying (NC) with parameters β_P, β_Q , and $c_P, c_Q \leq C$.

The following lower-bound in terms of ρ is obtained via information theoretic-arguments. In effect, given the VC class \mathcal{H} , we construct a set of distribution pairs $\{(P_i, Q_i)\}$ supported on $d_{\mathcal{H}}$ datapoints, which all belong to $\mathcal{F}_{(\text{NC})}(\rho, \beta_P, \beta_Q, C)$. All the distributions share the same marginals P_X, Q_X . Any two pairs are close to each other in the sense that Π_i, Π_j , where $\Pi_i \doteq P_i^{n_P} \times Q_i^{n_Q}$, are close in KL-divergence, while, however maintaining pairs $(P_i, Q_i), (P_j, Q_j)$ far in a pseudo-distance induced by Q_X . All the proofs from this section are in Appendix B.

Theorem 1 (ρ Lower-bound). *Suppose the hypothesis class \mathcal{H} has VC dimension $d_{\mathcal{H}} \geq 9$. Let $\hat{h} = \hat{h}(S_P, S_Q)$ denote any (possibly improper) classifier with access to two independent labeled samples $S_P \sim P^{n_P}$ and $S_Q \sim Q^{n_Q}$. Fix any $\rho \geq 1$, $0 \leq \beta_P, \beta_Q < 1$. Suppose either n_P or n_Q is sufficiently large so that*

$$\epsilon(n_P, n_Q) \doteq \min \left\{ \left(\frac{d_{\mathcal{H}}}{n_P} \right)^{1/(2-\beta_P)\rho}, \left(\frac{d_{\mathcal{H}}}{n_Q} \right)^{1/(2-\beta_Q)} \right\} \leq 1/2.$$

Then, for any \hat{h} , there exists $(P, Q) \in \mathcal{F}_{(\text{NC})}(\rho, \beta_P, \beta_Q, 1)$, and a universal constant c such that,

$$\mathbb{P}_{S_P, S_Q} \left(\mathcal{E}_Q(\hat{h}) > c \cdot \epsilon(n_P, n_Q) \right) \geq \frac{3 - 2\sqrt{2}}{8}.$$

As per Proposition 1 we can translate any upper-bound in terms of ρ to an upper-bound in terms of γ since $\rho \leq \gamma/\beta_P$. We investigate whether such upper-bounds in terms of γ are tight, i.e., given a class $\mathcal{F}_{(\text{NC})}(\rho, \beta_P, \beta_Q, C)$, are there distributions with $\rho = \gamma/\beta_P$ where the rate is realized.

The proof of the next result is similar to that of Theorem 1, however with the added difficulty that we need the construction to yield two forms of rates $\epsilon_1(n_P, n_Q), \epsilon_2(n_P, n_Q)$ over the data support (again $d_{\mathcal{H}}$ points). Combining these two rates matches the desired upper-bound. In effect, we follow the intuition that, to have $\rho = \gamma/\beta_P$ achieved on some subset $\mathcal{X}_1 \subset \mathcal{X}$, we need β_Q to behave as 1 locally on \mathcal{X}_1 , while matching the rate requires larger β_Q on the rest of the support (on $\mathcal{X} \setminus \mathcal{X}_1$).

Theorem 2 (γ Lower-bound). *Suppose the hypothesis class \mathcal{H} has VC dimension $d_{\mathcal{H}}, \lfloor d_{\mathcal{H}}/2 \rfloor \geq 9$. Let $\hat{h} = \hat{h}(S_P, S_Q)$ denote any (possibly improper) classifier with access to two independent labeled samples $S_P \sim P^{n_P}$ and $S_Q \sim Q^{n_Q}$. Fix any $0 < \beta_P, \beta_Q < 1$, $\rho \geq \max\{1/\beta_P, 1/\beta_Q\}$. Suppose either n_P or n_Q is sufficiently large so that*

$$\begin{aligned} \epsilon_1(n_P, n_Q) &\doteq \min \left\{ \left(\frac{d_{\mathcal{H}}}{n_P} \right)^{1/(2-\beta_P)\rho\beta_Q}, \left(\frac{d_{\mathcal{H}}}{n_Q} \right)^{1/(2-\beta_Q)} \right\} \leq 1/2, \text{ and} \\ \epsilon_2(n_P, n_Q) &\doteq \min \left\{ \left(\frac{d_{\mathcal{H}}}{n_P} \right)^{1/(2-\beta_P)\rho}, \left(\frac{d_{\mathcal{H}}}{n_Q} \right) \right\} \leq 1/2. \end{aligned}$$

Then, for any \hat{h} , there exists $(P, Q) \in \mathcal{F}_{(\text{NC})}(\rho, \beta_P, \beta_Q, 2)$, with **marginal-transfer-exponent** $\gamma = \rho \cdot \beta_P \geq 1$, with $C_\gamma \leq 2$, and a universal constant c such that,

$$\mathbb{E}_{S_P, S_Q} \mathcal{E}_Q(\hat{h}) \geq c \cdot \max\{\epsilon_1(n_P, n_Q), \epsilon_2(n_P, n_Q)\}.$$

Remark 1 (Tightness with upper-bound). *Write $\epsilon_1(n_P, n_Q) = \min\{\epsilon_1(n_P), \epsilon_1(n_Q)\}$, and similarly, $\epsilon_2(n_P, n_Q) = \min\{\epsilon_2(n_P), \epsilon_2(n_Q)\}$. Define $\epsilon_L \doteq \max\{\epsilon_1(n_P, n_Q), \epsilon_2(n_P, n_Q)\}$ as in the above lower-bound of Theorem 2. Next, define $\epsilon_H \doteq \min\{\epsilon_2(n_P), \epsilon_1(n_Q)\}$. It turns out that the*

best upper-bound we can show (as a function of γ) is in terms of ϵ_H so defined. It is therefore natural to ask whether or when ϵ_H and ϵ_L are of the same order.

Clearly, we have $\epsilon_1(n_P) \leq \epsilon_2(n_P)$ and $\epsilon_1(n_Q) \geq \epsilon_2(n_Q)$ so that $\epsilon_L \leq \epsilon_H$ (as to be expected).

Now, if $\beta_Q = 1$, we have $\epsilon_1(n_P) = \epsilon_2(n_P)$ and $\epsilon_1(n_Q) = \epsilon_2(n_Q)$, so that $\epsilon_L = \epsilon_H$. More generally, from the above inequalities, we see that $\epsilon_L = \epsilon_H$ in the two regimes where either $\epsilon_1(n_Q) \leq \epsilon_1(n_P)$ (in which case $\epsilon_L = \epsilon_H = \epsilon_1(n_Q)$), or $\epsilon_2(n_P) \leq \epsilon_2(n_Q)$ (in which case $\epsilon_L = \epsilon_H = \epsilon_2(n_P)$).

5 Upper-Bounds

The following lemma is due to [32].

Lemma 1. Let $A_n = \frac{d_{\mathcal{H}}}{n} \log \left(\frac{\max\{n, d_{\mathcal{H}}\}}{d_{\mathcal{H}}} \right) + \frac{1}{n} \log \left(\frac{1}{\delta} \right)$. With probability at least $1 - \frac{\delta}{3}$, $\forall h, h' \in \mathcal{H}$,

$$R(h) - R(h') \leq \hat{R}(h) - \hat{R}(h') + c\sqrt{\min\{P(h \neq h'), \hat{P}(h \neq h')\}A_n} + cA_n, \quad (4)$$

and

$$\frac{1}{2}P(h \neq h') - cA_n \leq \hat{P}(h \neq h') \leq 2P(h \neq h') + cA_n, \quad (5)$$

for a universal numerical constant $c \in (0, \infty)$, where \hat{R} denotes empirical risk on n iid samples.

Now consider the following algorithm. Let S_P be a sequence of n_P samples from P and S_Q a sequence of n_Q samples from Q . Also let $\hat{h}_{S_P} = \operatorname{argmin}_{h \in \mathcal{H}} \hat{R}_{S_P}(h)$ and $\hat{h}_{S_Q} = \operatorname{argmin}_{h \in \mathcal{H}} \hat{R}_{S_Q}(h)$. Choose \hat{h} as the solution to the following optimization problem.

Algorithm 1:

$$\begin{aligned} & \text{Minimize} && \hat{R}_{S_P}(h) \\ & \text{subject to} && \hat{R}_{S_Q}(h) - \hat{R}_{S_Q}(\hat{h}_{S_Q}) \leq c\sqrt{\hat{P}_{S_Q}(h \neq \hat{h}_{S_Q})A_{n_Q}} + cA_{n_Q} \quad (6) \\ & && h \in \mathcal{H}. \end{aligned}$$

The intuition is that, effectively, the constraint guarantees we maintain a near-optimal guarantee on $\mathcal{E}_Q(\hat{h})$ in terms of n_Q and the (NC) parameters for Q , while (as we show) still allowing the algorithm to select an h with a near-minimal value of $\hat{R}_{S_P}(h)$. The latter guarantee plugs into the transfer condition to obtain a term converging in n_P , while the former provides a term converging in n_Q , and altogether the procedure achieves a rate specified by the *min* of these two guarantees (which is in fact nearly minimax *optimal*, since it matches the lower bound up to logarithmic factors).

Formally, we have the following result for this learning rule; its proof is below.

Theorem 3 (Minimax Upper-Bounds). Assume (NC). Let \hat{h} be the solution from Algorithm 1. For a constant C depending on $\rho, C_\rho, \beta_P, c_{\beta_P}, \beta_Q, c_{\beta_Q}$, with probability at least $1 - \delta$,

$$\mathcal{E}_Q(\hat{h}) \leq C \min \left\{ A_{n_P}^{\frac{1}{(2-\beta_P)\rho}}, A_{n_Q}^{\frac{1}{2-\beta_Q}} \right\} = \tilde{O} \left(\min \left\{ \left(\frac{d_{\mathcal{H}}}{n_P} \right)^{\frac{1}{(2-\beta_P)\rho}}, \left(\frac{d_{\mathcal{H}}}{n_Q} \right)^{\frac{1}{2-\beta_Q}} \right\} \right).$$

Note that, by the lower bound of Theorem 1, this bound is optimal up to log factors.

Remark 2 (Effective Source Sample Size). From the above, we might view (ignoring $d_{\mathcal{H}}$) $\tilde{n}_P \doteq n_P^{\frac{(2-\beta_Q)/(2-\beta_P)\rho}{(2-\beta_P)\rho}}$ as the effective sample size contributed by P . In fact, the above minimax rate is of order $(\tilde{n}_P + n_Q)^{-1/(2-\beta_Q)}$, which yields added intuition into the combined effect of both samples. We have that, the effective source sample size \tilde{n}_P is smallest for large ρ , but also depends on $(2 - \beta_Q)/(2 - \beta_P)$, i.e., on whether P is noisier than Q .

Remark 3 (Rate in terms of γ). Note that, by Proposition 1, this also immediately implies a bound under the marginal transfer condition and RCS, simply taking $\rho \leq \gamma/\beta_P$. Furthermore, by the lower bound of Theorem 2, the resulting bound in terms of γ is tight in certain regimes up to log factors.

Proof of Theorem 3. In all the lines below, we let C serve as a generic constant (possibly depending on $\rho, \hat{C}_\rho, \beta_P, c_{\beta_P}, \beta_Q, c_{\beta_Q}$) which may be different in different appearances. Consider the event of probability at least $1 - \delta/3$ from Lemma 1 for the S_Q samples. In particular, on this event, if $\mathcal{E}_Q(h_P^*) = 0$, it holds that

$$\hat{R}_{S_Q}(h_P^*) - \hat{R}_{S_Q}(\hat{h}_{S_Q}) \leq c\sqrt{\hat{P}_{S_Q}(h_P^* \neq \hat{h}_{S_Q})A_{n_Q}} + cA_{n_Q}.$$

This means, under the (RCS) condition, h_P^* satisfies the constraint in the above optimization problem defining \hat{h} . Also, on this same event from Lemma 1 we have

$$\mathcal{E}_Q(\hat{h}_{S_Q}) \leq c\sqrt{Q(\hat{h}_{S_Q} \neq h_Q^*)A_{n_Q}} + cA_{n_Q},$$

so that (NC) implies

$$\mathcal{E}_Q(\hat{h}_{S_Q}) \leq C\sqrt{\mathcal{E}_Q(\hat{h}_{S_Q})^{\beta_Q} A_{n_Q}} + cA_{n_Q},$$

which implies the well-known fact from [28, 29] that

$$\mathcal{E}_Q(\hat{h}_{S_Q}) \leq C \left(\frac{d_{\mathcal{H}}}{n_Q} \log \left(\frac{n_Q}{d_{\mathcal{H}}} \right) + \frac{1}{n_Q} \log \left(\frac{1}{\delta} \right) \right)^{\frac{1}{2-\beta_Q}}. \quad (7)$$

Furthermore, following the analogous argument for S_P , it follows that for any set $\mathcal{G} \subseteq \mathcal{H}$ with $h_P^* \in \mathcal{G}$, with probability at least $1 - \delta/3$, the ERM $\hat{h}'_{S_P} = \operatorname{argmin}_{h \in \mathcal{G}} \hat{R}_{S_P}(h)$ satisfies

$$\mathcal{E}_P(\hat{h}'_{S_P}) \leq C \left(\frac{d_{\mathcal{H}}}{n_P} \log \left(\frac{n_P}{d_{\mathcal{H}}} \right) + \frac{1}{n_P} \log \left(\frac{1}{\delta} \right) \right)^{\frac{1}{2-\beta_P}}. \quad (8)$$

In particular, conditioned on the S_Q data, we can take the set \mathcal{G} as the set of $h \in \mathcal{H}$ satisfying the constraint in the optimization, and on the above event we have $h_P^* \in \mathcal{G}$ (assuming the (RCS) condition); furthermore, if $\mathcal{E}_Q(h_P^*) = 0$, then without loss we can simply define $h_Q^* = h_P^* = h^*$ (and it is easy to see that this does not affect the NC condition). We thereby establish the above inequality (8) for this choice of \mathcal{G} , in which case by definition $\hat{h}'_{S_P} = \hat{h}$. Altogether, by the union bound, all of these events hold simultaneously with probability at least $1 - \delta$. In particular, on this event, if the (RCS) condition holds then

$$\mathcal{E}_P(\hat{h}) \leq C \left(\frac{d_{\mathcal{H}}}{n_P} \log \left(\frac{n_P}{d_{\mathcal{H}}} \right) + \frac{1}{n_P} \log \left(\frac{1}{\delta} \right) \right)^{\frac{1}{2-\beta_P}}.$$

Applying the definition of ρ , this has the further implication that (again if (RCS) holds)

$$\mathcal{E}_Q(\hat{h}) \leq C \left(\frac{d_{\mathcal{H}}}{n_P} \log \left(\frac{n_P}{d_{\mathcal{H}}} \right) + \frac{1}{n_P} \log \left(\frac{1}{\delta} \right) \right)^{\frac{1}{(2-\beta_P)\rho}}.$$

Also note that, if $\rho = \infty$ this inequality trivially holds, whereas if $\rho < \infty$ then (RCS) necessarily holds so that the above implication is generally valid, without needing the (RCS) assumption explicitly. Moreover, again when the above events hold, using the event from Lemma 1 again, along with the constraint from the optimization, we have that

$$R_Q(\hat{h}) - R_Q(\hat{h}_{S_Q}) \leq 2c\sqrt{\hat{P}_{S_Q}(\hat{h} \neq \hat{h}_{S_Q})A_{n_Q}} + 2cA_{n_Q},$$

and (5) implies the right hand side is at most

$$C\sqrt{Q(\hat{h} \neq \hat{h}_{S_Q})A_{n_Q}} + cA_{n_Q} \leq C\sqrt{Q(\hat{h} \neq h_Q^*)A_{n_Q}} + C\sqrt{Q(\hat{h}_{S_Q} \neq h_Q^*)A_{n_Q}} + cA_{n_Q}.$$

Using the Bernstein class condition and (7), the second term is bounded by

$$C \left(\frac{d_{\mathcal{H}}}{n_Q} \log \left(\frac{n_Q}{d_{\mathcal{H}}} \right) + \frac{1}{n_Q} \log \left(\frac{1}{\delta} \right) \right)^{\frac{1}{2-\beta_Q}},$$

while the first term is bounded by

$$C\sqrt{\mathcal{E}_Q(\hat{h})^{\beta_Q} A_{n_Q}}.$$

Altogether, we have that

$$\begin{aligned}\mathcal{E}_Q(\hat{h}) &= R_Q(\hat{h}) - R_Q(\hat{h}_{S_Q}) + \mathcal{E}_Q(\hat{h}_{S_Q}) \\ &\leq C\sqrt{\mathcal{E}_Q(\hat{h})^{\beta_Q} A_{n_Q}} + C\left(\frac{d_{\mathcal{H}}}{n_Q} \log\left(\frac{n_Q}{d_{\mathcal{H}}}\right) + \frac{1}{n_Q} \log\left(\frac{1}{\delta}\right)\right)^{\frac{1}{2-\beta_Q}},\end{aligned}$$

which implies

$$\mathcal{E}_Q(\hat{h}) \leq C\left(\frac{d_{\mathcal{H}}}{n_Q} \log\left(\frac{n_Q}{d_{\mathcal{H}}}\right) + \frac{1}{n_Q} \log\left(\frac{1}{\delta}\right)\right)^{\frac{1}{2-\beta_Q}}. \quad \square$$

Remark 4. Note that the above Theorem 3 does not require (RCS): that is, it holds even when $\mathcal{E}_Q(h_P^*) > 0$, in which case $\rho = \infty$. However, for a related method we can also show a stronger result in terms of a modified definition of ρ :

Specifically, define $\mathcal{E}_Q(h, h_P^*) = \max\{R_Q(h) - R_Q(h_P^*), 0\}$, and suppose $\rho' > 0$, $C_{\rho'} > 0$ satisfy

$$\forall h \in \mathcal{H}, \quad C_{\rho'} \cdot \mathcal{E}_P(h) \geq \mathcal{E}_Q'(h, h_P^*).$$

This is clearly equivalent to ρ (Definition 2) under (RCS); however, unlike ρ , this ρ' can be finite even in cases where (RCS) fails. With this definition, we have the following result.

Proposition 2 (Beyond (RCS)). If $\hat{R}_{S_Q}(\hat{h}_{S_P}) - \hat{R}_{S_Q}(\hat{h}_{S_Q}) \leq c\sqrt{\hat{P}_{S_Q}(\hat{h}_{S_P} \neq \hat{h}_{S_Q})A_{n_Q}} + cA_{n_Q}$, that is, if \hat{h}_{S_P} satisfies (6), define $\hat{h} = \hat{h}_{S_P}$, and otherwise define $\hat{h} = \hat{h}_{S_Q}$. Assume (NC). For a constant C depending on $\rho', C_{\rho'}, \beta_P, c_{\beta_P}, \beta_Q, c_{\beta_Q}$, with probability at least $1 - \delta$,

$$\mathcal{E}_Q(\hat{h}) \leq \min\left\{\mathcal{E}_Q(h_P^*) + CA_{n_P}^{\frac{1}{(2-\beta_P)\rho'}}, CA_{n_Q}^{\frac{1}{2-\beta_Q}}\right\}.$$

The proof of this result is similar to that of Theorem 3, and as such is deferred to Appendix C.

An alternative procedure. Similar results as in Theorem 3 can be obtained for a method that swaps the roles of P and Q samples:

Algorithm 1' :

$$\begin{array}{ll}\text{Minimize} & \hat{R}_{S_Q}(h) \\ \text{subject to} & \hat{R}_{S_P}(h) - \hat{R}_{S_P}(\hat{h}_{S_P}) \leq c\sqrt{\hat{P}_{S_P}(h \neq \hat{h}_{S_P})A_{n_P}} + cA_{n_P} \\ & h \in \mathcal{H}.\end{array}$$

This version, more akin to so-called *hypothesis transfer* may have practical benefits in scenarios where the P data is accessible *before* the Q data, since then the feasible set might be calculated (or approximated) in advance, so that the P data itself would no longer be needed in order to execute the procedure. However this procedure presumes that h_P^* is not far from h_Q^* , i.e., that data S_P from P is not misleading, since it conditions on doing well on S_P . Hence we now require (RCS).

Proposition 3. Assume (NC) and (RCS). Let \hat{h} be the solution from Algorithm 1'. For a constant C depending on $\rho, C_{\rho}, \beta_P, c_{\beta_P}, \beta_Q, c_{\beta_Q}$, with probability at least $1 - \delta$,

$$\mathcal{E}_Q(\hat{h}) \leq C \min\left\{A_{n_P}^{\frac{1}{(2-\beta_P)\rho}}, A_{n_Q}^{\frac{1}{2-\beta_Q}}\right\} = \tilde{O}\left(\min\left\{\left(\frac{d_{\mathcal{H}}}{n_P}\right)^{\frac{1}{(2-\beta_P)\rho}}, \left(\frac{d_{\mathcal{H}}}{n_Q}\right)^{\frac{1}{2-\beta_Q}}\right\}\right).$$

The proof is very similar to that of Theorem 3, so is omitted for brevity.

6 Minimizing Sampling Cost

In this section (and continued in Appendix A.1), we discuss the value of having access to *unlabeled* data from Q . The idea is that unlabeled data can be obtained much more cheaply than labeled data,

so gaining access to unlabeled data can be realistic in many applications. Specifically, we begin by discussing an adaptive sampling scenario, where we are able to *draw* samples from P or Q , at different *costs*, and we are interested in optimizing the total cost of obtaining a given excess Q -risk.

Formally, consider the scenario where we have as input a value ϵ , and are tasked with producing a classifier \hat{h} with $\mathcal{E}_Q(\hat{h}) \leq \epsilon$. We are then allowed to *draw* samples from either P or Q toward achieving this goal, but at different costs. Suppose $c_P : \mathbb{N} \rightarrow [0, \infty)$ and $c_Q : \mathbb{N} \rightarrow [0, \infty)$ are *cost functions*, where $c_P(n)$ indicates the cost of sampling a batch of size n from P , and similarly define $c_Q(n)$. We suppose these functions are increasing, and concave, and unbounded.

Definition 5. Define $n_Q^* = d_{\mathcal{H}}/\epsilon^{2-\beta_Q}$, $n_P^* = d_{\mathcal{H}}/\epsilon^{(2-\beta_P)\gamma/\beta_P}$, and $c^* = \min\{c_Q(n_Q^*), c_P(n_P^*)\}$. We call $c^* = c^*(\epsilon; c_P, c_Q)$ the **minimax optimal cost** of sampling from P or Q to attain Q -error ϵ .

Note that the cost c^* is effectively the smallest possible, up to log factors, in the range of parameters given in Theorem 2. That is, in order to make the lower bound in Theorem 2 less than ϵ , either $n_Q = \tilde{\Omega}(n_Q^*)$ samples are needed from Q or $n_P = \tilde{\Omega}(n_P^*)$ samples are needed from P . We show that c^* is nearly achievable, adaptively with no knowledge of distributional parameters.

Procedure. We assume access to a large unlabeled data set U_Q sampled from Q_X . For our purposes, we will suppose this data set has size at least $\Theta(\frac{d_{\mathcal{H}}}{\epsilon} \log \frac{1}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta})$.

Let $A'_n = \frac{d_{\mathcal{H}}}{n} \log(\frac{\max\{n, d_{\mathcal{H}}\}}{d_{\mathcal{H}}}) + \frac{1}{n} \log(\frac{2n^2}{\delta})$. Then for any labeled data set S , define $\hat{h}_S = \operatorname{argmin}_{h \in \mathcal{H}} \hat{R}_S(h)$, and given an additional data set U (labeled or unlabeled) define a quantity

$$\hat{\delta}(S, U) = \sup \left\{ \hat{P}_U(h \neq \hat{h}_S) : h \in \mathcal{H}, \hat{R}_S(h) - \hat{R}_S(\hat{h}_S) \leq c \sqrt{\hat{P}_S(h \neq \hat{h}_S) A'_{|S|}} + c A'_{|S|} \right\},$$

where c is as in Lemma 1. Now we have the following procedure.

Algorithm 2:

0. $S_P \leftarrow \{\}, S_Q \leftarrow \{\}$
1. For $t = 1, 2, \dots$
2. Let $n_{t,P}$ be minimal such that $c_P(n_{t,P}) \geq 2^{t-1}$
3. Sample $n_{t,P}$ samples from P and add them to S_P
4. Let $n_{t,Q}$ be minimal such that $c_Q(n_{t,Q}) \geq 2^{t-1}$
5. Sample $n_{t,Q}$ samples from Q and add them to S_Q
6. If $c \sqrt{\hat{\delta}(S_Q, S_Q) A'_{|S_Q|}} + c A'_{|S_Q|} \leq \epsilon$, return \hat{h}_{S_Q}
7. If $\hat{\delta}(S_P, U_Q) \leq \epsilon/4$, return \hat{h}_{S_P}

The following theorem asserts that this procedure will find a classifier \hat{h} with $\mathcal{E}_Q(\hat{h}) \leq \epsilon$ while adaptively using a near-minimal cost associated with achieving this. The proof is in Appendix D.

Theorem 4 (Adapting to Sampling Costs). *Assume (NC) and (RCS). There exist a constant c' , depending on parameters $(C_\gamma, \gamma, c_{\beta_Q}, \beta_Q, c_{\beta_P}, \beta_P)$ but not on ϵ or δ , such that the following holds. Define sample sizes $\tilde{n}_Q = \frac{c'}{\epsilon^{2-\beta_Q}} (d_{\mathcal{H}} \log \frac{1}{\epsilon} + \log \frac{1}{\delta})$, and $\tilde{n}_P = \frac{c'}{\epsilon^{(2-\beta_P)\gamma/\beta_P}} (d_{\mathcal{H}} \log \frac{1}{\epsilon} + \log \frac{1}{\delta})$.*

Algorithm 2 outputs a classifier \hat{h} such that, with probability at least $1 - \delta$, we have $\mathcal{E}_Q(\hat{h}) \leq \epsilon$, and the total sampling cost incurred is at most $\min\{c_Q(\tilde{n}_Q), c_P(\tilde{n}_P)\} = \tilde{O}(c^)$.*

Thus, when c^* favors sampling from P , we end up sampling very few labeled Q data. These are scenarios where P samples are cheap relative to the cost of Q samples and w.r.t. parameters $(\beta_Q, \beta_P, \gamma)$ which determine the effective source sample size contributed for every target sample. Furthermore, we achieve this adaptively: without knowing (or even estimating) these relevant parameters.

Acknowledgments

We thank Mehryar Mohri for several very important discussions which helped crystallize many essential questions and directions on this topic.

References

- [1] Ilja Kuzborskij and Francesco Orabona. Stability and hypothesis transfer learning. In *Proceedings of the 30th International Conference on Machine Learning*, pages 942–950, 2013.
- [2] Simon S Du, Jayanth Koushik, Aarti Singh, and Barnabás Póczos. Hypothesis transfer learning via transformation functions. In *Advances in Neural Information Processing Systems*, pages 574–584, 2017.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [4] Shai Ben-David, Tyler Lu, Teresa Luu, and Dávid Pál. Impossibility theorems for domain adaptation. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 129–136, 2010.
- [5] Pascal Germain, Amaury Habrard, François Laviolette, and Emilie Morvant. A pac-bayesian approach for domain adaptation with specialization to linear classifiers. In *International Conference on Machine Learning*, pages 738–746, 2013.
- [6] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- [7] Mehryar Mohri and Andres Munoz Medina. New analysis and algorithm for learning with drifting distributions. In *International Conference on Algorithmic Learning Theory*, pages 124–138. Springer, 2012.
- [8] Corinna Cortes, Mehryar Mohri, and Andrés Munoz Medina. Adaptation based on generalized discrepancy. *Machine Learning Research*, forthcoming. URL <http://www.cs.nyu.edu/~mohri/pub/daj.pdf>.
- [9] Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. The MIT Press, 2009.
- [10] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- [11] Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul V Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in neural information processing systems*, pages 1433–1440, 2008.
- [12] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Multiple source adaptation and the rényi divergence. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 367–374. AUAI Press, 2009.
- [13] Corinna Cortes, Mehryar Mohri, Michael Riley, and Afshin Rostamizadeh. Sample selection bias correction theory. In *International conference on algorithmic learning theory*, pages 38–53. Springer, 2008.
- [14] Arthur Gretton, Alex Smola, Jiayuan Huang, Marcel Schmittfull, Karsten Borgwardt, and Bernhard Schölkopf. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4):5, 2009.
- [15] Jiayuan Huang, Arthur Gretton, Karsten M Borgwardt, Bernhard Schölkopf, and Alex J Smola. Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, pages 601–608, 2007.
- [16] Shai Ben-David and Ruth Uner. On the hardness of domain adaptation and the utility of unlabeled target samples. In *International Conference on Algorithmic Learning Theory*, pages 139–153. Springer, 2012.

- [17] Avishek Saha, Piyush Rai, Hal Daumé, Suresh Venkatasubramanian, and Scott L DuVall. Active supervised domain adaptation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 97–112. Springer, 2011.
- [18] Minmin Chen, Kilian Q Weinberger, and John Blitzer. Co-training for domain adaptation. In *Advances in neural information processing systems*, pages 2456–2464, 2011.
- [19] Rita Chattopadhyay, Wei Fan, Ian Davidson, Sethuraman Panchanathan, and Jieping Ye. Joint transfer and batch-mode active learning. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 253–261, 2013.
- [20] Liu Yang, Steve Hanneke, and Jaime Carbonell. A theory of transfer learning with applications to active learning. *Machine learning*, 90(2):161–189, 2013.
- [21] Christopher Berlind and Ruth Uner. Active nearest neighbors in changing environments. In *International Conference on Machine Learning*, pages 1870–1879, 2015.
- [22] Gilles Blanchard, Aniket Anand Deshmukh, Urun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *arXiv preprint arXiv:1711.07910*, 2017.
- [23] Clayton Scott. A generalized neyman-pearson criterion for optimal domain adaptation. In *Algorithmic Learning Theory*, pages 738–761, 2019.
- [24] T Tony Cai and Hongji Wei. Transfer learning for nonparametric classification: Minimax rate and adaptive classifier. *arXiv preprint arXiv:1906.02903*, 2019.
- [25] Samory Kpotufe and Guillaume Martinet. Marginal singularity, and the benefits of labels in covariate-shift. *arXiv preprint arXiv:1803.01833*, 2018.
- [26] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their expectation. *Theory of probability and its applications*, 16:264–280, 1971.
- [27] P. L. Bartlett and S. Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.
- [28] P. Massart and É. Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- [29] V. Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- [30] A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- [31] E. Mammen and A. B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- [32] V. Vapnik and A. Chervonenkis. *Theory of Pattern Recognition*. Nauka, 1974.
- [33] Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer, 2009.
- [34] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- [35] S. Hanneke and L. Yang. Surrogate losses in passive and active learning. *arXiv:1207.3772*, 2012.

A Additional Results

A.1 Reweighting the Source Data

In this section, we present a technique for using unlabeled data from Q to find a reweighting of the P data more suitable for transfer. This gives a technique for using the data effectively in a potentially practical way. As above, we again suppose access to the sample U_Q of unlabeled data from Q .

Additionally, we suppose we have access to a set \mathcal{P} of functions $f : \mathcal{X} \rightarrow [0, \infty)$, which we interpret as unnormalized *density* functions with respect to P_X . Let P_f denote the bounded measure whose marginal on \mathcal{X} has density f with respect to P_X , and the conditional $Y|X$ is the same as for P .

Now suppose $S_P = \{(x_i, y_i)\}_{i=1}^{n_P}$ is a sequence of n_P iid P -distributed samples. Continuing conventions from above $R_{P_f}(h) = \int \mathbb{1}[h(x) \neq y]f(x)dP(x, y)$ is a risk with respect to P_f , but now we also write $\hat{R}_{S_P, f}(h) = \frac{1}{n_P} \sum_{(x, y) \in S_P} \mathbb{1}[h(x) \neq y]f(x)$, and additionally we will use $P_{f^2}(h \neq h') = \int \mathbb{1}[h(x) \neq h'(x)]f^2(x)dP(x, y)$, and $\hat{P}_{S_P, f^2}(h \neq h') = \frac{1}{n_P} \sum_{(x, y) \in S_P} \mathbb{1}[h(x) \neq h'(x)]f^2(x)$; the reason f^2 is used instead of f is that this will represent a variance term in the bounds below. Other notations from above are defined analogously. In particular, also let $\hat{h}_{S_P, f} = \operatorname{argmin}_{h \in \mathcal{H}} \hat{R}_{S_P, f}(h)$. For simplicity, we will only present the case of \mathcal{P} having finite pseudo-dimension d_p (i.e., d_p is the VC dimension of the subgraph functions $\{(x, y) \mapsto \mathbb{1}[f(x) \leq y] : f \in \mathcal{P}\}$); extensions to general bracketing or empirical covering follow similarly.

For the remaining results in this section, we suppose the condition RCS holds for *all* P_f : that is, R_{P_f} is minimized in \mathcal{H} at a function $h_{P_f}^*$ having $\mathcal{E}_Q(h_{P_f}^*) = 0$. For instance, this would be the case if the Bayes optimal classifier is in the class \mathcal{H} .

Define $A''_n = \frac{d_{\mathcal{H}} + d_p}{n} \log\left(\frac{\max\{n, d_{\mathcal{H}} + d_p\}}{d_{\mathcal{H}} + d_p}\right) + \frac{1}{n} \log\left(\frac{1}{\delta}\right)$. Let us also extend the definition of $\hat{\delta}$ introduced above. Specifically, define $\hat{\delta}(S_P, f, U_Q)$ as

$$\sup \left\{ \hat{P}_{U_Q}(h \neq \hat{h}_{S_P, f}) : h \in \mathcal{H}, \hat{\mathcal{E}}_{S_P, f}(h) \leq c \sqrt{\hat{P}_{S_P, f^2}(h \neq \hat{h}_{S_P, f}) A''_{n_P}} + c \|f\|_{\infty} A''_{n_P} \right\}.$$

Now consider the following procedure.

Algorithm 3:

Choose \hat{f} to minimize $\hat{\delta}(S_P, f, U_Q)$ over $f \in \mathcal{P}$.

Choose \hat{h} to minimize $\hat{R}_{S_Q}(h)$ among $h \in \mathcal{H}$

subject to $\hat{\mathcal{E}}_{S_P, \hat{f}}(h) \leq c \sqrt{\hat{P}_{S_P, \hat{f}^2}(h \neq \hat{h}_{S_P, \hat{f}}) A''_{n_P}} + c \|\hat{f}\|_{\infty} A''_{n_P}$.

As we establish in the proof, \hat{f} is effectively being chosen to minimize an upper bound on the excess Q -risk of the resulting classifier \hat{h} . Toward analyzing the performance of this procedure, note that each f induces a marginal transfer exponent: that is, values $C_{\gamma, f}, \gamma_f$ such that $\forall h \in \mathcal{H}$, $C_{\gamma, f} P_{f^2}(h \neq h_{P_f}^*) \geq Q^{\gamma_f}(h \neq h_{P_f}^*)$. Similarly, each f induces a Bernstein Class Condition: there exist values $c_f > 0, \beta_f \in [0, 1]$ such that $P_{f^2}(h \neq h_{P_f}^*) \leq c_f \mathcal{E}_{P_f}^{\beta_f}(h)$.

The following theorem reveals that Algorithm 3 is able to perform nearly as well as applying the transfer technique from Theorem 3 directly under the measure in the family \mathcal{P} that would provide the best bound. The only losses compared to doing so are a dependence on d_p and the supremum of the density (which accounts for how different that measure is from P). The proof is in Appendix E.

Theorem 5. *Suppose $\beta_Q > 0$ and that (NC) and (RCS) hold for all $P_f, f \in \mathcal{P}$. There exist constants C_f depending on $\|f\|_{\infty}, C_{\gamma, f}, \gamma_f, c_f, \beta_f$, and a constant C depending on c_q, β_Q such that, for a sufficiently large $|U_Q|$, w.p. at least $1 - \delta$, the classifier \hat{h} chosen by Algorithm 3 satisfies*

$$\begin{aligned} \mathcal{E}_Q(\hat{h}) &\leq \inf_{f \in \mathcal{P}} C \min \left\{ C_f (A''_{n_P})^{\frac{\beta_f}{(2-\beta_f)\gamma_f}}, A''_{n_Q}^{\frac{1}{2-\beta_Q}} \right\} \\ &= \tilde{O} \left(\inf_{f \in \mathcal{P}} \min \left\{ C_f \left(\frac{d_{\mathcal{H}} + d_p}{n_P} \right)^{\frac{\beta_f}{(2-\beta_f)\gamma_f}}, \left(\frac{d_{\mathcal{H}}}{n_Q} \right)^{\frac{1}{2-\beta_Q}} \right\} \right). \end{aligned}$$

The utility of this theorem will of course depend largely on the family \mathcal{P} of densities. This class should contain a distribution with small γ_f marginal transfer exponent, while also small $\|f\|_\infty$ (which is captured by the C_f constant in the bound), and favorable noise conditions (i.e., large β_f).

A.2 Choice of Transfer from Multiple Sources

It is worth noting that all of the above analysis also applies to the case that, instead of a family of densities with respect to a single P , the set \mathcal{P} is a set of probability measures P_i , each with its own separate iid data set S_i of some size n_i . Lemma 1 can then be applied to all of these data sets, if we simply replace δ by $\delta/|\mathcal{P}|$ to accommodate a union bound; call the corresponding quantity A_n''' . Then, similarly to the above, we can use the following procedure.

Algorithm 4:
 Choose \hat{i} to minimize $\hat{\delta}(S_i, U_Q)$ over $P_i \in \mathcal{P}$.
 Choose \hat{h} to minimize $\hat{R}_{S_Q}(h)$ among $h \in \mathcal{H}$
 subject to $\hat{\mathcal{E}}_{S_i}(h) \leq c\sqrt{\hat{P}_{S_i}(h \neq \hat{h}_{S_i})}A_{n_i}''' + cA_{n_i}'''$.

To state a formal guarantee, let us suppose the conditions above hold for each of these distributions with respective values of C_{γ_i} , γ_i , c_i , β_i . We have the following theorem. Its proof is essentially identical to the proof of Theorem 5 (effectively just substituting notation), and is therefore omitted.

Theorem 6. *Suppose $\beta_Q > 0$ and that (NC) and (RCS) hold for all $P_i \in \mathcal{P}$. There exist constants C_i depending on C_{γ_i} , γ_i , c_i , β_i , and a constant C depending on c_q , β_Q such that, for a sufficiently large $|U_Q|$, with probability at least $1 - \delta$, the classifier \hat{h} chosen by Algorithm 4 satisfies*

$$\mathcal{E}_Q(\hat{h}) \leq \tilde{O} \left(\inf_{P_i \in \mathcal{P}} \min \left\{ C_i \left(\frac{d_{\mathcal{H}} + \log(|\mathcal{P}|)}{n_i} \right)^{\frac{\beta_i}{(2-\beta_i)\gamma_i}}, \left(\frac{d_{\mathcal{H}}}{n_Q} \right)^{\frac{1}{2-\beta_Q}} \right\} \right).$$

B Lower-Bounds Proofs

Our lower-bounds rely on the following extensions of Fano inequality.

Proposition 4 (Thm 2.5 of [33]). *Let $\{\Pi_h\}_{h \in \mathcal{H}}$ be a family of distributions indexed over a subset \mathcal{H} of a semi-metric $(\mathcal{F}, \text{dist})$. Suppose $\exists h_0, \dots, h_M \in \mathcal{H}$, where $M \geq 2$, such that:*

- (i) $\text{dist}(h_i, h_j) \geq 2s > 0$, $\forall 0 \leq i < j \leq M$,
- (ii) $\Pi_{h_i} \ll \Pi_{h_0} \quad \forall i \in [M]$, and the average KL-divergence to Π_{h_0} satisfies

$$\frac{1}{M} \sum_{i=1}^M \mathcal{D}_{kl}(\Pi_{h_i} | \Pi_{h_0}) \leq \alpha \log M, \text{ where } 0 < \alpha < 1/8.$$

Let $Z \sim \Pi_h$, and let $\hat{h} : Z \mapsto \mathcal{F}$ denote any improper learner of $h \in \mathcal{H}$. We have for any \hat{h} :

$$\sup_{h \in \mathcal{H}} \Pi_h \left(\text{dist}(\hat{h}(Z), h) \geq s \right) \geq \frac{\sqrt{M}}{1 + \sqrt{M}} \left(1 - 2\alpha - \sqrt{\frac{2\alpha}{\log(M)}} \right) \geq \frac{3 - 2\sqrt{2}}{8}.$$

The following proposition would be needed to construct packings (of spaces of distributions) of the appropriate size.

Proposition 5 (Varshamov-Gilbert bound). *Let $d \geq 8$. Then there exists a subset $\{\sigma_0, \dots, \sigma_M\}$ of $\{-1, 1\}^d$ such that $\sigma_0 = (1, \dots, 1)$,*

$$\text{dist}(\sigma_i, \sigma_j) \geq \frac{d}{8}, \quad \forall 0 \leq i < j \leq M, \quad \text{and} \quad M \geq 2^{d/8},$$

where $\text{dist}(\sigma, \sigma') \doteq \text{card}(\{i \in [m] : \sigma(i) \neq \sigma'(i)\})$ is the Hamming distance.

Results similar to the following lemma are known.

Lemma 2 (A basic KL upper-bound). *For any $0 < p, q < 1$, we let $\mathcal{D}_{kl}(p|q)$ denote $\mathcal{D}_{kl}(\text{Ber}(p)|\text{Ber}(q))$. Now let $0 < \epsilon < 1/2$ and let $z \in \{-1, 1\}$. We have*

$$\mathcal{D}_{kl}(1/2 + (z/2) \cdot \epsilon | 1/2 - (z/2) \cdot \epsilon) \leq c_0 \cdot \epsilon^2, \text{ for some } c_0 \text{ independent of } \epsilon.$$

Proof. Write $\frac{p}{q} \doteq \frac{1/2 + (z/2)\epsilon}{1/2 - (z/2)\epsilon} = 1 + \frac{2z\epsilon}{1 - z\epsilon}$, and use the fact that

$$\mathcal{D}_{kl}(p|q) \leq \chi^2(p|q) = q \left(1 - \frac{p}{q}\right)^2 + (1-q) \left(1 - \frac{1-p}{1-q}\right)^2 = q \left(\frac{2z\epsilon}{1 - z\epsilon}\right)^2 + (1-q) \left(\frac{-2z\epsilon}{1 + z\epsilon}\right)^2.$$

□

Proof of Theorem 1. Let $d = d_{\mathcal{H}} - 1$. Pick $x_0, x_1, x_2, \dots, x_d$ a shatterable subset of \mathcal{X} under \mathcal{H} . These will form the support of marginals P_X, Q_X . Furthermore, let $\tilde{\mathcal{H}}$ denote the *projection* of \mathcal{H} onto $\{x_i\}_{i=0}^d$ (i.e., the quotient space of equivalences $h \equiv h'$ on $\{x_i\}$), with the additional constraint that all $h \in \tilde{\mathcal{H}}$ classify x_0 as 1. We can now restrict attention to $\tilde{\mathcal{H}}$ as the *effective* hypothesis class.

Let $\sigma \in \{-1, 1\}^d$. We will construct a family of distribution pairs (P_σ, Q_σ) indexed by σ to which we then apply Proposition 4 above. For any P_σ, Q_σ , we let $\eta_{P,\sigma}, \eta_{Q,\sigma}$ denote the corresponding regression functions (i.e., $\mathbb{E}_{P_\sigma}[Y|x]$, and $\mathbb{E}_{Q_\sigma}[Y|x]$). To proceed, fix $\epsilon = c_1 \cdot \epsilon(n_P, n_Q) \leq 1/2$, for a constant $c_1 < 1$ to be determined, where $\epsilon(n_P, n_Q)$ is as defined in the theorem's statement.

- *Distribution Q_σ .* We have that $Q_\sigma = Q_X \times Q_{Y|X}^\sigma$, where $Q_X(x_0) = 1 - \epsilon^{\beta_Q}$, while $Q_X(x_i) = \frac{1}{d}\epsilon^{\beta_Q}$, $i \geq 1$. Now, the conditional $Q_{Y|X}^\sigma$ is fully determined by $\eta_{Q,\sigma_i}(x_0) = 1$, and $\eta_{Q,\sigma}(x_i) = 1/2 + (\sigma_i/2) \cdot \epsilon^{1-\beta_Q}$, $i \geq 1$.

- *Distribution P_σ .* We have that $P_\sigma = P_X \times P_{Y|X}^\sigma$, $P_X(x_0) = 1 - \epsilon^{\beta_P}$, while $P_X(x_i) = \frac{1}{d}\epsilon^{\beta_P}$, $i \geq 1$. Now, the conditional $P_{Y|X}^\sigma$ is fully determined by $\eta_{P,\sigma}(x_0) = 1$, and $\eta_{P,\sigma}(x_i) = 1/2 + (\sigma_i/2) \cdot \epsilon^{\rho(1-\beta_P)}$, $i \geq 1$.

- *Verifying that $(P_\sigma, Q_\sigma) \in \mathcal{F}_{(NC)}(\rho, \beta_P, \beta_Q, 1)$.* For any $\sigma \in \{-1, 1\}^d$, let $h_\sigma \in \tilde{\mathcal{H}}$ denote the corresponding Bayes classifier (remark that the Bayes is the same for both P_σ and Q_σ). Now, pick any other $h_{\sigma'} \in \tilde{\mathcal{H}}$, and let $\text{dist}(\sigma, \sigma')$ denote the Hamming distance between σ, σ' (as in Proposition 5). We then have that

$$\mathcal{E}_{Q_\sigma}(h_{\sigma'}) = \text{dist}(\sigma, \sigma') \cdot \frac{1}{d}\epsilon^{\beta_Q} \cdot \epsilon^{1-\beta_Q} = \frac{\text{dist}(\sigma, \sigma')}{d} \cdot \epsilon,$$

$$\text{while } Q_X(h_{\sigma'} \neq h_\sigma) = \frac{\text{dist}(\sigma, \sigma')}{d} \cdot \epsilon^{\beta_Q},$$

$$\text{and similarly, } \mathcal{E}_{P_\sigma}(h_{\sigma'}) = \frac{\text{dist}(\sigma, \sigma')}{d} \cdot \epsilon^\rho, \text{ while } P_X(h_{\sigma'} \neq h_\sigma) = \frac{\text{dist}(\sigma, \sigma')}{d} \cdot \epsilon^{\rho\beta_P}.$$

The condition is also easily verified for classifiers not labeling x_0 as 1. Since $(\text{dist}(\sigma, \sigma')/d) \leq 1$, it follows that (1) holds with exponents β_P and β_Q for any P_σ and Q_σ respectively (with $C_{P_\sigma} = 1$, $C_{Q_\sigma} = 1$), and that any P_σ admits a transfer-exponent ρ w.r.t. Q_σ , with $C_\rho = 1$.

- *Reduction to a packing.* Now apply Proposition 5 to identify a subset Σ of $\{-1, 1\}^d$, where $|\Sigma| = M \geq 2^{d/8}$, and $\forall \sigma, \sigma' \in \Sigma$, we have $\text{dist}(\sigma, \sigma') \geq d/8$. It should be clear then that for any $\sigma, \sigma' \in \Sigma$,

$$\mathcal{E}_{Q_\sigma}(h_{\sigma'}) \geq \frac{d}{8} \cdot \frac{1}{d}\epsilon^{\beta_Q} \cdot \epsilon^{1-\beta_Q} = \epsilon/8.$$

Furthermore, by construction, any classifier $\hat{h} : \{x_i\} \mapsto \{0, 1\}$ can be reduced to a decision on σ , and we henceforth view $\text{dist}(\sigma, \sigma')$ as the semi-metric referenced in Proposition 4, with effective indexing set Σ .

- KL bounds in terms of n_P and n_Q . Define $\Pi_\sigma = P_\sigma^{n_P} \times Q_\sigma^{n_Q}$. We can now verify that all $\Pi_\sigma, \Pi_{\sigma'}$ are close in KL-divergence. First notice that, for any $\sigma, \sigma' \in \Sigma$ (in fact in $\{-1, 1\}^d$)

$$\begin{aligned} \mathcal{D}_{\text{kl}}(\Pi_\sigma | \Pi_{\sigma'}) &= n_P \cdot \mathcal{D}_{\text{kl}}(P_\sigma | P_{\sigma'}) + n_Q \cdot \mathcal{D}_{\text{kl}}(Q_\sigma | Q_{\sigma'}) \\ &= n_P \cdot \mathbb{E}_{P_X} \mathcal{D}_{\text{kl}}\left(P_{Y|X}^\sigma | P_{Y|X}^{\sigma'}\right) + n_Q \cdot \mathbb{E}_{Q_X} \mathcal{D}_{\text{kl}}\left(Q_{Y|X}^\sigma | Q_{Y|X}^{\sigma'}\right) \\ &= n_P \cdot \sum_{i=1}^d \frac{\epsilon^{\rho\beta_P}}{d} \mathcal{D}_{\text{kl}}\left(P_{Y|x_i}^\sigma | P_{Y|x_i}^{\sigma'}\right) + n_Q \cdot \sum_{i=1}^d \frac{\epsilon^{\beta_Q}}{d} \mathcal{D}_{\text{kl}}\left(Q_{Y|x_i}^\sigma | Q_{Y|x_i}^{\sigma'}\right) \\ &\leq c_0 \left(n_P \cdot \epsilon^{\rho(2-\beta_P)} + n_Q \cdot \epsilon^{(2-\beta_Q)} \right) \end{aligned} \quad (9)$$

$$\leq c_0 d (c_1^{\rho(2-\beta_P)} + c_1^{2-\beta_Q}) \leq 2c_0 c_1 d. \quad (10)$$

where, for inequality (9), we used Lemma 2 to upper-bound the divergence terms. It follows that, for c_1 sufficiently small so that $2c_0 c_1 \leq 1/16$, we get that (10) is upper bounded by $(1/8) \log M$. Now apply Proposition 4 and conclude. \square

We need the following lemma for the next result.

Lemma 3. *Let $\epsilon_1, \epsilon_2, \alpha, \alpha_1, \alpha_2 \geq 0$, and $\alpha_1 + \alpha_2 \leq 1$. We then have that*

$$\begin{aligned} \text{For } \alpha \geq 1, \quad & \alpha_1 \epsilon_1^\alpha + \alpha_2 \epsilon_2^\alpha \geq (\alpha_1 \epsilon_1 + \alpha_2 \epsilon_2)^\alpha, \text{ and} \\ \text{for } \alpha \leq 1, \quad & \alpha_1 \epsilon_1^\alpha + \alpha_2 \epsilon_2^\alpha \leq (\alpha_1 \epsilon_1 + \alpha_2 \epsilon_2)^\alpha. \end{aligned}$$

Proof. W.l.o.g., let $\alpha_1 + \alpha_2 > 0$, and normalize the l.h.s. of each of the above inequalities by $(\alpha_1 + \alpha_2)^{-1} \geq 1$. The results follows by Jensen's inequality and the convexity of $z \mapsto z^\alpha$ for $\alpha \geq 1$, and concavity of $z \mapsto z^\alpha$ for $\alpha \leq 1$. \square

We can now show Theorem 2.

Proof of Theorem 2. We proceed similarly (as far as high-level arguments) as for the proof of Theorem 1, but with a different construction where distributions now all satisfy $\gamma = \rho \cdot \beta_P$, and are broken into two subfamilies (corresponding to the rates ϵ_1 and ϵ_2), and the final result holds by considering the intersection of these subfamilies. For simplicity, in what follows, assume d is even, otherwise, the arguments hold by just replacing d by $d - 1$. First, define $x_0, x_1, x_2, \dots, x_d, \tilde{\mathcal{H}}$ as in that proof.

Let $\sigma \in \{-1, 1\}^d$. Next we construct distribution pairs P_σ, Q_σ indexed by σ , with corresponding regression functions $\eta_{P,\sigma}, \eta_{Q,\sigma}$. Fix $\epsilon_1 = c_1 \cdot \epsilon_1(n_P, n_Q) \leq 1/2$, and $\epsilon_2 = c_2 \cdot \epsilon_2(n_P, n_Q) \leq 1/2$, for some $c_1, c_2 < 1$ to be determined.

The construction is now broken up over $I_1 \doteq \{1, \dots, \frac{d}{2}\}$, and $I_2 \doteq \{\frac{d}{2} + 1, \dots, d\}$. Fix a constant $\frac{1}{2} \leq \tau < 1$; this ensures that $\epsilon_2/\tau \leq 1$. We will later impose further conditions on τ .

- *Distribution Q_σ .* We let $Q_\sigma = Q_X \times Q_{Y|X}^\sigma$, where $Q_X(x_0) = 1 - \frac{1}{2} \left(\epsilon_1^{\beta_Q} + (\epsilon_2/\tau) \right)$, while $Q_X(x_i) = \frac{1}{d} \epsilon_1^{\beta_Q}$ for $i \in I_1$, and $Q_X(x_i) = \frac{1}{d} (\epsilon_2/\tau)$ for $i \in I_2$. Now, the conditional $Q_{Y|X}^\sigma$ is fully determined by $\eta_{Q,\sigma}(x_0) = 1$, and $\eta_{Q,\sigma}(x_i) = 1/2 + (\sigma_i/2) \cdot \epsilon_1^{1-\beta_Q}$ for $i \in I_1$, and $\eta_{Q,\sigma}(x_i) = 1/2 + (\sigma_i/2) \cdot \tau$ for $i \in I_2$.

- *Distribution P_σ .* We let $P_\sigma = P_X \times P_{Y|X}^\sigma$, where $P_X(x_0) = 1 - \frac{1}{2} \left(\epsilon_1^{\gamma\beta_Q} + \epsilon_2^\gamma \right)$, while $P_X(x_i) = \frac{1}{d} \epsilon_1^{\gamma\beta_Q}$ for $i \in I_1$, and $P_X(x_i) = \frac{1}{d} \epsilon_2^\gamma$ for $i \in I_2$. Now, the conditional $P_{Y|X}^\sigma$ is fully determined by $\eta_{P,\sigma}(x_0) = 1$, and $\eta_{P,\sigma}(x_i) = 1/2 + (\sigma_i/2) \cdot \epsilon_1^{(1-\beta_P)\rho\beta_Q}$ for $i \in I_1$, and $\eta_{P,\sigma}(x_i) = 1/2 + (\sigma_i/2) \cdot \epsilon_2^{(1-\beta_P)\rho}$ for $i \in I_2$.

- *Verifying that $(P_\sigma, Q_\sigma) \in \mathcal{F}_{(NC)}(\rho, \beta_P, \beta_Q, 2)$.* For any $\sigma \in \{-1, 1\}^d$, define $h_\sigma \in \tilde{\mathcal{H}}$ as in the proof of Theorem 1. Now, pick any other $h_{\sigma'} \in \tilde{\mathcal{H}}$, and let $\text{dist}_I(\sigma, \sigma')$ denote the Hamming distance

between σ, σ' , restricted to indices in I (that is the Hamming distance between subvectors σ_I and σ'_I). We then have that

$$\begin{aligned}\mathcal{E}_{Q_\sigma}(h_{\sigma'}) &= \text{dist}_{I_1}(\sigma, \sigma') \cdot \frac{1}{d} \epsilon_1^{\beta_Q} \cdot \epsilon_1^{1-\beta_Q} + \text{dist}_{I_2}(\sigma, \sigma') \cdot \frac{1}{d} (\epsilon_2/\tau) \tau \\ &= \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} \epsilon_1 + \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} \epsilon_2,\end{aligned}$$

$$\text{while } Q_X(h_{\sigma'} \neq h_\sigma) = \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} \epsilon_1^{\beta_Q} + \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} (\epsilon_2/\tau).$$

$$\begin{aligned}\text{Similarly, } \mathcal{E}_{P_\sigma}(h_{\sigma'}) &= \text{dist}_{I_1}(\sigma, \sigma') \cdot \frac{1}{d} \epsilon_1^{\gamma \beta_Q} \cdot \epsilon_1^{(1-\beta_P)\rho \beta_Q} + \text{dist}_{I_2}(\sigma, \sigma') \cdot \frac{1}{d} \epsilon_2^\gamma \cdot \epsilon_2^{(1-\beta_P)\rho}, \\ &= \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} \epsilon_1^{\rho \beta_Q} + \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} \epsilon_2^\rho,\end{aligned}$$

$$\text{while } P_X(h_{\sigma'} \neq h_\sigma) = \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} \epsilon_1^{\gamma \beta_Q} + \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} \epsilon_2^\gamma.$$

The condition is also easily verified for classifiers not labeling x_0 as 1. We apply Lemma 3 repeatedly in what follows. First, by the above, we have that

$$Q_X(h_{\sigma'} \neq h_\sigma) \leq \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} \epsilon_1^{\beta_Q} + 2 \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} \epsilon_2^{\beta_Q} \leq 2 \mathcal{E}_{Q_\sigma}^{\beta_Q}(h_{\sigma'}).$$

On the other hand,

$$P_X(h_{\sigma'} \neq h_\sigma) = \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} (\epsilon_1^{\rho \beta_Q})^{\beta_P} + \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} (\epsilon_2^\rho)^{\beta_P} \leq \mathcal{E}_{P_\sigma}^{\beta_P}(h_{\sigma'}),$$

Finally we have that

$$\mathcal{E}_{P_\sigma}(h_{\sigma'}) \geq \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} \epsilon_1^\rho + \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} \epsilon_2^\rho \geq \mathcal{E}_{Q_\sigma}^\rho(h_{\sigma'}).$$

- *Verifying that γ is a marginal-transfer-exponent P_X to Q_X .* Using the above derivations, the condition that $\gamma \geq 1$, and further imposing the condition that $\tau \geq (1/2)^{1/\gamma}$, we have

$$P_X(h_{\sigma'} \neq h_\sigma) \geq \frac{\text{dist}_{I_1}(\sigma, \sigma')}{d} (\epsilon_1^{\beta_Q})^\gamma + \frac{1}{2} \frac{\text{dist}_{I_2}(\sigma, \sigma')}{d} (\epsilon_2/\tau)^\gamma \geq \frac{1}{2} Q_X^\gamma(h_{\sigma'} \neq h_\sigma).$$

where we again used Lemma 3.

- *Reduction to sub-Packings.* Now, in a slight deviation from the proof of Theorem 1, we define two separate packings (in Hamming distance), indexed by some ς as follows. Fix any $\varsigma \in \{-1, 1\}^{d/2}$, and applying Proposition 2, let $\Sigma_1(\varsigma) \subset \{\sigma \in \{-1, 1\}^d : \sigma_{I_2} = \varsigma\}$, and $\Sigma_2(\varsigma) \subset \{\sigma \in \{-1, 1\}^d : \sigma_{I_1} = \varsigma\}$ denote m -packings of $\{-1, 1\}^{d/2}$, $m \geq d/16$, of size $M + 1$, $M \geq 2^{d/16}$.

Clearly, for any $\sigma, \sigma' \in \Sigma_1(\varsigma)$ we have $\mathcal{E}_{Q_\sigma}(h_{\sigma'}) \geq \epsilon_1/16$, while for any $\sigma, \sigma' \in \Sigma_2(\varsigma)$ we have $\mathcal{E}_{Q_\sigma}(h_{\sigma'}) \geq \epsilon_2/16$.

- *KL Bounds in terms of n_P and n_Q .* Again, define $\Pi_\sigma = P_\sigma^{n_P} \times Q_\sigma^{n_Q}$. First, for any ς fixed, let $\sigma, \sigma' \in \Sigma_1(\varsigma)$. As in the proof of Theorem 1, we apply Lemma 2 to get that

$$\begin{aligned}\mathcal{D}_{\text{kl}}(\Pi_\sigma | \Pi_{\sigma'}) &= n_P \cdot \mathbb{E}_{P_X} \mathcal{D}_{\text{kl}}(P_{Y|X}^\sigma | P_{Y|X}^{\sigma'}) + n_Q \cdot \mathbb{E}_{Q_X} \mathcal{D}_{\text{kl}}(Q_{Y|X}^\sigma | Q_{Y|X}^{\sigma'}) \\ &= n_P \cdot \sum_{i \in I_1} \frac{\epsilon_1^{\gamma \beta_Q}}{d} \mathcal{D}_{\text{kl}}(P_{Y|x_i}^\sigma | P_{Y|x_i}^{\sigma'}) + n_Q \cdot \sum_{i \in I_1} \frac{\epsilon_1^{\beta_Q}}{d} \mathcal{D}_{\text{kl}}(Q_{Y|x_i}^\sigma | Q_{Y|x_i}^{\sigma'}) \\ &\leq n_P \cdot c_0 \frac{1}{2} \epsilon_1^{(2-\beta_P)\rho \beta_Q} + n_Q \cdot c_0 \frac{1}{2} \epsilon_1^{(2-\beta_Q)} \\ &\leq c_0 \frac{d}{2} (c_1^{(2-\beta_P)\rho \beta_Q} + c_1^{2-\beta_Q}) \leq c_0 c_1 d.\end{aligned}$$

Similarly, for any ς fixed, let $\sigma, \sigma' \in \Sigma_2(\varsigma)$; expanding over I_2 , we have:

$$\mathcal{D}_{\text{kl}}(\Pi_\sigma | \Pi_{\sigma'}) \leq n_P \cdot c_0 \frac{1}{2} \epsilon_2^{(2-\beta_P)\rho} + n_Q \cdot c_0 \frac{1}{2} \epsilon_2 \cdot \tau \leq c_0 c_1 d.$$

It follows that, for c_1 sufficiently small so that $c_0 c_1 \leq 1/16$, we can apply Proposition 4 twice, to get that for all ς , there exist σ_{I_1} and σ_{I_2} , such that for some constant c , we have

$$\mathbb{E}_{\Pi_\sigma} \left(\mathcal{E}_{Q_\sigma}(\hat{h}) \right) \geq c \cdot \epsilon_1, \text{ where } \sigma = [\sigma_{I_1}, \varsigma], \text{ and } \mathbb{E}_{\Pi_\sigma} \left(\mathcal{E}_{Q_\sigma}(\hat{h}) \right) \geq c \cdot \epsilon_2, \text{ where } \sigma = [\varsigma, \sigma_{I_2}].$$

It follows that $c \cdot \max\{\epsilon_1, \epsilon_2\}$ is a lower-bound for either $\sigma = [\sigma_{I_1}, \varsigma]$ or $\sigma = [\varsigma, \sigma_{I_2}]$. \square

C Upper Bounds Proofs

Proof of Proposition 2. To reduce redundancy, we refer to arguments presented in the proof of Theorem 3, rather than repeating them here. As in the proof of Theorem 3, we let C serve as a generic constant (possibly depending on $\rho', C_{\rho'}, \beta_P, c_{\beta_P}, \beta_Q, c_{\beta_Q}$) which may be different in different appearances. Define a set

$$\mathcal{G} = \left\{ h \in \mathcal{H} : \hat{R}_{S_Q}(h) - \hat{R}_{S_Q}(\hat{h}_{S_Q}) \leq c \sqrt{\hat{P}_{S_Q}(h \neq \hat{h}_{S_Q}) A_{n_Q}} + c A_{n_Q} \right\}.$$

We can rephrase the definition of \hat{h} as saying $\hat{h} = \hat{h}_{S_P}$ when $\hat{h}_{S_P} \in \mathcal{G}$, and otherwise $\hat{h} = \hat{h}_{S_Q}$.

We suppose the event from Lemma 1 holds for both S_Q and S_P ; by the union bound, this happens with probability at least $1 - \delta$. In particular, as in (8) from the proof of Theorem 3, we have

$$\mathcal{E}_P(\hat{h}_{S_P}) \leq C A_{n_P}^{\frac{1}{2-\beta_P}}.$$

Together with the definition of ρ' , this implies

$$\mathcal{E}_Q(\hat{h}_{S_P}, h_P^*) \leq C A_{n_P}^{\frac{1}{(2-\beta_P)\rho'}},$$

which means

$$\mathcal{E}_Q(\hat{h}_{S_P}) \leq \mathcal{E}_Q(h_P^*) + \mathcal{E}_Q(\hat{h}_{S_P}, h_P^*) \leq \mathcal{E}_Q(h_P^*) + C A_{n_P}^{\frac{1}{(2-\beta_P)\rho'}}. \quad (11)$$

Now, if $R_Q(\hat{h}_{S_P}) \leq R_Q(\hat{h}_{S_Q})$, then (due to the event from Lemma 1) we have $\hat{h}_{S_P} \in \mathcal{G}$, so that $\hat{h} = \hat{h}_{S_P}$, and thus the rightmost expression in (11) bounds $\mathcal{E}_Q(\hat{h})$. On the other hand, if $R_Q(\hat{h}_{S_P}) > R_Q(\hat{h}_{S_Q})$, then regardless of whether $\hat{h} = \hat{h}_{S_P}$ or $\hat{h} = \hat{h}_{S_Q}$, we have $\mathcal{E}_Q(\hat{h}) \leq \mathcal{E}_Q(\hat{h}_{S_P})$, so that again the rightmost expression in (11) bounds $\mathcal{E}_Q(\hat{h})$. Thus, in either case,

$$\mathcal{E}_Q(\hat{h}) \leq \mathcal{E}_Q(h_P^*) + C A_{n_P}^{\frac{1}{(2-\beta_P)\rho'}}.$$

Furthermore, as in the proof of Theorem 3, every $h \in \mathcal{G}$ satisfies $\mathcal{E}_Q(h) \leq C A_{n_Q}^{\frac{1}{2-\beta_Q}}$. Since the algorithm only picks $\hat{h} = \hat{h}_{S_P}$ if $\hat{h}_{S_P} \in \mathcal{G}$, and otherwise picks $\hat{h} = \hat{h}_{S_Q}$, which is clearly in \mathcal{G} , we may note that we always have $\hat{h} \in \mathcal{G}$. We therefore conclude that

$$\mathcal{E}_Q(\hat{h}) \leq C A_{n_Q}^{\frac{1}{2-\beta_Q}},$$

which completes the proof. \square

D Proofs for Adaptive Sampling Costs

Proof of Theorem 4. First note that since $\sum_n \frac{1}{2n^2} < 1$, by the union bound and Lemma 1, with probability at least $1 - \delta$, for every $h, h' \in \mathcal{H}$, every set S_P in the algorithm has

$$R_P(h) - R_P(h') \leq \hat{R}_{S_P}(h) - \hat{R}_{S_P}(h') + c \sqrt{\min\{P(h \neq h'), \hat{P}_{S_P}(h \neq h')\} A'_{|S_P|}} + c A'_{|S_P|}$$

and

$$\hat{P}_{S_P}(h \neq h') \leq 2P(h \neq h') + cA'_{|S_P|}$$

every set S_Q in the algorithm has

$$R_Q(h) - R_Q(h') \leq \hat{R}_{S_Q}(h) - \hat{R}_{S_Q}(h') + c\sqrt{\min\{Q(h \neq h'), \hat{P}_{S_Q}(h \neq h')\}A'_{|S_Q|}} + cA'_{|S_Q|}$$

and

$$\hat{P}_{S_Q}(h \neq h') \leq 2Q(h \neq h') + cA'_{|S_Q|},$$

and we also have for the set U_Q that

$$\frac{1}{2}Q(h \neq h') - cA'_{|U_Q|} \leq \hat{P}_{U_Q}(h \neq h') \leq 2Q(h \neq h') + cA'_{|U_Q|},$$

which by our choice of the size of U_Q implies

$$\frac{1}{2}Q(h \neq h') - \frac{\epsilon}{8} \leq \hat{P}_{U_Q}(h \neq h') \leq 2Q(h \neq h') + \frac{\epsilon}{8}.$$

For the remainder of this proof, we suppose these inequalities hold.

In particular, these imply

$$R_Q(\hat{h}_{S_Q}) - R_Q(h^*) \leq c\sqrt{\hat{P}_{S_Q}(\hat{h}_{S_Q} \neq h^*)A'_{|S_Q|}} + cA'_{|S_Q|}.$$

Furthermore,

$$\hat{R}_{S_Q}(h^*) - \hat{R}_{S_Q}(\hat{h}_{S_Q}) \leq c\sqrt{\hat{P}_{S_Q}(h^* \neq \hat{h}_{S_Q})A'_{|S_Q|}} + cA'_{|S_Q|},$$

so that $h = h^*$ is included in the supremum in the definition of $\hat{\delta}(S_Q, S_Q)$. Together these imply

$$\mathcal{E}_Q(\hat{h}_{S_Q}) \leq R_Q(\hat{h}_{S_Q}) - R_Q(h^*) \leq c\sqrt{\hat{\delta}(S_Q, S_Q)A'_{|S_Q|}} + cA'_{|S_Q|}.$$

Thus, if the algorithm returns \hat{h}_{S_Q} in Step 6, then $\mathcal{E}_Q(\hat{h}_{S_Q}) \leq \epsilon$.

Also by the above inequalities, we have

$$\hat{R}_{S_P}(h^*) - \hat{R}_{S_P}(\hat{h}_{S_P}) \leq c\sqrt{\hat{P}_{S_P}(h^* \neq \hat{h}_{S_P})A'_{|S_P|}} + cA'_{|S_P|},$$

so that h^* is included in the supremum in the definition of $\hat{\delta}(S_P, U_Q)$. Thus,

$$\mathcal{E}_Q(\hat{h}_{S_P}) \leq Q(\hat{h}_{S_P} \neq h^*) \leq 2\hat{P}_{U_Q}(\hat{h}_{S_P} \neq h^*) + \frac{\epsilon}{2} \leq 2\hat{\delta}(S_P, U_Q) + \frac{\epsilon}{2},$$

and hence if the algorithm returns \hat{h}_{S_P} in Step 7 we have $\mathcal{E}_Q(\hat{h}_{S_P}) \leq \epsilon$ as well. Furthermore, the algorithm will definitely return at some point, since the bound in Step 6 approaches 0 as the sample size grows. Altogether, this establishes that, on the above event, the \hat{h} returned by the algorithm satisfies $\mathcal{E}_Q(\hat{h}) \leq \epsilon$, as claimed.

It remains to show that the cost satisfies the stated bound. For this, first note that since the costs incurred by the algorithm grow as a function that is upper and lower bounded by a geometric series, it suffices to argue that, for an appropriate choice of the constant c' , the algorithm would halt if ever it reached a set S_P of size at least n_P^* or a set S_Q of size at least n_Q^* (which ever were to happen first); the result would then follow by choosing the actual constant c' in the theorem slightly larger than this, to account for the algorithm slightly “overshooting” this target (by at most a numerical constant factor).

First suppose it reaches S_Q of size at least n_Q^* . Now, as in the proof of Theorem 3, on the above event, every $h \in \mathcal{H}$ included in the supremum in the definition of $\hat{\delta}(S_Q, S_Q)$ has

$$\mathcal{E}_Q(h) \leq C \left(A'_{|S_Q|} \right)^{\frac{1}{2-\beta_Q}},$$

which further implies

$$Q(h \neq h^*) \leq C \left(A'_{|S_Q|} \right)^{\frac{\beta_Q}{2-\beta_Q}},$$

so that (by the triangle inequality and the above inequalities)

$$\hat{P}_{S_Q}(h \neq \hat{h}_{S_Q}) \leq C \left(A'_{|S_Q|} \right)^{\frac{\beta_Q}{2-\beta_Q}}.$$

Thus, in Step 6,

$$c\sqrt{\hat{\delta}(S_Q, S_Q)A_{|S_Q|}} + cA_{|S_Q|} \leq C \left(A'_{|S_Q|} \right)^{\frac{1}{2-\beta_Q}},$$

which, by our choice of n_Q^* is at most ϵ . Hence, in this case, the algorithm will return in Step 6 (or else would have returned on some previous round).

On the other hand, suppose S_P reaches a size at least n_P^* . In this case, again by the same argument used in the proof of Theorem 3, every $h \in \mathcal{H}$ included in the supremum in the definition of $\hat{\delta}(S_P, U_Q)$ has

$$\mathcal{E}_P(h) \leq C \left(A'_{|S_P|} \right)^{\frac{1}{2-\beta_P}},$$

which implies

$$P(h \neq h^*) \leq C \left(A'_{|S_P|} \right)^{\frac{\beta_P}{2-\beta_P}},$$

and hence

$$Q(h \neq h^*) \leq C \left(A'_{|S_P|} \right)^{\frac{\beta_P}{(2-\beta_P)\gamma}}.$$

By the above inequalities and the triangle inequality (since \hat{h}_{S_P} is clearly also included as an h in that supremum), this implies

$$\hat{P}_{U_Q}(h \neq \hat{h}_{S_P}) \leq C \left(A'_{|S_P|} \right)^{\frac{\beta_P}{(2-\beta_P)\gamma}} + \frac{\epsilon}{8}.$$

Altogether we get that

$$\hat{\delta}(S_P, U_Q) \leq C \left(A'_{|S_P|} \right)^{\frac{\beta_P}{(2-\beta_P)\gamma}} + \frac{\epsilon}{8}.$$

By our choice of n_P^* (for an appropriate choice of constant factors), the right hand side is at most $\epsilon/4$. Therefore, in this case the algorithm will return in Step 7 (if it had not already returned in some previous round). This completes the proof. \square

E Proofs for Reweighting Results

The following lemma is known (see [34, 35]), following from the general form of Bernstein's inequality and standard VC arguments, in combination with the well-known fact that, since the VC dimension of $\{(x, y) \mapsto \mathbb{1}[h(x) \neq y] : h \in \mathcal{H}\}$ is $d_{\mathcal{H}}$, and pseudo-dimension of \mathcal{P} is d_p , it follows that the pseudo-dimension of $\{(x, y) \mapsto \mathbb{1}[h(x) \neq y]f(x) : h \in \mathcal{H}, f \in \mathcal{P}\}$ is at most $\propto d_{\mathcal{H}} + d_p$.

Lemma 4. *With probability at least $1 - \frac{\delta}{3}$, $\forall f \in \mathcal{P}, \forall h, h' \in \mathcal{H}$,*

$$R_{P_f}(h) - R_{P_f}(h') \leq \hat{R}_{S_P, f}(h) - \hat{R}_{S_P, f}(h') + c\sqrt{\min\{P_{f^2}(h \neq h'), \hat{P}_{S_P, f^2}(h \neq h')\}A''_{n_P}} + c\|f\|_{\infty}A''_{n_P}$$

and $\frac{1}{2}P_{f^2}(h \neq h') - c\|f\|_{\infty}A''_{n_P} \leq \hat{P}_{S_P, f^2}(h \neq h') \leq 2P_{f^2}(h \neq h') + c\|f\|_{\infty}A''_{n_P}$, for a universal numerical constant $c \in (0, \infty)$.

Proof of Theorem 5. Let us suppose the event from Lemma 4 holds, as well as the event from Lemma 1 for S_Q , and also the part (5) from the event in Lemma 1 holds for U_Q . The union bound implies all of these hold simultaneously with probability at least $1 - \delta$. For simplicity, and without loss of generality, we will suppose the constants c in these two lemmas are the same. Regarding the

sufficient size of $|U_Q|$, for this result it suffices to have $|U_Q| \geq n_P^{\frac{\beta_f}{(2-\beta_f)\gamma_f}}$ for all $f \in \mathcal{P}$; for instance, in the typical case where $\gamma_f \geq 1$ for all $f \in \mathcal{P}$, it would suffice to simply have $|U_Q| \geq n_P$.

First note that, exactly as in the proof of Theorem 3, since the event in Lemma 4 implies $h_{P_f}^*$ satisfies the constraint in the optimization defining \hat{h} , and the RCS assumption implies $\mathcal{E}_Q(h_{P_f}^*) = 0$, and hence by (NC) that $Q(h_{P_f}^* \neq h_Q^*) = 0$, we immediately get that

$$\mathcal{E}_Q(\hat{h}) \leq CA_{n_Q}^{\frac{1}{2-\beta_Q}}.$$

Thus, it only remains to establish the other term in the minimum as a bound.

Similarly to the proofs above, we let C_f be a general f -dependent constant (with the same restrictions on dependences mentioned in the theorem statement), which may be different in each appearance below. For each $f \in \mathcal{P}$, denote by \hat{h}_f the $h \in \mathcal{H}$ that minimizes $\hat{R}_{S_Q}(h)$ among $h \in \mathcal{H}$ subject to $\hat{\mathcal{E}}_{S_P,f}(h) \leq c\sqrt{\hat{P}_{S_P,f^2}(h \neq \hat{h}_{S_P,f})A''_{n_P}} + c\|f\|_\infty A''_{n_P}$. Also note that $\hat{h}_{S_P,f}$ certainly satisfies the constraint in the set defining $\hat{\delta}(S_P, f, U_Q)$, and that the event from Lemma 4 implies $h_{P_f}^*$ also satisfies this same constraint. Therefore, the event for U_Q from Lemma 1, and the triangle inequality, imply

$$\mathcal{E}_Q(\hat{h}_f) \leq Q(\hat{h}_f \neq h_{P_f}^*) \leq Q(\hat{h}_f \neq \hat{h}_{S_P,f}) + Q(h_{P_f}^* \neq \hat{h}_{S_P,f}) \leq 4\hat{\delta}(S_P, f, U_Q) + 4cA_{|U_Q|}.$$

Thus, \hat{f} is being chosen to minimize an upper bound on the excess Q -risk of the resulting classifier.

Next we relax this expression to match that in the theorem statement. Again using (5), we get that

$$\begin{aligned} \hat{\delta}(S_P, f, U_Q) &\leq cA_{|U_Q|} + \\ &2 \sup \left\{ Q(h \neq \hat{h}_{S_P,f}) : h \in \mathcal{H}, \hat{\mathcal{E}}_{S_P,f} \leq c\sqrt{\hat{P}_{S_P,f^2}(h \neq \hat{h}_{S_P,f})A''_{n_P}} + c\|f\|_\infty A''_{n_P} \right\}. \end{aligned}$$

Again since $h_{P_f}^*$ and $\hat{h}_{S_P,f}$ both satisfy the constraint in this set, the supremum on the right hand side is at most

$$2 \sup \left\{ Q(h \neq h_{P_f}^*) : h \in \mathcal{H}, \hat{\mathcal{E}}_{S_P,f} \leq c\sqrt{\hat{P}_{S_P,f^2}(h \neq \hat{h}_{S_P,f})A''_{n_P}} + c\|f\|_\infty A''_{n_P} \right\}.$$

Then using the marginal transfer condition, this is at most

$$C_f \sup \left\{ P_{f^2}(h \neq h_{P_f}^*)^{\frac{1}{\gamma_f}} : h \in \mathcal{H}, \hat{\mathcal{E}}_{S_P,f} \leq c\sqrt{\hat{P}_{S_P,f^2}(h \neq \hat{h}_{S_P,f})A''_{n_P}} + c\|f\|_\infty A''_{n_P} \right\},$$

and the Bernstein Class condition further bounds this as

$$C_f \sup \left\{ \mathcal{E}_{P_f}^{\frac{\beta_f}{\gamma_f}} : h \in \mathcal{H}, \hat{\mathcal{E}}_{S_P,f} \leq c\sqrt{\hat{P}_{S_P,f^2}(h \neq \hat{h}_{S_P,f})A''_{n_P}} + c\|f\|_\infty A''_{n_P} \right\}.$$

Finally, by essentially the same argument as in the proof of Theorem 3 above, every $h \in \mathcal{H}$ with $\hat{\mathcal{E}}_{S_P,f} \leq c\sqrt{\hat{P}_{S_P,f^2}(h \neq \hat{h}_{S_P,f})A''_{n_P}} + c\|f\|_\infty A''_{n_P}$ satisfies

$$\mathcal{E}_{P_f}(h) \leq C_f(A''_{n_P})^{\frac{1}{2-\beta_f}},$$

so that the above supremum is at most $C_f(A''_{n_P})^{\frac{\beta_f}{(2-\beta_f)\gamma_f}}$ for a (different) appropriate choice of C_f . Altogether we have established that

$$\hat{\delta}(S_P, f, U_Q) \leq cA_{|U_Q|} + C_f(A''_{n_P})^{\frac{\beta_f}{(2-\beta_f)\gamma_f}}.$$

By our condition on $|U_Q|$ specified above, this implies

$$\hat{\delta}(S_P, f, U_Q) \leq C_f(A''_{n_P})^{\frac{\beta_f}{(2-\beta_f)\gamma_f}}.$$

We therefore have that

$$\begin{aligned}\mathcal{E}_Q(\hat{h}) &= \mathcal{E}_Q(\hat{h}_{\hat{f}}) \leq 4\hat{\delta}(S_P, \hat{f}, U_Q) + 4cA_{|U_Q|} = \inf_{f \in \mathcal{P}} 4\hat{\delta}(S_P, \hat{f}, U_Q) + 4cA_{|U_Q|} \\ &\leq \inf_{f \in \mathcal{P}} C_f(A''_{n_P})^{\frac{\beta_f}{(2-\beta_f)\gamma_f}},\end{aligned}$$

where we have again used the condition on $|U_Q|$. This completes the proof. \square