

---

# Complementary Benefits of Contrastive Learning and Self-Training Under Distribution Shift

---

**Saurabh Garg\***

Carnegie Mellon University  
sgarg2@andrew.cmu.edu

**Amrith Setlur\***

Carnegie Mellon University  
asetlur@andrew.cmu.edu

**Zachary C. Lipton**

Carnegie Mellon University  
zlipton@andrew.cmu.edu

**Sivaraman Balakrishnan**

Carnegie Mellon University  
sbalakri@andrew.cmu.edu

**Virginia Smith**

Carnegie Mellon University  
smithv@andrew.cmu.edu

**Aditi Raghunathan**

Carnegie Mellon University  
aditirag@andrew.cmu.edu

## Abstract

Self-training and contrastive learning have emerged as leading techniques for incorporating unlabeled data, both under distribution shift (unsupervised domain adaptation) and when it is absent (semi-supervised learning). However, despite the popularity and compatibility of these techniques, their efficacy in combination remains surprisingly unexplored. In this paper, we first undertake a systematic empirical investigation of this combination, finding (i) that in domain adaptation settings, self-training and contrastive learning offer significant complementary gains; and (ii) that in semi-supervised learning settings, surprisingly, the benefits are not synergistic. Across eight distribution shift datasets (*e.g.*, BREEDs, WILDS), we demonstrate that the combined method obtains 3–8% higher accuracy than either approach independently. Finally, we theoretically analyze these techniques in a simplified model of distribution shift demonstrating scenarios under which the features produced by contrastive learning can yield a good initialization for self-training to further amplify gains and achieve optimal performance, even when either method alone would fail.

## 1 Introduction

Even under natural, non-adversarial distribution shifts, the performance of machine learning models often drops [65, 83, 47, 29]. While we might hope to retrain our models on labeled samples from the new distribution, this option is often unavailable due to the expense or impracticality of collecting new labels. Consequently, researchers have investigated the Unsupervised Domain Adaptation (UDA) setting. Here, given labeled source data and unlabeled out-of-distribution (OOD) target data, the goal is to produce a classifier that performs well on the target. Because UDA is generally underspecified [7], researchers have focused on two main paths: (i) domain adaptation papers that explore heuristics for incorporating the unlabeled target data, relying on benchmark datasets ostensibly representative of “real-world shifts” to adjudicate progress [72, 85, 47]; and (ii) theoretically motivated papers that explore structural assumptions under which UDA problems are well posed [77, 74]. This work engages with the former focusing on two popular methods: self-training and contrastive pretraining.

Self-training [75, 52, 79, 90, 86] and contrastive pretraining [13, 16, 93] were both proposed, initially, for traditional Semi-Supervised Learning (SSL) problems, where the labeled and unlabeled data are drawn from the same distribution. Here the central challenge is statistical: to exploit the unlabeled data to learn a better predictor than one would get by training on the (small) labeled data alone. More

---

\*Equal contribution.

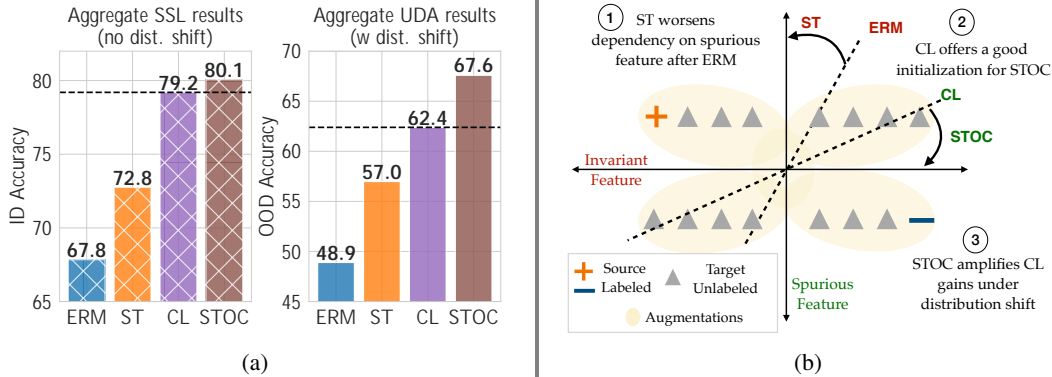


Figure 1: *Self-training over Contrastive learning (STOC) improves over Contrastive Learning (CL) under distribution shift.* (a) We observe that in SSL settings, where labeled and unlabeled data are drawn from the same distribution, STOC offers negligible improvements over CL. In contrast, in UDA settings where there is distribution shift between labeled and unlabeled data, STOC offers gains over CL. Results aggregated across 8 benchmarks. Results on individual data in Table 1 and 2. (b) 2-D illustration of our simplified distribution setup, depicting decision boundaries learned by ERM and CL and how Self-Training (ST) updates those. 1, 2, and 3 summarize our theoretical results in Sec. 4.

recently, these methods have emerged as favored empirical approaches for UDA, demonstrating efficacy on many popular benchmarks [70, 30, 12, 76]. In self-training, one first learns a predictor using source labeled data. The predictor then produces pseudolabels for the unlabeled target data, and a new predictor is trained on the pseudolabeled data. Contrastive pretraining learns representations from unlabeled data by enforcing invariance to specified augmentations. These representations are subsequently used to learn a classifier. In UDA, the representations are trained on the union of the source and target data. In attempts to explain their strong empirical performance, several researchers have attempted analyses under various assumptions on the data, task, and inductive biases of the function class [87, 39, 73, 76, 12, 40, 38, 11]. Despite the strong results, there has been surprisingly little work (both empirically and theoretically) exploring when either might be expected to perform best and whether the benefits might be complementary.

In this paper, we investigate the complementary benefits of self-training and contrastive pretraining. Interestingly, we find that the combination yields significant gains in UDA despite producing negligible gains in SSL. In experiments across eight distribution shift benchmarks (e.g. BREEDs [72], FMoW [47], Visda [63]), we observe that re-using unlabeled data for self-training (with Fix-Match [79]) after learning contrastive representations (with SwAV [13]), yields a 5% average improvement on OOD accuracy in UDA as compared to a 0.8% average improvement in SSL (Fig. 1).

Next, we address the question *why the combination of self-training and contrastive learning proves synergistic in distribution shift scenarios*. To facilitate our analysis, we consider a simplified distribution shift setting that includes two types of features: (i) invariant features that perfectly predict the label; and (ii) domain-dependent features that are predictive of the label in just source. Our theoretical analysis reveals that self-training can achieve optimal target performance but requires a “good” enough classifier to start with. We observe that source-only ERM fails to provide a “good” initialization. On the other hand, contrastive pretraining on unlabeled data performs better than ERM but is still sub-optimal. This implies that contrastive pretraining ends up decreasing reliance on domain-dependent features (as compared to ERM) but doesn’t completely eliminate them. Nevertheless, contrastive pretraining does provide a “good” initialization for self-training, i.e., “good” initial pseudolabels on the target unlabeled data. As a result, self-training on top of contrastive learned features effectively unlearns the reliance on domain-dependent features and generalizes perfectly OOD. In contrast, for SSL settings (i.e., in distribution), our analysis highlights that contrastive pretraining already acquires sufficient predictive features such that linear probing with (a small amount of) labeled data picks up those features and attains near-optimal ID generalization.

Finally, we connect our theoretical understanding of “good” representations from contrastive learning and improved linear transferability from self-training back to observed empirical gains. We linearly probe representations (fix representations and train only the linear head) learned by contrastive pretraining vs. no pretraining and find: (i) contrastive pretraining substantially improves the ceiling on

the target accuracy (performance of optimal linear probe) compared to ERM; (ii) self-training mainly improves linear transfer, *i.e.* OOD performance for the linear probe trained with source labeled data.

## 2 Setup and Preliminaries

**Task.** Our goal is to learn a predictor that maps inputs  $x \in \mathcal{X} \subseteq \mathbb{R}^d$  to outputs  $y \in \mathcal{Y}$ . We parameterize predictors  $f = h \circ \Phi : \mathbb{R}^d \rightarrow \mathcal{Y}$ , where  $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^k$  is a feature map and  $h \in \mathbb{R}^k$  is a classifier that maps the representation to the final scores or logits. Let  $P_S, P_T$  be the source and target joint probability measures over  $\mathcal{X} \times \mathcal{Y}$  with  $p_S$  and  $p_T$  as the corresponding probability density (or mass) functions. The distribution over unlabeled samples from both the union of source and target is denoted as  $P_U = \rho_1 P_S + \rho_2 P_T$ .

We study two particular scenarios: (i) Unsupervised Domain Adaptation (UDA); and (ii) Semi-Supervised Learning (SSL). In UDA, we assume that the source and target distributions have the same label marginals  $P_S p_S(y) = P_T p_T(y)$  (*i.e.*, no label proportion shift) and the same Bayes optimal predictor, *i.e.*,  $\arg \max_y p_S p(y | x) = \arg \max_y p_T p(y | x)$ . We are given labeled samples from the source, and unlabeled pool from the target. In contrast in SSL, there is no distribution shift, *i.e.*,  $P_S = P_T = P_U$ . Here, we are given a small number of labeled examples and a comparatively large amount of unlabeled examples, both drawn from the same distribution, which we denote as  $P_T$ .

Unlabeled data is typically much cheaper to obtain, and our goal in both these settings is to leverage this along with labeled data to achieve good performance on the target distribution. In the UDA scenario, the challenge lies in generalizing out-of-distribution, while in SSL, the challenge is to generalize in-distribution despite the paucity of labeled examples. A predictor  $f$  is evaluated on distribution  $P$  via its accuracy, *i.e.*,  $\mathbb{E}_P \arg \max_y f(x) = y$ .

**Methods.** We now introduce the algorithms used for learning from labeled and unlabeled data.

1. *Source-only ERM (ERM)*: A standard approach is to simply perform supervised learning on the labeled data by minimizing the empirical risk  $\frac{1}{n} \sum_{i=1}^n \ell(h \circ \Phi(x_i), y_i)$ , for some classification loss  $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$  (*e.g.*, softmax cross-entropy) and labeled points  $(x_i, y_i)_{i=1}^n$ .
2. *Contrastive Learning (CL)*: We first use the unlabeled data to learn a feature extractor. In particular, the objective is to learn a feature extractor  $\Phi_{cl}$  that maps augmentations (for *e.g.* crops or rotations) of the same input close to each other and far from augmentations of random other inputs [13, 16, 93]. Once we have  $\Phi_{cl}$ , we learn a linear classifier  $h$  on top to minimize a classification loss on the labeled source data. We could either keep  $\Phi_{cl}$  fixed or propagate gradients through. When clear from context, we also use CL to refer to just the contrastively pretrained backbone without training for downstream classification.
3. *Self-training (ST)*: This is a two-stage procedure, where the first stage performs source-only ERM by just looking at source-labeled data. In the second stage, we iteratively apply the current classifier on the unlabeled data to generate “pseudo-labels” and then update the classifier by minimizing a classification loss on the pseudolabeled data [52].

## 3 Self-Training Improves Contrastive Pretraining Under Distribution Shift

**Self-Training Over Contrastive learning (STOC).** Finally, rather than starting with a source-only ERM classifier, we propose to initialize self-training with a CL classifier, that was pretrained on unlabeled source and target data. ST uses that same unlabeled data again for pseudolabeling. As we demonstrate experimentally and theoretically, this combination of methods improves substantially over each independently.

**Datasets.** For both UDA and SSL, we conduct experiments across eight benchmark datasets: four BREEDs datasets [72]—Entity13, Entity30, Nonliving26, Living17; FMoW [47, 18] from WILDS benchmark; Officehome [85]; Visda [64, 63]; and CIFAR-10 [48]. Each of these datasets consists of domains, enabling us to construct source-target pairs (*e.g.*, CIFAR10, we consider CIFAR10  $\rightarrow$  CINIC shift [22]). In the UDA setup, we adopt the source and target domains standard to previous studies (details in App. C.2). Because the SSL setting lacks distribution shift, we do not need to worry about domain designations and default to using source alone. To simulate limited supervision in SSL, we sub-sample the original labeled training set to 10%.

Table 1: *Results in the UDA setup.* We report accuracy on target (OOD) data from which we only observe unlabeled examples during training. For benchmarks with multiple target distributions (e.g., OH, Visda), we report avg accuracy on those targets. Results with source performance, individual target performance, and standard deviation numbers are in App. C.4.

Method	Living17	Nonliv26	Entity13	Entity30	FMoW (2 tgts)	Visda (2 tgts)	OH (3 tgts)	CIFAR $\tilde{N}$ CINIC	Avg
ERM	60.31	45.54	68.32	55.75	56.50	20.91	9.51	74.33	48.90
ST	71.29	56.79	77.93	66.37	56.79	38.03	10.47	78.19	56.98
CL	74.14	57.02	76.58	66.01	61.78	63.49	22.63	77.51	62.39
STOC (ours)	<b>82.22</b>	<b>62.23</b>	<b>81.84</b>	<b>72.00</b>	<b>65.25</b>	<b>70.08</b>	<b>27.12</b>	<b>79.94</b>	<b>67.59</b>

Table 2: *Results in the SSL setup.* We report accuracy on hold-out ID data. Recall that SSL uses labeled and unlabeled data from the same distribution during training. Refer to App. C.5 for ERM and ST.

Method	Living17	Nonliv26	Entity13	Entity30	FMoW	Visda	OH	CIFAR	Avg
CL	91.15	84.58	90.73	85.47	43.05	97.67	49.73	91.78	79.27
STOC (ours)	92.00	85.95	91.27	86.14	44.43	97.70	49.95	93.06	80.06

**Experimental Setup and Protocols.** SwAV [13] is the specific algorithm that we use for contrastive pretraining. In all UDA settings, unless otherwise specified, we pool all the (unlabeled) data from the source and target to perform SwAV. For self-training, we apply FixMatch [79], where the loss on source labeled data and on pseudolabeled target data are minimized simultaneously. For both methods, we fix the algorithm-specific hyperparameters to the original recommendations. For SSL settings, we perform SwAV and FixMatch on in-distribution unlabeled data. We experiment with Resnet18, Resnet50 [42] trained from scratch (*i.e.* random initialization). We do not consider off-the-shelf pretrained models (e.g., on Imagenet [68]) to avoid confounding our conclusions about contrastive pretraining. However, we note that our results on most datasets tend to be comparable to and sometimes exceed those obtained with ImageNet-pretrained models. For source-only ERM, as with other methods (FixMatch, SwAV), we default to using strong augmentation techniques: random horizontal flips, random crops, augmentation with Cutout [24], and RandAugment [21]. Moreover, unless otherwise specified, we default to full finetuning with source-only ERM, both from scratch and after contrastive pretraining, and for ST with FixMatch. For UDA, given that the setup precludes access to labeled data from the target distribution, we use source hold-out performance to pick the best hyperparameters. During pretraining, early stopping is done according to lower values of pretraining loss. For more details on datasets, model architectures, and experimental protocols, see App. C.

**Results on UDA setup.** Both ST and CL individually improve over ERM across all datasets, with CL significantly performing better than ST on 5 out of 8 benchmarks (see Table 1). Even on datasets where ST is better than CL, their performance remains close. Combining ST and CL with STOC shows a 3–8% improvement over the best alternative, yielding an absolute improvement in average accuracy of 5.2%.

Note that by default, we train with CL on the combined unlabeled data from source and target. However, to better understand the significance of unlabeled target data in contrastive pretraining, we perform an ablation where the CL model was trained solely on unlabeled source data (refer to this as CL (source only); see App. C.4). We observe that ST on top of CL (source only) improves over ST (from scratch). However, the average performance of ST over CL (source only) is similar to that of standalone CL, maintaining an approximate 6% performance gap observed between CL and ST. This brings two key insights to the fore: (i) the observed benefit is not merely a result of the contrastive pretraining objective alone, but specifically CL with unlabeled target data helps; and (ii) both CL and ST leverage using target unlabeled data in a complementary nature.

**Results on SSL setup.** While CL improves over ST (as in UDA), unlike UDA, STOC doesn’t offer any significant improvements over CL (see Table 2; ERM and ST results (refer to App. C.5)). We conduct ablation studies with varying proportions of labeled data used for SSL, illustrating that there’s considerable potential for improvement (see App. C.5). These findings highlight that the complementary nature of STOC over CL and ST individually is an artifact of distribution shift.

## 4 Theoretical Analysis and Intuitions

Our results on real-world datasets suggest that although self-training may offer little to no improvement over contrastive pretraining for in-distribution (*i.e.*, SSL) settings, it leads to substantial improvements when facing distribution shifts in UDA (Sec. 3). Why do these methods offer complementary gains, but only under distribution shifts? In this section, we seek to answer this question by first replicating all the empirical trends of interest in a simple data distribution with an intuitive story (Sec. 4.1). In this toy model, we formally characterize the gains afforded by contrastive pretraining and self-training both individually (Secs. 4.2, 4.3) and when used together (Sec. 4.4).

**Data distribution** We consider binary classification and model the inputs as consisting of two kinds of features:  $x = [x_{\text{in}}, x_{\text{sp}}]$ , where  $x_{\text{in}} \sim \mathcal{P}^{\mathcal{R}^{d_{\text{in}}}}$  is the invariant feature that is predictive of the label across both source  $\mathcal{P}_{\text{S}}$  and target  $\mathcal{P}_{\text{T}}$  and  $x_{\text{sp}} \sim \mathcal{P}^{\mathcal{R}^{d_{\text{sp}}}}$  is the spurious feature that is correlated with the label  $y$  only on the source domain  $\mathcal{P}_{\text{S}}$  but uncorrelated with label  $y$  in  $\mathcal{P}_{\text{T}}$ . Formally, we sample  $y \sim \text{Unif}\{1, -1\}$  and generate inputs  $x$  conditioned on  $y$  as follows:

$$\begin{aligned} \mathcal{P}_{\text{S}} : x_{\text{in}} &\sim \mathcal{N}(\gamma y w^*, \Sigma_{\text{in}}) & x_{\text{sp}} &\sim \mathcal{N}(y \mathbf{1}_{d_{\text{sp}}}) \\ \mathcal{P}_{\text{T}} : x_{\text{in}} &\sim \mathcal{N}(\gamma y w^*, \Sigma_{\text{in}}) & x_{\text{sp}} &\sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{sp}}), \end{aligned} \quad (1)$$

where  $\gamma$  is the margin afforded by the invariant feature<sup>2</sup>. We set the covariance of the invariant features  $\Sigma_{\text{in}} = \sigma_{\text{in}}^2 \rho \mathbf{I}_{d_{\text{in}}} + w^* w^{*\top}$ . This makes the variance along the latent predictive direction  $w^*$  to be zero. Note that the spurious feature is also completely predictive of the label in the source data. In fact, when  $d_{\text{sp}}$  is sufficiently large,  $x_{\text{sp}}$  is more predictive (than  $x_{\text{in}}$ ) of  $y$  in the source. In the target,  $x_{\text{sp}}$  is distributed as a Gaussian with  $\Sigma_{\text{sp}} = \sigma_{\text{sp}}^2 \mathbf{I}_{d_{\text{sp}}}$ . We use  $w_{\text{in}} = [w^*, 0, \dots, 0]^{\top}$  to refer to the invariant direction/feature, and  $w_{\text{sp}} = [0, \dots, 0, \frac{1}{d_{\text{sp}}} \mathbf{1}_{d_{\text{sp}}}^{\top}]^{\top}$  for the spurious direction.

**Data for UDA vs. SSL** For convenience, we assume access to infinite unlabeled data and replace their empirical quantities with population counterparts. For SSL, we sample both finite labeled and infinite unlabeled data from the same distribution  $\mathcal{P}_{\text{T}}$ , where spurious features are absent (to exclude easy-to-generalize features). For UDA, we assume infinite labeled data from  $\mathcal{P}_{\text{S}}$  and infinite unlabeled from  $\mathcal{P}_{\text{T}}$ . Importantly, note that due to distribution shift, population access of  $\mathcal{P}_{\text{S}}$  doesn’t trivialize the problem as “ERM” on infinite labeled source data *does not* achieve optimal performance on target.

**Methods and objectives** Recall from Section 2 that we learn linear classifiers  $h$  over feature extractor  $\Phi$ . For our toy setup, we consider linear feature extractors *i.e.*  $\Phi$  is a matrix in  $\mathbb{R}^{d \times k}$  and the prediction is given by  $\text{sgn}(\rho h^{\top} \Phi x)$ . We use the exponential loss  $\ell(\rho f(x), y) = \exp(\rho y f(x))$ .

*Self-training.* ST performs ERM in the first stage using labeled data from the source, and then subsequently updates the head  $h$  by iteratively generating pseudolabels on the unlabeled target:

$$\begin{aligned} L_{\text{st}}(\rho h; \Phi) &: \mathbb{E}_{\mathcal{P}_{\text{T}}(x)} \ell(\rho h^{\top} \Phi x, \text{sgn}(\rho h^{\top} \Phi x)) \\ \text{Update: } h^{t+1} &: \frac{h^t - \eta \nabla_{h^t} L_{\text{st}}(\rho h^t; \Phi)}{\|h^t - \eta \nabla_{h^t} L_{\text{st}}(\rho h^t; \Phi)\|_2} \end{aligned} \quad (2)$$

For ERM and ST, we train both  $h$  and  $\Phi$  (equivalent to  $\Phi$  being identity and training a linear head).

*Contrastive pretraining.* We obtain  $\Phi_{\text{cl}} = \arg \min_{\Phi} L_{\text{cl}}(\rho \Phi)$  by minimizing the Barlow Twins objective [93], which prior works have shown is also equivalent to spectral contrastive and non-contrastive objectives [33, 11]. Given probability distribution  $\mathcal{P}_{\text{A}}(a | x)$  for input  $x$ , and marginal  $\mathcal{P}_{\text{A}}$ , we consider a constrained form of Barlow Twins in (3) which enforces features of “positive pairs”  $a_1, a_2$  to be close while ensuring feature diversity. We assume a strict regularization  $\rho > 0$  for the theory arguments in the rest of the paper, and in App. D.2 we prove that all our claims hold for small  $\rho$  as well. For augmentations, we scale the magnitude of each co-ordinate uniformly by an independent amount, *i.e.*,  $a \sim \mathcal{P}_{\text{A}}(a | x) \sim \mathcal{C}(c, x)$ , where  $c \sim \text{Unif}(0, 1)^d$ . We try to mirror practical settings where the augmentations are fairly “generic”, not encoding information about which features are invariant or spurious, and hence perturb all features symmetrically.

$$\begin{aligned} L_{\text{cl}}(\rho \Phi) &: \mathbb{E}_{x \sim \mathcal{P}_{\text{U}}} \mathbb{E}_{a_1, a_2 \sim \mathcal{P}_{\text{A}}(\cdot | x)} \|\Phi \rho a_1 - \Phi \rho a_2\|_2^2 \\ \text{s.t. } & \mathbb{E}_{a \sim \mathcal{P}_{\text{A}}} \Phi \rho a \Phi \rho a^{\top} = \mathbf{I}_k \quad \text{tr} \Phi \rho a \Phi \rho a^{\top} \propto \rho \end{aligned} \quad (3)$$

<sup>2</sup>See App. D.1 for similarities and differences of our setup with prior works.

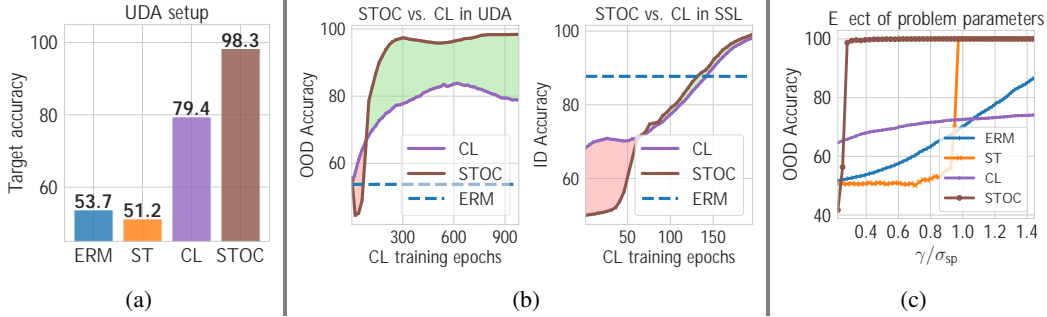


Figure 2: *Our simplified model of shift captures real-world trends and theoretical behaviors:* (a) Target (OOD) accuracy separation in the UDA setup (for problem parameters in Example 1). (b) Comparison of the benefits of STOC (ST over CL) over just CL in UDA and SSL settings, done across training iterations for contrastive pretraining. (c) Comparison between different methods in UDA setting, as we vary problem parameters  $\gamma$  and  $\sigma_{sp}$ , connecting our theory results in Sec. 4.

Keeping the  $\Phi_{cl}$  fixed, we then learn a linear classifier  $h_{cl}$  over  $\Phi_{cl}$  to minimize the exponential loss on labeled source data (refer to as *linear probing*). For STOC, keeping the  $\Phi_{cl}$  fixed and initializing the linear head with the CL linear probe (instead of source only ERM), we perform ST with (2).

**Example 1.** For the setup in (1), we chose  $\gamma = 0.5$ ,  $\sigma_{sp}^2 = 1.$ , and  $\sigma_{in}^2 = 0.05$  with  $d_{in} = 5$  and  $d_{sp} = 20$  for our running example.  $\gamma/\sigma_{sp}$  controls signal to noise ratio in the source such that spurious feature is easy-to-learn and the invariant feature is harder-to-learn.  $\sigma_{sp}$  controls the noise in target which we show later is critical in unlearning the spurious feature with CL.

#### 4.1 Simulations and Intuitive Story: A Comparative Study Between SSL and DA

**Our setup captures real-world trends in UDA setting.** Our toy setup (in Example 1) accentuates the behaviors observed on real-world datasets (Fig. 2(a)): (i) both ERM and ST yield close to random performance (though ST performs slightly worse than ERM); (ii) CL improves over ERM but still yields sub-optimal target performance; (iii) STOC then further improves over CL, achieving near-optimal target performance. Note that, a linear predictor can improve target performance only by reducing its dependence on spurious feature  $x_{sp}$ , and increasing it on invariant feature  $x_{in}$  (along  $w^*$ ). Given this, we can explain our trends if we understand the following: (i) how ST reduces dependence on spurious feature when done after CL; (ii) why CL helps reduce but not completely eliminate the reliance of linear head on spurious features. Before we present intuitions, we ablate over a key problem parameter that affects both the target performance and conditions for ST to work.

**Effect of  $\gamma/\sigma_{sp}$  on success of ST.** By increasing the ratio of margin  $\gamma$  and variance of spurious feature on target  $\sigma_{sp}$  (keeping others constant), the problem becomes easier because  $\gamma$  directly affects the signal on  $x_{in}$  and reducing  $\sigma_{sp}$  helps ST to unlearn  $x_{sp}$  (see App. D.3). In Fig. 2(c), we see that a phase transition occurs for ST, *i.e.*, after a certain threshold of  $\gamma/\sigma_{sp}$ , ST successfully recovers the optimal target predictor. This hints that ST has a binary effect, where beyond a certain magnitude of  $\gamma/\sigma_{sp}$ , ST can amplify the signal on domain invariant feature to obtain optimal target predictor. On the other hand, the performance of CL and ERM improve gradually where CL achieves high performance even at small ratios of  $\gamma/\sigma_{sp}$ . One way of viewing this trend with CL is that it magnifies the effective  $\gamma/\sigma_{sp}$  in its representation space, because of which a linear head trained these representations have a good performance at low values of the ratio. Consequently, the *phase transition* of STOC occurs much sooner than that of ST. Finally, we note that for CL the rate of performance increase diminishes at high values of  $\gamma/\sigma_{sp}$  because CL fails to reduce dependency along  $x_{sp}$  beyond a certain point.

**An intuitive story.** We return to the question of why self-training improves over contrastive learning under distribution shift in our Example 1. When the classifier at initialization of ST relies more on spurious features, ST aggravates this dependency. However, as the problem becomes easier (with increasing  $\gamma/\sigma_{sp}$ ), the source-only ERM classifier will start relying more on invariant rather than spurious feature. Once this ERM classifier is sufficiently accurate on the target, ST unlearns any dependency on spurious features achieving optimal target performance. In contrast, we observe that CL performs better than ERM but is still sub-optimal. This implies that CL ends up decreasing reliance on spurious features (as compared to ERM) but doesn't completely eliminate them. Combining ST and CL, a natural hypothesis explaining our trends is that CL provides a "favorable" initialization

for ST by sufficiently increasing signal on invariant features. Our empirical findings emphasize an intriguing contrast suggesting that ST and CL improve target performance in complementary ways.

**Why disparate behaviors for out-of-distribution vs. in distribution?** In the SSL setup, recall, there is no distribution shift. In Example 1, we sample  $50k$  unlabeled data and 100 labeled data from the same (target) distribution to simulate SSL setup. Substantiating our findings on real-world data, we observe that STOC provides a small to negligible improvement over CL (refer to App. D). To understand why such disparate behaviors emerge, recall that in the UDA setting, the main benefit of STOC lies in picking up reliance on “good” features for OOD data, facilitated by CL initialization. While contrastive pretraining uncovers features that are “good” for OOD data, it also learns more predictive source-only features (which are not predictive at all on target). As a result, linear probing with source-labeled data picks up these source-only features, leaving considerable room for improvement on OOD data with further self-training. On the other hand, in the SSL setting, the limited ID labeled data might provide enough signal to pick up features predictive on ID data, leaving little to no room for improvement for further self-training. Corroborating our intuitions, throughout the CL training in the toy setup, when CL doesn’t achieve near-perfect generalization, the improvements provided by STOC for each checkpoint remain minimal. On the other hand, for UDA setup, after reaching a certain training checkpoint in CL, STOC yields significant improvement (Fig. 2(b)).

In the next sections, we formalize our intuitions and analyze why ST and CL offer complementary benefits when dealing with distribution shifts. Formal statements and proofs are in App. E.

#### 4.2 Conditions for Success and Failure of Self-training over ERM from Scratch

In our results on Example 1, we observe that performing ST after ERM yields a classifier with near-random target accuracy. In Theorem 2, we characterize conditions under which ST fails and succeeds.

**Theorem 2** (Informal; Conditions for success and failure of ST over ERM). *The target accuracy of ERM classifier, is given by  $0.5 + \text{erfc}^{-1}(\gamma^2 \sqrt{\frac{1}{2d_{sp} \cdot \sigma_{sp}^2}})$ . Then ST performed in the second stage yields: (i) a classifier with  $0.5$  target accuracy when  $\gamma \leq 2\sigma_{sp}$  and  $\sigma_{sp} \neq 1$ ; and (ii) a classifier with near-perfect target accuracy when  $\gamma \neq \sigma_{sp}$ .*

The informal theorem above abstracts the exact dependency of  $\gamma$ ,  $\sigma_{sp}$ , and  $d_{sp}$  for the success and failure of ST over ERM. Our analysis highlights that while ERM learns a perfect predictor along  $w_{in}$  (with norm  $\gamma$ ), it also learns to depend on  $w_{sp}$  (with norm  $\frac{\gamma}{d_{sp}}$ ) because of the perfect correlation of  $x_{sp}$  with labels on the source. Our conditions depict that when the  $\gamma/\sigma_{sp}$  is sufficiently small, then ST continues to erroneously enhance its reliance on the  $x_{sp}$  feature for target prediction, resulting in near-random target performance. Conversely, when  $\gamma/\sigma_{sp}$  is larger than 1, the signal in  $x_{in}$  is correctly used for predictor on the majority of target points, and ST eliminates the  $x_{sp}$  dependency, converging to an optimal target classifier.

Our proof analysis shows that if the ratio of the norm of the classifier along in the direction of  $w^*$  is smaller than  $w_{sp}$  by a certain ratio then the generated pseudolabels (incorrectly) use  $x_{sp}$  for its prediction further increasing the component along  $w_{sp}$ . Moreover, normalization further diminishes the reliance along  $w^*$ , culminating in a near-random performance. The opposite occurs when the ERM classifier achieves a signal along  $w^*$  that is sufficiently stronger than along  $w_{sp}$ . Upon substituting the parameters used in Example 1, the ERM and ST performances as determined by Theorem 2 align with our empirical results, notably, ST performance on target being near-random.

#### 4.3 CL Captures Both Features But Amplifies Invariant Over Spurious Features

Here we show that minimizing the contrastive loss (3) on unlabeled data from both  $P_S$  and  $P_T$  gives us a feature extractor  $\Phi_{cl}$  that has a higher inner product with the invariant feature over the spurious feature. First, we derive a closed form expression for  $\Phi_{cl}$  that holds for any linear backbone and augmentation distribution. Then, we introduce assumptions on the augmentation distribution (or equivalently on  $w^*$ ) and other problem parameters, that are sufficient to prove amplification.

**Proposition 3** (Barlow Twins solution). *The solution for (3) is  $U_k^T \Sigma_A^{-1/2}$  where  $U_k$  are the top  $k$  eigenvectors of  $\Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2}$ . Here,  $\Sigma_A : E_{a \sim P_A} [aa^T]$  is the covariance over augmentations, and  $\Sigma : E_{x \sim P_U} [xx^T]$  is the covariance matrix of mean augmentations  $\mathbb{E}_{P_A(a|x)} [aa^T]$ .*

The above result captures the effect of augmentations through the matrix  $U_k$ . If there were no augmentations, then  $\Sigma_A = \Sigma$ , implying that  $U_k$  could then be any random orthonormal matrix. On

the other hand if augmentation distributions change prevalent covariances in the data, *i.e.*,  $\Sigma_A$  is very different from  $\Sigma$ , the matrix  $U_k$  would bias the CL solution towards directions that capture significant variance in marginal distribution on augmented data, but have low conditional variance, when conditioned on original point  $x$ —precisely the directions with low invariance loss. Hence, we can expect that CL would learn components along both invariant  $w_{in}$  and spurious  $w_{sp}$  because: (i) these directions explain a large fraction of variance in the raw data; (ii) augmentations that randomly scale down dimensions would add little variance along  $w_{sp}$  and  $w_{in}$  compared to noise directions in their null space. On the other hand it is unclear which of these directions is amplified more in  $\Phi_{cl}$ . The following assumption and amplification result conveys that when the noise in target  $p\sigma_{sp}q$  is sufficiently large, the CL solution amplifies the invariant feature over the spurious feature.

**Assumption 4** (Informal; Alignment of  $w^*$  with augmentations). *We assume that  $w^*$  aligns with  $P_{Ap} \mid xq$ , *i.e.*,  $\langle w^*, E_{a|x} r a^\top w^* \rangle \gg \langle w^*, x^\top \text{diag} p \mathbb{1}_a q w^* \rangle$  is high. Hence, we assume  $w^* \approx \mathbb{1}_{d_{in}} \{\sqrt{d_{in}}\}$ .*

One implication of Assumption 4 is that when  $w^* \approx \mathbb{1}_{d_{in}} \{\sqrt{d_{in}}\}$ , only the top two eigenvectors lie in the space spanned by  $w_{in}$  and  $w_{sp}$ . To analyze our amplification with fewer eigenvectors from Proposition 3 while retaining all relevant phenomena, we assume  $w^* \approx \mathbb{1}_{d_{in}} \{\sqrt{d_{in}}\}$  for mathematical convenience. While Assumption 4 permits a tighter theoretical analysis, our empirical results in Sec. 4.1 hold more generally for  $w^* \sim N(p0, \mathbf{I}_{d_{in}})q$ .

**Theorem 5** (Informal; CL recovers both  $w_{in}$  and  $w_{sp}$  but amplifies  $w_{in}$ ). *Under Assumption 4, the CL solution  $\Phi_{cl} = [\phi_1, \phi_2, \dots, \phi_k]^\top$  satisfies  $\phi_j^\top w_{in} \approx \phi_j^\top w_{sp} = 0 \ @j \neq 3, \phi_1 \approx c_1 w_{in} + c_3 w_{sp}$  and  $\phi_2 \approx c_2 w_{in} + c_4 w_{sp}$ . For constants  $K_1, K_2 \geq 0, \gamma \in [K_1 K_2 \{\sigma_{sp}, d_{sp}\}^{-2}, \infty[$ ,  $\sigma_{sp} \geq \sigma_{sp0}$ , such that for  $\sigma_{sp} \neq \sigma_{sp0}, |c_3 - K_1 K_2^2 d_{in} \{2L\sigma_{in}^2 (d_{in} - 1)\}| \leq \epsilon$ , and  $|c_4 - L \frac{d_{sp}}{d_{in}} \{\gamma\}| \leq \epsilon$ , where  $L = 1 + K_2^2$ .*

We analyze the amplification of  $w_{in}$  with contrastive learning in the regime where  $\sigma_{sp}$  is large enough. In other words, if the target distribution has sufficient noise along the spurious feature, the augmentations prevent the CL solution from extracting components along  $w_{sp}$ . Thus, in our analysis, we first analyze the amplification factors asymptotically  $p\sigma_{sp} \gg \sigma_{sp0}$ , and then use the asymptotic behavior to draw conclusions for the regime where  $\sigma_{sp}$  is large but finite.

Theorem 5 conveys two results: (i) CL recovers components along both  $w_{in}$  and  $w_{sp}$  through  $\phi_1, \phi_2$ ; and (ii) it increases the norm along  $w_{in}$  more than  $w_{sp}$ . The latter is evident because the margin separating labeled points along  $w_{in}$  is now amplified by a factor of  $|c_2 c_4| \approx \Omega(p^L \frac{d_{sp}}{d_{in}} \{\gamma\}q$  in  $\phi_2$ . Naturally, this will improve the target performance of a linear predictor trained over CL representations. At the same time, we also see that in  $\phi_1$ , the component along  $w_{sp}$  is still significant ( $c_3 \approx \text{Op}\{L\sigma_{in}^2\}q$ ). Intuitively, CL prefers the invariant feature since augmentations amplify the noise along  $w_{sp}$  in the target domain. At the same time, the variance induced by augmentations along  $w_{sp}$  in source is still very small due to which the dependence on  $w_{sp}$  is not completely alleviated. Due to the remaining components along  $w_{sp}$ , the target performance for CL can remain less than ideal. Both the above arguments on target performance are captured in Corollary 6.

**Corollary 6** (Informal; CL improves OOD error over ERM but is still imperfect). *For  $\gamma, \sigma_{sp}, d_{sp}$  defined as in Theorem 5,  $\sigma_{sp} \geq \sigma_{sp1}$  such that for all  $\sigma_{sp} \neq \sigma_{sp1}$ , the target accuracy of CL (linear predictor on  $\Phi_{cl}$ ) is  $\geq 0.5 \text{erfc} p \frac{L'}{\gamma \sqrt{2\sigma_{sp}q}}$  and  $\leq 0.5 \text{erfc} p \frac{4L'}{\gamma \sqrt{2\sigma_{sp}q}}$ , where  $L' = K_1^2 K_2 \{\sigma_{in}^2 (1 - 1/d_{in})\}$ . When  $\sigma_{sp1} \leq \sigma_{in} \frac{1}{\sqrt{d_{in}}}$ , the lower bound on accuracy is strictly better than ERM from scratch.*

While  $\Phi_{cl}$  is still not ideal for linear probing, in the next part we will see how  $\Phi_{cl}$  can instead be sufficient for subsequent self-training to unlearn the remaining components along spurious features.

#### 4.4 Improvements with Self-training Over Contrastive Learning

The result in the previous section highlights that while CL may improve over ERM, the linear probe continues to depend on the spurious feature. Next, we characterize the behavior STOC. Recall, in the ST stage, we iteratively update the linear head with (2) starting with the CL backbone and head.

**Theorem 7** (Informal; ST improves over CL). *Under the conditions of Theorem 5 and  $d_{sp} \geq K_1^2 + K_2^{2/3}$ , the target accuracy of ST over CL is lower bounded by  $0.5 \text{erfc} p \frac{c_2 c_4}{\gamma \sqrt{2\sigma_2}q}$  and  $\leq 0.5 \text{erfc} p \frac{L}{d_{sp} \{\sqrt{2\sigma_{sp}}\}}$  where  $c_2$  and  $c_4$  are the coefficients of feature  $\phi_2$  along  $w_{in}$  and  $w_{sp}$  learned by BT.*



The above theorem states that when  $\frac{1}{\sigma_{sp}} \ll 1$  the target accuracy of ST over CL is close to 1. In Example 1, the lower bound of the accuracy of ST over CL is  $\text{erfc}\left(\frac{1}{\sqrt{10}}\right) \approx 0.2$  showing near-perfect target generalization. Recall that Theorem 6 shows that CL yields a linear head that mainly depends on both the invariant direction  $w_{in}$  and the spurious direction  $w_{sp}$ . At initialization, the linear head trained on the CL backbone has negligible dependence on  $\phi_2$  (under conditions in Theorem 6). Building on that, the analysis in Theorem 7 captures that ST gradually reduces the dependence on  $w_{sp}$  by learning a linear head that has a larger reliance on  $\phi_2$ , which has a higher “effective” margin on the target, thus increasing overall dependency on  $w_{in}$ .

**Theoretical comparison with SSL.** Our analysis until now shows that linear probing with source labeled data during CL picks up features that are more predictive of source label under distribution shift, leaving a significant room for improvement on OOD data when self-trained further. In UDA, the primary benefit of ST lies in picking up the features with a high “effective” margin on target data that are not picked up by linear head trained during CL. In contrast, in the SSL setting, the limited ID labeled data may provide enough signal in picking up high-margin features that are predictive on ID data, leaving little to no room for improvement for further ST. We formalize this intuition in App. E.

#### 4.5 Reconciling Practice: Implications for Deep Non-Linear Networks

In this section, we experiment with deep non-linear backbone (*i.e.*,  $\Phi_{cl}$ ). When we continue to fix  $\Phi_{cl}$  during CL and STOC, the trends we observed with linear networks in Sec. 4.1 continue to hold. We then perform full fine-tuning with CL and STOC, *i.e.*, propagate gradients even to  $\Phi_{cl}$ , as commonly done in practice. We present key takeaways here but detailed experiments are in App. D.4.

**Benefits of augmentation for self-training.** ST while updating  $\Phi_{cl}$  can hurt due to overfitting issues when training with the finite sample of labeled and unlabeled data (drop by >10% over CL). This is due to the ability of deep networks to overfit on confident but incorrect pseudolabels on target data [94]. This exacerbates components along  $w_{sp}$  and we find that augmentations (and other heuristics) typically used in practice (*e.g.* in FixMatch [79]) help avoid overfitting on incorrect pseudolabels.

**Can ERM and ST over contrastive pretraining improve features?** We find that self-training can also slightly improve features when we update the backbone with the second stage of STOC and when the CL backbone is early stopped sub-optimally (*i.e.* at an earlier checkpoint in Fig. 2(b)). This feature finetuning can now widen the gap between STOC and CL in SSL settings, as compared to the linear probing gap (as in 2). This is because STOC can now improve performance beyond just recovering the generalization gap for the linear head (which is typically small). However, STOC benefits are negligible when CL is not early stopped sub-optimally, *i.e.*, trained till convergence. Thus, it remains unclear if STOC and CL have complementary benefits for feature learning in UDA or SSL settings. Investigating this is an interesting avenue for future work.

### 5 Connecting Experimental Gains with Theoretical Insights

Our theory emphasizes that under distribution shift contrastive pretraining improves the representations for target data, while self-training primarily improves linear classifiers learned on top. To investigate different methods in our UDA setup, we study the representations learned by each of them. We fix the representations and train linear heads over them to answer two questions: (i) How good are the representations in terms of their *ceiling* of target accuracy (performance of the optimal linear probe)?—we evaluate this by training the classifier head on target labeled data (*i.e.*, target linear probe); and (ii) How well do heads trained on source generalize to target?—we assess this by training a head on source labeled data (source linear probe) and evaluate its difference with target linear probe. For both, we plot target accuracy. We make two *intriguing* observations Fig. 3):

**Does CL improve representations over ERM features?** Yes. We observe a substantial difference in accuracy (14% gap) of target linear probes on backbones trained




Figure 3: *Target accuracy with source and target linear probes*, which freezes backbones trained with various objectives and trains only the head in UDA setup. Avg. accuracy across all datasets. We observe that: (i) ST improves the linear transferability of source probes, and (ii) CL improves representations.

with contrastive pretraining (*i.e.* CL, STOC) and without it (*i.e.*, ERM, ST) highlighting that CL significantly pushes the performance ceiling over non-contrastive features. As a side, our findings also stand in contrast to recent studies suggesting that ERM features might be “good enough” for OOD generalization [67, 46]. Instead, the observed gains with contrastively pretrained backbones (*i.e.* CL, STOC) demonstrate that target unlabeled data can be leveraged to further improve over ERM features.

**Do CL features yield *perfect* linear transferability from source to target?** Recent works [40, 76] conjecture that under certain conditions CL representations, linear probes learned with source labeled data may transfer perfectly from source to target. However, we observe that this doesn’t hold strictly in practice, and in fact, the linear transferability can be further improved with ST. We first note a significant gap between the performance of source linear probes and target linear probes illustrating that linear transferability is not perfect in practice. Moreover, while the accuracy of target linear probes doesn’t change substantially between CL and STOC, the accuracy of the source linear probe improves significantly. Similar observations hold for ERM and ST, methods trained without contrastive pretraining. This highlights that ST performs “feature refinement” to improve source to target linear transfer (with relatively small improvements in their respective target probe performance). *The findings highlight the complementary nature of benefits on real-world data: ST improves linear transferability while CL improves representations.*

## 6 Connections to Prior Work

Our empirical results and our analyses offer a perspective that contrasts with the prior literature that argues for the individual optimality of contrastive pretraining and self-training. We outline the key differences from existing studies here, and delve into other related works in App. A.

**Limitations of prior work analyzing contrastive learning** Prior works [40, 76] analyzing CL first make assumptions on the consistency of augmentations with labels [39, 11, 73, 44], and specifically for UDA make stronger ones on the augmentation graph connecting examples from same domain or class more than cross-class/cross-domain ones. While this is sufficient to prove linear transferability, it is unclear if this holds in practice when augmentations are imperfect, *i.e.* if they fail to mask the spurious features completely—as corroborated by our findings in Sec. 5. We show why this also fails in our simplified setup in App. F.1.

**Limitations of prior work analyzing self-training** Prior research views self-training as consistency regularization, ensuring pseudolabels for original samples align with their augmentations [12, 87, 79]. This approach abstracts the role played by the optimization algorithm and instead evaluates the global minimizer of a population objective promoting pseudolabel consistency. It also relies on specific assumptions about class-conditional distributions to guarantee pseudolabel accuracy across domains. However, this framework doesn’t address issues in iterative label propagation. For example, when augmentation distribution has long tails, the consistency of pseudolabels depends on the sampling frequency of “favorable” augmentations (for more discussion see App. F.2). Our analysis thus follows the iterative examination of self-training [17].

## 7 Conclusion

In this study, we highlight the synergistic behavior of self-training and contrastive pretraining under distribution shift. Shifts in distribution are commonplace in real-world applications of machine learning, and even under natural, non-adversarial distribution shifts, the performance of machine learning models often drops. By simply combining existing techniques in self-training and contrastive learning, we find that we can improve accuracy by 3–8% rather than using either approach independently. Despite these significant improvements, we note that one limitation of this combined approach is that performing self-training sequentially after contrastive pretraining increases the computation cost for UDA. The potential for integrating these benefits into one unified training paradigm is yet unclear, presenting an interesting direction for future exploration.

Beyond this, we note that our theoretical framework primarily confines the analysis to training the backbone and linear network independently during the pretraining and fine-tuning/self-training phases. Although our empirical observations apply to deep networks with full fine-tuning, we leave a more rigorous theoretical study of full fine-tuning for future work. Our theory also relies on a covariate shift assumption (where we assume that label distribution also doesn’t shift). Investigating the complementary nature of self-training and contrastive pretraining beyond the covariate shift assumption would be another interesting direction for future work.

## Acknowledgements

SG acknowledges the JP Morgan AI Ph.D. Fellowship and Bloomberg Ph.D. Fellowship for their support.

## References

- [1] Alexandari, A., Kundaje, A., and Shrikumar, A. (2021). Adapting to label shift with bias-corrected calibration. In *International Conference on Machine Learning (ICML)*.
- [2] Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. (2019). A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*.
- [3] Azizzadenesheli, K., Liu, A., Yang, F., and Anandkumar, A. (2019). Regularized learning for domain adaptation under label shifts. In *International Conference on Learning Representations (ICLR)*.
- [4] Bardes, A., Ponce, J., and LeCun, Y. (2021). Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*.
- [5] Baricz, Á. (2008). Mills’ ratio: Monotonicity patterns and functional inequalities. *Journal of Mathematical Analysis and Applications*, 340(2):1362–1370.
- [6] Bekker, J. and Davis, J. (2020). Learning from positive and unlabeled data: a survey. *Machine Learning*.
- [7] Ben-David, S., Lu, T., Luu, T., and Pál, D. (2010). Impossibility Theorems for Domain Adaptation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [8] Berthelot, D., Carlini, N., Cubuk, E. D., Kurakin, A., Sohn, K., Zhang, H., and Raffel, C. (2019). Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*.
- [9] Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- [10] Blanchard, G., Lee, G., and Scott, C. (2011). Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems*, 24.
- [11] Cabannes, V., Kiani, B. T., Balestrieri, R., LeCun, Y., and Bietti, A. (2023). The ssl interplay: Augmentations, inductive bias, and generalization. *arXiv preprint arXiv:2302.02774*.
- [12] Cai, T., Gao, R., Lee, J., and Lei, Q. (2021). A theory of label propagation for subpopulation shift. In *International Conference on Machine Learning*, pages 1170–1182. PMLR.
- [13] Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. (2020). Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924.
- [14] Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., and Joulin, A. (2021). Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660.
- [15] Chapelle, O., Scholkopf, B., and Zien, A. (2006). Semi-supervised learning. 2006. *Cambridge, Massachusetts: The MIT Press View Article*, 2.
- [16] Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. (2020a). A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- [17] Chen, X., Chen, W., Chen, T., Yuan, Y., Gong, C., Chen, K., and Wang, Z. (2020b). Self-pu: Self boosted and calibrated positive-unlabeled training. In *International Conference on Machine Learning*, pages 1510–1519. PMLR.

- [18] Christie, G., Fendley, N., Wilson, J., and Mukherjee, R. (2018). Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [19] Cortes, C., Mansour, Y., and Mohri, M. (2010). Learning Bounds for Importance Weighting. In *Advances in Neural Information Processing Systems (NIPS)*.
- [20] Cortes, C. and Mohri, M. (2014). Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519.
- [21] Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. (2020). Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703.
- [22] Darlow, L. N., Crowley, E. J., Antoniou, A., and Storkey, A. J. (2018). Cinic-10 is not imagenet or cifar-10. *arXiv preprint arXiv:1810.03505*.
- [23] Deledalle, C.-A., Denis, L., Tabti, S., and Tupin, F. (2017). *Closed-form expressions of the eigen decomposition of  $2 \times 2$  and  $3 \times 3$  Hermitian matrices*. PhD thesis, Université de Lyon.
- [24] DeVries, T. and Taylor, G. W. (2017). Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- [25] Elkan, C. and Noto, K. (2008). Learning classifiers from only positive and unlabeled data. In *International Conference Knowledge Discovery and Data Mining (KDD)*, pages 213–220.
- [26] Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. (2016). Domain-adversarial training of neural networks. *The journal of machine learning research*.
- [27] Gardner, J., Pleiss, G., Weinberger, K. Q., Bindel, D., and Wilson, A. G. (2018). Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [28] Garg, S., Balakrishnan, S., and Lipton, Z. (2022a). Domain adaptation under open set label shift. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [29] Garg, S., Balakrishnan, S., Lipton, Z., Neyshabur, B., and Sedghi, H. (2022b). Leveraging unlabeled data to predict out-of-distribution performance. In *International Conference on Learning Representations (ICLR)*.
- [30] Garg, S., Erickson, N., Sharpnack, J., Smola, A., Balakrishnan, S., and Lipton, Z. (2023). Rlsbench: A large-scale empirical study of domain adaptation under relaxed label shift. In *International Conference on Machine Learning (ICML)*.
- [31] Garg, S., Wu, Y., Balakrishnan, S., and Lipton, Z. (2020). A unified view of label shift estimation. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [32] Garg, S., Wu, Y., Smola, A., Balakrishnan, S., and Lipton, Z. (2021). Mixture proportion estimation and PU learning: A modern approach. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [33] Garrido, Q., Chen, Y., Bardes, A., Najman, L., and Lecun, Y. (2022). On the duality between contrastive and non-contrastive self-supervised learning. *arXiv preprint arXiv:2206.02574*.
- [34] Grandvalet, Y. and Bengio, Y. (2006). Entropy regularization.
- [35] Gretton, A., Smola, A. J., Huang, J., Schmittfull, M., Borgwardt, K. M., and Schölkopf, B. (2009). Covariate Shift by Kernel Mean Matching. *Journal of Machine Learning Research (JMLR)*.
- [36] Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al. (2020). Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284.

- [37] Gulrajani, I. and Lopez-Paz, D. (2020). In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*.
- [38] HaoChen, J. Z. and Ma, T. (2022). A theoretical study of inductive biases in contrastive learning. *arXiv preprint arXiv:2211.14699*.
- [39] HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. (2021). Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34:5000–5011.
- [40] HaoChen, J. Z., Wei, C., Kumar, A., and Ma, T. (2022). Beyond separability: Analyzing the linear transferability of contrastive representations to related subpopulations. *arXiv preprint arXiv:2204.02683*.
- [41] He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738.
- [42] He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep Residual Learning for Image Recognition. In *Computer Vision and Pattern Recognition (CVPR)*.
- [43] Joachims, T. et al. (1999). Transductive inference for text classification using support vector machines. In *Icml*, volume 99, pages 200–209.
- [44] Johnson, D. D., Hanchi, A. E., and Maddison, C. J. (2022). Contrastive learning can find an optimal basis for approximately view-invariant functions. *arXiv preprint arXiv:2210.01883*.
- [45] Kakade, S. M., Sridharan, K., and Tewari, A. (2008). On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. *Advances in neural information processing systems*, 21.
- [46] Kirichenko, P., Izmailov, P., and Wilson, A. G. (2022). Last layer re-training is sufficient for robustness to spurious correlations. *arXiv preprint arXiv:2204.02937*.
- [47] Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Gao, I., Lee, T., David, E., Stavness, I., Guo, W., Earnshaw, B. A., Haque, I. S., Beery, S., Leskovec, J., Kundaje, A., Pierson, E., Levine, S., Finn, C., and Liang, P. (2021). WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning (ICML)*.
- [48] Krizhevsky, A. and Hinton, G. (2009). Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer.
- [49] Kschischang, F. R. (2017). The complementary error function. *Online, April*.
- [50] Kumar, A., Ma, T., and Liang, P. (2020). Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*, pages 5468–5479. PMLR.
- [51] Kumar, A., Raghunathan, A., Jones, R. M., Ma, T., and Liang, P. (2022). Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*.
- [52] Lee, D.-H. et al. (2013). Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896.
- [53] Lipton, Z. C., Wang, Y.-X., and Smola, A. (2018). Detecting and Correcting for Label Shift with Black Box Predictors. In *International Conference on Machine Learning (ICML)*.
- [54] Long, M., Cao, Y., Wang, J., and Jordan, M. (2015). Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR.
- [55] Long, M., Zhu, H., Wang, J., and Jordan, M. I. (2017). Deep transfer learning with joint adaptation networks. In *International conference on machine learning*. PMLR.

- [56] Loshchilov, I. and Hutter, F. (2016). Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- [57] Ma, M. Q., Tsai, Y.-H. H., Liang, P. P., Zhao, H., Zhang, K., Salakhutdinov, R., and Morency, L.-P. (2021). Conditional contrastive learning for improving fairness in self-supervised learning. *arXiv preprint arXiv:2106.02866*.
- [58] Mishra, S., Saenko, K., and Saligrama, V. (2021). Surprisingly simple semi-supervised domain adaptation with pretraining and consistency. *arXiv preprint arXiv:2101.12727*.
- [59] Muandet, K., Balduzzi, D., and Schölkopf, B. (2013). Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18. PMLR.
- [60] Nagarajan, V., Andreassen, A., and Neyshabur, B. (2020). Understanding the failure modes of out-of-distribution generalization. *arXiv preprint arXiv:2010.15775*.
- [61] Oord, A. v. d., Li, Y., and Vinyals, O. (2018). Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- [62] Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., and Wang, B. (2019). Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415.
- [63] Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. (2017). Visda: The visual domain adaptation challenge.
- [64] Peng, X., Usman, B., Saito, K., Kaushik, N., Hoffman, J., and Saenko, K. (2018). Syn2real: A new benchmark for synthetic-to-real visual domain adaptation. *arXiv preprint arXiv:1806.09755*.
- [65] Quinero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. D. (2008). *Dataset shift in machine learning*. Mit Press.
- [66] Roberts, M., Mani, P., Garg, S., and Lipton, Z. (2022). Unsupervised learning under latent label shift. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [67] Rosenfeld, E., Ravikumar, P., and Risteski, A. (2022). Domain-adjusted regression or: Erm may already learn features sufficient for out-of-distribution generalization. *arXiv preprint arXiv:2202.06856*.
- [68] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252.
- [69] Saerens, M., Latinne, P., and Decaestecker, C. (2002). Adjusting the Outputs of a Classifier to New a Priori Probabilities: A Simple Procedure. *Neural Computation*.
- [70] Sagawa, S., Koh, P. W., Lee, T., Gao, I., Xie, S. M., Shen, K., Kumar, A., Hu, W., Yasunaga, M., Marklund, H., Beery, S., David, E., Stavness, I., Guo, W., Leskovec, J., Saenko, K., Hashimoto, T., Levine, S., Finn, C., and Liang, P. (2021). Extending the wilds benchmark for unsupervised adaptation. In *NeurIPS Workshop on Distribution Shifts*.
- [71] Sagawa, S., Raghunathan, A., Koh, P. W., and Liang, P. (2020). An investigation of why overparameterization exacerbates spurious correlations. In *International Conference on Machine Learning*, pages 8346–8356. PMLR.
- [72] Santurkar, S., Tsipras, D., and Madry, A. (2021). Breeds: Benchmarks for subpopulation shift. In *International Conference on Learning Representations (ICLR)*.
- [73] Saunshi, N., Ash, J., Goel, S., Misra, D., Zhang, C., Arora, S., Kakade, S., and Krishnamurthy, A. (2022). Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning*, pages 19250–19286. PMLR.
- [74] Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., and Mooij, J. (2012). On Causal and Anticausal Learning. In *International Conference on Machine Learning (ICML)*.

- [75] Scudder, H. (1965). Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371.
- [76] Shen, K., Jones, R. M., Kumar, A., Xie, S. M., HaoChen, J. Z., Ma, T., and Liang, P. (2022). Connect, not collapse: Explaining contrastive learning for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 19847–19878. PMLR.
- [77] Shimodaira, H. (2000). Improving Predictive Inference Under Covariate Shift by Weighting the Log-Likelihood Function. *Journal of Statistical Planning and Inference*.
- [78] Shu, R., Bui, H. H., Narui, H., and Ermon, S. (2018). A dirt-t approach to unsupervised domain adaptation. *arXiv preprint arXiv:1802.08735*.
- [79] Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C. A., Cubuk, E. D., Kurakin, A., and Li, C.-L. (2020). Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33.
- [80] Stewart, G. W. (1993). On the early history of the singular value decomposition. *SIAM review*, 35(4):551–566.
- [81] Sun, B., Feng, J., and Saenko, K. (2017). Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*. Springer.
- [82] Sun, B. and Saenko, K. (2016). Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*. Springer.
- [83] Torralba, A. and Efros, A. A. (2011). Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528. IEEE.
- [84] Van Engelen, J. E. and Hoos, H. H. (2020). A survey on semi-supervised learning. *Machine learning*, 109(2):373–440.
- [85] Venkateswara, H., Eusebio, J., Chakraborty, S., and Panchanathan, S. (2017). Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5018–5027.
- [86] Wang, D., Shelhamer, E., Liu, S., Olshausen, B., and Darrell, T. (2021). Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*.
- [87] Wei, C., Shen, K., Chen, Y., and Ma, T. (2020). Theoretical analysis of self-training with deep networks on unlabeled data. *arXiv preprint arXiv:2010.03622*.
- [88] Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. (2018). Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742.
- [89] Xie, Q., Luong, M.-T., Hovy, E., and Le, Q. V. (2020a). Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10687–10698.
- [90] Xie, X., Chen, J., Li, Y., Shen, L., Ma, K., and Zheng, Y. (2020b). Self-supervised cyclegan for object-preserving image-to-image domain adaptation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 498–513. Springer.
- [91] Yang, X., Song, Z., King, I., and Xu, Z. (2022). A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering*.
- [92] Zadrozny, B. (2004). Learning and Evaluating Classifiers Under Sample Selection Bias. In *International Conference on Machine Learning (ICML)*.
- [93] Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. (2021). Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR.

- [94] Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR)*.
- [95] Zhang, J., Menon, A., Veit, A., Bhojanapalli, S., Kumar, S., and Sra, S. (2021). Coping with label shift via distributionally robust optimisation. In *International Conference on Learning Representations (ICLR)*.
- [96] Zhang, K., Schölkopf, B., Muandet, K., and Wang, Z. (2013). Domain Adaptation Under Target and Conditional Shift. In *International Conference on Machine Learning (ICML)*.
- [97] Zhang, R. (2019). Making convolutional networks shift-invariant again. In *ICML*.
- [98] Zhang, W., Ouyang, W., Li, W., and Xu, D. (2018). Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [99] Zhang, Y., Liu, T., Long, M., and Jordan, M. (2019). Bridging theory and algorithm for domain adaptation. In *International Conference on Machine Learning*. PMLR.
- [100] Zhu, X. and Ghahramani, Z. (2003). Learning from labeled and unlabeled data with label propagation. *CMU CALD tech report CMU-CALD-02-107, 2002*.



## Appendix

### Appendix Table of Contents

<b>A Other Related Works</b>	<b>17</b>
<b>B More Details on Problem Setup</b>	<b>18</b>
<b>C Additional Experiments and Details</b>	<b>19</b>
C.1 Additional setup and notation . . . . .	19
C.2 Dataset details . . . . .	19
C.3 Method details . . . . .	20
C.4 Additional UDA experiments . . . . .	22
C.5 Additional SSL experiments . . . . .	23
C.6 Other experimental details . . . . .	23
<b>D Additional Results in Toy Setup</b>	<b>23</b>
D.1 Detailed description of our simplified setup . . . . .	23
D.2 Discussion on self-training and contrastive learning objectives . . . . .	25
D.3 Additional empirical results in our simplified setup . . . . .	26
D.4 Reconciling Practice: Experiments with deep networks in toy setup . . . . .	27
<b>E Formal Statements from Sec. 4</b>	<b>29</b>
E.1 Analysis of ERM and ST: Formal Statement of Theorem 2 . . . . .	29
E.1.1 Proof of Proposition 3 . . . . .	34
E.1.2 Analysis with $\rho \neq 0$ in Contrastive Pretraining Objective (3) . . . . .	36
E.1.3 Proof of Theorem 5 . . . . .	38
E.1.4 Proof of Corollary 6 . . . . .	41
E.2 Analysis of STOC: Formal Statement of Theorem 7 . . . . .	42
E.3 Analysis for SSL . . . . .	48
<b>F Limitations of Prior Work</b>	<b>48</b>
F.1 Contrastive learning analysis . . . . .	48
F.2 Self-training analysis . . . . .	48
<b>G Additional Lemmas</b>	<b>49</b>

### A Other Related Works

**Unsupervised domain adaption.** Without assumption on the nature of shift, UDA is underspecified [7]. This challenge has been addressed in various ways by researchers. One approach is to investigate additional structural assumptions under which UDA problems are well posed [77, 74]. Popular settings for which DA is well-posed include (i) *covariate shift* [96, 92, 19, 20, 35] where  $p(x)$  can change from source to target but  $p(y|x)$  remains invariant; and (ii) *label shift* [69, 53, 3, 1, 31, 95, 66, 30] where the label marginal  $p(y)$  can change but  $p(x|y)$  is shared across source and target. Principled methods with strong theoretical guarantees exists for adaptation under these settings when target

distribution’s support is a subset of the source support. Other works [25, 6, 32, 28] extend the label shift setting to scenarios where previously unseen classes may appear in the target and  $p_{\mathcal{X}}|_{\mathcal{Y}}q$  remains invariant among seen classes. A complementary line of research focuses on constructing benchmarks to develop heuristics for incorporating the unlabeled target data, relying on benchmark datasets ostensibly representative of “real-world shifts” to adjudicate progress [72, 85, 70, 62, 63]. As a result, various benchmark-driven heuristics have been proposed [54, 55, 82, 81, 99, 98, 26, 79]. Our work engages with the latter, focusing on two popular methods: self-training and contrastive pretraining.

**Domain generalization.** In domain generalization, the model is given access to data from multiple different domains and the goal is to generalize to a previously unseen domain at test time [10, 59]. For a survey of different algorithms for domain generalization, we refer the reader to Gulrajani and Lopez-Paz [37]. A crucial distinction here is that unlike the domain generalization setting, in DA problems, we have access to unlabeled examples from the test domain.

**Semi-supervised learning.** To learn from a small amount of labeled supervision, semi-supervised learning methods leverage unlabeled data alongside to improve learning models. One of the seminal works in SSL is the pseudolabeling method [75], where a classifier is trained on the labeled data and then used to classify the unlabeled data, which are then added to the training set. The work of Zhu and Ghahramani [100] built on this by introducing graph-based methods, and the transductive SVMs [43] presented an SVM-based approach. More recent works have focused on deep learning techniques, and similar to UDA, self-training and contrastive pretraining have emerged as two prominent choices. We delve into these methods in greater detail in the following paragraphs. For a discussion on other SSL methods, we refer interested readers to [15, 84, 91].

**Self-training.** Two popular forms of self-training are pseudolabeling [52] and conditional entropy minimization [34], which have been observed to be closely connected [8, 52, 79, 78]. Motivated by its strong performance in SSL and UDA settings [79, 89, 30, 78], several theoretical works have made attempts to understand its behavior [50, 87, 17]. [87, 12] aims to understand the behavior of the global minimizer of self-training objective by studying input consistency regularization, which enforces stability of the prediction for different augmentations of the unlabeled data. Our analysis of self-training is motivated by the work of Chen et al. [17] which explores the iterative behavior of self-training to unlearn spurious features. The setting of spurious features is of particular interest, since prior works have specifically analyzed the failures of out-of-distribution generalization in the presence of spurious features [60, 71].

**Contrastive learning.** An alternate line of work that uses unlabeled data for learning representations in the pretraining stage is contrastive learning [36, 61, 13, 16, 88]. Given an augmentation distribution, the main goal of contrastive objectives is to map augmentations drawn from the same input (positive pairs) to similar features, and force apart features corresponding to augmentations of different inputs (negative pairs) [13, 14, 41]. Prior works [11, 44, 38] have also shown a close relationship between contrastive [16, 39] and non-contrastive objectives [4, 93]. Consequently, in our analysis pertaining to the toy setup we focus on the mathematically non-contrastive objective Barlow Twins [93]. Using this pretrained backbone (either as an initialization or as a fixed feature extractor) a downstream predictor is learned using labeled examples. Several works [39, 73, 38, 2, 44] have analyzed the in-distribution generalization of the downstream predictor via label consistency arguments on the graph of positive pairs (augmentation graph). In contrast, we study the impact of contrastive learning under distribution shifts in the UDA setup. Other works [76, 40] that examine contrastive learning for UDA also conform to the augmentation graph view point, making additional assumptions that guarantee linear transferability. In our simplified setup involving spurious correlations, these abstract assumptions break easily when the augmentations are of a generic nature, akin to practice. Finally, some empirical works [58, 57] have found self-supervised objectives like contrastive pretraining to reduce dependence on spurious correlations. Corroborating their findings, we extensively evaluate the complementary benefits of contrastive learning and self-training on real-world datasets. Finding differing results in SSL and UDA settings, we further examine their behavior theoretically in our toy setup.

## B More Details on Problem Setup

In this section, we elaborate on our setup and methods studied in our work.

**Unsupervised Domain Adaptation (UDA).** We assume that we are given labeled data from the *source* distribution and unlabeled data from a shifted, *target* distribution, with the goal of performing well on target data. We assume that the source and target distributions have the same label marginals  $P_S(p|y) = P_T(p|y)$  (i.e., no label proportion shift) and the same Bayes optimal predictor, i.e.,  $\arg \max_y p_S(p|y | x) = \arg \max_y p_T(p|y | x)$ . Here, even with infinite labeled source data, the challenge lies in generalizing out-of-distribution. In experiments, we assume access to finite data but in theory, we assume population access to labeled source and unlabeled target.

**Semi-Supervised Learning (SSL).** Here, there is no distribution shift, i.e.,  $P_S = P_T = P_U$ . We are given a small number of labeled examples and a comparatively large amount of unlabeled examples, both drawn from the same distribution. Without loss of generality, we denote this distribution with  $P_T$ . The goal in SSL is to generalize in-distribution. The challenge is primarily due to limited access to labeled data. Here, in experiments, we assume limited access to labeled data but a comparatively larger amount of unlabeled in-distribution data. In theory, we assume population access to unlabeled data but limited labeled examples.

**Methods.** As discussed in the main paper, we compare four methods for learning with labeled and unlabeled data. Table 8 summarizes the main methods and key differences between those methods in UDA and SSL setup. For exact implementation in our experiments, we refer reader to App. C.3.

## C Additional Experiments and Details

### C.1 Additional setup and notation

Recall, our goal is to learn a predictor that maps inputs  $x \in \mathcal{X} \subseteq \mathbb{R}^d$  to outputs  $y \in \mathcal{Y}$ . We parameterize predictors  $f = h \circ \Phi : \mathbb{R}^d \rightarrow \mathcal{Y}$ , where  $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^k$  is a feature map and  $h \in \mathbb{R}^k$  is a classifier that maps the representation to the final scores or logits. With  $A : \mathcal{X} \rightarrow \mathcal{A}$ , we denote the augmentation function that takes in an input  $x$  and outputs an augmented view of the input  $A(x)$ . Unless specified otherwise, we perform full-finetuning in all of our experiments on real-world data. That is, we backpropagate gradients in both the linear head  $h$  and the backbone  $\phi$ . For UDA, we denote source labeled points as  $\{x_i, y_i\}_{i=1}^n$  and target unlabeled points as  $\{x'_i\}_{i=1}^m$ . For SSL, we use the same notation for labeled and unlabeled in-distribution data.

### C.2 Dataset details

For both UDA and SSL, we conduct experiments across eight benchmark datasets. Each of these datasets consists of domains, enabling us to construct source-target pairs for UDA. The adopted source and target domains are standard to previous studies [76, 30, 70]. Because the SSL setting lacks distribution shift, we do not need to worry about domain designations and default to using source alone. To simulate limited supervision in SSL, we sub-sample the original labeled training set to 10%. Below provide exact details about the datasets used in our benchmark study.

- **CIFAR10** We use the original CIFAR10 dataset [48] as the source dataset. For target domains, we consider CINIC10 [22] which is a subset of Imagenet restricted to CIFAR10 classes and downsampled to 32/32.
- **FMoW** In order to consider distribution shifts faced in the wild, we consider FMoW-WILDs [47, 18] from WILDS benchmark, which contains satellite images taken in different geographical regions and at different times. We use the original train as source and OOD val and OOD test splits as target domains as they are collected over different time-period. Overall, we obtain 3 different domains (1 source and 2 targets).
- **BREEDs** We also consider BREEDs benchmark [72] in our setup to assess robustness to subpopulation shifts. BREEDs leverage class hierarchy in ImageNet [68] to re-purpose original classes to be the subpopulations and defines a classification task on superclasses. We consider distribution shift due to subpopulation shift which is induced by directly making the subpopulations present in the training and test distributions disjoint. BREEDs benchmark contains 4 datasets **Entity-13**, **Entity-30**, **Living-17**, and **Non-living-26**, each focusing on different subtrees and levels in the hierarchy. Overall, for each of the 4 BREEDs datasets (i.e., Entity-13, Entity-30, Living-

Dataset	Domains			
CIFAR10	CIFAR10v1	CINIC10		
FMoW	Year 2002-'13	Year '13-'16	Year '16-'18	
Visda	Synthetic	Real-1	Real-2	
Entity13	Sub-pop. 1	Sub-pop. 2		
Entity30	Sub-pop. 1	Sub-pop. 2		
Living17	Sub-pop. 1	Sub-pop. 2		
Nonliving26	Sub-pop. 1	Sub-pop. 2		
Officehome	Product	Real	ClipArt	Art

Figure 4: Examples from all the domains in each dataset.

17, and Non-living-26), we obtain one different domain which we consider as target. We refer to source and target as follows: BREEDs sub-population 1, BREEDs sub-population 2.

- **OfficeHome** We use four domains (art, clipart, product and real) from OfficeHome dataset [85]. We use the product domain as source and the other domains as target.
- **Visda** We use three domains (train, val and test) from the Visda dataset [64, 63]. While ‘train’ domain contains synthetic renditions of the objects, ‘val’ and ‘test’ domains contain real world images. To avoid confusing, the domain names with their roles as splits, we rename them as ‘synthetic’, ‘Real-1’ and ‘Real-2’. We use the synthetic (original train set) as the source domain and use the other domains as target.

We summarize the information about source and target domains in Table 3.

**Train-test splits** We partition each source and target dataset into 80% and 20% i.i.d. splits. We use 80% splits for training and 20% splits for evaluation (or validation). We throw away labels for the 80% target split and only use labels in the 20% target split for final evaluation. The rationale behind splitting the target data is to use a completely unseen batch of data for evaluation. This avoids evaluating on examples where a model potentially could have overfit. over-fitting to unlabeled examples for evaluation. In practice, if the aim is to make predictions on all the target data (i.e., transduction), we can simply use the (full) target set for training and evaluation.

**Simulating SSL settings and limited supervision.** For SSL settings, we choose the in-distribution domain as the source domain. To simulate limited supervision in SSL, we sub-sample the original labeled training set to 10% and use all the original dataset as unlabeled data. For evaluation, we further split the original holdout set into two partitions (one for validation and the other to report final accuracy numbers).

### C.3 Method details

For implementation, we build on top of WILDs [70] and RLSbench [30] open source libraries.

Dataset	Source	Target
CIFAR10	CIFAR10v1	CINIC10
FMoW	FMoW (2002–’13)	FMoW (2013–’16), FMoW (2016–’18)
Entity13	Entity13 (sub-population 1)	Entity13 (sub-population 2)
Entity30	Entity30 (sub-population 1)	Entity30 (sub-population 2),
Living17	Living17 (sub-population 1)	Living17 (sub-population 2),
Nonliving26	Nonliving26 (sub-population 1)	Nonliving26 (sub-population 2),
Officehome	Product	Product, Art, ClipArt, Real
Visda	Synthetic (originally referred to as train)	Synthetic, Real-1 (originally referred to as val), Real-2 (originally referred to as test)

Table 3: Details of source and target sets in each dataset considered in our testbed.

**ERM (Source only) training.** We consider Empirical Risk Minimization (ERM) on the labeled source data as a baseline. Since this simply ignores the unlabeled target data, we call this as source only training. As mentioned in the main paper, we perform source only training with data augmentations. Formally, we minimize the following ERM loss:

$$L_{\text{source only}} = \frac{1}{n} \sum_{i=1}^n \ell(Ax_i, y_i), \quad (4)$$

where  $A$  is the stochastic data augmentation operation and  $\ell$  is a loss function. For SSL, the ERM baseline only uses the small of labeled data available.

**Contrastive Learning (CL).** We perform contrastive pretraining on the unlabeled dataset to obtain the backbone  $\phi_{cl}$ . And then we perform full fine-tuning with source labeled data by initializing the backbone with  $\phi_{cl}$ . We use SwAV [13] for contrastive pretraining. The main idea behind SwAV is to train a model to identify different views of the same image as similar, while also ensuring that it finds different images to be distinct. This is accomplished through a *swapped* prediction mechanism, where the goal is to compute a code from an augmented version of the image and predict this code from other augmented versions of the same image. In particular, given two image features  $\phi(x'_{a1})$  and  $\phi(x'_{a2})$  from two different augmentations of the same image  $x'$ , i.e.,  $x'_{a1}, x'_{a2} = Ax'$ , SwAV computes their codes  $z_{a1}$  and  $z_{a2}$  by matching the features to a set of  $K$  prototypes  $\{c_1, \dots, c_K\}$ . Then SwAV minimizes the following loss such that  $\phi(x'_{a1})$  can compute codes  $z_{a2}$  and  $\phi(x'_{a2})$  can compute codes  $z_{a1}$ :

$$L_{\text{SwAV}} = \sum_{i=1}^m \mathbb{E}_{x'_{i,a1}, x'_{i,a2} \sim A(x'_i)} \ell(\phi(x'_{i,a1}), z_{i,a2}) + \ell(\phi(x'_{i,a2}), z_{i,a1}), \quad (5)$$

where  $\ell$  computes KL-divergence between codes computed with features (e.g.  $\phi(x_{a1})$ ) and the code computed by another view (e.g.  $z_{a2}$ ). For more details about the algorithm, we refer the reader to Caron et al. [13]. In all UDA settings, unless otherwise specified, we pool all the (unlabeled) data from the source and target to perform SwAV. For SSL, we leverage in-distribution unlabeled data.

We employ SimCLR [16] for the CIFAR10 dataset, aligning with previous studies that have utilized contrastive pretraining on the same dataset [51, 76]. The reason for this choice is that SwAV relies on augmentations that involve cropping images to a smaller resolution, making it more suitable for datasets with larger resolutions beyond 32 × 32.

**Self-Training (ST).** For self-training, we apply FixMatch [79], where the loss on labeled data and on pseudolabeled unlabeled data are minimized simultaneously. Sohn et al. [79] proposed FixMatch as a variant of the simpler Pseudo-label method [52]. This algorithm dynamically generates pseudolabels and overfits on them in each batch. FixMatch employs consistency regularization on the unlabeled data. In particular, while pseudolabels are generated on a weakly augmented view of

the unlabeled examples, the loss is computed with respect to predictions on a strongly augmented view. The intuition behind such an update is to encourage a model to make predictions on weakly augmented data consistent with the strongly augmented example. Moreover, FixMatch only overfits to the assigned labeled with weak augmentation if the confidence of the prediction with strong augmentation is greater than some threshold  $\tau$ . Refer to  $A_{\text{weak}}$  as the weak-augmentation and  $A_{\text{strong}}$  as the strong-augmentation function. Then, FixMatch uses the following loss function:

$$L_{\text{FixMatch}}(f, \theta) = \frac{1}{n} \sum_{i=1}^n \ell(f(A_{\text{strong}}(x_i), y_i)) + \frac{\lambda}{m} \sum_{i=1}^m \ell(f(A_{\text{strong}}(x'_i), \hat{y}_i)) \mid \max_y f_y(A_{\text{strong}}(x'_i)) \geq \tau,$$

where  $\hat{y}_i = \arg \max_y f_y(A_{\text{weak}}(x_i))$ . For UDA, our unlabeled data is the union of source and target unlabeled data. For SSL, we only leverage in-distribution unlabeled data.

We adapted our implementation from Sagawa et al. [70] which matches the implementation of Sohn et al. [79] except for one detail. While Sohn et al. [79] augments labeled examples with weak augmentation, Sagawa et al. [70] proposed to strongly augment the labeled source examples.

**Self-Training Over Contrastive learning (STOC).** Finally, rather than performing FixMatch from a randomly initialized backbone, we initialize FixMatch with a contrastive pretrained backbone.

#### C.4 Additional UDA experiments

Table 4: *Results in the UDA setup.* We report accuracy on target (OOD) data from which we only observe unlabeled examples during training. For benchmarks with multiple target distributions (e.g., OH, Visda), we report average accuracy on those targets.

Method	Living17	Nonliv26	Entity13	Entity30	FMoW (2 tgts)	Visda (2 tgts)	OH (3 tgts)	CIFAR-100
ERM	60.2±0.1	45.4±0.2	68.6±0.1	55.7±0.0	56.5±0.1	20.8±0.2	9.5±0.2	74.3±0.1
ST	71.1±0.2	56.8±0.1	78.0±0.3	66.7±0.1	56.9±0.4	39.1±0.1	11.1±0.1	78.3±0.3
CL	74.1±0.2	57.4±0.3	76.9±0.2	66.6±0.3	61.5±0.5	63.2±0.2	22.8±0.1	77.5±0.1
STOC (ours)	<b>82.6±0.1</b>	<b>62.1±0.2</b>	<b>81.9±0.2</b>	<b>72.0±0.2</b>	<b>65.3±0.1</b>	<b>70.1±0.2</b>	<b>27.1±0.3</b>	<b>79.9±0.3</b>

Table 5: *Results in the UDA setup with source only contrastive pretraining.* We report accuracy on target (OOD) data from which we only observe unlabeled examples during training. For benchmarks with multiple target distributions (e.g., OH, Visda), we report average accuracy on those targets.

Method	Living17	Nonliv26	Entity13	Entity30	FMoW (2 tgts)	Visda (2 tgts)	OH (3 tgts)	CIFAR-100
CL (source only)	67.3±0.1	49.1±0.2	71.5±0.1	58.5±0.3	53.9±0.1	33.3±0.2	21.7±0.1	77.7±0.1
STOC (source only)	75.0±0.2	58.4±0.1	79.8±0.3	67.5±0.1	56.3±0.4	42.7±0.1	25.7±0.1	77.8±0.1
CL	74.1±0.2	57.4±0.3	76.9±0.2	66.6±0.3	61.5±0.5	63.2±0.2	22.8±0.1	77.5±0.1
STOC	<b>82.6±0.1</b>	<b>62.1±0.2</b>	<b>81.9±0.2</b>	<b>72.0±0.2</b>	<b>65.3±0.1</b>	<b>70.1±0.2</b>	<b>27.1±0.3</b>	<b>79.9±0.3</b>

## C.5 Additional SSL experiments

Table 6: *Results in the SSL setup.* We report accuracy on hold-out ID data. Recall that SSL uses labeled and unlabeled data from the same distribution during training.

Method	Living17	Nonliv26	Entity13	Entity30	FMoW	Visda	OH	CIFAR
ERM	76.8 $\pm$ 0.1	64.9 $\pm$ 0.2	80.1 $\pm$ 0.0	70.4 $\pm$ 0.3	33.6 $\pm$ 0.4	99.2 $\pm$ 0.0	32.0 $\pm$ 0.2	85.5 $\pm$ 0.1
ST	85.4 $\pm$ 0.1	75.7 $\pm$ 0.2	85.4 $\pm$ 0.2	77.3 $\pm$ 0.1	33.6 $\pm$ 0.3	99.2 $\pm$ 0.1	32.0 $\pm$ 0.1	93.1 $\pm$ 0.1
CL	91.1 $\pm$ 0.5	84.6 $\pm$ 0.6	90.7 $\pm$ 0.4	85.5 $\pm$ 0.3	43.1 $\pm$ 0.2	97.6 $\pm$ 0.3	49.7 $\pm$ 0.2	91.7 $\pm$ 0.2
STOC (ours)	92.0 $\pm$ 0.1	85.8 $\pm$ 0.2	91.3 $\pm$ 0.3	86.1 $\pm$ 0.2	44.4 $\pm$ 0.1	97.7 $\pm$ 0.2	49.9 $\pm$ 0.2	93.06 $\pm$ 0.3

## C.6 Other experimental details

**Augmentations.** For weak augmentation, we leverage random horizontal flips and random crops of pre-defined size. For SwAV, we also perform multicrop augmentation as proposed in Caron et al. [13]. For strong augmentation, we apply the following transformations sequentially: random horizontal flips, random crops of pre-defined size, augmentation with Cutout [24], and RandAugment [21]. For the exact implementation of RandAugment, we directly use the implementation of Sohn et al. [79]. Unless specified otherwise, for all methods, we default to using strong augmentation techniques.

**Architectures.** In our work, we experiment with Resnet18, Resnet50 [42] trained from scratch (*i.e.* random initialization). We do not consider off-the-shelf pretrained models (*e.g.*, on Imagenet [68]) to avoid confounding our conclusions about contrastive pretraining. However, we note that our results on most datasets tend to be comparable to and sometimes exceed those obtained with ImageNet pretrained models. For BREEDs datasets, we employ Resnet18 architecture. For other datasets, we train a Resnet50 architecture.

Except for Resnets on CIFAR dataset, we used the standard pytorch implementation [27]. For Resnet on Cifar, we refer to the implementation here: <https://github.com/kuangliu/pytorch-cifar>. For all the architectures, whenever applicable, we add antialiasing [97]. We use the official library released with the paper.

**Hyperparameters.** For all the methods, we fix the algorithm-specific hyperparameters to the original recommendations. For UDA, given that the setup precludes access to labeled data from the target distribution, we use source hold-out performance to pick the best hyperparameters. During pretraining, early stopping is done according to lower values of pretraining loss.

We tune the learning rate and  $\ell_2$  regularization parameter by fixing the batch size for each dataset that corresponds to the maximum we can fit to 15GB GPU memory. We default to using cosine learning rate schedule [56]. We set the number of epochs for training as per the suggestions of the authors of respective benchmarks. For SSL, we run both ERM and FixMatch for approximately 2000 epochs. Note that we define the number of epochs as a full pass over the labeled training source data. We summarize the learning rate, batch size, number of epochs, and  $\ell_2$  regularization parameter used in our study in Table 7.

**Compute infrastructure.** Our experiments were performed across a combination of Nvidia T4, A6000, and V100 GPUs.

## D Additional Results in Toy Setup

In this section we will first give more details on our simplified setup that captures both contrastive pretraining and self-training in the same framework. Then, we provide some additional empirical results that are not captured theoretically but mimic behaviors observed in real world settings, highlighting the richness of our setup.

### D.1 Detailed description of our simplified setup

In this subsection, we will first re-iterate the problem setup in Sec. 4 and provide some comparisons between our setup and those in closely related works. We will then describe the four methods: ERM,

Dataset	Batch size	$\ell_2$ regularization set	Learning rate set
CIFAR10	200	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.2, 0.1, 0.05, 0.01, 0.003, 0.001u
FMoW	64	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.01, 0.003, 0.001, 0.0003, 0.0001u
Entity13	256	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.4, 0.2, 0.1, 0.05, 0.02, 0.01, 0.005u
Entity30	256	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.4, 0.2, 0.1, 0.05, 0.02, 0.01, 0.005u
Entity30	256	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.4, 0.2, 0.1, 0.05, 0.02, 0.01, 0.005u
Nonliving26	256	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.4, 0.2, 0.1, 0.05, 0.02, 0.01, 0.005u
Officehome	96	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.01, 0.003, 0.001, 0.0003, 0.0001u
Visda	96	t0.001, 0.0001, $10^{-5}$ , 0.0u	t0.03, 0.01, 0.003, 0.001, 0.0003u

Table 7: Details of the batch size, learning rate set and  $\ell_2$  regularization set considered in our testbed.

ST, CL, and STOC, providing details on the exact estimates returned by these algorithms in the SSL and UDA settings.

**Data distribution.** We consider binary classification and model the inputs as consisting of two kinds of features:  $x = [x_{in}, x_{sp}]$  where  $x_{in} \sim \mathcal{P} \mathcal{R}^{d_{in}}$  is the invariant feature that is predictive of the label across both source  $\mathcal{P}_S$  and target  $\mathcal{P}_T$  and  $x_{sp} \sim \mathcal{P} \mathcal{R}^{d_{sp}}$  is the spurious feature that is correlated with the label  $y$  only on the source domain  $\mathcal{P}_S$  but uncorrelated with label  $y$  in  $\mathcal{P}_T$ . Here,  $x_{in} \sim \mathcal{P} \mathcal{R}^{d_{in}}$  determines the label using the ground truth classifier  $w^* \sim \text{Unif}(\mathcal{S}^{d_{in}-1})$ , and  $x_{sp} \sim \mathcal{P} \mathcal{R}^{d_{sp}}$  is strongly correlated with the label on source but random noise on target. Formally, we sample  $y \sim \text{Unif}\{1, -1\}$  and generate inputs  $x$  conditioned on  $y$  as follows

$$\begin{aligned} \mathcal{P}_S : x_{in} &\sim \mathcal{N}(\gamma y w^*, \Sigma_{in}) \quad x_{sp} \sim y \mathbf{1}_{d_{sp}} \\ \mathcal{P}_T : x_{in} &\sim \mathcal{N}(\gamma y w^*, \Sigma_{in}) \quad x_{sp} \sim \mathcal{N}(\mathbf{0}, \Sigma_{sp}), \end{aligned} \quad (6)$$

where  $\gamma$  is the margin afforded by the invariant feature. We set covariance of the invariant features  $\Sigma_{in} = \sigma_{in}^2 \rho \mathbf{I}_{d_{in}} - w^* w^{*\top}$  to capture structure in the invariant feature that the variance is less along the latent predictive direction  $w^*$ . Note that the spurious feature is completely predictive of the label in the source data, and is distributed as spherical Gaussian in the target data with  $\Sigma_{sp} = \sigma_{sp}^2 \mathbf{I}_{d_{sp}}$ .

**Why is our simplified setup interesting?** In our setup,  $x_{in}$  is the hard to learn feature that generalizes from source to target. The hardness of learning this feature is determined by the value of the margin  $\gamma$  and how it compares with size of the spurious feature ( $d_{sp}$ ). Since,  $\gamma \ll d_{sp}$  is small in our setup,  $x_{in}$  is much harder to learn on source data (even with population access) compared to the spurious feature  $x_{sp}$  which generalizes poorly from source to target. These two types of features have been captured in similar analysis on spurious correlations [71, 60] since it imitates pitfalls emanating from the presence of spurious features in real world datasets (e.g., the easy to learn background feature in image classification problems). While this setup is simple, it is also expressive enough to elucidate both self-training and contrastive learning behaviors we observe in real world settings. Specifically, it captures the separation results we observe in Sec. 3.

**Differences of our setup with prior works.** While our distribution shift settings bears the above similarities it also has important differences with works analyzing self-training and contrastive pretraining individually. Chen et al. [17] analyze the iterative nature of self-training algorithm, where the premise is that we are given a classifier that not only has good performance on source data but in addition does not rely too much on the spurious feature. Under the strong condition of small norms along the spurious feature, they show that self-training can provably unlearn this small dependence when the target data along the spurious feature is random noise. This assumption is clearly violated in setups where the spurious correlation is strong (as in our toy setup), i.e., the dependence on the spurious feature is rather large (much larger than that on the invariant feature) for any classifier that is trained directly on source data. Consequently, we show the need for “good” pretrained representations from contrastive pretraining over which if we train a linear predictor (using source labeled data), it will provably have a reduced “effective” dependence on the spurious feature.

Using an augmentation distribution similar to ours, Saunshi et al. [73] carried out contrastive pretraining analysis with the backbone belonging to a capacity constrained function class (similar analysis also in [40]). Our setup differs from this in two key ways: (i) we specifically consider a distribution



shift from source to target. Unlike their setting, it is not sufficient to make augmentations consistent with ground truth labels, since the predictor that uses just the spurious feature also assigns labels consistent with both ground truth predictions and augmentations on the source data; and (ii) our augmentation distribution assumes no knowledge of the invariant feature, which is why we augment all dimensions uniformly, as opposed to selectively augmenting a set of dimensions. In other words, we assume no knowledge of the structure of the optimal target predictor. For *e.g.*, if we had knowledge of the spurious dimensions we could have just selectively augmented those. Assuming knowledge of these perfect augmentations is not ideal for two reasons: (a) it makes the problem so easy that just training an ERM model on source data with these augmentations would already yield a good target predictor (which rarely happens in practice); and (b) in real-world datasets perfect augmentations for the downstream task are not known. Hence, we stick to generic augmentations in our setup.

## D.2 Discussion on self-training and contrastive learning objectives

Method	UDA Setup	SSL Setup
<b>ERM:</b>	$h_{\text{erm}} \arg \min_h \mathbb{E}_{P_S} \ell_{\text{ph}}(x, y)$	$h_{\text{erm}} \arg \min_h \frac{1}{n} \sum_{i=1}^n \ell_{\text{ph}}(x_i, y_i)$
<b>ST:</b>	Starting from $h_{\text{erm}}$ optimize over $h$ (to get $h_{\text{st}}$ ): $\mathbb{E}_{P_T(x)} \ell_{\text{ph}}(x, \text{sgn}(\text{ph}(x)))$	Starting from $h_{\text{erm}}$ optimize over $h$ (to get $h_{\text{st}}$ ): $\mathbb{E}_{P_T(x)} \ell_{\text{ph}}(x, \text{sgn}(\text{ph}(x)))$
<b>CL:</b>	$\arg \min_{\phi} \mathcal{L}_{\text{cl}}(\phi)$ Use $\rho \mathbb{P}_S \rho$ for $\mathcal{L}_{\text{cl}}(\phi)$ $h_{\text{cl}} \arg \min_h \mathbb{E}_{P_S} \ell_{\text{ph}}(x, y)$	$\arg \min_{\phi} \mathcal{L}_{\text{cl}}(\phi)$ Use $P_T \rho$ for $\mathcal{L}_{\text{cl}}(\phi)$ $h_{\text{cl}} \arg \min_h \frac{1}{n} \sum_{i=1}^n \ell_{\text{ph}}(x_i, y_i)$
<b>STOC:</b>	Starting from $h_{\text{cl}}$ optimize over $h$ (to get $h_{\text{stoc}}$ ): $\mathbb{E}_{P_T(x)} \ell_{\text{ph}}(x, \text{sgn}(\text{ph}(x)))$	Starting from $h_{\text{cl}}$ optimize over $h$ (to get $h_{\text{stoc}}$ ): $\mathbb{E}_{P_T(x)} \ell_{\text{ph}}(x, \text{sgn}(\text{ph}(x)))$

Table 8: **Description of methods for SSL vs. UDA:** For each method we provide exact objectives used for experiments and analysis in the SSL and UDA setups (pertaining to Sec. 4).

In text we will describe our objectives and methods for the UDA setup. In Table 8 we contrast the differences in the methods and objectives for SSL and UDA setups. Recall from Section 2 that we learn linear classifiers  $h$  over features extractors  $\Phi$ . We consider linear feature extractor i.e.  $\Phi$  is a matrix in  $\mathbb{R}^{k \times d}$ . For mathematical convenience, we assume access to infinite unlabeled data and hence replace the empirical quantities over unlabeled data with their population counterpart. In the UDA setting, we further assume access to infinite labeled data from the source. Note that due to distribution shift between source and target, ‘‘ERM’’ on infinite labeled data from the source does not necessarily achieve optimal performance on the target. For binary classification, we assume that the linear layer  $h$  maps features to a scalar in  $\mathbb{R}$  such that the prediction is  $\text{sgn}(\text{ph}^\top \Phi x)$ . We use the exponential loss  $\ell_{\text{ph}}(x, y) = \exp(-y \text{ph}(x))$  as the classification loss.

*Contrastive pretraining.* We obtain  $\Phi_{\text{cl}} : \arg \min_{\Phi} \mathcal{L}_{\text{cl}}(\Phi)$  by minimizing the Barlow Twins objective [93], which prior works have shown is also equivalent to spectral contrastive and non-contrastive objectives [33, 11]. In Sec. 4, we consider a constrained form of Barlow Twins in (3) which enforces representations of different augmentations  $a_1, a_2$  of the same input  $x$  to be close in representation space, while ensuring feature diversity by staying in the constraint set. We assume a strict constraint on regularization  $\rho > 0$  for the theoretical arguments in the rest of the main paper. In App. E.1.2 we prove that all our claims hold for small  $\rho$  as well. In (7), we redefine the pretraining objective with a regularization term (instead of a constraint set) where  $\kappa$  controls the strength of the regularization term, with higher values of  $\kappa$  corresponding to stronger constraints on feature diversity. We then learn a linear classifier  $h_{\text{cl}}$  over  $\Phi_{\text{cl}}$  to minimize the exponential loss on labeled source data.

$$\mathcal{L}_{\text{cl}}(\Phi) : \mathbb{E}_{x \sim P_U} \mathbb{E}_{a_1, a_2 \sim P_A(\cdot|x)} \left\{ \Phi \rho a_1 - \Phi \rho a_2 \right\}_2^2 + \kappa \mathbb{E}_{a \sim P_A} \Phi \rho a \Phi \rho a^\top - \mathbf{I}_k \quad (7)$$

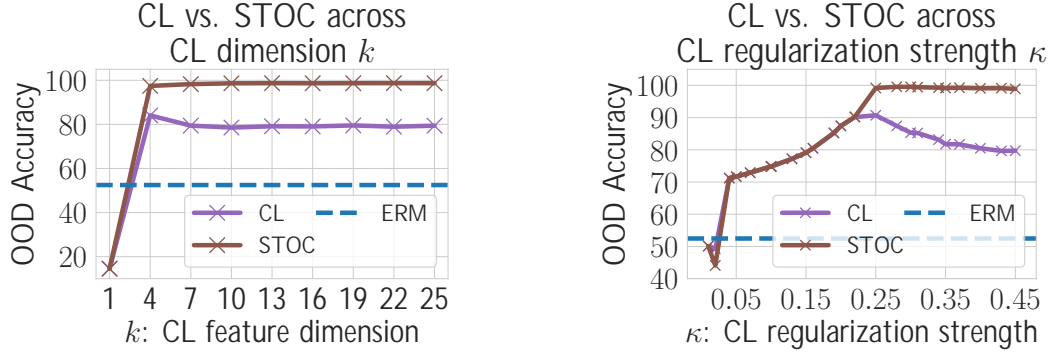


Figure 5: **Ablations on pretraining hyperparameters:** In the UDA setup we plot the performance of CL and STOC as we vary two pretraining hyper-parameters: (*left*) the output dimension  $k$  of the feature extractor  $\Phi$ ; and (*right*) the strength  $\rho\kappa\eta$  of the regularizer in the Barlow Twins objective in (7). While ablating on  $k$  we fix  $\kappa = 0.5$ , and while ablating on  $\kappa$  we fix  $k = 10$ . Other problem parameters are taken from Example 1.

*Augmentations.* Data augmentations play a key role in contrastive pre-training (and also as we see later, state-of-the-art self-training variants like FixMatch). Given input  $x \in \mathcal{X}$ , let  $\mathbb{P}_{\mathcal{A}} \rho a \mid x \mathbb{q}$  denote the distribution over its augmentations, and  $\mathbb{P}_{\mathcal{A}}$  denote the marginal distribution over all possible augmentations. We use the following simple augmentations where we scale the magnitude of each co-ordinate by a uniformly independent amount, *i.e.*,

$$a \sim \mathbb{P}_{\mathcal{A}} \rho c \mid x \mathbb{q} \quad c \in \mathcal{D} \text{ where, } c \sim \text{Unif}(0, 1)^d. \quad (8)$$

The performance of different methods heavily depends on the assumptions we make on augmentations. We try to mirror practical settings where the augmentations are fairly “generic”, not encoding any information about which features are invariant or spurious, and hence perturb all features symmetrically.

*Self-training.* ST performs ERM in the first stage using labeled data from the source, and then subsequently updates the head  $h$  by iteratively generating pseudolabels on the unlabeled target:

$$\mathcal{L}_{\text{st}} \rho h; \Phi \mathbb{q} : \mathbb{E}_{\mathbb{P}_{\mathcal{T}}(x)} \ell \rho h^{\top} \Phi x, \text{sgn} \rho h^{\top} \Phi x \mathbb{q} \mathbb{q} \quad \text{Update: } h^{t+1} = \frac{h^t}{\|h^t\|_2} \frac{\eta \gamma \rho h^{\top} \Phi x; \Phi \mathbb{q}}{\eta \gamma \rho h^{\top} \Phi x; \Phi \mathbb{q}} \quad (9)$$

For convenience, we keep the feature backbone  $\Phi$  fixed across the self-training iterations and only update the linear head on the pseudolabels.

*STOC(Self-training after contrastive learning).* Finally, we can combine the two unsupervised objectives where we do the self-training updates (2) with  $h_0 = h_{\text{cl}}$  and  $\Phi_0 = \Phi_{\text{cl}}$  starting with the contrastive learning model rather than just source-only ERM. Here, we only update  $h$  and fix  $\Phi_{\text{cl}}$ .

### D.3 Additional empirical results in our simplified setup

We conduct two ablations on the hyperparameters for contrastive pretraining. First, we vary the dimensionality  $k$  of the linear feature extractor  $\Phi \in \mathbb{R}^{k \times d}$ . Second, we vary the regularization strength  $\kappa$  that enforces feature diversity in the Barlow Twins objective (7). In Figure 5 we plot these ablations in the UDA setup.

**Varying feature dimension.** We find that CL recovers the full set of predictive features (*i.e.* both spurious and invariant) only when  $k$  is large enough (Figure 5(*left*)). Since the dimensionality of the true feature is 5 in our Example 1, reducing  $k$  below the true feature dimension hurts CL. Once  $k$  crosses a certain threshold, CL features completely capture the projection of the invariant feature  $w_{\text{in}}$ . After this point, it amplifies the component along  $w_{\text{in}}$ . It retains the amplification over the spurious feature  $w_{\text{sp}}$  even as we increase  $k$ . This is confirmed by our finding that further increasing  $k$  does not hurt CL performance. This is also inline with our theoretical observations, where we find that for




Figure 6: **Results with linear backbone:** We plot the OOD accuracy for ERM, CL, ST and STOC in the UDA setup and ID accuracy in the SSL setup when the feature extractor  $\Phi$  is a linear network. Note, that the feature extractor is still fixed during CL and STOC.

suitable  $w^*$ , the subspace spanned by  $w_{in}$  and  $w_{sp}$  are contained in a low rank space (as low as rank 2) of the contrastive representations (Theorem 5). Once CL has amplified the dependence along  $w_{in}$  STOC improves over CL by unlearning any remaining dependence on the spurious  $w_{sp}$ . The above arguments for the CL trend also explain why the performance of STOC continues to remain 100% as we vary  $k$ .

**Varying regularization strength.** In our main theoretical arguments we consider the constrained form of the Barlow Twins objective (3) with a strict constraint of  $\rho = 0$  (we relax this theoretically as well, see E.1.2). For our experiments, we optimize the regularized version of this objective (7), where the constraint term now appears as a regularizer which enforces feature diversity, *i.e.* the features learned through contrastive pretraining span orthogonal parts of the input space (as governed under the metric defined by augmentation covariance matrix  $\Sigma_A$ ). If  $\kappa$  is very low, then trivial solutions exist for the Barlow Twins objective. For *e.g.*,  $\phi = \mathbf{0}$  (zero vector) achieves very low invariance loss. When  $\kappa = 0.05$ , we find that CL recovers these trivial solutions (Figure 5(right)). Hence, both CL and STOC perform poorly. As we increase  $\kappa$  the performance of both CL and STOC improve, mainly because the features returned by  $\Phi_{cl}$  now comprise of the predictive directions  $w_{in}$  and  $w_{sp}$ , as predictive by our theoretical arguments for  $\rho = 0$  (which corresponds to large  $\kappa$ ). On the other hand, when  $\kappa$  is too high optimization becomes hard since  $\kappa$  directly effects the Lipschitz constant of the loss function. Hence, the performance of CL drops by some value. Note that this does not effect the performance of STOC since CL continues to amplify  $w_{in}$  over  $w_{sp}$  even if it is returning suboptimal solutions with respect to the optimization loss of the pretraining objective.

#### D.4 Reconciling Practice: Experiments with deep networks in toy setup

In this section we delve into the details of Sec. 4.5, *i.e.*, we analyze performance of different methods when we make some design choices that imitate practice. First, we look at experiments involving a deep non-linear backbone  $\Phi$ . Here, the non-linear  $\Phi$  is learned during contrastive pretraining and




Figure 7: **Results with non-linear backbone:** We plot the OOD accuracy for ERM, CL, ST and STOC in the UDA setup and ID accuracy in the SSL setup when the feature extractor  $\Phi$  is a non-linear one-hidden layer network with ReLU activations. Note, that the feature extractor is still fixed during CL and STOC.




Figure 8: **Finetuning the contrastive representations during STOC:** We propagate gradients to the feature backbone  $\Phi$  when running STOC algorithm. Note that CL still fixes the contrastive representations when learning a fixed linear head over it. On the (left) we show results in UDA setup where we compare the performance of STOC with and without augmentations (along with other practical design choices like confidence thresholds and continuing to optimize source loss as done in FixMatch) when the feature backbone is non-linear. On the (right) we show results for STOC and CL in the SSL setup when the feature backbone is linear.

fixed for CL and STOC. Then, we investigate trends when we continue to propagate gradients onto  $\Phi$  during STOC (we call this full-finetuning). Unlike previous cases, this allows features to be updated.

**Results with non-linear feature extractor  $\Phi$ .** In Fig. 7 we plot the performance of the four methods when we use a non-linear feature extractor during contrastive pretraining. This feature extractor is a one-hidden layer neural network (hidden dimension is 500) with ReLU activations. We find that the trends observed with linear backbones in Fig. 6 are also replicated with the non-linear one. Specifically, we note that STOC improves over CL under distribution shifts, whereas CL is already close to optimal when there are no distribution shifts. We also see that CL and ST individually are subpar. In SSL, we see a huge drop in the performance of ST (over ERM) mainly because we only fit on pseudolabels during ST. This is different from practice where we continue to optimize loss on labeled data points while fitting the pseudolabels. Consequently, when we continue to optimize performance on source labeled data the performance of ST in SSL setup is improves from 51.1%  $\bar{N}$  72.6%.

**Results with full fine-tuning.** Up till this point, we have only considered the case (for both SSL and UDA) where we fix the contrastive learned features when running CL and STOC, *i.e.*, we only optimized the linear head  $h$ . Now, we shall consider the setting where gradients are propagated to  $\Phi$  during STOC. Note that we still fix the representations for training the linear head during CL. Results for this setting are in Figure 8. We show two interesting trends that imitate real world behaviors.

*STOC benefits from augmentations during full-finetuning:* In the UDA setup we find that ST while updating  $\Phi_{cl}$  can hurt due to overfitting issues when training with the finite sample of labeled and unlabeled data (drop by  $\downarrow$  7% over CL). This is due to overfitting on confident but incorrect pseudolabels on target data. This can exacerbate components along spurious feature  $w_{sp}$  from source. One reasoning behind this is that deep neural networks can perfectly memorize them on finite unlabeled target data [94]. Heuristics typically used in practice (*e.g.* in FixMatch [79]) help avoid overfitting on incorrect pseudolabels: (i) confidence thresholding; to pick confident pseudolabel examples; (ii) pseudolabel a different augmented input than the one on which the self-training loss is optimized; and (iii) optimize source loss with labeled data simultaneously when fitting pseudolabels. Intuitively, thresholding introduces a curriculum where we only learn confident examples in the beginning whose pseudolabels are mainly determined by component along the invariant feature  $w_{in}$ . Augmentations prevent the neural network from memorizing incorrect pseudolabels and optimizing source loss prevents forgetting of features learned during CL. When we implement these during full-finetuning in STOC we see that STOC now improves over CL (by  $\downarrow$  20%).

*Can we improve contrastive pretraining features during STOC?* We find that self-training can also improve features learned during contrastive pretraining when we update the full backbone during STOC (see Figure 8(right)). Specifically, in the SSL setup we find that STOC can now improve substantially over CL. Recall, that when we fixed  $\Phi_{\text{cl}}$  this was not possible (see E.3 and Fig. 2(b)). This is mainly because STOC can now improve performance beyond just recovering the generalization gap for the linear head (which is typically small). This feature improvement is observed even when we fully finetune a linear feature extractor. Similar trends are also observed with the non-linear backbone. But, it becomes harder to identify a good stopping criterion for CL training. Thus, it remains unclear if STOC and CL have complementary benefits for feature learning in UDA or SSL settings. Investigating this is an interesting avenue for future work.

## E Formal Statements from Sec. 4

Recall from Section 2 that we learn linear classifiers  $h$  over features extractors  $\Phi$ . We consider linear feature extractor i.e.  $\Phi$  is a matrix in  $\mathbb{R}^{d \times k}$  and the linear layer  $h : \mathbb{R}^k \rightarrow \mathbb{R}$  with a prediction as  $\text{sgn}(h^\top \Phi x)$ . We use the exponential loss  $\ell(p(x), y) = \exp(-y f(p(x)))$ .

### E.1 Analysis of ERM and ST: Formal Statement of Theorem 2

For ERM and ST, we train both  $h$  and  $\Phi$ . This is equivalent to  $\Phi = I_{d \times d}$  being identity and training a linear head  $h$ . Recall that the ERM classifier is obtained by minimizing the population loss on labeled source data:

$$h_{\text{ERM}} = \arg \min_h \mathbb{E}_{(x,y) \sim P_S} \ell(p(x), y). \quad (10)$$

We split Theorem 2 into Theorem 8 and Theorem 9. Before we characterize the ERM solution, we recall some additional notation. Define  $w_{\text{in}} = (w^*, 0, \dots, 0)^\top$ , and  $w_{\text{sp}} = (0, \dots, 0, \frac{1}{d_{\text{sp}}})^\top$ . The following proposition characterizes  $h_{\text{ERM}}$  and 0-1 error of the classifier on target:

**Theorem 8** (ERM classifier and its error on target). *ERM classifier obtained as in (10) is given by*

$$\frac{h_{\text{ERM}}}{\|h_{\text{ERM}}\|_2} = \frac{\gamma w_{\text{in}} + \frac{\alpha}{d_{\text{sp}}} w_{\text{sp}}}{\gamma^2 + \frac{\alpha}{d_{\text{sp}}}}.$$

The target accuracy of  $h_{\text{ERM}}$  is given by  $0.5 + \text{erfc} \left( \frac{\gamma^2 \alpha}{2 d_{\text{sp}} \sigma_{\text{sp}}} \right)$ .

*Proof.* To prove this theorem, we first derive a closed-form expression for the ERM classifier and then use Lemma 29 to derive its 0-1 error on target. For Gaussian data with the same covariance matrices for class conditional  $P_S(p|x) = 1$  and  $P_S(p|x) = 0$ , Bayes decision rule is given by the Fisher's linear discriminant direction (Chapter 4; Bishop [9]):

$$h(p|x) = \begin{cases} 1, & \text{if } h^\top x \geq 0 \\ 0, & \text{otherwise} \end{cases}$$

where  $h = \frac{2}{\gamma^2 + \frac{\alpha}{d_{\text{sp}}}} (\gamma w_{\text{in}} + \frac{\alpha}{d_{\text{sp}}} w_{\text{sp}})$ . Plugging  $h$  in Lemma 29 we get the desired result.  $\square$

ST performs ERM in the first stage using labeled data from the source, and then subsequently updates the head  $h$  by iteratively generating pseudolabels on the unlabeled target:

$$L_{\text{st}}(h) : \mathbb{E}_{P_T(x)} \ell(p(x), \text{sgn}(h^\top x)). \quad (11)$$

Starting with  $h_{\text{ST}}^0 = h_{\text{ERM}} / \|h_{\text{ERM}}\|_2$  (the classifier obtained with ERM) we perform the following iterative procedure for self-training:

$$h_{\text{ST}}^{t+1} = \frac{h_{\text{ST}}^t - \eta \nabla_h L_{\text{st}}(h_{\text{ST}}^t)}{\|h_{\text{ST}}^t - \eta \nabla_h L_{\text{st}}(h_{\text{ST}}^t)\|_2} \quad (12)$$

Next, we characterize ST solution:

**Theorem 9** (ST classifier and its error on target). *Starting with ERM solution, ST will lead to:*

- (i) (Necessary condition)  $h_{ST}^t \approx w_{sp}$  as  $t \rightarrow \infty$ , such that the target accuracy is 50% for all  $\sigma_{sp} \neq 1$  and  $\gamma \propto \frac{1}{2\sqrt{\sigma_{sp}}}$ .
- (ii) (Sufficient condition)  $h_{ST}^t \approx w_{in}$  as  $t \rightarrow \infty$ , such that the target accuracy is 100% when the problem parameters  $\gamma, \sigma_{sp}$  satisfy:  $\gamma \neq \sigma_{sp}$ .

*Proof.* The proof can be divided into two parts: (i) deriving closed-form expressions for updates on  $h_{ST}^t$  in terms of  $h_{ST}^{t-1}$  and (ii) obtaining conditions under which the component along  $w_{in}$  monotonically increases or decreases with  $t$  after re-normalizing the norm of updated  $h$ . For notation convenience, we denote  $h_{ST}$  with  $h$  in the rest of the proof.

**Part-1.** First, the loss of self-training with classifier  $h : \mathbb{R}^{d_{in}}, h_{sp} \in \mathbb{R}^{d_{sp}}$  where  $h_{in} \in \mathbb{R}^{d_{in}}$  and  $h_{sp} \in \mathbb{R}^{d_{sp}}$  is given by:

$$L_{st} \mathbb{P} h^T x, \text{sgn} \mathbb{P} h^T x \mathbb{Q} \quad (13)$$

$$\mathbb{E}_{\mathbb{P}(x)} \exp \left[ \text{sign} \mathbb{P} h^T x \mathbb{Q} \mathbb{P} h^T x \mathbb{Q} \right] \quad (14)$$

$$\mathbb{E}_{\mathbb{P}(x)} \exp \left[ |h^T x| \right] \quad (15)$$

$$\mathbb{E}_{\mathbb{P}(x)} \exp \left[ |h_{in}^T x_{in} - h_{sp}^T x_{sp}| \right] \quad (16)$$

$$\mathbb{E}_{y \sim \mathcal{U}\{-1,1\}, z \sim \mathcal{N}(0,1)} \exp \left[ \gamma y h_{in}^T w^* \right. \\ \left. \sigma_{in} \mathbb{P} \|h_{in}\|_2^2 \mathbb{P} h_{in}^T w^* \mathbb{Q}^2 \mathbb{Q} \sigma_{sp} \|h_{sp}\|_2 \right] z \quad (17)$$

$$\mathbb{E}_{z \sim \mathcal{N}(0,1)} \exp \left[ \left| \gamma h_{in}^T w^* \sigma_{in} \mathbb{P} \|h_{in}\|_2^2 \mathbb{P} h_{in}^T w^* \mathbb{Q}^2 \mathbb{Q} \sigma_{sp} \|h_{sp}\|_2 \right| z \right], \quad (18)$$

where (16) to (17) is implied by simply replacing the definition of target distribution and (17) to (18) is implied by the symmetry of the function with respect to  $y$  and  $-y$  due to the symmetry of the absolute function and Gaussian distribution. For a classifier  $h^t$ , we denote  $\mu_t = \gamma h_{in}^T w^*$  and  $\sigma_t = \sigma_{in} \mathbb{P} \|h_{in}^t\|_2^2 \mathbb{P} h_{in}^t w^* \mathbb{Q}^2 \mathbb{Q} \sigma_{sp} \|h_{sp}^t\|_2$ . With this notation, we can re-write the loss in (18) as  $L_{st} \mathbb{P} h^t \mathbb{Q} = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma_t^2)} \left[ \exp \left[ \mu_t z / \sigma_t \right] \right]$ .

Now we derive a closed-form expression of  $L_{st} \mathbb{P} h^t \mathbb{Q}$  in Lemma 30:

$$L_{st} \mathbb{P} h^t \mathbb{Q} = \frac{1}{2} \exp \left[ \frac{\sigma_t^2}{2} \mu_t \text{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{G}_t}{2} \right] \exp \left[ \frac{\sigma_t^2}{2} \mu_t \text{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{G}_t}{2} \right] \quad (19)$$

Define the Mill's ratio as  $\mathbb{P} x \mathbb{Q} = \exp \left[ -x^2 / 2 \right] \text{erfc} \left[ x / \sqrt{2} \right] / \sqrt{\pi}$  as in Baricz [5]. We will frequently use standard properties of the Mill's ratio. We list them in Lemma 21 for completeness. Define:

$$\alpha_1 \mathbb{P} \mu_t, \sigma_t \mathbb{Q} = \exp \left[ \frac{\sigma_t^2}{2} \mu_t \text{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{G}_t}{2} \right] \exp \left[ \frac{\sigma_t^2}{2} \mu_t \text{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{G}_t}{2} \right], \\ \frac{C}{\pi} \exp \left[ \frac{\mu_t^2}{2\sigma_t^2} \right] \mathbb{P} \sigma_t \frac{\mu_t}{\sigma_t} \mathbb{Q} \mathbb{P} \sigma_t \frac{\mu_t}{\sigma_t} \mathbb{Q} \quad (20)$$

$$\alpha_2 \mathbb{P} \mu_t, \sigma_t \mathbb{Q} = \exp \left[ \frac{\sigma_t^2}{2} \mu_t \text{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{G}_t}{2} \right] \exp \left[ \frac{\sigma_t^2}{2} \mu_t \text{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{G}_t}{2} \right] \\ \frac{C}{\pi} \exp \left[ \frac{\mu_t^2}{2\sigma_t^2} \right] \mathbb{P} \sigma_t \frac{\mu_t}{\sigma_t} \mathbb{Q} \mathbb{P} \sigma_t \frac{\mu_t}{\sigma_t} \mathbb{Q} \quad (21)$$

Let  $\tilde{h}^{t+1}$  denote the un-normalized gradient descent update at iterate  $t + 1$ . We have:

$$\tilde{h}^{t+1} = h^t - \eta \frac{\mathbb{B} L_{st} \mathbb{P} h^t \mathbb{Q}}{\mathbb{B} h} \quad (22)$$

Now we will individually argue about the update of  $\mathbf{h}^{t+1}$  along the first  $d_{\text{in}}$  dimensions and the last  $d_{\text{sp}}$  dimensions. First, we have:

$$\begin{aligned}
\mathbf{h}_{\text{in}}^{t+1} &= \mathbf{h}_{\text{in}}^t + \eta \frac{\mathbf{B} \mathcal{L}_{\text{st}} \mathbf{p} \mathbf{h}^t \mathbf{q}}{\mathbf{B} h_{\text{in}}} \\
&= \mathbf{h}_{\text{in}}^t + \frac{\eta}{2} \exp\left\{-\frac{\sigma_t^2}{2}\right\} \mu_t \operatorname{erfc}\left\{\frac{\mu_t}{2\sigma_t}\right\} \frac{\mathcal{G}^t}{2} \\
&\quad + \exp\left\{-\frac{\sigma_t^2}{2}\right\} \mu_t \operatorname{erfc}\left\{\frac{\mu_t}{2\sigma_t}\right\} \frac{\mathcal{G}^t}{2} \gamma w^* \\
&= \frac{\eta}{2} \exp\left\{-\frac{\sigma_t^2}{2}\right\} \mu_t \operatorname{erfc}\left\{\frac{\mu_t}{2\sigma_t}\right\} \frac{\mathcal{G}^t}{2} \\
&\quad + \exp\left\{-\frac{\sigma_t^2}{2}\right\} \mu_t \operatorname{erfc}\left\{\frac{\mu_t}{2\sigma_t}\right\} \frac{\mathcal{G}^t}{2} \\
&\quad + \frac{2}{\sigma_t} \frac{\mu_t^2}{\pi} \exp\left\{-\frac{\mu_t^2}{2\sigma_t^2}\right\} \mathbf{p} 2h_{\text{in}}^t \mathbf{p} 2h_{\text{in}}^{t\top} w^* \mathbf{q} w^* \mathbf{q} \sigma_{\text{in}}^2 \\
\mathbf{h}_{\text{in}}^t &= \frac{\eta}{2} \alpha_1 \mathbf{p} \mu_t, \sigma_t \mathbf{q} \gamma w^* + \frac{\eta}{2} \alpha_2 \mathbf{p} \mu_t, \sigma_t \mathbf{q} \mathbf{p} 2h_{\text{in}}^t \mathbf{p} 2h_{\text{in}}^{t\top} w^* \mathbf{q} w^* \mathbf{q} \sigma_{\text{in}}^2. \quad (23)
\end{aligned}$$

Notice that the update of  $\mathbf{h}_{\text{in}}^{t+1}$  is split into two components, one along  $w^*$  and the other along the orthogonal component  $2h_{\text{in}}^t \mathbf{p} 2h_{\text{in}}^{t\top} w^* \mathbf{q} w^*$ . We will now argue that since at initialization, the component along  $\mathbf{p} I - w^* w^{*\top} \mathbf{q}$  is zero then it will remain zero. In particular, we have:

$$\mathbf{h}_{\text{in}}^0 \top \mathbf{p} I - w^* w^{*\top} \mathbf{q} \mathcal{G} w^{*\top} \mathbf{p} I - w^* w^{*\top} \mathbf{q} = 0. \quad (24)$$

With (23), we can argue that if  $\mathbf{p} I - w^* w^{*\top} \mathbf{q} h_{\text{in}}^t = 0$ , then  $\mathbf{p} I - w^* w^{*\top} \mathbf{q} \mathbf{h}_{\text{in}}^{t+1} = 0$  implying that  $\mathbf{p} I - w^* w^{*\top} \mathbf{q} \mathbf{h}_{\text{in}}^t = 0$  for all  $t \geq 0$ . Hence, we have:

$$\begin{aligned}
\mathbf{h}_{\text{inv}}^{t+1} &= \mathbf{h}_{\text{in}}^t + \eta \frac{\mathbf{B} \mathcal{L}_{\text{st}} \mathbf{p} \mathbf{h}^t \mathbf{q}}{\mathbf{B} h_{\text{in}}} \\
&= \mathbf{h}_{\text{in}}^t + \frac{\eta}{2} \alpha_1 \mathbf{p} \mu_t, \sigma_t \mathbf{q} \gamma w^*. \quad (25)
\end{aligned}$$

Second, we have the update  $\mathbf{h}_{\text{sp}}^{t+1}$  given by:

$$\begin{aligned}
\mathbf{h}_{\text{sp}}^{t+1} &= \mathbf{h}_{\text{sp}}^t + \eta \frac{\mathbf{B} \mathcal{L}_{\text{st}} \mathbf{p} \mathbf{h}^t \mathbf{q}}{\mathbf{B} h_{\text{sp}}} \\
&= \mathbf{h}_{\text{sp}}^t + \frac{\eta}{2} \exp\left\{-\frac{\sigma_t^2}{2}\right\} \mu_t \operatorname{erfc}\left\{\frac{\mu_t}{2\sigma_t}\right\} \frac{\mathcal{G}^t}{2} \\
&\quad + \exp\left\{-\frac{\sigma_t^2}{2}\right\} \mu_t \operatorname{erfc}\left\{\frac{\mu_t}{2\sigma_t}\right\} \frac{\mathcal{G}^t}{2} \frac{2}{\sigma_t} \frac{\mu_t^2}{\pi} \exp\left\{-\frac{\mu_t^2}{2\sigma_t^2}\right\} h_{\text{sp}}^t \sigma_{\text{sp}}^2 \\
\mathbf{h}_{\text{sp}}^t &= \frac{\eta}{2} \alpha_2 \mathbf{p} \mu_t, \sigma_t \mathbf{q} h_{\text{sp}}^t \sigma_{\text{sp}}^2. \quad (26)
\end{aligned}$$

Re-writing the expressions (25) and (26) for the update of  $\mathbf{h}^{t+1}$ , we have:

$$\mathbf{h}_{\text{in}}^{t+1} = \mathbf{h}_{\text{in}}^t \mathbf{p} I - \frac{\eta}{2} \alpha_1 \mathbf{p} \mu_t, \sigma_t \mathbf{q} \gamma^2 \{ \mu_t \mathbf{q} \}. \quad (27)$$

$$\mathbf{h}_{\text{sp}}^{t+1} = \mathbf{h}_{\text{sp}}^t \mathbf{p} I - \frac{\eta}{2} \alpha_2 \mathbf{p} \mu_t, \sigma_t \mathbf{q} \sigma_{\text{sp}}^2 \mathbf{q}. \quad (28)$$

Here, we replace  $\mathbf{h}_{\text{sp}}^t = \mu_t w^* \{ \gamma$  in (25) to get (27). Updates in (27) and (28) show that  $\mathbf{h}_{\text{in}}^{t+1}$  remains in the direction of  $\mathbf{h}_{\text{in}}^t$  and  $\mathbf{h}_{\text{sp}}^{t+1}$  remains in the direction of  $\mathbf{h}_{\text{sp}}^t$ .

**Part-2.** Now we will derive conditions under which  $h_{\text{in}}^t$  and  $h_{\text{sp}}^t$  will show monotonic behavior for necessary and sufficient conditions. We will first argue the condition under which ST will provably fail and converge to a classifier with a random target performance. For this, at every  $t$ , if we have:

$$\frac{\|h_{\text{sp}}^{t+1}\|_2}{\|h_{\text{sp}}^t\|_2} \leq \|h_{\text{sp}}^t\|_2, \quad (29)$$

then we can argue that as  $t \rightarrow \infty$ , we have  $\|h_{\text{sp}}^t\|_2 \rightarrow 1$  and hence, the ST classifier will have random target performance. Thus, we will focus on conditions, under which the norm on  $\|h_{\text{sp}}^t\|_2$  increases with  $t$ . Re-writing (29), we have:

$$\|h_{\text{sp}}^{t+1}\|_2 \leq \|h_{\text{sp}}^t\|_2 \|h_{\text{sp}}^t\|_2 \quad (30)$$

$$\|h_{\text{sp}}^{t+1}\|_2 \leq \|h_{\text{sp}}^t\|_2 \|h_{\text{in}}^{t+1}\|_2 \|h_{\text{sp}}^t\|_2 \quad (31)$$

$$\|h_{\text{sp}}^{t+1}\|_2 \leq \|h_{\text{sp}}^t\|_2 \|h_{\text{in}}^{t+1}\|_2 \|h_{\text{sp}}^t\|_2 \quad (32)$$

$$\frac{\|h_{\text{sp}}^{t+1}\|_2}{\|h_{\text{sp}}^t\|_2} \leq \frac{\|h_{\text{in}}^{t+1}\|_2}{\|h_{\text{in}}^t\|_2}. \quad (33)$$

Plugging in (27) and (28) into (33), we get:

$$\left| 1 - \frac{\eta}{2} \alpha_2 \rho(\mu_t, \sigma_t) \sigma_{\text{sp}}^2 \right| \leq \left| 1 - \frac{\eta}{2} \alpha_1 \rho(\mu_t, \sigma_t) \gamma^2 \mu_t \right|. \quad (34)$$

For small enough  $\eta$ , we have the necessary condition for the failure of ST as:

$$\alpha_2 \rho(\mu_t, \sigma_t) \sigma_{\text{sp}}^2 \leq \alpha_1 \rho(\mu_t, \sigma_t) \gamma^2 \mu_t. \quad (35)$$

Now we show in Lemma 11 and Lemma 10 that if the conditions assumed in the theorem continue to hold, then we can success and failure respectively.  $\square$

**Lemma 10** (Necessary conditions for ST). *Define  $\alpha_1$  and  $\alpha_2$  as in (20) and (21) respectively. If  $\sigma_{\text{sp}} \ll 1$  and  $\gamma \propto \frac{1}{2\sqrt{\sigma_{\text{sp}}}}$ , then we have for all  $t$ :*

$$\alpha_2 \rho(\mu_t, \sigma_t) \frac{\sigma_{\text{sp}}^2 \mu_t}{\gamma^2} \propto \alpha_1 \rho(\mu_t, \sigma_t). \quad (36)$$

*Proof.* We upper bound and lower bound  $\alpha_1$  and  $\alpha_2$  by using the properties of  $\rho$ . Recall:

$$\alpha_1 \rho(\mu_t, \sigma_t) \leq \frac{C}{\pi} \exp\left(-\frac{\mu_t^2}{2\sigma_t^2}\right) \leq \frac{C}{\pi} \exp\left(-\frac{\mu_t}{\sigma_t}\right) \leq \frac{C}{\pi} \exp\left(-\frac{\mu_t}{\sigma_t}\right). \quad (37)$$

and

$$\alpha_2 \rho(\mu_t, \sigma_t) \leq \frac{C}{\pi} \exp\left(-\frac{\mu_t^2}{2\sigma_t^2}\right) \leq \frac{C}{\pi} \exp\left(-\frac{\mu_t}{\sigma_t}\right) \leq \frac{C}{\pi} \exp\left(-\frac{\mu_t}{\sigma_t}\right) \frac{2}{\sigma_t}. \quad (38)$$

We now use Taylor's expansion on  $\rho$  and we get:

$$\rho(\sigma_t) \leq \rho'(\sigma_t) \frac{\mu_t}{\sigma_t} \leq \rho(\sigma_t) \frac{\mu_t}{\sigma_t} \leq \rho(\sigma_t) \frac{\mu_t}{\sigma_t} \leq \rho'(\sigma_t) \frac{\mu_t}{\sigma_t} \leq R'' \frac{\mu_t}{\sigma_t}^2 \quad (39)$$

and similarly, we get:

$$\rho(\sigma_t) \leq \rho'(\sigma_t) \frac{\mu_t}{\sigma_t} \leq \rho(\sigma_t) \frac{\mu_t}{\sigma_t} \leq \rho(\sigma_t) \frac{\mu_t}{\sigma_t} \leq \rho'(\sigma_t) \frac{\mu_t}{\sigma_t} \leq R'' \frac{\mu_t}{\sigma_t}^2 \quad (40)$$



where  $R'' = r'' \rho \sigma_0 q$ . This is because  $r'' \rho q$  takes positive values and is a decreasing function in  $\sigma_t$  (refer to Lemma 21). We now lower bound  $\alpha_1 \rho \mu_t, \sigma_t q$  and upper bound  $\alpha_2 \rho \mu_t, \sigma_t q$ :

$$\frac{\alpha_1 \rho \mu_t, \sigma_t q}{\frac{2}{\pi} \exp \frac{\mu_t^2}{2\sigma_t^2}} \geq 2r' \rho \sigma_t q \frac{\mu_t}{\sigma_t} R'' \frac{\mu_t}{\sigma_t}^2 \quad (41)$$

$$\frac{\alpha_2 \rho \mu_t, \sigma_t q}{\frac{2}{\pi} \exp \frac{\mu_t^2}{2\sigma_t^2}} \leq 2r \rho \sigma_t q \frac{2}{2} R'' \frac{\mu_t}{\sigma_t}^2 \quad (42)$$

Substituting the lower bound and upper bound in (36) gives us the following as stricter a necessary condition (i.e., (43) implies (36)):

$$2r \rho \sigma_t q \frac{2}{2} R'' \frac{\mu_t}{\sigma_t}^2 \frac{2}{\sigma_t} \frac{\sigma_{sp}^2}{\gamma^2} \mu_t \leq 2r' \rho \sigma_t q \frac{\mu_t}{\sigma_t} R'' \frac{\mu_t}{\sigma_t}^2 \quad (43)$$

$$\bar{\theta} \bar{n} \frac{2r \rho \sigma_t q}{2} R'' \frac{\mu_t}{\sigma_t}^2 \frac{2}{\sigma_t} \frac{\sigma_{sp}^2}{\gamma^2} \leq 2r' \rho \sigma_t q \frac{1}{\sigma_t} R'' \frac{\mu_t}{\sigma_t}^2 \quad (44)$$

$$\bar{\theta} \bar{n} \frac{r \rho \sigma_t q}{R''} \frac{\mu_t}{\sigma_t}^2 \frac{1}{\sigma_t} \frac{\sigma_{sp}^2}{\gamma^2} \leq r \rho \sigma_t q \frac{1}{\sigma_t} \frac{R''}{2} \frac{\mu_t}{\sigma_t}^2 \quad (45)$$

$$\bar{\theta} \bar{n} R'' \frac{\mu_t}{\sigma_t}^2 \frac{\sigma_{sp}^2}{\gamma^2} \frac{R''}{2} \frac{\mu_t}{\sigma_t} \leq r \rho \sigma_t q \frac{1}{\sigma_t} \frac{1}{1} \frac{\sigma_{sp}^2}{\gamma^2} \quad (46)$$

$$\bar{\theta} \bar{n} R'' \frac{\mu_t^2}{\sigma_t} \frac{\sigma_{sp}^2}{\gamma^2} \frac{R''}{2} \frac{\mu_t}{\sigma_t} \leq \rho \sigma_t r \rho \sigma_t q \frac{1}{1} \frac{1}{1} \frac{\sigma_{sp}^2}{\gamma^2} \quad (47)$$

Now, we will argue the monotonicity of LHS and RHS in (47). Observe that LHS is increasing in  $\mu_t$  and decreasing in  $\sigma_t$  and RHS is decreasing in  $\sigma_t$  as  $\rho \sigma_t r \rho \sigma_t q \frac{1}{1} q$  is increasing (and the multiplier is negative). Moreover, if (47) holds true for maximum value of RHS and minimum of LHS, then we would have (36). Thus substituting  $\mu_t = \gamma$  and  $\sigma_t = \sigma_0$  in LHS and  $\sigma_t = \sigma_{sp}$  in RHS, we get:

$$R'' \frac{\gamma^2}{\sigma_0} \frac{\sigma_{sp}^2}{\gamma^2} \frac{R''}{2} \frac{\gamma}{\sigma_0} \leq \rho \sigma_{sp} r \rho \sigma_{sp} q \frac{1}{1} \frac{1}{1} \frac{\sigma_{sp}^2}{\gamma^2} \quad (48)$$

$$\bar{\theta} \bar{n} R'' \frac{\sigma_{sp}^2}{\sigma_0} \frac{R''}{2} \frac{\gamma}{\sigma_0} \leq \rho \sigma_{sp} r \rho \sigma_{sp} q \frac{1}{1} \frac{1}{1} \frac{\sigma_{sp}^2}{\gamma^2} \quad (49)$$

$$(50)$$

Taking  $\gamma = \frac{1}{2\sqrt{\sigma_{sp}}}$  and substituting  $R'' = r'' \rho \sigma_0 q$ :

$$\rho \sigma_{sp} \{4q r'' \rho \sigma_0 q \sigma_{sp} \leq \rho \sigma_{sp} r \rho \sigma_{sp} q \frac{1}{1} \frac{1}{1} \frac{4}{4} \sigma_{sp}^3\} \quad (51)$$

Analytically solving the above expression, we get that (51) is satisfied for all values of  $\sigma_{sp} \neq 1$  when  $d_{sp} \neq 1$ . For example, the expression in (51) is also satisfied for the problem parameter used in the running example of the main paper.

□

As a remark, we note that in the proof of Lemma 10, the conditions derived are loose because of the relaxations made to simplify the proof. In principle, the proof (and hence the conditions) can be tightened by carefully propagating second-order terms (which depend on  $\sigma_t$ ) in (40).

**Lemma 11** (Sufficiency conditions for ST). *Define  $\alpha_1$  and  $\alpha_2$  as in (20) and (21) respectively. If  $\sigma_{sp} \leq \gamma$ , then we have for all  $t$ :*

$$\alpha_2 \rho \mu_t, \sigma_t q \frac{\sigma_{sp}^2}{\gamma^2} \mu_t \geq \alpha_1 \rho \mu_t, \sigma_t q. \quad (52)$$

*Proof.* We upper bound and lower bound  $\alpha_1$  and  $\alpha_2$  by using the properties of  $r$   $p$   $q$ . Recall:

$$\alpha_1 p \mu_t, \sigma_t q \leq \frac{C}{\pi} \exp \left( -\frac{\mu_t^2}{2\sigma_t^2} \right) r(\sigma_t) \frac{\mu_t}{\sigma_t} r(\sigma_t) \frac{\mu_t}{\sigma_t} . \quad (53)$$

and

$$\alpha_2 p \mu_t, \sigma_t q \leq \frac{C}{\pi} \exp \left( -\frac{\mu_t^2}{2\sigma_t^2} \right) r(\sigma_t) \frac{\mu_t}{\sigma_t} r(\sigma_t) \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} . \quad (54)$$

We now use Taylor's expansion on  $r$   $p$   $q$  and we get:

$$r p \sigma_t q \approx r' p \sigma_t q \frac{\mu_t}{\sigma_t} \approx r(\sigma_t) \frac{\mu_t}{\sigma_t} \approx r p \sigma_t q + r' p \sigma_t q \frac{\mu_t}{\sigma_t} + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \quad (55)$$

and similarly, we get:

$$r p \sigma_t q \approx r' p \sigma_t q \frac{\mu_t}{\sigma_t} + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \approx r(\sigma_t) \frac{\mu_t}{\sigma_t} \approx r p \sigma_t q + r' p \sigma_t q \frac{\mu_t}{\sigma_t} + R'' \frac{\mu_t^2}{\sigma_t^2} \quad (56)$$

where  $R'' = r'' p \sigma_t q$ . This is because  $r' p q$  takes positive values and is a decreasing function in  $\sigma_t$  (refer to Lemma 21). We now lower bound  $\alpha_1 p \mu_t, \sigma_t q$  and upper bound  $\alpha_2 p \mu_t, \sigma_t q$ :

$$\alpha_1 p \mu_t, \sigma_t q \geq \frac{2}{\pi} \exp \left( -\frac{\mu_t^2}{2\sigma_t^2} \right) 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \quad (57)$$

$$\alpha_2 p \mu_t, \sigma_t q \leq \frac{2}{\pi} \exp \left( -\frac{\mu_t^2}{2\sigma_t^2} \right) 2r p \sigma_t q + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \frac{2}{\sigma_t} \quad (58)$$

Substituting the lower bound and upper bound in (52) gives us the following as stricter a sufficient condition (i.e., (59) implies (52)):

$$2r p \sigma_t q + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \frac{2}{\sigma_t} \frac{\sigma_{sp}^2 \mu_t}{\gamma^2} \geq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \quad (59)$$

$$\bar{\delta} \bar{\eta} \left( 2r p \sigma_t q + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \frac{2}{\sigma_t} \right) \geq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{\gamma^2}{\sigma_{sp}^2 \mu_t} \quad (60)$$

$$\bar{\delta} \bar{\eta} \left( 2r p \sigma_t q + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \frac{2}{\sigma_t} \right) - 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{\gamma^2}{\sigma_{sp}^2 \mu_t} \geq 0 \quad (61)$$

$$\bar{\delta} \bar{\eta} \left( 2r p \sigma_t q + \sigma_t r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \right) - 2r' p \sigma_t q \frac{\gamma^2}{\sigma_{sp}^2} \geq 0 \quad (62)$$

$$\bar{\delta} \bar{\eta} \left( 2r' p \sigma_t q + r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} \right) - 2r' p \sigma_t q \frac{\gamma^2}{\sigma_{sp}^2} \geq 0 \quad (63)$$

$$\bar{\delta} \bar{\eta} \left( r'' p \sigma_t q \frac{\mu_t^2}{\sigma_t^2} - 2r' p \sigma_t q \right) - 1 \frac{\gamma^2}{\sigma_{sp}^2} \geq 0 \quad (64)$$

Hence, when  $1 - \frac{\gamma^2}{\sigma_{sp}^2} \geq 0$ , we have condition in (64) hold true as  $r' p \sigma_t q$  is always negative. Hence, the condition  $\gamma \geq \sigma_{sp}$  gives us the necessary condition.  $\square$

### E.1.1 Proof of Proposition 3

For convenience, we first restate the Proposition 3 which gives us a closed form solution for (3) when  $\rho = 0$ . Then, we provide the proof, focusing first on the case of  $k = 1$ , and then showing that extension to  $k \geq 1$  is straightforward and renders the final form in the proposition that follows.

**Proposition 12** (Barlow Twins solution). *The solution for (3) is  $U_k^\top \Sigma_A^{-1/2}$  where  $U_k$  are the top  $k$  eigenvectors of  $\Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2}$ . Here,  $\Sigma_A : E_{a \sim P_A} \Gamma a a^\top$  is the covariance over augmentations, and  $\Sigma : E_{x \sim P_U} \Gamma \bar{x} \bar{x}^\top$  is the covariance matrix of mean augmentations  $\bar{x} = E_{P_A(a|x)} \Gamma a$ .*

*Proof.* We will use  $\phi^\top x$  to denote  $\phi^\top x$  where  $\phi \in \mathbb{R}^d$ . Throughout the proof, we use  $a$  to denote augmentation and  $x$  to denote the input. We will use  $P_A \Gamma a | x$  as the probability measure over the space of augmentations  $A$ , given some input  $x \in X$  (with corresponding density)  $p_A \Gamma a | x$ . Next, we use  $p_A \Gamma a$  to denote the density associate with the marginal probability measure over augmentations:  $P_A = \int_X P_A \Gamma a | x dP_U$ . Finally, the joint distribution over positive pairs  $A_+ \Gamma a_1, a_2$   $\times P_A \Gamma a_1 | x$   $\times P_A \Gamma a_2 | x$   $dP_U$ , gives us the positive pair graph over augmentations.

Before we solve the optimization problem in (3) for  $\Phi \in \mathbb{R}^{k \times d}$  for any general  $k$ , let us first consider the case where  $k = 1$ , i.e. we only want to find a single linear projection  $\phi$ . The constraint  $\rho = 0$ , transfers onto  $\phi$  in the following way:

$$E_{a \sim P_A} \Gamma \phi^\top a \phi^\top a = 1 \quad \phi^\top \Sigma_A \phi = 1 \quad (65)$$

Under the above constraint we want to minimize the invariance loss, which according to Lemma 22 is given by  $2 \int_A \phi^\top a \Gamma L \phi^\top a dP_A$ , where  $L \Gamma \phi^\top a \Gamma a$  is the following linear operator.

$$L \Gamma \phi^\top a \Gamma a = \int_A \frac{A_+ \Gamma a, a'}{P_A \Gamma a} \phi^\top a' \Gamma a da' \quad (66)$$

Based on the definition of the operator, we can reformulate the constrained optimization for contrastive pretraining as:

$$\arg \min_{\phi: \phi^\top \Sigma_A \phi = 1} \int_A \phi^\top a \Gamma L \phi^\top a dP_A \quad (67)$$

$$\Leftrightarrow \arg \min_{\phi: \phi^\top \Sigma_A \phi = 1} E_{a \sim P_A} \Gamma \phi^\top a \phi^\top a = \int_A \phi^\top a \Gamma \phi^\top a' \Gamma A_+ \Gamma a, a' da da' \quad (68)$$

$$\Leftrightarrow \arg \min_{\phi: \phi^\top \Sigma_A \phi = 1} E_{a \sim P_A} \Gamma \phi^\top a \phi^\top a = \int_X \int_A \phi^\top a \Gamma \phi^\top a' \Gamma p_A \Gamma a | x \Gamma p_A \Gamma a' | x dP_U \quad (69)$$

$$\Leftrightarrow \arg \min_{\phi: \phi^\top \Sigma_A \phi = 1} E_{a \sim P_A} \Gamma \phi^\top a \phi^\top a = \int_X \Gamma \bar{\phi}^\top \bar{x} \Gamma \bar{x}^\top dP_U, \quad (70)$$

where  $\bar{\phi}^\top \bar{x} = E_{a \sim P_A(\cdot|x)} \Gamma \phi^\top a \Gamma x = E_{c \sim \text{Unif}[0,1]} \Gamma r \Gamma \phi^\top p c \Gamma x$ . Note that,

$$\bar{\phi}^\top \bar{x}^2 = E_{c \sim \text{Unif}[0,1]} \Gamma r \Gamma \phi^\top p c \Gamma x^2 \quad (71)$$

$$\phi^\top \rho E_{c \sim \text{Unif}[0,1]} \Gamma r c \Gamma x \Gamma \phi^\top x = E_{c \sim \text{Unif}[0,1]} \Gamma r c \Gamma x \Gamma \phi^\top x \quad (72)$$

$$\Leftrightarrow \int_X \Gamma \bar{\phi}^\top \bar{x} \Gamma \bar{x}^\top dP_U = \phi^\top \Sigma \phi \quad (73)$$

Further, since  $E_{a \sim P_A} \Gamma \phi^\top a \phi^\top a = \phi^\top \Sigma \phi$  we can now rewrite our main optimization problem for  $k = 1$  as:

$$\arg \min_{\phi: \phi^\top \Sigma_A \phi = 1} \phi^\top \Sigma \phi \quad (74)$$

$$\arg \max_{\phi: \phi^\top \Sigma_A \phi = 1} \phi^\top \Sigma \phi \quad (75)$$

Recall that in our setup both  $\Sigma$  and  $\Sigma_A$  are positive definite and invertible matrices. To solve the above problem, let's consider a re-parameterization:  $\phi' = \Sigma_A^{-1/2} \phi$ , thus  $\phi^\top \Sigma_A \phi = 1$ , is equivalent to the constraint  $\|\phi'\|_2^2 = 1$ . Based on this re-parameterization we are now solving:

$$\arg \max_{\|\phi'\|_2=1} \phi'^\top \Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2} \phi', \quad (76)$$

which is nothing but the top eigenvector for  $\Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2}$ .

Now, to extend the above argument from  $k = 1$  to  $k \leq 1$ , we need to care of one additional form of constraint in the form of feature diversity:  $\phi_i^\top \Sigma_A \phi_j = 0$  when  $i \neq j$ . But, we can easily redo the reformulations above and arrive at the following optimization problem:

$$\arg \max_{\substack{\|\phi'_i\|_2^2 = 1, \text{ @ } i \\ \phi_i'^\top \phi_j' = 0, \text{ @ } i \neq j}} \phi'_1, \phi'_2, \dots, \phi'_k \top \Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2} \phi'_1, \phi'_2, \dots, \phi'_k, \quad (77)$$

where  $\phi'_i = \Sigma_A^{-1/2} \phi_i$ . The above is nothing but the top  $k$  eigenvectors for the matrix  $\Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2}$ . This completes the proof of Proposition 12.  $\square$

### E.1.2 Analysis with $\rho \leq 0$ in Contrastive Pretraining Objective (3)

In (3) we considered the strict version of the optimization problem where  $\rho = 0$ . Here, we will consider the following optimization problem that we optimize for our experiments in the simplified setup:

$$L_{\text{clp}}(\Phi, \kappa) : \mathbb{E}_{x \sim P_U} \mathbb{E}_{a_1, a_2 \sim P_A(\cdot|x)} \|\Phi p a_1 - \Phi p a_2\|_2^2 - \kappa \mathbb{E}_{a \sim P_A} \Phi p a q \Phi p a q^\top - \mathbf{I}_k \frac{2}{F}, \quad (78)$$

where  $\kappa \geq 0$  is some finite constant (note that every  $\rho$  corresponds to some  $\kappa$  and particularly  $\rho = 0$ , corresponds to  $\kappa = 8$ ). Let  $\Phi^*$  be the solution for (3) with  $\rho = 0$ , i.e. the solution described in Proposition 3. Now, we will show that in practice we can provably recover something close to  $\Phi^*$  when  $\kappa$  is large enough.

**Theorem 13** (Solution for (78) is approximately equal to  $\Phi^*$ ). *If  $\Phi$  is some solution that achieves low values of the objective  $L_{\text{clp}}(\Phi, \kappa)$  in (78), i.e.,  $L_{\text{clp}}(\Phi, \kappa) \leq \epsilon$ , then there exists matrix  $W \in \mathbb{R}^{k \times k}$  such that:*

$$\mathbb{E}_{a \sim P_A} \|W \Phi^* p a - \Phi p a\|_2^2 \leq \frac{k\epsilon}{2\gamma_{k+1}},$$

$$\text{where, } \gamma_{k+1} \leq \frac{2\gamma_1^2}{k\epsilon} \leq 1 - \frac{\epsilon}{\kappa} \leq \frac{\gamma_1}{k},$$

where  $\gamma_{k+1}$  is the  $(k+1)^{\text{th}}$  eigenvalue for  $\mathbf{I}_d - \Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2}$ . Here,  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ .

*Proof.* Since we know that  $L_{\text{clp}}(\Phi, \kappa) \leq \epsilon$ , we can individually bound the invariance loss and the regularization term:

$$\mathbb{E}_{x \sim P_U} \mathbb{E}_{a_1, a_2 \sim P_A(\cdot|x)} \|\Phi p a_1 - \Phi p a_2\|_2^2 \leq \epsilon \quad (79)$$

$$\mathbb{E}_{a \sim P_A} \Phi p a q \Phi p a q^\top - \mathbf{I}_k \frac{2}{F} \leq \frac{\epsilon}{\kappa} \quad (80)$$

Thus,

$$\text{@ } i \text{ P rks : } 1 - \frac{\epsilon}{\kappa} \leq \Phi_i^\top \Sigma_A \Phi_i \leq 1 - \frac{\epsilon}{\kappa} \quad (81)$$

$$\text{@ } i \text{ P rks : } \mathbb{E}_{x \sim P_U} \mathbb{E}_{a_1, a_2 \sim P_A(\cdot|x)} \Phi_i^\top a_1 - \Phi_i^\top a_2 \|^2 \leq \epsilon \quad (82)$$

Let  $\phi_1^*, \phi_2^*, \phi_3^*, \dots, \phi_d^*$  be the solution returned by the analytical solution for  $\rho = 0$ , i.e. the solution in Proposition 3. Now, since  $\Phi^*$  would span  $\mathbb{R}^d$  when  $\Sigma_A$  is full rank, we can denote:

$$\Phi_i = \sum_{j=1}^d \eta_i^{(j)} \phi_j^* \quad (83)$$

Now from Lemma 22, the invariance loss for  $\Phi_i$  can be written using the operator  $Lp\phi qpaq$   $\phi paq$   $\int_A \frac{A(a, a')}{p_A(a)} \phi pa' q da'$ :

$$\text{Invariance Loss}_{p\Phi_i q} : E_{x \sim P_U} E_{a_1, a_2 \sim P_A(\cdot | x)} p\Phi_i^T a_1 \Phi_i^T a_2 q^2 \quad (84)$$

$$= 2 E_{a \sim P_A} r\Phi_i paq Lp\Phi_i qpaqs \quad (85)$$

$$= 2 E_{a \sim P_A} \sum_{j=1}^d \eta_i^{(j)} \phi_j^* L \sum_{j=1}^d \eta_i^{(j)} \phi_j^* paq \quad (86)$$

$$= 2 E_{a \sim P_A} \sum_{j=1}^d \eta_i^{(j)} \phi_j^* \sum_{j=1}^d \eta_i^{(j)} Lp\phi_j^* qpaq \quad (87)$$

$$= 2 \sum_{j=1}^d \eta_i^{(j) 2} E_{a \sim P_A} \phi_j^* paq Lp\phi_j^* qpaq \quad (88)$$

$$= 2 \sum_{m=1, n=1, m \neq n}^d \eta_i^{(m)} \eta_i^{(n)} E_{a \sim P_A} r\phi_m^* paq Lp\phi_n^* qpaqs \quad (89)$$

Since,  $\phi_i^* p q$  are eigenfunctions of the operator  $L$  [38], we can conclude that:

$$\sum_{m=1, n=1, m \neq n}^d \eta_i^{(m)} \eta_i^{(n)} E_{a \sim P_A} r\phi_m^* paq Lp\phi_n^* qpaqs = 0,$$

and if  $\gamma_1 \leq \gamma_2 \leq \gamma_3 \dots \leq \gamma_d$  are the eigenvalues for  $\phi_1^*, \phi_2^*, \phi_3^*, \dots, \phi_d^*$  under the decomposition of  $Lp\phi qpaq$  then:

$$E_{x \sim P_U} E_{a_1, a_2 \sim P_A(\cdot | x)} p\Phi_i^T a_1 \Phi_i^T a_2 q^2 = 2 \sum_{j=1}^d \gamma_j \eta_i^{(j) 2} \quad (90)$$

Recall, we are also aware of a condition on the regularization term:  $1 - \frac{\epsilon}{\kappa} \leq \Phi_i^T \Sigma_A \Phi_i \leq 1 + \frac{\epsilon}{\kappa}$ .

$$\Phi_i^T \Sigma_A \Phi_i = \sum_{j=1}^d \eta_i^{(j)} \phi_j^* \Sigma_A \sum_{j=1}^d \eta_i^{(j)} \phi_j^* = \sum_{j=1}^d \eta_i^{(j) 2} \quad (91)$$

$$\Rightarrow 1 - \frac{\epsilon}{\kappa} \leq \sum_{j=1}^d \eta_i^{(j) 2} \leq 1 + \frac{\epsilon}{\kappa} \quad (92)$$

In order to show that the projection of  $\Phi_i$  on  $\Phi^*$  is significant, we need to argue that the term  $\sum_{j=k+1}^d \eta_i^{(j) 2}$  is small. The argument for this begins with the condition on invariance loss, and the fact that  $\gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_k \leq \gamma_{k+1} \leq \dots \leq \gamma_d$ :

$$\frac{\epsilon}{2} \leq \sum_{j=k+1}^d \eta_i^{(j) 2} \gamma_j \leq \gamma_{k+1} \sum_{j=k+1}^d \eta_i^{(j) 2} \quad (93)$$

$$\Rightarrow \sum_{j=k+1}^d \eta_i^{(j) 2} \leq \frac{\epsilon}{2\gamma_{k+1}} \quad (94)$$

Extending the above result @i by simply adding the bounds completes the claim of our first result in Theorem 13. Next, we will lower bound the eigenvalue  $\gamma_{k+1}$ . Recall that,  $\sum_{j=1}^k \eta_i^{(j)2} \leq \frac{\epsilon}{\kappa} \frac{\epsilon}{2\gamma_{k+1}}$ . Thus,

$$\gamma_{k+1} \geq 1 - \frac{C}{\kappa} \frac{\epsilon}{2\gamma_{k+1}} \leq \sum_{j=1}^k \gamma_j \eta_i^{(j)2} \leq k\gamma_{k+1} \frac{\epsilon}{2\gamma_1} \quad (95)$$

We assume that all eigenvalues are strictly positive, which is true under our augmentation distribution. Given,  $\gamma_{k+1} \leq \gamma_1$ , we can rearrange the above to get:

$$\gamma_{k+1} \leq \frac{2\gamma_1^2}{k\epsilon} \leq 1 - \frac{C}{\kappa} \frac{\gamma_1}{k} \quad (96)$$

This completes the claim of our second result in Theorem 13.  $\square$

### E.1.3 Proof of Theorem 5

In this section, we prove our main theorem about the recovery of both spurious  $w_{\text{sp}}$ , invariant  $w_{\text{in}}$  features by the contrastive learning feature backbone, and also the amplification of the invariant over the spurious feature (where amplification is defined relatively with respect to what is observed in the data distribution alone). We begin by defining some quantities needed for analysis, that are fully determined by the choice of problem parameters for the model in (6).

From Section 4, we recall the definitions of  $w_{\text{in}} : \{r w^*, 0, \dots, 0\}$  and  $w_{\text{sp}} : \{r_0, \dots, 0, w\}$  where  $w' = \frac{1}{d_{\text{sp}}} \{ \sqrt{d_{\text{in}}}\mathbf{q} \}$ . Let us now define  $u_1, u_2$  as the top two eigenvectors of  $\Sigma_A$  with eigenvalues  $\lambda_1, \lambda_2$  (note that in our problem setup both  $\Sigma_A$  and  $\Sigma$  are full rank positive definite matrices), and  $\tau = \frac{\lambda_1}{\lambda_1 + \lambda_2}$ . Next we define  $\alpha$  as the angle between  $u_1$  and  $w_{\text{in}}$ , i.e.,  $\cos \alpha = u_1^\top w_{\text{in}}$ . Based on the definitions of  $\alpha$  and  $\tau$ , both of which are fully determined by the eigen decomposition of the post-augmentation feature covariance matrix  $\Sigma_A$ , we now restate Theorem 5:

**Theorem 14** (Formal; CL recovers both invariant  $w_{\text{in}}$  and spurious  $w_{\text{sp}}$  but amplifies  $w_{\text{in}}$ ). *Under Assumption 4  $w^* = \frac{1}{d_{\text{in}}} \{ \sqrt{d_{\text{in}}}\mathbf{q} \}$ , the CL solution  $\Phi_{\text{cl}} = \{ \phi_1, \phi_2, \dots, \phi_k \}$  satisfies  $\phi_j^\top w_{\text{in}} = \phi_j^\top w_{\text{sp}} = 0$  @  $j \neq 3$ . For  $\tau, \alpha$  as defined above, the solution for  $\phi_1, \phi_2$  is:*

$$\begin{aligned} w^* &= \cot \alpha \tau, & w^* &= \cos \theta, & \sin \theta \\ w' &= 1/\tau, & w' &= \cot \alpha \tau, & \sin \theta, & \cos \theta \end{aligned}$$

where  $0 \leq \alpha, \theta \leq \pi/2$ . Let us redefine  $\phi_1 = c_1 w_{\text{in}} + c_3 w_{\text{sp}}$  and  $\phi_2 = c_2 w_{\text{in}} + c_4 w_{\text{sp}}$ .

For constants  $K_1, K_2 \geq 0, \gamma = K_1 K_2 \{ \sigma_{\text{sp}}, d_{\text{sp}} \} \sigma_{\text{sp}}^2 \{ K_2^2, \epsilon \} \geq 0, D \sigma_{\text{sp}0}$ , such that for  $\sigma_{\text{sp}} \leq \sigma_{\text{sp}0}$ :

$$\begin{aligned} \frac{K_1 K_2^2 d_{\text{in}}}{2L \sigma_{\text{in}}^2 p d_{\text{in}} - 1} \epsilon &\leq \frac{c_1}{c_3} \leq \frac{K_1 K_2^2 d_{\text{in}}}{2L \sigma_{\text{in}}^2 p d_{\text{in}} - 1} \epsilon \\ \frac{L}{\gamma} \frac{d_{\text{sp}}}{d_{\text{sp}}} &\leq \left| \frac{c_2}{c_4} \right| \leq \frac{L}{\gamma} \frac{d_{\text{sp}}}{d_{\text{sp}}} \epsilon, \end{aligned}$$

where  $L = 1 + K_2^2$ .

*Proof.* We will first show that the only components of interest are  $\phi_1, \phi_2$ . Then, we will prove conditions on the amplification of  $w_{\text{in}}$  over  $w_{\text{sp}}$  in  $\phi_1, \phi_2$ . Following is the proof overview:

- I. When  $w^* = \frac{1}{d_{\text{in}}} \{ \sqrt{d_{\text{in}}}\mathbf{q} \}$ , from the closed form expressions for  $\Sigma_A$  and  $\Sigma$ , show that the solution returned by solving the Barlow Twins objective depends on  $w_{\text{in}}$  and  $w_{\text{sp}}$  only through the first two components  $\phi_1, \phi_2$ .

II. For the components  $\phi_1, \phi_2$ , we will show that the dependence along  $w_{in}$  is amplified compared to  $w_{sp}$  when the target data sufficiently denoises the spurious feature (*i.e.*,  $\sigma_{sp}$  is sufficiently large).

**Part-I:**

We can divide the space  $\mathbb{R}^d$  into two subspaces that are perpendicular to each other. The first subspace is  $\mathcal{W} = \{b_1 w_{in} + b_2 w_{sp} : b_1, b_2 \in \mathbb{R}, i.e. \text{ the rank 2 subspace spanned by } w_{in} \text{ and } w_{sp}\}$ . The second subspace is  $\mathcal{W}_\perp$  where  $\mathcal{W}_\perp = \{u \in \mathbb{R}^d : u^\top w_{in} = 0, u^\top w_{sp} = 0\}$ . Then, from Lemma 23 we can conclude that the matrix  $\Sigma_A$  can be written as:

$$\Sigma_A = \Sigma_{A\mathcal{W}} + \Sigma_{A\mathcal{W}_\perp} \quad (97)$$

where  $\Sigma_{A\mathcal{W}_\perp} = E_{a \sim P_A} \Pi_{\mathcal{W}_\perp} \text{paqp} \Pi_{\mathcal{W}_\perp} \text{paqq}^\top$  is the covariance matrix in the null space of  $\mathcal{W}$ , and  $\Pi_{\mathcal{W}_\perp} \text{paq}$  is the projection of augmentation  $a$  into the null space of  $\mathcal{W}$ , *i.e.* the covariance matrix in the space of non-predictive (noise) features. Similarly we can define:

$$\Sigma = \Sigma_{\mathcal{W}} + \Sigma_{\mathcal{W}_\perp} \quad (98)$$

Here again  $\Sigma_{\mathcal{W}_\perp} = E_{x \sim P_U} \Pi_{\mathcal{W}_\perp} pE_{c \sim \text{Unif}[0,1]^d} p c \text{d} x \text{qqp} \Pi_{\mathcal{W}_\perp} pE_{c \sim \text{Unif}[0,1]^d} p c \text{d} x \text{qqq}^\top$  is the covariance matrix of mean augmentations after they are projected onto the null space of predictive features. The above decomposition also follows from result in Lemma 23.

From Proposition 3, the closed form expression for the solution returned by optimizing the Barlow Twins objective in (3) is  $U^\top \Sigma_A^{-1/2}$  where  $U$  are the top-k eigenvectors of:

$$\Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2} \quad (99)$$

When  $w^* = \frac{1}{d_{in}} \{ \frac{?}{d_{in}} \}$ , then  $\Sigma_{A\mathcal{W}_\perp} = \Sigma_{\mathcal{W}_\perp} B$  where  $B$  is a diagonal matrix with diagonal given by  $\frac{1}{3} \text{diag} \Sigma_{\mathcal{W}_\perp} q$ . Further, since  $\text{diag} \Sigma_{\mathcal{W}_\perp} q = p \mathbb{1}_d$  for some constant  $p > 0$ , the eigenvectors of  $\Sigma_{\mathcal{W}_\perp}$  and  $\Sigma_{A\mathcal{W}_\perp}$  are exactly the same. Hence, when we consider the SVD of the expression  $\Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2}$ , the matrices  $\Sigma_{A\mathcal{W}_\perp}$  and  $\Sigma_{\mathcal{W}_\perp}$  have no effect on the SVD components that lie along the span of the predictive features. In fact, we only need to consider two rank 2 matrices (first terms in (98), (97)) and only do the SVD of  $\Sigma_{A\mathcal{W}}^{-1/2} \Sigma_{\mathcal{W}} \Sigma_{A\mathcal{W}}^{-1/2}$ .

There are only two eigenvectors of  $\Sigma_{A\mathcal{W}}^{-1/2} \Sigma_{\mathcal{W}} \Sigma_{A\mathcal{W}}^{-1/2}$ . We use  $\lambda_1, \lambda_2$  to denote the eigenvalues of  $\Sigma_{A\mathcal{W}}$ , and  $\text{rcosp} \alpha q w^*, \text{rsinp} \alpha q w'^s^\top, \text{rsinp} \alpha q w^*, \text{rcosp} \alpha q w'^s^\top$  for the corresponding eigenvectors. Similarly, we use  $\lambda_1, \lambda_2$  to denote the eigenvalues of  $\Sigma_{\mathcal{W}}$ , and  $\text{rcosp} \beta q w^*, \text{rsinp} \beta q w'^s^\top, \text{rsinp} \beta q w^*, \text{rcosp} \beta q w'^s^\top$  for the corresponding eigenvectors. Let  $\text{SVD}_U p q$  denote the operation of obtaining the singular vectors of a matrix. Then, to compute the components of the final expression:  $\text{SVD}_U p \Sigma_A^{-1/2} \Sigma \Sigma_A^{-1/2} q^\top \Sigma_A^{-1/2}$  that lies along the span of predictive features (in  $\mathcal{W}$ ), we need only look at the decomposition of the following matrix:

$$\begin{pmatrix} \cos \theta & , & \text{sinp} \theta q \\ \text{sin} \theta & , & \text{cosp} \theta q \end{pmatrix} \text{SVD}_U \begin{pmatrix} \sqrt{\lambda_1} & , & 0 \\ 0 & , & \sqrt{\lambda_2} \end{pmatrix} \begin{pmatrix} \text{cosp} \alpha & \beta q & , & \text{sinp} \alpha & \beta q \\ \text{sinp} \alpha & \beta q & , & \text{cosp} \alpha & \beta q \end{pmatrix} \begin{pmatrix} \lambda_1 & & 0 \\ & \lambda_2 & \\ 0 & & \lambda_2 \end{pmatrix} \quad (100)$$

Based on the above definitions of  $\theta, \alpha, \lambda_1, \lambda_2$ , we can then formulate  $\phi_1$  and  $\phi_2$  in the following way:

$$\begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix} = \begin{pmatrix} w^* \frac{\cos(\alpha)}{\sqrt{\lambda_1}}, & w^* \frac{\sin(\alpha)}{\sqrt{\lambda_2}} \\ w' \frac{\sin(\alpha)}{\sqrt{\lambda_1}}, & w' \frac{-\cos(\alpha)}{\sqrt{\lambda_2}} \end{pmatrix} \begin{pmatrix} \cos \theta & \sin \rho \theta q \\ \sin \theta & \cos \rho \theta q \end{pmatrix} \quad (101)$$

To summarize, using arguments in Lemma 23 and the fact that  $w^* = \frac{1}{\sqrt{d_{\text{in}}}} \frac{?}{d_{\text{in}}}$ , we can afford to focus on just two rank two matrices  $\Sigma_{A_W}, \Sigma_{W}$  in the operation:  $\text{SVD}_U \rho \Sigma_A^{-1/2} q \Sigma \Sigma_A^{-1/2}$ . The other singular vectors from the SVD only impact directions that span  $W_{\perp}$ , and the singular vectors obtained by considering only the rank 2 matrices lie only in the space of  $W$ .

## Part-II:

From the previous part we obtained forms of  $\phi_1, \phi_2$  in terms of:  $\lambda_1, \lambda_2, \alpha, \theta$ , all of which are fully specified by the SVD of  $\Sigma_{A_W}$  and  $\Sigma_W$ . If we define  $\tau = \frac{\sqrt{\lambda_1}}{\sqrt{\lambda_2}}$ , we can evaluate  $c_1, c_2, c_3, c_4$  as:

$$c_1 = \frac{\cot \rho \alpha q}{\tau} \tan \rho \theta q \quad (102)$$

$$c_2 = 1 - \frac{\cot \rho \alpha q \tan \rho \theta q}{\tau} \quad (103)$$

$$c_3 = \frac{1}{\tau} \cot \rho \alpha q \tan \rho \theta q \quad (104)$$

$$c_4 = \frac{\tan \rho \theta q}{\tau} \cot \rho \alpha q \quad (105)$$

Now, we are ready to begin proofs for our claims on the amplification factors, *i.e.* on the ratios  $c_1\{c_3, |c_2\{c_4|$ .

We will first prove some limiting conditions for  $c_1\{c_3$ , followed by those on  $|c_2\{c_4|$ . For each of these conditions we will rely on the forms for  $c_1, c_2, c_3, c_4$  derived in the previous part, in terms of  $\alpha, \theta, \tau$  (where  $0 < \alpha, \theta < \pi/2$ ). We will also rely on some lemmas that characterize the asymptotic behavior of  $\alpha, \theta$  and  $\tau$  as we increase  $\sigma_{\text{sp}}$ . We defer the full proof of these helper lemmas to later sections.

### Asymptotic behavior of $c_1\{c_3$ .

From Lemma 25 and Lemma 26, when  $\gamma = K_1 \sqrt{z}$  and  $\sigma_{\text{sp}} = K_2 \sqrt{z}$ , then:

$$\lim_{z \rightarrow \infty} \frac{c_1}{c_3} = \frac{\cot \alpha}{1} \frac{\tau \tan \theta}{\tau \cot \alpha \tan \theta} = \lim_{z \rightarrow \infty} \tau \tan \theta = \frac{K_1 K_2^2}{\rho_1 K_2^2 q 2 \sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}} q}}, \quad (106)$$

where we apply Moore-Osgood when applying limits on intermediate forms. We can do this since  $\tau \tan \theta$  approaches a constant, and each of  $\cot \alpha, \tau$  and  $\tan \theta$  are continuous and smooth functions of  $z$  (see Lemma 24).

### Asymptotic behavior of $|c_2\{c_4|$ .

When we consider the limiting behavior of  $|c_2\{c_4|$ , as we increase  $z$  or equivalently  $\sigma_{\text{sp}}$  when  $\gamma = K_1 \sqrt{z}$  and  $\sigma_{\text{sp}} = K_2 \sqrt{z}$ , then we get:

$$\lim_{z \rightarrow \infty} \left| \frac{c_2}{c_4 z} \right| = \left| \frac{1 - \cot \rho \alpha q \tan \rho \theta q}{\frac{\tan(\theta) z}{\tau} \cot \rho \alpha q z} \right|. \quad (107)$$

From Lemma 26,  $\cot \alpha \tan \theta \rightarrow 0$ . Next, if we consider  $\lim_{z \rightarrow \infty} z \tan \theta = \lim_{z \rightarrow \infty} \tau \tan \theta = z \tau^2$ . For  $z \tau^2$ , we invoke Lemma 28, which states that when  $\gamma = K_1 \sqrt{z}$  and  $\sigma_{\text{sp}} = K_2 \sqrt{z}$ , then:

$$\lim_{z \rightarrow \infty} \frac{z}{\tau^2} = \frac{2 \sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}} q}}{1 - 4 \rho_1 K_2^2}. \quad (108)$$

Further, in our bound on  $c_1\{c_3$ , we derived that  $\tau \tan \theta \rightarrow \frac{K_1 K_2^2}{(1 + K_2^2) 2 \sigma_{\text{in}}^2 (1 - 1/d_{\text{in}})}$ . Once again using Moore-Osgood we can plug this along with (108) to get:



$$\lim_{z \rightarrow \infty} \frac{\tan \rho \theta q z}{\tau} = \frac{K_1 K_2^2}{\rho 1 \quad K_2^2 q \rho 3 \quad 4 K_2^2 q}. \quad (109)$$

Finally, from Lemma 27, when  $\gamma = K_1 \sqrt{z}$  and  $\sigma_{sp} = K_2 \sqrt{z}$ , then:

$$\lim_{z \rightarrow \infty} \frac{z}{\tan \alpha} = \frac{K_1}{\rho 1 \quad 4 \{3 K_2^2 q\}}. \quad (110)$$

Plugging, 109 and 110 into 107 we get the following limit:

$$\lim_{z \rightarrow \infty} \left| \frac{c_2}{c_4 z} \right| = \frac{1}{K_1} \frac{K_2^2}{4 \{3 K_2^2 q\}}. \quad (111)$$

Since  $z = K_1 \sqrt{d_{sp} \gamma}$ ,

$$\lim_{z \rightarrow \infty} \left| \frac{c_2 \gamma}{c_4 K_1 \sqrt{d_{sp} \gamma}} \right| = \frac{1}{K_1} \frac{K_2^2}{4 \{3 K_2^2 q\}} \quad \text{and} \quad \lim_{z \rightarrow \infty} \left| \frac{c_2 \gamma}{c_4 \sqrt{d_{sp} \gamma}} \right| = \frac{1}{K_1} \frac{K_2^2}{4 \{3 K_2^2 q\}}. \quad (112)$$

Since both  $c_1 \{c_3\}$  and  $j^{c_2} \{c_4\}$  are continuous functions of  $z$ , with  $\liminf_{z \rightarrow \infty}$  and  $\limsup_{z \rightarrow \infty}$  converging to the limits in 106 and 107 for both quantities respectively, we conclude that  $\forall \epsilon > 0$  there exists  $\sigma_{sp0}$  such that for all  $\sigma_{sp} \neq \sigma_{sp0}$ , the following is true:

$$\frac{K_1 K_2^2 d_{in}}{2 L \sigma_{in}^2 \rho d_{in} \quad 1 q} \in \neq \frac{c_1}{c_3} \neq \frac{K_1 K_2^2 d_{in}}{2 L \sigma_{in}^2 \rho d_{in} \quad 1 q} \in \quad (113)$$

$$\frac{\rho 1 \quad K_2^2 q \quad d_{sp}}{\gamma} \in \neq \left| \frac{c_2}{c_4} \right| \neq \frac{\rho 1 \quad K_2^2 q \quad d_{sp}}{\gamma} \in, \quad (114)$$

This completes both Part-I and Part-II of the proof for Theorem 5. □

### E.1.4 Proof of Corollary 6

**Corollary 15** (CL improves OOD error over ERM but is still imperfect). *For  $\gamma, \sigma_{sp}, d_{sp}$  defined as in Theorem 5,  $D_{\sigma_{sp1}}$  such that  $\sigma_{sp} \neq \sigma_{sp1}$ , the target accuracy of CL (linear predictor on  $\Phi_{cl}$ ) is  $\neq 0.5 \operatorname{erfc} \left[ \frac{L'}{\gamma \sqrt{2} \sigma_{sp} q} \right]$  and  $\neq 0.5 \operatorname{erfc} \left[ \frac{4L'}{\gamma \sqrt{2} \sigma_{sp} q} \right]$ , where  $L' = K_2^2 K_1 \{\sigma_{in}^2 (1 - 1/d_{in})\}$ . When  $\sigma_{sp1} \leq \sigma_{in} \leq 1/d_{in}$ , the lower bound on accuracy is strictly better than ERM from scratch.*

*Proof.* Recall from Theorem 14, all  $\phi_j$ , for  $j \neq 3$ , lie in the null space of  $w_{in}$  and  $w_{sp}$ . Since, the predictive features are strictly contained in the rank two space spanned by  $w_{in}$  and  $w_{sp}$ , without loss of generality we can restrict ourselves to the case where  $k = 2$ , and when doing training a head  $h = [h_1, h_2]^T \in \mathbb{R}^2$  over contrastive pretrained representations using source labeled data, we get the following max margin solution:

$$\begin{aligned} h_1 &= \frac{c_1 \gamma}{c_3} \frac{a}{d_{sp}} \\ h_2 &= \frac{c_2 \gamma}{c_4} \frac{a}{d_{sp}} \end{aligned} \quad (115)$$

Without loss of generality we can divide both  $h_1$  and  $h_2$  by  $h_1$  and get the final classifier to be  $\phi_1 = \frac{h_2}{h_1} \phi_2$ :

$$\begin{aligned} \rho c_1 w_{in} \quad c_3 w_{sp} q & \quad \frac{h_2}{h_1} \quad \rho c_2 w_{in} \quad c_4 w_{sp} q \\ & \quad \rho c_1 w_{in} \quad c_3 w_{sp} q \quad \frac{\rho c_2 \gamma}{\rho c_1 \gamma} \frac{c_4}{c_3} \frac{a}{d_{sp} q} \quad \rho c_2 w_{in} \quad c_4 w_{sp} q \end{aligned} \quad (116)$$

From Lemma 29, we can derive the target accuracy of the classifier  $h$  on top of CL representations to be the following:

$$0.5 \operatorname{erfc} \frac{c_1 \beta c_2}{c_3 \beta c_4} \frac{\gamma}{2\sigma_{\text{sp}}} \quad (117)$$

where  $\beta = \frac{c_2\gamma + c_4}{c_3 + \beta c_4} \frac{\gamma}{d_{\text{sp}}}$ .

Substituting  $\beta$  into the expression  $\frac{c_1 + \beta c_2}{c_3 + \beta c_4}$  we get:

$$\frac{c_1^2 \gamma}{c_1 c_3 \gamma} \frac{c_1 c_3}{c_3^2} \frac{a}{d_{\text{sp}}} \frac{c_2^2 \gamma}{c_2 c_4 \gamma} \frac{c_2 c_4}{c_4^2} \frac{a}{d_{\text{sp}}} \quad (118)$$

We first substitute expressions for  $c_1, c_2, c_3, c_4$  from (102), (103), (104) and (105) in the above expression. Then for  $\gamma = K_1 \frac{z}{\sigma_{\text{sp}}}, K_2 \frac{z}{\sigma_{\text{sp}}}$ , we substitute the expressions for  $\cot \alpha, \tan \theta$ , and  $\tau = \lambda_1 \lambda_2$  with their corresponding closed form expressions (as functions of  $z$ ) from Lemma 24. On the resulting expression we apply do repeated applications of L'Hôpital's rule to get the following result:

$$\lim_{z \rightarrow \infty} \frac{c_1^2 \gamma}{c_1 c_3 \gamma} \frac{c_1 c_3}{c_3^2} \frac{a}{d_{\text{sp}}} \frac{c_2^2 \gamma}{c_2 c_4 \gamma} \frac{c_2 c_4}{c_4^2} \frac{a}{d_{\text{sp}}} = \frac{2K_2^2 K_1}{\sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}}}} \quad (119)$$

Based on  $\gamma, d_{\text{sp}}, \sigma_{\text{sp}}$  defined in Theorem 5, and (119) we can conclude that  $D\sigma_{\text{sp}1}$  such that for all  $\sigma_{\text{sp}} \neq \sigma_{\text{sp}1}$ :

$$\frac{4K_2^2 K_1}{\sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}}}} \neq \frac{c_1^2 \gamma}{c_1 c_3 \gamma} \frac{c_1 c_3}{c_3^2} \frac{a}{d_{\text{sp}}} \frac{c_2^2 \gamma}{c_2 c_4 \gamma} \frac{c_2 c_4}{c_4^2} \frac{a}{d_{\text{sp}}} \neq \frac{K_2^2 K_1}{\sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}}}} \quad (120)$$

Finally, applying (120) to Lemma 29, we conclude the following: When  $\gamma = K_1 K_2 \frac{z}{\sigma_{\text{sp}}}, d_{\text{sp}} = \sigma_{\text{sp}}^2 K_2^2$ , there exists  $\sigma_{\text{sp}1}$ , such that for any  $\sigma_{\text{sp}} \neq \sigma_{\text{sp}1}$ , target accuracy of CL is at least  $0.5 \operatorname{erfc} \frac{L'}{\sqrt{2}\sigma_{\text{sp}}}$  and at most  $0.5 \operatorname{erfc} \frac{4L'}{\sqrt{2}\sigma_{\text{sp}}}$ , where  $L' = \frac{K_2^2 K_1}{\sigma_{\text{in}}^2 (1 - 1/d_{\text{in}})}$ .

**Comparison with ERM.** Recall from Theorem 8 the performance of ERM classifier (trained from scratch) is  $0.5 \operatorname{erfc} \frac{\gamma^2}{2d_{\text{sp}}\sigma_{\text{sp}}}$ . The lower bound on the performance of classifier over CL representations is strictly better than ERM when:

$$\frac{a\gamma}{d_{\text{sp}}} > L'$$

$$\frac{\partial \dot{u}}{\partial \dot{u}} \frac{K_2^2 K_1}{\sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}}}} \frac{a\gamma}{d_{\text{sp}}} > \frac{\partial \dot{u}}{\partial \dot{u}} \frac{K_2^2 K_1}{\sigma_{\text{in}}^2 \rho_1 \sqrt{d_{\text{in}}}} \frac{K_1 K_2^2}{\sigma_{\text{sp}}^2}$$

$$\frac{\partial \dot{u}}{\partial \dot{u}} \sigma_{\text{sp}} \frac{a}{\sigma_{\text{in}} \sqrt{1 - 1/d_{\text{in}}}} > \frac{\partial \dot{u}}{\partial \dot{u}} \sigma_{\text{sp}1} \frac{a}{\sigma_{\text{in}} \sqrt{1 - 1/d_{\text{in}}}}$$

This completes our proof of Corollary 6. □

## E.2 Analysis of STOC: Formal Statement of Theorem 7

Recall ERM solution over contrastive pretraining. We showed that without loss of generality when  $k$  (the output dimensionality of  $\Phi$ ) is greater than 2, we can restrict  $k$  to 2 and the  $\Phi$  can be denoted as  $\Gamma \phi_1, \phi_2 S^\top$  where  $\phi_1 = c_1 w^* + c_3 w_{\text{sp}}$  and  $\phi_2 = c_2 w^* + c_4 w_{\text{sp}}$ . The ERM solution of the linear head is then given by  $h_1, h_2 \in \mathbb{R}$ :

$$h_1 = c_1 \gamma + c_3 \frac{a}{d_{\text{sp}}}, \text{ and } h_2 = c_2 \gamma + c_4 \frac{a}{d_{\text{sp}}}. \quad (121)$$

STOC performs self-training of the linear head over the CL solution. Before introducing the result, we need some additional notation. Let  $h^t$  denote the solution of the linear head at iterate  $t$ . Without loss of generality, assume that the coefficients in  $\phi_1 = c_1 w_{\text{in}} + c_3 w_{\text{sp}}$  and  $\phi_2 = c_2 w_{\text{in}} + c_4 w_{\text{sp}}$  are such that  $c_2$  is positive and  $c_1, c_3$ , and  $c_4$  are negative. Moreover, for simplicity of exposition, assume that  $j_{c_4} \leq j_{c_3}$ .

**Theorem 16.** Under the conditions of Corollary 15 and when  $\frac{\gamma^2}{\sigma_{\text{sp}}} \asymp \frac{-c_3 - c_4}{(c_2 + c_1) \cdot j c_1 j} - \frac{c_4}{c_1 \cdot c_2}$ , the target accuracy of ST over CL is lower bounded by  $0.5 \operatorname{erfc} \left( \frac{c_4}{c_1 \cdot c_2} \right)$  with  $L \asymp 1$ .

Before proving Theorem 16, we first connect the condition  $\frac{\gamma^2}{\sigma_{\text{sp}}} \asymp \frac{-c_3 - c_4}{(c_2 + c_1) \cdot j c_1 j} - \frac{c_4}{c_1 \cdot c_2}$  with the result obtained with contrastive learning.

**Remark 1.** We first argue that  $\frac{-c_3 - c_4}{(c_2 + c_1) \cdot j c_1 j}$  term dominates and hence, if we have  $\frac{\gamma^2}{\sigma_{\text{sp}}} \asymp \frac{-c_3 - c_4}{(c_2 + c_1) \cdot j c_1 j}$ , then we get the result in Theorem 16. First, recall that as  $\sigma_{\text{sp}}$  increases, we have  $\left| \frac{c_3}{c_1} \right|$  converge to  $\frac{2L\sigma_{\text{in}}^2(d_{\text{in}}-1)}{K_1 K_2^2 d_{\text{in}}}$ ,  $c_2 \tilde{\rightarrow} 1$  and  $\frac{c_1}{c_2} \tilde{\rightarrow} 0$ . Using these limits, we get:

$$\frac{\gamma^2}{\sigma_{\text{sp}}} \asymp \frac{K_1^2}{K_2^2 z^{3/2}} \asymp \frac{2L\sigma_{\text{in}}^2 d_{\text{in}}}{K_1 K_2^2 d_{\text{in}}} \cdot \frac{1}{z^{3/2}}. \quad (122)$$

which reduces the following condition:  $d_{\text{sp}} \asymp K_1^2 K_2^{2/3} \frac{d_{\text{in}}}{2L\sigma_{\text{in}}^2(d_{\text{in}}-1)} \cdot \frac{2}{3}$ .

*Proof.* First, we create an outline of the proof. We argue about the updates of  $h^t$  showing that both  $h_1^t$  and  $h_2^t$  increase with  $j h_2^t j$  becoming greater than  $j h_1^t j$  for some large  $t$ . Then we show that  $j h_2^t j \asymp j h_1^t j$  is sufficient to obtain near-perfect target generalization.

**Part 1.** Recall the loss of used for self-training of  $h$ :

$$L_{\text{st}} \rho h^t q = \mathbb{E}_{P_T(x)} \ell \rho h^T \Phi x, \operatorname{sgn} \rho h^T \Phi x q q \quad (123)$$

$$\mathbb{E}_{P_T(x)} \exp \left[ -\rho h^T \Phi x \right] \quad (124)$$

$$\mathbb{E}_{z \sim N(0,1)} \operatorname{r exp} \rho \left[ j c_1 \gamma h_1 - c_2 \gamma h_2 - \rho c_3 \sigma_{\text{sp}} h_1 - c_4 \sigma_{\text{sp}} h_2 \right] z j q s. \quad (125)$$

Define  $\mu_t = c_1 \gamma h_1^t - c_2 \gamma h_2^t$  and  $\sigma_t = c_3 \sigma_{\text{sp}} h_1^t + c_4 \sigma_{\text{sp}} h_2^t$ . With this notation, we can re-write the loss in (125) as  $L_{\text{st}} \rho h^t q = \mathbb{E}_{z \sim N(0, \sigma_t^2)} \operatorname{r exp} \rho \left[ j \mu_t - z j q s$ .

Similar to the the treatment in Theorem 9, we now derive a closed-form expression of  $L_{\text{st}} \rho h^t q$  in Lemma 30:

$$L_{\text{st}} \rho h^t q = \frac{1}{2} \exp \left[ \frac{\sigma_t^2}{2} \mu_t \right] \operatorname{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{F}_t}{2} \exp \left[ \frac{\sigma_t^2}{2} \mu_t \right] \operatorname{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{F}_t}{2}. \quad (126)$$

Define:

$$A_1 \rho \mu_t, \sigma_t q = \exp \left[ \frac{\sigma_t^2}{2} \mu_t \right] \operatorname{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{F}_t}{2} \frac{c}{\pi} \exp \left[ \frac{\mu_t^2}{2\sigma_t^2} \right] \operatorname{r} \left[ \sigma_t \frac{\mu_t}{\sigma_t} \right], \quad (127)$$

$$A_2 \rho \mu_t, \sigma_t q = \exp \left[ \frac{\sigma_t^2}{2} \mu_t \right] \operatorname{erfc} \left[ \frac{\mu_t}{2\sigma_t} \right] \frac{\mathcal{F}_t}{2} \frac{c}{\pi} \exp \left[ \frac{\mu_t^2}{2\sigma_t^2} \right] \operatorname{r} \left[ \sigma_t \frac{\mu_t}{\sigma_t} \right], \quad (128)$$

$$A_3 \rho \mu_t, \sigma_t q = \frac{2}{\pi} \exp \left[ \frac{\mu_t^2}{2\sigma_t^2} \right]. \quad (129)$$

Let  $\mathfrak{h}^{t+1}$  denote the un-normalized gradient descent update at iterate  $t + 1$ . We have:

$$\mathfrak{h}^{t+1} = h^t - \eta \frac{\text{B}L_{\text{st}}\rho h^t \mathbf{q}}{\text{B}h}. \quad (130)$$

Now we will individually argue about the update of  $\mathfrak{h}^{t+1}$ . First, we have:

$$\begin{aligned} \mathfrak{h}_1^{t+1} &= h_1^t - \eta \frac{\text{B}L_{\text{st}}\rho h^t \mathbf{q}}{\text{B}h_1} \\ \mathfrak{h}_1^{t+1} &= h_1^t - \eta \underbrace{\left[ \frac{A_1 \rho \sigma_t c_3 \sigma_{\text{sp}}}{j} - \frac{c_1 \gamma \mathbf{q}}{h_1} - \frac{A_2 \rho \sigma_t c_3 \sigma_{\text{sp}}}{j} + \frac{c_1 \gamma \mathbf{q}}{h_1} - \frac{A_3 c_3 \sigma_{\text{sp}} \mathbf{S}}{\bar{n}} \right]}_{\delta_1}. \end{aligned} \quad (131)$$

and second, we have:

$$\begin{aligned} \mathfrak{h}_2^{t+1} &= h_2^t - \eta \frac{\text{B}L_{\text{st}}\rho h^t \mathbf{q}}{\text{B}h_2} \\ \mathfrak{h}_2^{t+1} &= h_2^t - \eta \underbrace{\left[ \frac{A_1 \rho \sigma_t c_4 \sigma_{\text{sp}}}{j} - \frac{c_2 \gamma \mathbf{q}}{h_2} - \frac{A_2 \rho \sigma_t c_4 \sigma_{\text{sp}}}{j} + \frac{c_2 \gamma \mathbf{q}}{h_2} - \frac{A_3 c_4 \sigma_{\text{sp}} \mathbf{S}}{\bar{n}} \right]}_{\delta_2}. \end{aligned} \quad (132)$$

We will now argue the conditions under which  $h_2^{t+1}$  increases till its value reaches  $1/\sqrt{2}$ . In particular, we will argue that when  $h_2^t$  is negative, the norm  $|h_2^t|$  decreases and when  $h_2^t$  becomes positive, then its norm increases. We show that the following three conditions are sufficient to argue the increasing value of  $h_2^t$ : for all  $t$ , we have (i)  $\mu_t \neq \mu_c$  and  $j\sigma_t/j = \sigma_c$  for constant  $\mu_c = j c_1 - \gamma j$ ; (ii)  $\delta_2 = 0$ ; (iii)  $j\delta_2/j \neq \delta_1$ . In Lemma 18, we argue that our assumption on the initialization of the backbone learned with BT implies the previous three conditions.

**Case-1.** When  $h_2^t$  is negative (and after the update, it remains negative). Then we want to argue the following:

$$\frac{\rho h_2^t}{\rho h_2^t} \frac{\eta \delta_2 \mathbf{q}^2}{\eta \delta_2 \mathbf{q}^2} \frac{\eta \delta_2 \mathbf{q}^2}{\rho h_1^t \eta \delta_1 \mathbf{q}^2} \preceq \rho h_2^t \mathbf{q}^2 \quad (133)$$

$$\bar{n} \frac{\rho h_2^t \eta \delta_2 \mathbf{q}^2}{\rho h_2^t \mathbf{q}^2} \preceq \rho h_2^t \eta \delta_2 \mathbf{q}^2 - \rho h_1^t \eta \delta_1 \mathbf{q}^2 \quad (134)$$

$$\bar{n} \frac{h_2^{t^2}}{\rho h_2^t \mathbf{q}^2} \frac{\eta^2 \delta_2^2 - 2\eta \delta_2 h_2^t}{\rho h_2^t \mathbf{q}^2} \preceq h_2^{t^2} \eta^2 \delta_2^2 - 2\eta h_2^t \delta_2 - h_1^{t^2} \eta^2 \delta_1^2 - 2\eta h_1^t \delta_1 \quad (135)$$

$$\bar{n} \left( 1 - \frac{\eta^2 \delta_2^2 - 2\eta \delta_2 h_2^t}{\rho h_2^t \mathbf{q}^2} \right) \preceq 1 - \eta^2 \delta_2^2 - 2\eta h_2^t \delta_2 - \eta^2 \delta_1^2 - 2\eta h_1^t \delta_1 \quad (136)$$

$$\bar{n} \eta^2 \delta_2^2 - 2\eta \delta_2 h_2^t \preceq \eta^2 \delta_2^2 - 2\eta h_2^t \delta_2 - \eta^2 \delta_1^2 - 2\eta h_1^t \delta_1 - \rho h_2^t \mathbf{q}^2 \quad (137)$$

$$\bar{n} \eta^2 \delta_2^2 \rho h_1^t \mathbf{q}^2 - 2\eta \delta_2 h_2^t \rho h_1^t \mathbf{q}^2 \preceq \eta^2 \delta_1^2 \rho h_2^t \mathbf{q}^2 - 2\eta h_1^t \delta_1 \rho h_2^t \mathbf{q}^2 \quad (138)$$

$$\bar{n} \eta^2 \delta_2^2 \rho h_1^t \mathbf{q}^2 - \eta^2 \delta_1^2 \rho h_2^t \mathbf{q}^2 \preceq 2\eta \delta_2 h_2^t \rho h_1^t \mathbf{q}^2 - 2\eta h_1^t \delta_1 \rho h_2^t \mathbf{q}^2 \quad (139)$$

$$\bar{n} \eta \delta_2 \rho h_1^t \mathbf{q} - \eta \delta_1 \rho h_2^t \mathbf{q} \preceq \eta \delta_2 \rho h_1^t \mathbf{q} - \eta \delta_1 \rho h_2^t \mathbf{q} \preceq 2h_2^t h_1^t \eta \delta_2 \rho h_1^t \mathbf{q} - \eta \delta_1 \rho h_2^t \mathbf{q} \quad (140)$$

$$\bar{n} \eta \delta_2 \rho h_1^t \mathbf{q} - \eta \delta_1 \rho h_2^t \mathbf{q} \preceq 2h_2^t h_1^t \eta \delta_2 \rho h_1^t \mathbf{q} \quad (141)$$

Since  $\delta_2 = 0$ ,  $j\delta_2/j \neq j\delta_1/j$  and  $h_2^t = h_1^t = 0$ , we have  $\eta \delta_2 \rho h_1^t \mathbf{q} - \eta \delta_1 \rho h_2^t \mathbf{q} \mathbf{S}$  as positive. This implies inequality (140) to (141) and for small enough  $\eta$ , (141) will continue to hold true.

**Case-2.** When  $h_2^t$  is positive but less than  $1\{\frac{?}{2}$ . Then we want to argue the following:

$$\frac{\rho h_2^t}{\rho h_2^t} \frac{\eta \delta_2 q^2}{\eta \delta_2 q^2} \neq \rho h_2^t q^2 \quad (142)$$

$$\bar{n} \quad \frac{\rho h_2^t}{\rho h_2^t q^2} \neq \rho h_2^t \quad \eta \delta_2 q^2 \quad \rho h_1^t \quad \eta \delta_1 q^2 \quad (143)$$

$$\bar{n} \quad \frac{h_2^{t^2}}{\rho h_2^t q^2} \neq h_2^{t^2} \quad \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \quad h_1^{t^2} \quad \eta^2 \delta_1^2 \quad 2\eta h_1^t \delta_1 \quad (144)$$

$$\bar{n} \quad 1 \quad \frac{\eta^2 \delta_2^2}{\rho h_2^t q^2} \neq 1 \quad \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \quad \eta^2 \delta_1^2 \quad 2\eta h_1^t \delta_1 \quad (145)$$

$$\bar{n} \quad \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \neq \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \quad \eta^2 \delta_1^2 \quad 2\eta h_1^t \delta_1 \quad \rho h_2^t q^2 \quad (146)$$

$$\bar{n} \quad \eta^2 \delta_2^2 \rho h_1^t q^2 \quad 2\eta \delta_2 h_2^t \rho h_1^t q^2 \neq \eta^2 \delta_1^2 \rho h_2^t q^2 \quad 2\eta h_1^t \delta_1 \rho h_2^t q^2 \quad (147)$$

$$\bar{n} \quad \eta^2 \delta_2^2 \rho h_1^t q^2 \quad \eta^2 \delta_1^2 \rho h_2^t q^2 \neq 2\eta \delta_2 h_2^t \rho h_1^t q^2 \quad 2\eta h_1^t \delta_1 \rho h_2^t q^2 \quad (148)$$

$$\bar{n} \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \neq 2h_2^t h_1^t \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad (149)$$

$$\bar{n} \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \neq 2h_2^t h_1^t \quad (150)$$

Since  $\delta_2 > 0, j\delta_2 j \neq j\delta_1 j, h_1^t \propto 1\{\frac{?}{2}$  and  $0 < h_2^t < 1\{\frac{?}{2}$ , we have  $\eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q$  as positive. This implies inequality (149) to (150). Focusing on (150), we note that  $h_1^t \delta_2$  is positive and greater in magnitude than  $h_2^t \delta_1$ . Moreover, since  $h_2^t h_1^t$  is negative, (150) will continue to hold true.

Now, when  $h_2^t$  is positive and greater than  $1\{\frac{?}{2}$ , then  $h_2^t$  will stay in that region. Convergence of STOC together with conditions of convergence as in Lemma 17 will imply that the at convergence  $h_2^t$  will remain greater than  $1\{\frac{?}{2}$ , such that  $\frac{h_1^{t_c}}{h_2^{t_c}} > \frac{\delta_1}{\delta_2}$ . Now we bound the target error of STOC.

**Part 2.** To bound the accuracy at any iterate  $t$  when  $h_2^t \neq 1\{\frac{?}{2}$ , we have from Lemma 29:

$$\mathbb{E}_{P_T} y \quad h^{t \top} \phi_{cl} x \quad i \quad 0 \quad \mathbb{E}_{z \sim N(0,1)} z \quad i \quad \frac{c_1 \gamma h_1^t}{j c_3 \sigma_{sp} h_1^t} \quad \frac{c_2 \gamma h_2^t}{c_4 \sigma_{sp} h_2^t j} \quad . \quad (151)$$

We now upper bound and lower bound the fraction  $\frac{c_1 \gamma h_1^t + c_2 \gamma h_2^t}{j c_3 \sigma_{sp} h_1^t + c_4 \sigma_{sp} h_2^t j}$  in RHS in (151): (i)  $c_1 \gamma h_1^t \quad c_2 \gamma h_2^t \neq c_2 \gamma h_2^t$  since both  $c_1 \gamma h_1^t$  and  $c_2 \gamma h_2^t$  have same sign; (ii)  $j c_3 \sigma_{sp} h_1^t \quad c_4 \sigma_{sp} h_2^t j \propto j c_4 \sigma_{sp} h_2^t j$  because  $j c_4 \sigma_{sp} h_2^t j \neq j c_3 \sigma_{sp} h_1^t j$  and they have opposite signs. Hence, from (151), we have:

$$\mathbb{E}_{P_T} y \quad h^{t \top} \phi_{cl} x \quad i \quad 0 \quad \mathbb{E}_{z \sim N(0,1)} z \quad i \quad \frac{c_2 \gamma h_2^t}{j c_4 \sigma_{sp} h_2^t j} \quad \mathbb{E}_{z \sim N(0,1)} z \quad i \quad \frac{c_2 \gamma}{j c_4 \sigma_{sp} j} \quad . \quad (152)$$

Substituting the definition of  $\text{erfc}$ , the expression (152) gives us the required lower bound on the target accuracy. □

**Lemma 17** (Convergence of STOC). *Assume the gradient updates as in (131) and (132). Then STOC converges at  $t = t_c$  when  $\frac{h_1^{t_c}}{h_2^{t_c}} > \frac{\delta_1}{\delta_2}$ . For  $t \leq t_c$ , (131) and (132) make no updates to the linear  $h$ .*

*Proof.* When the gradient updates  $\delta_1$  and  $\delta_2$  are such that  $h_1^{t+1}$  matches  $h_1^t$ , we have convergence of STOC.

$$\frac{\rho h_2^t}{\rho h_2^t} \frac{\eta \delta_2 q^2}{\eta \delta_2 q^2} \frac{\eta \delta_2 q^2}{\rho h_1^t} \frac{\eta \delta_1 q^2}{\eta \delta_1 q^2} \rho h_2^t q^2 \quad (153)$$

$$\bar{n} \quad \frac{\rho h_2^t}{\rho h_2^t} \frac{\eta \delta_2 q^2}{\rho h_2^t q^2} \rho h_2^t \quad \eta \delta_2 q^2 \quad \rho h_1^t \quad \eta \delta_1 q^2 \quad (154)$$

$$\bar{n} \quad \frac{h_2^{t^2}}{\rho h_2^t q^2} \frac{\eta^2 \delta_2^2}{\eta^2 \delta_2^2} \frac{2\eta \delta_2 h_2^t}{2\eta \delta_2 h_2^t} \quad h_2^{t^2} \quad \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \quad h_1^{t^2} \quad \eta^2 \delta_1^2 \quad 2\eta h_1^t \delta_1 \quad (155)$$

$$\bar{n} \quad 1 \quad \frac{\eta^2 \delta_2^2}{\rho h_2^t q^2} \frac{2\eta \delta_2 h_2^t}{2\eta \delta_2 h_2^t} \quad 1 \quad \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \quad \eta^2 \delta_1^2 \quad 2\eta h_1^t \delta_1 \quad (156)$$

$$\bar{n} \quad \eta^2 \delta_2^2 \quad 2\eta \delta_2 h_2^t \quad \eta^2 \delta_2^2 \quad 2\eta h_2^t \delta_2 \quad \eta^2 \delta_1^2 \quad 2\eta h_1^t \delta_1 \quad \rho h_2^t q^2 \quad (157)$$

$$\bar{n} \quad \eta^2 \delta_2^2 \rho h_1^t q^2 \quad 2\eta \delta_2 h_2^t \rho h_1^t q^2 \quad \eta^2 \delta_1^2 \rho h_2^t q^2 \quad 2\eta h_1^t \delta_1 \rho h_2^t q^2 \quad (158)$$

$$\bar{n} \quad \eta^2 \delta_2^2 \rho h_1^t q^2 \quad \eta^2 \delta_1^2 \rho h_2^t q^2 \quad 2\eta \delta_2 h_2^t \rho h_1^t q^2 \quad 2\eta h_1^t \delta_1 \rho h_2^t q^2 \quad (159)$$

$$\bar{n} \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad 2h_2^t h_1^t \quad \eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad (160)$$

Thus either  $\eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad 0$  or  $\eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad 2h_2^t h_1^t$ . Since  $\eta$  is such that  $h_1 \quad \eta \delta_1 \quad 0$ ,  $\eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad 2h_2^t h_1^t$  implying that  $\eta \delta_2 \rho h_1^t q \quad \eta \delta_1 \rho h_2^t q \quad 0$  giving us the required condition.  $\square$

**Lemma 18.** Under the initialization conditions assumed in Theorem 16, for all  $t$ , we have: (i)  $\mu_t \neq \mu_c$  and  $j\sigma_{tj} \propto \sigma_c$  for constant  $\mu_c \quad j \in \{1, 2\}$  and  $\sigma_c \quad j \in \{c_3, c_4\}$ ; (ii)  $\delta_2 \neq 0$ ; (iii)  $j\delta_{2j} \neq \delta_1$ , where  $\delta_1 \quad A_1 \quad \rho \sigma_t c_3 \sigma_{sp} \quad c_1 \gamma q \quad A_2 \quad \rho \sigma_t c_3 \sigma_{sp} \quad c_1 \gamma q \quad A_3 c_3 \sigma_{sp}$  and  $\delta_2 \quad A_1 \quad \rho \sigma_t c_4 \sigma_{sp} \quad c_2 \gamma q \quad A_2 \quad \rho \sigma_t c_4 \sigma_{sp} \quad c_2 \gamma q \quad A_3 c_4 \sigma_{sp}$  for  $A_1, A_2$  and  $A_3$  defined in (127), (128), and (129).

*Proof.* Recall,  $\mu_t \quad c_1 \gamma h_1^t \quad c_2 \gamma h_2^t$  and  $\sigma_t \quad c_3 \sigma_{sp} h_1^t \quad c_4 \sigma_{sp} h_2^t$ . First, we argue that  $\mu_t$  increases from the initialization value. Notice that  $\mu_0 \quad c_1 \gamma h_1^0 \quad c_2 \gamma h_2^0$ . Due to Corollary 15, we have  $h_2^0 \neq 0$ . And since  $j \in \{1, 2\}$ , we get  $\mu_0 \neq j \mu_c$  as both  $c_1$  and  $h_1^0$  are of same sign. Moreover, as training progresses with  $h_1^t$  remaining negative and  $h_2^t$  remaining positive, we have  $\mu_t$  stays greater than  $\mu_0$ .

Recall the definition of  $A_1, A_2$ , and  $A_3$  in (127), (128), and (129). Moreover, recall the definition of  $\alpha_1 \rho \mu_t, \sigma_t q$  and  $\alpha_2 \rho \mu_t, \sigma_t q$ :

$$\alpha_1 \rho \mu_t, \sigma_t q \quad \frac{C}{2\pi} \exp \left( -\frac{\mu_t^2}{2\sigma_t^2} \right) \quad \rho \quad \sigma_t \quad \frac{\mu_t}{\sigma_t} \quad \rho \quad \sigma_t \quad \frac{\mu_t}{\sigma_t} \quad . \quad (161)$$

and

$$\alpha_2 \rho \mu_t, \sigma_t q \quad \frac{C}{\pi} \exp \left( -\frac{\mu_t^2}{2\sigma_t^2} \right) \quad \rho \quad \sigma_t \quad \frac{\mu_t}{\sigma_t} \quad \rho \quad \sigma_t \quad \frac{\mu_t}{\sigma_t} \quad \frac{2}{\sigma_t} \quad . \quad (162)$$

Thus, we have  $\alpha_1 \rho \mu_t, \sigma_t q \quad A_3 \quad A_1 \quad \sigma_t$  and  $\alpha_2 \rho \mu_t, \sigma_t q \quad A_3 \quad \sigma_t \quad A_2 \quad \frac{2}{\sigma_t} A_3$ . Replacing the definition of  $A_1, A_2$ , and  $A_3$  in  $\delta_1$  and  $\delta_2$ , we get:

$$\delta_1 \quad \sigma_t c_3 \sigma_{sp} \quad \alpha_2 \rho \mu_t, \sigma_t q \quad c_1 \gamma \alpha_1 \rho \mu_t, \sigma_t q \quad \text{and} \quad \delta_2 \quad \sigma_t c_4 \sigma_{sp} \quad \alpha_2 \rho \mu_t, \sigma_t q \quad c_2 \gamma \alpha_1 \rho \mu_t, \sigma_t q \quad (163)$$

We now upper bound and lower bound  $\alpha_1$  and  $\alpha_2$  by using the properties of  $\rho \mu_t, \sigma_t q$ . We use Taylor's expansion on  $\rho \mu_t, \sigma_t q$  and we get:

$$\rho \mu_t, \sigma_t q \quad \rho' \mu_t, \sigma_t q \quad \frac{\mu_t}{\sigma_t} \quad \propto \quad \rho \quad \sigma_t \quad \frac{\mu_t}{\sigma_t} \quad \propto \quad \rho \mu_t, \sigma_t q \quad \rho' \mu_t, \sigma_t q \quad \frac{\mu_t}{\sigma_t} \quad \rho'' \mu_t, \sigma_t q \quad \frac{\mu_t^2}{\sigma_t^2} \quad (164)$$

and similarly, we get:

$$\rho \mu_t, \sigma_t q \quad \rho' \mu_t, \sigma_t q \quad \frac{\mu_t}{\sigma_t} \quad \rho'' \mu_t, \sigma_t q \quad \frac{\mu_t^2}{\sigma_t^2} \quad \propto \quad \rho \quad \sigma_t \quad \frac{\mu_t}{\sigma_t} \quad \propto \quad \rho \mu_t, \sigma_t q \quad \rho' \mu_t, \sigma_t q \quad \frac{\mu_t}{\sigma_t} \quad R'' \quad \frac{\mu_t^2}{\sigma_t^2} \quad (165)$$

where  $R'' = r'' p \sigma_t q$ . This is because  $r'' p q$  takes positive values and is a decreasing function in  $\sigma_t$  (refer to Lemma 21). We now lower bound  $\alpha_1 p \mu_t, \sigma_t q$  and upper bound  $\alpha_2 p \mu_t, \sigma_t q$ :

$$\frac{\alpha_1 p \mu_t, \sigma_t q}{\frac{2}{\pi} \exp \frac{\mu_t^2}{2\sigma_t^2}} \geq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \quad (166)$$

$$\frac{\alpha_2 p \mu_t, \sigma_t q}{\frac{2}{\pi} \exp \frac{\mu_t^2}{2\sigma_t^2}} \leq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} \quad (167)$$

**Part-1.** We first prove that  $\delta_2 \geq 0$ . Substituting the lower bound and upper bound in (163) gives us the following as stricter a sufficient condition (i.e., (168) implies  $\delta_2 \geq 0$ ):

$$2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} \frac{\sigma_{sp} p c_4 q}{\gamma c_2} \geq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \quad (168)$$

$$\bar{\delta} \bar{n} \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} \leq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{\gamma c_2}{\sigma_{sp} p c_4 q} \quad (169)$$

$$\bar{\delta} \bar{n} \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} - 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{\gamma c_2}{\sigma_{sp} p c_4 q} \geq 0 \quad (170)$$

$$\bar{\delta} \bar{n} \frac{\mu_t^2}{\sigma_t} - 2r' p \sigma_t q \mu_t \frac{\gamma c_2}{\sigma_{sp} p c_4 q} \geq 0 \quad (171)$$

$$\bar{\delta} \bar{n} \frac{\mu_t^2}{\sigma_t} - 2r' p \sigma_t q \mu_t \frac{\gamma c_2}{\sigma_{sp} p c_4 q} \geq 0 \quad (172)$$

$$\bar{\delta} \bar{n} \frac{\mu_t^2}{\sigma_t} - 2r' p \sigma_t q \mu_t \frac{\gamma c_2}{\sigma_{sp} p c_4 q} \geq 0 \quad (173)$$

Thus, if we have  $\mu_t \geq \frac{\sigma_{sp} p c_4}{\gamma c_2}$ , then (168) holds true.

**Part-2.** Next, we prove that  $j\delta_2 j \geq \delta_1$ . Substituting the lower bound and upper bound in (163) gives us the following as stricter a sufficient condition (i.e., (174) implies  $j\delta_2 j \geq \delta_1$ ):

$$2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} \frac{\sigma_{sp} p c_4 c_3 q}{\gamma p c_2 c_1 q} \geq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \quad (174)$$

$$\bar{\delta} \bar{n} \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} \leq 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{\gamma p c_2 c_1 q}{\sigma_{sp} p c_4 c_3 q} \quad (175)$$

$$\bar{\delta} \bar{n} \frac{\mu_t}{\sigma_t} \frac{2}{\sigma_t} - 2r' p \sigma_t q \frac{\mu_t}{\sigma_t} \frac{\gamma p c_2 c_1 q}{\sigma_{sp} p c_4 c_3 q} \geq 0 \quad (176)$$

$$\bar{\delta} \bar{n} \frac{\mu_t^2}{\sigma_t} - 2r' p \sigma_t q \mu_t \frac{\gamma p c_2 c_1 q}{\sigma_{sp} p c_4 c_3 q} \geq 0 \quad (177)$$

$$\bar{\delta} \bar{n} \frac{\mu_t^2}{\sigma_t} - 2r' p \sigma_t q \mu_t \frac{\gamma p c_2 c_1 q}{\sigma_{sp} p c_4 c_3 q} \geq 0 \quad (178)$$

$$\bar{\delta} \bar{n} \frac{\mu_t^2}{\sigma_t} - 2r' p \sigma_t q \mu_t \frac{\gamma p c_2 c_1 q}{\sigma_{sp} p c_4 c_3 q} \geq 0 \quad (179)$$

Thus, if we have  $\mu_t \geq \frac{\sigma_{sp} p c_4 c_3}{\gamma p c_2 c_1}$ , then (174) holds true which in-turn implies  $j\delta_2 j \geq \delta_1$ . Plugging in  $\mu_t \geq \mu_0$ , we get the required condition. □

### E.3 Analysis for SSL

For SSL analysis, we argue that the projection learned by contrastive pretraining can significantly improve the generalization of the linear head learned on top, leaving little to no room for improvement for self-training. Our analysis leverages the margin-based bound for linear models from Kakade et al. [45]. Before introducing the result, we present some additional notation. Let  $\text{Err}_D^{\gamma} p w q$  denote 0-1 error of a classifier on a distribution  $D$ . Define 0-1 error with margin  $\gamma$  as  $\mathbb{E}_{\text{Err}}^{\gamma} p w q = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\gamma y_i w^J x_i \leq \gamma S}$ .

**Theorem 19** (Corollary 6 in Kakade et al. [45]). *For all classifiers  $w$  and margin  $\gamma$ , we have with probability at least  $1 - \delta$ :*

$$\text{Err}_T p w q \leq \mathbb{E}_{\text{Err}}^{\gamma} p w q + 4 \frac{B}{\gamma} \frac{1}{n} \frac{C}{\log p 1 \{ \delta q \}} \frac{C}{\log p \log_2 p 4 B \{ \gamma q q \}}, \quad (180)$$

where  $B$  is an upper bound on the  $\ell_2$  norm of the input points  $x$ .

When  $\mathbb{E}_{\text{Err}}^{\gamma} p w q$  is close to zero, the denominating term in RHS of (180) is  $4 \frac{B}{\gamma} \frac{1}{n}$ . With Proposition 3, CL solution  $\phi$  obtained on the target domain alone (for SSL setup) is  $w_{in}$  when  $k = 1$ . For larger  $k$ 's the CL solution is dominated by the  $\phi_1 = w_{in}$ . Thus, SSL mainly reduces the  $B$  on the projected data by reducing the dependency from order  $\frac{1}{d}$  to 1 where 1 is the dimensionality of the output of  $\phi$  without altering the margin. Thus, we get a tighter upper bound for linear probing performed on top CL features when compared with linear probing done on inputs directly.

Intuitively, since the target data has only one predictive feature (along  $w_{in}$ ), CL directly recovers this predictive feature as it is the predominant direction that minimizes invariance loss.

## F Limitations of Prior Work

### F.1 Contrastive learning analysis

Prior works that analyze contrastive learning show that minimizers of the CL objective recover clusters in the augmentation graph, which weights pairs of augmentations with their probability of being sampled as a positive pair [39, 11, 73, 44]. When there is no distribution shift in the downstream task, assumptions made on the graph in the form of consistency of augmentations with downstream labels, is sufficient to ensure that a linear probed head has good ID generalization. Under distribution shift, these assumptions are not sufficient and stronger ones are needed. *E.g.*, some works assume that same-domain/class examples are weighted higher than cross-class cross-domain pairs [40, 76].

Using notation defined in [76], the assumption on the augmentation graph requires cross-class and same-domain weights ( $\beta$ ) to be higher than cross-class and cross-domain weights ( $\gamma$ ). It is unclear if examples from different classes in the same domain will be “connected” if strong spurious features exist in the source domain and augmentations fail to mask them completely (*e.g.*, image background may not be completely masked by augmentations but it maybe perfectly predictive of the label on source domain). In such cases, the linear predictor learnt over CL would fail to generalize OOD. In our toy setup as well, the connectivity assumption fails since on source  $x_{sp}$  is perfectly predictive of the label and the augmentations are imperfect, *i.e.*, augmentations do not mask  $x_{sp}$  and examples of different classes do not overlap in source (*i.e.*,  $\beta > 0$ ). On the other hand, since  $x_{sp}$  is now random on target, augmentations of different classes may overlap, *i.e.*,  $\gamma > 0$ , thus breaking the connectivity assumption. This is also highlighted in our empirical findings of CL furnishing representations that do not fully enable linear transferability from source to target (see Sec. 5). These empirical findings also call into question existing assumptions on data augmentations, highlighting that perfect linear transferability may not typically hold in practice. It is in this setting that we believe self-training can improve over contrastive learning by unlearning source-only features and improving linear transferability.

### F.2 Self-training analysis

Some prior works on self-training view it as consistency regularization that constrain pseudolabels of original samples to be consistent with all their augmentations [12, 87, 79]. This framework



abstracts the role played by the optimization algorithm and instead evaluates the global minimizer of a population objective that enforces consistency of pseudolabels. In addition, certain expansion assumptions on class-conditional distributions are needed to ensure that pseudolabels have good accuracy on source and target domains. This framework does not account for challenges involved in propagating labels iteratively. For *e.g.*, when augmentation distribution has long tails, the consistency of pseudolabels depends on the sampling frequency of “favorable” augmentations. As an illustration, consider our augmentation distribution in the toy setup in Sec. 4. If it were not uniform over dimensions, but instead something that was highly skewed, then a large number of augmentations need to be sampled for every data point to propagate pseudolabels successfully from source labeled samples to target unlabeled samples during self-training. This might hurt the performance of ST when we are optimizing for only finitely many iterations and over finitely many datapoints. This is why in our analysis we instead adopt the iterative analysis of self-training [17].

## G Additional Lemmas

In this section we define some additional lemmas that we use in our theoretical analysis in E.

**Lemma 20** (Upper bound and lower bounds on  $\text{erfc}$ ; Kschischang [49]). *Define*  $\text{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty \exp(-z^2) dz$ . *Then we have:*

$$\frac{2}{\pi} \frac{\exp(-x^2)}{x} \leq \text{erfc}(x) \leq \frac{2}{\pi} \frac{\exp(-x^2)}{x^2 + 4\pi}$$

**Lemma 21** (Properties of Mill’s ratio [5]). *Define the Mill’s ratio as*  $r(x) = \frac{\exp(-x^2)}{\sqrt{2\pi}} \text{erfc}(x)$ . *Then following assertions are true: (i)  $r(x)$  is a strictly decreasing log-convex function; (ii)  $r'(x) = x r(x) - 1$  is an increasing function with  $r'(x) \leq 0$  for all  $x$ ; (iii)  $r''(x) = r(x) - x^2 r'(x)$  is a decreasing function with  $r''(x) \geq 0$  for all  $x$ ; (iv)  $x^2 r'(x)$  is a decreasing function of  $x$ .*

**Lemma 22** (invariance loss as product with operator  $L$ ). *The invariance loss for some  $\phi \in \mathbb{R}^d$  is given as:*  $2 \int_A \phi(a) L\phi(a) dP_A$  *where the operator*  $L$  *is defined as:*

$$L\phi(a) = \int_A \frac{A + a, a'}{p_A(a)} \phi(a') da'$$

*Proof.* The invariance loss for  $\phi$  is given by:

$$\mathbb{E}_{x \sim P_U} \mathbb{E}_{a_1, a_2 \sim P_A(\cdot|x)} \phi(a_1)^\top \phi(a_2) = 2 \mathbb{E}_{x \sim P_U} \mathbb{E}_{a \sim P_A(\cdot|x)} \phi(a)^2 = 2 \mathbb{E}_{a_1, a_2 \sim A} \phi(a_1)^\top \phi(a_2) \quad (181)$$

$$\geq 2 \int_A \phi(a)^2 dP_A \geq 2 \int_A \phi(a) \int_A \frac{A + a, a_2}{p_A(a)} \phi(a_2) da_2 dP_A \quad (182)$$

$$\geq 2 \int_A \phi(a) L\phi(a) dP_A \quad (183)$$

□

**Lemma 23.** *If  $W$  is the space spanned by  $w_{in}$  and  $w_{sp}$ , and  $W_\perp$  is the null space for  $W$ , then for any  $u \in W$  and any  $v \in W_\perp$ , the covariance along these directions  $\mathbb{E}_{a \sim P_A} r a^\top u v^\top a$  is 0.*

*Proof:* We can write the covariance over augmentations after we break down the augmentation  $a$  into two projections:  $a = \Pi_W a + \Pi_{W_\perp} a$

$$\mathbb{E}_{a \sim P_A} r a^\top u v^\top a = \mathbb{E}_{a \sim P_A} u^\top \Pi_W r a + \mathbb{E}_{a \sim P_A} v^\top \Pi_{W_\perp} r a \quad (184)$$

$$= \mathbb{E}_{a \sim P_A} u^\top \Pi_W r a + \mathbb{E}_{a \sim P_A} v^\top \Pi_{W_\perp} r a \quad (185)$$

$$= u^\top \mathbb{E}_{a \sim P_A} \Pi_W r a \Pi_{W_\perp} r a^\top v = 0 \quad (186)$$

where the last inequality follows from the fact that  $\mathbb{E}_{a \sim P_A} \Pi_W r a \Pi_{W_\perp} r a^\top = \mathbb{E}_{a \sim P_A} r \Pi_W r a \mathbb{E}_{a \sim P_A} r \Pi_{W_\perp} r a^\top$ , since the noise in the null space of  $W$  is drawn independent of the component along  $W$ , and furthermore the individual expectations evaluate to zero.

**Lemma 24** (closed-form expressions for eigenvalues and eigenvectors of  $\Sigma_A, \Sigma$ ). For a  $2 \times 2$  real symmetric matrix  $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$  the eigenvalues  $\lambda_1, \lambda_2$  are given by the following expressions:

$$\lambda_1 = \frac{\rho a + b + \delta q}{2}, \quad \lambda_2 = \frac{\rho a + b - \delta q}{2},$$

where  $\delta = \sqrt{4c^2 - \rho a - b q^2}$ . Further, the eigenvectors are given by  $U = \begin{pmatrix} \cos \rho \theta q & \sin \rho \theta q \\ \sin \rho \theta q & \cos \rho \theta q \end{pmatrix}$ , where:

$$\tan \rho \theta q = \frac{b - a - \delta}{2c}.$$

For full proof of these statements see [23]. Here, we will use these statements to arrive at closed form expressions for the eigenvalues and eigenvectors of  $\Sigma_A, \Sigma$ .

*Proof.* We can now substitute the above formulae with  $a, b, c, d$  taken from the expressions of  $\Sigma_A$  and  $\Sigma$ , to get the following values:  $\lambda_1, \lambda_2$  are the eigenvalues of  $\Sigma_A$ , with  $\alpha$  determining the corresponding eigenvectors  $(\cos \alpha q, \sin \alpha q, \sin \alpha q, -\cos \alpha q)$ ; and  $\kappa_1, \kappa_2$  are the eigenvalues of  $\Sigma$ , with  $\beta$  determining the corresponding eigenvectors:  $(\cos \beta q, \sin \beta q, \sin \beta q, -\cos \beta q)$ .

$$\lambda_1 = \frac{\frac{1}{8} \gamma^2 + 1 + \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} + 1 + \frac{1}{d_{in}} \frac{d_{sp}}{2} + \frac{2\sigma_{sp}^2}{3} + \frac{1}{6}}{d \sqrt{\gamma^2 d_{sp} + \gamma^2 + 1 + \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} + 1 + \frac{1}{d_{in}} \frac{d_{sp}}{2} + \frac{2\sigma_{sp}^2}{3} + \frac{1}{6}}} \quad (187)$$

$$\lambda_2 = \frac{\frac{1}{8} \gamma^2 + 1 + \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} + 1 + \frac{1}{d_{in}} \frac{d_{sp}}{2} + \frac{2\sigma_{sp}^2}{3} + \frac{1}{6}}{d \sqrt{\gamma^2 d_{sp} + \gamma^2 + 1 + \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} + 1 + \frac{1}{d_{in}} \frac{d_{sp}}{2} + \frac{2\sigma_{sp}^2}{3} + \frac{1}{6}}} \quad (188)$$

$$\kappa_1 = \frac{\frac{1}{8} \gamma^2 + \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2}}{d \sqrt{\gamma^2 d_{sp} + \gamma^2 + \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2}}} \quad (189)$$

$$\kappa_2 = \frac{\frac{1}{8} \gamma^2 + \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2}}{d \sqrt{\gamma^2 d_{sp} + \gamma^2 + \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2}}} \quad (190)$$

$$\tan \rho \alpha q = \frac{\frac{1}{\gamma} \frac{d_{sp}}{2} + \frac{2\sigma_{sp}^2}{3} + \frac{1}{6} + \gamma^2 + 1 + \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} + 1 + \frac{1}{d_{in}}}{d \sqrt{\gamma^2 d_{sp} + \gamma^2 + 1 + \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} + 1 + \frac{1}{d_{in}} \frac{d_{sp}}{2} + \frac{2\sigma_{sp}^2}{3} + \frac{1}{6}}} \quad (191)$$

$$\tan \rho \beta q = \frac{\frac{1}{\gamma} \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2} + \gamma^2 + \gamma^2 d_{sp} + \gamma^2 + \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2}}{d \sqrt{\gamma^2 d_{sp} + \gamma^2 + \frac{d_{sp}}{2} + \frac{\sigma_{sp}^2}{2}}} \quad (192)$$

Consider the subclass of problem parameters,  $d_{sp} = z, \gamma = K_1 \sqrt{z}$  and  $\sigma_{sp} = K_2 \sqrt{z}$  for fixed constants  $K_1, K_2 \geq 0$  and some variable  $z \geq 0$ , which we can vary to give us different problem instances for our toy model in (6).

$$\lambda_1 = \frac{1}{8} \frac{K_1^2}{z} \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} \frac{1}{d_{in}} \frac{z}{2} \frac{2K_2^2 z}{3} \frac{1}{6} \frac{d}{K_1^2 \frac{K_1^2}{z} \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} \frac{1}{d_{in}} \frac{z}{2} \frac{2K_2^2 z}{3} \frac{1}{6}}^2 \quad (193)$$

$$\lambda_2 = \frac{1}{8} \frac{K_1^2}{z} \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} \frac{1}{d_{in}} \frac{z}{2} \frac{2K_2^2 z}{3} \frac{1}{6} \frac{d}{K_1^2 \frac{K_1^2}{z} \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} \frac{1}{d_{in}} \frac{z}{2} \frac{2K_2^2 z}{3} \frac{1}{6}}^2 \quad (194)$$

$$\kappa_1 = \frac{1}{8} \frac{K_1^2}{z} \frac{z}{2} \frac{K_2^2 z}{2} \frac{d}{K_1^2 \frac{K_1^2}{z} \frac{z}{2} \frac{K_2^2 z}{2}}^2 \quad (195)$$

$$\kappa_2 = \frac{1}{8} \frac{K_1^2}{z} \frac{z}{2} \frac{K_2^2 z}{2} \frac{d}{K_1^2 \frac{K_1^2}{z} \frac{z}{2} \frac{K_2^2 z}{2}}^2 \quad (196)$$

$$\tan \rho \alpha q = \frac{1}{K_1} \frac{z}{2} \frac{2K_2^2 z}{3} \frac{1}{6} \frac{K_1^2}{z} \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} \frac{1}{d_{in}} \frac{d}{K_1^2 \frac{K_1^2}{z} \frac{1}{3d_{in}} \frac{\sigma_{in}^2}{3} \frac{1}{d_{in}} \frac{z}{2} \frac{2K_2^2 z}{3} \frac{1}{6}}^2 \quad (197)$$

$$\tan \rho \beta q = \frac{1}{K_1} \frac{z}{2} \frac{K_2^2 z}{2} \frac{K_1^2}{z} \frac{d}{K_1^2 \frac{K_1^2}{z} \frac{z}{2} \frac{K_2^2 z}{2}}^2 \quad (198)$$

From Stewart [80], we can use the closed form expression for the singular vectors of a  $2 \times 2$  full rank asymmetric matrix  $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ . The singular vectors are given by

$$\begin{pmatrix} \cos \theta & \sin \theta \\ \sin \theta & \cos \theta \end{pmatrix},$$

where,  $\tan 2\theta q$  is given by:

$$\tan 2\theta q = \frac{2ac - 2bd}{a^2 - b^2 - c^2 + d^2}.$$

Now, substituting the values in the expression from (100), we get singular vectors of the above form where  $\theta \in (0, \pi/2)$  satisfies:

$$\theta = \frac{1}{2} \tan^{-1} \frac{2 \tan \rho \beta q - \rho \kappa_1 \kappa_2 q \sqrt{\lambda_1 \lambda_2}}{\rho \lambda_2 \kappa_1 - \lambda_1 \kappa_2 q \tan^2 \rho \alpha q} \quad (199)$$

□

**Lemma 25** (asymptotic behavior of  $\tau \tan \theta$ ). For  $\gamma = K_1 \sqrt{z}, \sigma_{sp} = K_2 \sqrt{z}$ ,

$$\lim_{z \rightarrow \infty} \tau \tan \theta = \frac{K_1 K_2^2}{\rho \lambda_2^2 q^2 \sigma_{in}^2 \rho \lambda_1 \sqrt{d_{in} q}}$$

*Proof.* In order to determine the asymptotic nature of  $\tan \rho \theta q$  as  $z \rightarrow \infty$ , we take the limit of a slightly different term first, since we have the closed form expression of  $\tan 2\theta q$ .

$$\lim_{z \rightarrow \infty} \tau \tan p 2\theta q = \frac{\frac{C}{\lambda_1}}{\lambda_2} \frac{2 \tan p \alpha - \beta q \frac{\rho^{\lambda_1} \{\kappa_2\} - 1}{\lambda_1 \{\lambda_2 q - \rho^{\lambda_1} \lambda_1 \{\lambda_2 \kappa_2\} - 1} q \tan^2 p \alpha - \beta q}}{2 \tan p \alpha - \beta q \frac{\kappa_1 \{\kappa_2\} - 1}{\kappa_1 \{\kappa_2\} - \lambda_2 \{\lambda_1 - 1}}},$$

since it is easy to see that  $\lim_{z \rightarrow \infty} \tan^2 p \alpha - \beta q \frac{\lambda_1 \kappa_1}{\lambda_2 \kappa_2} - 1 = 0$ .

If we use  $\tan p \alpha - \beta q \frac{\tan \alpha - \tan \beta}{1 + \tan \alpha \tan \beta}$ , and substitute the functions of  $z$ , for all the quantities in the above expression using Lemma 24, we derive:  $\lim_{z \rightarrow \infty} \tau \tan 2\theta = 2K_1 K_2^2 \{(1 + K_2^2) 2\sigma_{in}^2 (1 - 1/d_{in})\}$ .

Since  $\tau \tilde{N} \mathcal{B}$ ,  $\tan p 2\theta q \tilde{N} 0$ , and further from Taylor approximation of  $\tan p 2\theta q$ ,  $\tan p 2\theta q \tilde{N} 2\theta$ . We can use this to derive the limit for  $\tau \tan \theta$ , which would just be  $\frac{1}{2} \{2 - 2K_1 K_2^2 \{(1 + K_2^2) 2\sigma_{in}^2 (1 - 1/d_{in})\} - K_1 K_2^2 \{(1 + K_2^2) 2\sigma_{in}^2 (1 - 1/d_{in})\}\}$ .

□

**Lemma 26** (asymptotic behaviors of  $\cot \alpha, \tan \theta$ ). For  $\gamma = K_1 \{\sqrt{z}, \sigma_{sp} = K_2 \frac{?}{z}$  following the expressions in Lemma 24,

$$\lim_{z \rightarrow \infty} \cot \alpha = 0, \quad \lim_{z \rightarrow \infty} \tan \theta = 0.$$

*Proof.* For  $\tan \theta$ , since  $\tau \tilde{N} \mathcal{B}$ , and  $\tau \tan \theta$  approaches a constant (from Lemma 25), we conclude  $\lim_{z \rightarrow \infty} \tan \theta = 0$ . For  $\cot \alpha$ ,

$$\lim_{z \rightarrow \infty} \frac{z}{2} = \frac{2K_2^2 z}{3} = \frac{1}{6} = \frac{K_1^2}{z} - 1 = \frac{1}{3d_{in}} = \frac{\sigma_{in}^2}{3} - 1 = \frac{1}{d_{in}} = \mathcal{B},$$

and,

$$\lim_{z \rightarrow \infty} \frac{d}{K_1^2} = \frac{\frac{K_1^2}{z} - 1}{\frac{1}{3d_{in}} - \frac{\sigma_{in}^2}{3} + 1} = \frac{\frac{1}{d_{in}} - \frac{z}{2} - \frac{2K_2^2 z}{3} - \frac{1}{6}}{2} = \mathcal{B}.$$

Thus,  $\cot \alpha \tilde{N} 0$ .

□

**Lemma 27** (asymptotic behavior of  $z \cot \alpha$ ). For  $\gamma = K_1 \{\sqrt{z}, \sigma_{sp} = K_2 \frac{?}{z}$  following the expressions in Lemma 24,

$$\lim_{z \rightarrow \infty} z \cot \alpha = \frac{K_1}{1 - 4\{3K_2^2\}}.$$

*Proof.* The expression for  $z \cot \alpha$  or  $z \{\tan \alpha$  follows from Lemma 24:

$$\lim_{z \rightarrow \infty} z \cot \alpha = \frac{z K_1}{p \frac{a}{p^2} - K_1^2},$$

where  $p = \frac{K_1^2}{z} - 1 = \frac{1}{3d_{in}} - \frac{\sigma_{in}^2}{3} + 1 = \frac{1}{d_{in}} - \frac{z}{2} - \frac{2K_2^2 z}{3} - \frac{1}{6}$ . Applying L'Hôpital's (relevant expressions are continuous in  $z$ ) rule we get:  $\lim_{z \rightarrow \infty} z \cot \alpha = \frac{K_1}{1 + 4/3 K_2^2}$ .

□

**Lemma 28** (asymptotic behavior of  $z \{\tau^2$ ). For  $\gamma = K_1 \{\sqrt{z}, \sigma_{sp} = K_2 \frac{?}{z}$  following the expressions in Lemma 24,

$$\lim_{z \rightarrow \infty} z \{\tau^2 = \frac{2\sigma_{in}^2 \{3p - 1\} d_{in} q}{1 - 4\{3K_2^2\}}.$$

*Proof.* For  $\tau = \lambda_1 \{\lambda_2$ , substituting the relevant expressions from Lemma 24, we get:

$$z \{\tau^2 = \frac{z \lambda_2}{\lambda_1} = \frac{2K_1^2 \{z p - 1\} \{3d_{in} q - 2\sigma_{in}^2 p - 1\} d_{in} q}{2K_1^2 \{z p - 1\} \{3d_{in} q - 2\sigma_{in}^2 p - 1\} d_{in} q} = \frac{p}{\frac{a}{K_1^2} - \frac{p^2}{K_1^2}},$$

where  $p = z\{2 - 2K_2^2 z\}^3 - 1\{6$ . Applying L'Hôpital's (relevant expressions are continuous in  $z$ ) rule we get:  $\lim_{z \rightarrow \infty} z\{\tau^2 - \frac{2\sigma_{in}^2/3(1-1/d_{in})}{1+4/3K_2^2}\}$ .  $\square$

**Lemma 29** (0-1 error of a classifier on target). Assume a classifier of the form  $w = l_1 w_{in} + l_2 w_{sp}$  where  $l_1, l_2 \in \mathbb{R}$  and  $w_{in} = (w^*, 0, \dots, 0)^T$ , and  $w_{sp} = (0, \dots, 0, \frac{1}{d_{sp}} \mathbf{1}_{d_{sp}})^T$ . Then the target accuracy of this classifier is given by  $0.5 + \operatorname{erfc} \frac{l_1 \gamma}{\sqrt{2} \cdot l_2 \cdot \sigma_{sp}}$ .

*Proof.* Assume  $(x, y) \in \mathcal{P}_T$ . Accuracy of  $w$  is given by  $\mathbb{E}_{\mathcal{P}_T} \mathbb{P}(\operatorname{sign}(w^T x) = y)$ .

$$\begin{aligned} \mathbb{E}_{\mathcal{P}_T} \mathbb{P}(\operatorname{sign}(w^T x) = y) &= \mathbb{E}_{\mathcal{P}_T} \mathbb{P}(y = \operatorname{sign}(w^T x)) \\ &= \mathbb{E}_{\mathcal{P}_T} \mathbb{P}(y = \mathbb{1}_{\{w^T x \geq 0\}}) \\ &= \mathbb{E}_{\mathcal{P}_T} \mathbb{P}(y = \mathbb{1}_{\{l_1 w_{in}^T x + l_2 w_{sp}^T x \geq 0\}}) \\ &= \mathbb{E}_{\mathcal{P}_T} \mathbb{P}(y = \mathbb{1}_{\{l_1 y + l_2 \sigma_{sp} z \geq 0\}}) \\ &= \mathbb{E}_{z \sim N(0,1)} \mathbb{P}(y = \mathbb{1}_{\{l_1 y + l_2 \sigma_{sp} z \geq 0\}}) \\ &= \mathbb{E}_{z \sim N(0,1)} \mathbb{P}(y = \mathbb{1}_{\{z \geq -\frac{l_1 y}{l_2 \sigma_{sp}}\}}) \\ &= \mathbb{E}_{z \sim N(0,1)} \mathbb{P}(z \geq -\frac{l_1 y}{l_2 \sigma_{sp}}) \end{aligned}$$

Using the definition of  $\operatorname{erfc}$  function, we get the aforementioned accuracy expression.  $\square$

**Lemma 30.** For  $\sigma > 0$  and  $\mu \in \mathbb{R}$ , we have

$$\begin{aligned} g(\mu, \sigma) &= \mathbb{E}_{z \sim N(0, \sigma)} \mathbb{P}(\exp(\mu z) \geq z) \tag{200} \\ &= \frac{1}{2} \exp(\sigma^2 \{2 - \mu\}) \operatorname{erfc}(\mu \sqrt{2} \sigma) + \exp(\sigma^2 \{2 - \mu\}) \operatorname{erfc}(\mu \sqrt{2} \sigma) \tag{201} \end{aligned}$$

*Proof.* The proof uses simple algebra and the definition of  $\operatorname{erfc}$  function.

$$\begin{aligned} g(\mu, \sigma) &= \mathbb{E}_{z \sim N(0, \sigma)} \mathbb{P}(\exp(\mu z) \geq z) \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(\mu z) \mathbb{1}_{\{\exp(\mu z) \geq z\}} \exp(-\frac{z^2}{2\sigma^2}) dz \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(\mu z) \mathbb{1}_{\{\mu z \geq z\}} \exp(-\frac{z^2}{2\sigma^2}) dz \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(\mu z) \mathbb{1}_{\{\mu \geq 1\}} \exp(-\frac{z^2}{2\sigma^2}) dz + \frac{1}{2\pi} \int_{-\infty}^{-\mu} \exp(\mu z) \mathbb{1}_{\{\mu z \geq z\}} \exp(-\frac{z^2}{2\sigma^2}) dz \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(\mu z) \mathbb{1}_{\{\mu \geq 1\}} \exp(-\frac{z^2}{2\sigma^2}) dz + \frac{1}{2\pi} \int_{-\infty}^{-\mu} \exp(\mu z) \mathbb{1}_{\{\mu z \geq z\}} \exp(-\frac{z^2}{2\sigma^2}) dz \\ &= \frac{1}{2} \exp(\sigma^2 \{2 - \mu\}) \operatorname{erfc}(\mu \sqrt{2} \sigma) + \exp(\sigma^2 \{2 - \mu\}) \operatorname{erfc}(\mu \sqrt{2} \sigma) \end{aligned}$$

$\square$