

Surprisingly Simple Semi-Supervised Domain Adaptation with Pretraining and Consistency

Samarth Mishra¹
samarthm@bu.edu

Kate Saenko^{1,2}
saenko@bu.edu

Venkatesh Saligrama¹
srv@bu.edu

¹ Boston University

² MIT-IBM Watson AI Lab

Abstract

Most modern unsupervised domain adaptation (UDA) approaches are rooted in domain alignment, *i.e.*, learning to align source and target features to learn a target domain classifier using source labels. In semi-supervised domain adaptation (SSDA), when the learner can access few target domain labels, prior approaches have followed UDA theory to use domain alignment for learning. We show that the case of SSDA is different and a good target classifier can be learned without needing alignment. We use self-supervised pretraining (via rotation prediction) and consistency regularization to achieve well separated target clusters, aiding in learning a low error target classifier. With our Pretraining and Consistency (PAC) approach, we achieve state of the art target accuracy on this semi-supervised domain adaptation task, surpassing multiple adversarial domain alignment methods, across multiple datasets. PAC, while using simple techniques, performs remarkably well on large and challenging SSDA benchmarks like DomainNet and Visda-17, often outperforming recent state of the art by sizeable margins. Code for our experiments can be found at <https://github.com/venkatesh-saligrama/PAC>.

1 Introduction

The problem of visual domain adaptation arises when a learner must leverage labeled source domain data to classify instances in the target domain, where it has limited access to ground-truth annotated labels. An example of this is the problem of learning to classify real-world images based on hand-sketched depictions. The problem is challenging because discriminative features that are learnt while training to classify source domain instances may not be meaningful or sufficiently discriminative in the target domain. As described in prior works, this situation can be viewed as arising from a “domain-shift”, where the joint distribution of features and labels in the source domain does not follow the same law in the target domain.

We propose a novel method for semi-supervised domain adaptation (SSDA), where the learner, in addition to unlabeled target examples, is granted access to a few labeled target domain examples for training. Our method is based on enhancing clusterability of target

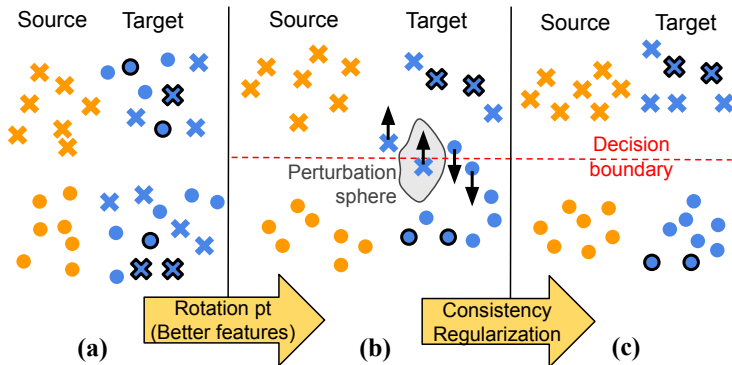


Figure 1: Feature space illustration of our approach. Shapes of points represent classes and colors represent different domains. Target domain points with a black border are labelled. Initial feature space might not be discriminative on target causing subsequent errors with domain alignment. Pretraining with rotation prediction can help fix initial features to some extent and cluster them better by class (a \rightarrow b). Consistency regularization helps cluster features by making the classifier invariant to input perturbations, keeping decision boundaries away from data-dense regions. Image space perturbations can help correctly cluster points that initially lie on the wrong side of the boundary (b \rightarrow c).

domain features independent of the source domain. Prior SSDA methods [19, 20, 35] draw upon approaches developed within the context of visual domain adaptation using labelled source images and unlabelled target images [0, 25, 34, 39, 44]. These works in turn are based on adversarial domain alignment, and draw upon Ben-David *et al.* [8]’s elegant theory of domain adaptation. Generalization bounds in Ben-David *et al.* [8] suggest that, good adaptation is possible if the divergence between induced source and target domain feature distributions is small.

While domain alignment is a meaningful goal in the absence of labels, we believe that the situation is dramatically different in the presence of few target domain labels. For our work, we draw inspiration from Castelli and Cover [4], who show that, in binary classification problems, a learner with no knowledge of underlying feature distributions, can benefit exponentially from labelled examples, and the learner’s error probability approaches Bayes error. As such, our key insight is that we can forego domain alignment if we learn representations that result in compact clusters in the target domain. The identity of these clusters can then be deduced from few-shot labels.

Our proposed approach, PAC (pretraining and consistency), is based on combining well-known methods and novel objectives, which together lead to enhanced clusterability. In particular, we use rotation prediction pre-training [15], label consistency [37], and cross-entropy losses on labelled source and target data, which together result in more compact and pure target domain clusters. Cross-entropy loss on target-domain labelled examples implies good feature separation for different target classes. Consistency loss on perturbed examples and pretraining result in improving similarity between labelled target examples and unlabelled examples. Together these objectives are effective in realizing compact clusters, and result in improved prediction. Notably, PAC achieves better target accuracy than most comparable state-of-the-art methods on the large and challenging DomainNet [31] benchmark by 3-5 %, and on VisDA-17, it is better by 4-10%.

Pretraining. Following recent domain adaptation methods we also utilize an Imagenet [12] pretrained backbone as a starting point. We found that these features are somewhat biased to-

wards Imagenet images and to enhance them, we use self-supervision via rotation prediction. Our approach is informed by Gidaris *et al.* [15]’s work for learning semantically meaningful features from unlabelled images. In addition, it was recently found in a study by Wallace *et al.* [22] to be more semantically meaningful than an array of other self-supervision objectives. Most importantly for us, we found that self-supervision via rotation prediction also resulted in better target clusterability as compared to Imagenet pretraining. Fig 1 (a), shows a hypothetical scenario where simply aligning features can lead to poor generalization, and how pretraining can somewhat remedy this situation.

Label Consistency. Also key to our approach is label consistency using image space perturbations. Label consistency or *consistency regularization* enforces invariance of classifier predictions to small input perturbations, and thus to inputs from neighborhoods that may form clusters¹. In our approach, we do this using image augmentation methods like RandAugment [16] and color jittering, and the model is trained to produce the same output for both a perturbed and an unperturbed version of the image resulting in meaningful neighborhood for images².

Our contributions in this paper are two-fold:

- We propose a novel semi-supervised domain adaptation method PAC, based on label consistency and rotation prediction for pretraining, which performs comparably or better than state of the art on SSDA across multiple datasets. In contrast to prior works, we forego domain alignment, and pose objectives that improve target clusterability.
- We perform ablative analysis on individual components of our method, illustrating their behavior, and their impact on performance. Our analysis provides an understanding of these components and shows how they can be combined with other techniques.

2 Background

Adversarial Domain Alignment. The generalization bound for target error for unsupervised domain adaptation scenario in Ben-David *et al.* [3] is a sum of three terms: the empirical source error; the divergence between source and target distributions X_S and X_T respectively; and the smallest error of a jointly trained classifier. In the absence of labels, this bound is somewhat intuitive, namely that a joint classifier is required, and that such a classifier in general would suffer from distributional differences. Motivated by this theory, a range of visual domain adaptation approaches attempt to find a feature space ϕ such that the distributions $\phi(X_S)$ and $\phi(X_T)$ have low divergence, so that the distributional differences are mitigated. [3]’s distributional difference is measured in terms of the worst-case differences in disagreement of any two classifiers on source and target spaces. As such the goal of minimizing distributional alignment results in a minimax objective, which, along with the minimization of a classifier error on source labels is broadly the approach adopted by a range of recent domain adaptation methods [11, 13, 25, 52, 59, 42].

The Exponential Value of Labels. Castelli & Cover [5, 7] showed that for a conventional binary classification problem with u unlabelled examples and ℓ labelled examples, the learner’s

¹Fig 1 depicts consistency regularization helping correctly cluster points that may initially lie on the wrong side of the decision boundary, and that may result in errors with a domain alignment strategy.

²label consistency is related to cluster assumption [9], which suggests similar labels for examples within a cluster, as well as the concept of *low-density separation assumption*[10], which requires classifier decision boundaries to lie in low-density regions of feature space.

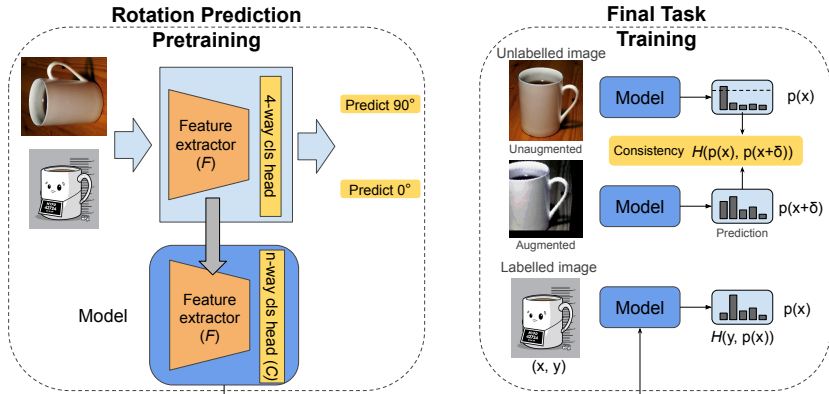


Figure 2: A diagram of our Pretraining and Consistency (PAC) approach. We first train our backbone for the self-supervised task of predicting rotations (left). This backbone is then used as a warm start for our classification model, which uses labelled data with the cross entropy criterion and consistency regularization for the unlabelled data (right).

probability of error approaches the Bayes error as

$$P_{err} - P_{bayes} = O\left(\frac{1}{u}\right) + \exp(-D(q_1, q_2)\ell + o(\ell)) \quad (1)$$

where q_1 and q_2 are feature distributions conditioned on class-labels, and unknown to the learner. $D(q_1, q_2)$ is the Bhattacharyya Distance between the distributions. A key insight is that the error probability decays exponentially with the product of inter-class distance and number of labels. Thus low error is achievable with only a few labels for a feature distribution that is well separated across classes. This is the key motivation behind our approach.

For learning a good target classifier, we formulate an objective that, in addition to penalizing source and target domain mistakes on labelled data, improves separability of features belonging to different target classes. Informally, our objective is to optimize *Source Loss + Target Loss + Target Clustering*.

The first term leverages source data to realize good feature representation. On the other hand, the second and third terms, together attempt to improve target domain representation. As such target loss promotes good feature separation for labelled examples between different target classes. We utilize label consistency to improve target clustering by enforcing target invariance of classifier predictions to small input perturbations, and thus to inputs from neighborhoods that may form clusters. Target clustering is also aided by our initialization with self-supervised pretraining. Together, the two terms result in operationalizing [1]’s concept resulting in well-separated clusters, and allow for generalization from few labels.

3 Pretraining and Consistency (PAC)

Before describing our approach, we introduce some notation for ease of precise exposition. Available to the model are two sets of labelled images : $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, the labelled source images and $\mathcal{D}_t = \{(x_i^t, y_i^t)\}_{i=1}^{n_t}$, the few labelled target images, and additionally the set of unlabelled target images $\mathcal{D}_u = \{x_i^u\}_{i=1}^{n_u}$. The goal is to predict labels for these images in \mathcal{D}_u . The final classification model consists of two components : the feature extractor F and the classifier C . F generates features $F(x)$ for an input image x which the classifier operates

on to produce output class scores $C(F(x)) \in \mathbb{R}^K$, where K is the number of categories that the images in the dataset could belong to. In our experiments, F is a convolutional network and produces features with unit ℓ_2 -norm, *i.e.* $\|F(x)\|_2 = 1$ (following [65]). C consists of one or two fully connected layers.

An overview of PAC is shown in Fig 2. Our final model is trained in two stages:

Pretraining with Rotation Prediction. We first train our feature extractor F with the self-supervised task of predicting image rotations (Fig 2 (left)) on both the source and target datasets, *i.e.*, all images in $\mathcal{D}_s, \mathcal{D}_t$ and \mathcal{D}_u . Without using image category labels, we train a 4-way classifier to predict one out of 4 possible angles ($0^\circ, 90^\circ, 180^\circ, 270^\circ$) that an input image has been rotated by. We follow the procedure of Gidaris *et al.* [65] and in each minibatch, we introduce all 4 rotations of a given image to the classifier. This backbone is then used as the initialization for training the final classifier in the next stage.

Consistency Regularized Classifier Training. Consistency regularization promotes the final model $C \circ F$ to produce the same output for both an input image x and a perturbed version $x + \delta$. We introduce these perturbations using image level augmentations — RandAugment [66] along with additional color jittering. Given an unlabelled image $x \in \mathcal{D}_u$, we first compute the model’s predicted class distributions

$$\begin{aligned} p_x &= p(y|x; F, C) = \text{softmax}(C(F(x))) \\ q_x &= p(y|x + \delta; F, C) = \text{softmax}(C(F(x + \delta))) \end{aligned}$$

p_x is then confidence thresholded using a threshold τ and the following is used as the consistency regularization loss.

$$\mathcal{L}_{CR}(x) = \mathbb{1}[\max_{k \in K} p_x(k) \geq \tau] H(p_x, q_x) \quad (2)$$

where $\mathbb{1}$ is an indicator function and $H(p_x, q_x) = \sum_{k \in K} -p_x(k) \log(q_x(k))$ is cross-entropy. Note that $p_x(\cdot)$ has been used to index into the K elements of p_x . Intuitively, an unperturbed version of image x is used to compute *pseudo-targets* for the perturbed version $x + \delta$, which is only used when the pseudo-target has high confidence ($\max_k p_x(k) \geq \tau$). We also note here that the target p_x is treated as a constant for gradient computation with respect to the network parameters. For the labelled examples from \mathcal{D}_s and \mathcal{D}_t , we use the same perturbations but with ground truth labels as targets.

The model is optimized using minibatch-SGD, with minibatches M_s, M_t and M_u sampled from $\mathcal{D}_s, \mathcal{D}_t$ and \mathcal{D}_u respectively. The final optimization criterion used is

$$\begin{aligned} \mathcal{L} &= \frac{1}{|M_s|} \sum_{(x,y) \in M_s} H(\bar{y}, x) + \frac{1}{|M_t|} \sum_{(x,y) \in M_t} H(\bar{y}, x) \\ &\quad + \frac{1}{|M_u|} \sum_{x \in M_u} \mathcal{L}_{CR}(x) \end{aligned}$$

where $\bar{y} \in \mathbb{R}^K$ is the one-hot representation of $y \in [K]$ or $\bar{y}(i) = \mathbb{1}[i = y]$ and $H(\bar{y}, x)$ has been overloaded to mean $H(\bar{y}, C(F(x)))$.

4 Experiments

Datasets. We evaluated PAC on four different datasets: DomainNet [67], VisDA-17 [68], Office-Home [69] and Office [62]. DomainNet [67] is a recent large scale domain adaptation

Net	Method	R to C		R to P		P to C		C to S		S to P		R to S		P to R		MEAN	
		1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot
Alexnet	S+T	43.3	47.1	42.4	45.0	40.1	44.9	33.6	36.4	35.7	38.4	29.1	33.3	55.8	58.7	40.0	43.4
	DANN	43.3	46.1	41.6	43.8	39.1	41.0	35.9	36.5	36.9	38.9	32.5	33.4	53.6	57.3	40.4	42.4
	ADR	43.1	46.2	41.4	44.4	39.3	43.6	32.8	36.4	33.1	38.9	29.1	32.4	55.9	57.3	39.2	42.7
	CDAN	46.3	46.8	45.7	45.0	38.3	42.3	27.5	29.5	30.2	33.7	28.8	31.3	56.7	58.7	39.1	41.0
	MME	48.9	55.6	48.0	49.0	46.7	51.7	36.3	39.4	39.4	43.0	33.3	37.9	56.8	60.7	44.2	48.2
	Meta-MME	-	56.4	-	50.2	-	51.9	-	39.6	-	43.7	-	38.7	-	60.7	-	48.7
	APE	47.7	54.6	49.0	50.5	46.9	52.1	38.5	42.6	38.5	42.2	33.8	38.7	57.5	61.4	44.6	48.9
	BiAT	54.2	58.6	49.2	50.6	44.0	52.0	37.7	41.9	39.6	42.1	37.2	42.0	56.9	58.8	45.5	49.4
	CDAC	56.9	61.4	55.9	57.5	51.6	58.9	44.8	50.7	48.1	51.7	44.1	46.7	63.8	66.8	52.1	56.2
	PAC	55.4	61.7	54.6	57.9	47.0	59.8	46.9	52.9	38.6	43.9	38.7	48.2	56.7	59.7	48.3	54.7
Resnet-34	S+T	55.6	60.0	60.6	62.2	56.8	59.4	50.8	55.0	56.0	59.5	46.3	50.1	71.8	73.9	56.8	60.0
	DANN	58.2	59.8	61.4	62.8	56.3	59.6	52.8	55.4	57.4	59.9	52.2	54.9	70.3	72.2	58.4	60.7
	ADR	57.1	60.7	61.3	61.9	57.0	60.7	51.0	54.4	56.0	59.9	49.0	51.1	72.0	74.2	57.6	60.4
	CDAN	65.0	69.0	64.9	67.3	63.7	68.4	53.1	57.8	63.4	65.3	54.5	59.0	73.2	78.5	62.5	66.5
	MME	70.0	72.2	67.7	69.7	69.0	71.7	56.3	61.8	64.8	66.8	61.0	61.9	76.1	78.5	66.4	68.9
	Meta-MME	-	73.5	-	70.3	-	72.8	-	62.8	-	68.0	-	63.8	-	79.2	-	70.1
	APE	70.4	76.6	70.8	72.1	72.9	76.7	56.7	63.1	64.5	66.1	63.0	67.8	76.6	79.4	67.8	71.7
	BiAT	73.0	74.9	68.0	68.8	71.6	74.6	57.9	61.5	63.9	67.5	58.5	62.1	77.0	78.6	67.1	69.7
	CDAC	77.4	79.6	74.2	75.1	75.5	79.3	67.6	69.9	71.0	73.4	69.2	72.5	80.4	81.9	73.6	76.0
	PAC	74.9	78.6	73.0	74.3	72.6	76.0	65.8	69.6	67.9	69.4	68.7	70.2	76.7	79.3	71.4	73.9

Table 1: Accuracy on the DomainNet dataset (%) for one-shot and three-shot settings on 4 domains, R: Real, C: Clipart, P: Painting, S: Sketch. PAC, though simple, is strong enough to be competitive with or outperform other state of the art approaches on most scenarios. Top 2 accuracies in each column are highlighted in bold

benchmark with 6 different visual domains and 345 classes. We use a subset of 4 domains (Clipart, Paintings, Real, and Sketch) and 126 classes for our experiments. This subset had close to 36500 images per domain. A total of 7 different scenarios out of the possible 12 were used for evaluation. VisDA-17 is another large benchmark with a single adaptation scenario : the source domain (the VisDA "train" split) consists of 152,398 synthetic images from 12 categories, and the target domain (the VisDA "validation" split) consists of 55,388 real images. Office-Home [40] and Office [32] datasets are relatively smaller in size with images of objects typically found in office and home environments. More details and the performance comparison for these datasets are in the supp.

For each adaptation scenario, following [55], we use 1-shot and 3-shot settings for evaluation, where 1 and 3 target labels per class are available to the learner respectively (See supp. for an analysis varying the number of target shots). For each scenario, 3 examples in the target domain are held out for validation, except in VisDA-17 where 20 examples per class were held out because of the fewer number of categories.

Implementation Details. All our experiments were implemented in PyTorch [29] using W&B [9] for tracking experiments. For evaluations on the DomainNet dataset, we used an Alexnet and a Resnet-34 [17] backbone, while on VisDA-17 we evaluated our method with a ResNet-34 backbone. 1 or 2 fully connected layers we used for the classifier C . Optimization was done via SGD with a momentum 0.9. Complete details of all experiments are in supp.

4.1 Results

Comparison to other approaches. We compare PAC with different recent SSDA approaches : MME [65], BiAT [19], Meta-MME [23], APE [20], LIRR [22], and CDAC [24] using results reported by these papers. Besides this, we also include in the tables, baseline approaches using adversarial domain alignment—DANN [24], ADR [33] and CDAN [25], that were evaluated by Saito *et al.* [65]. The baseline “S+T” is a method that simply uses all labelled data available to it to train the network using cross-entropy loss. Note that PAC

Method	Overall Accuracy			
	1-shot	3-shot	1-pct	5-pct
S+T	57.7	59.9	76.2	82.9
MME	69.7	70.7	80.5	84.1
LIRR	-	-	81.7	84.5
LIRR+CosC	-	-	82.3	85.1
PAC	75.2	80.4	86.0	88.9

Table 2: Results on VisDA-17. PAC outperforms MME and LIRR, both current state of the art approaches, by a sizeable margin. CosC stands for cosine classifier and “1-pct” and “5-pct” denote scenarios where 1% and 5% of the target images are labelled respectively.

can be construed as “S+T” along with additional consistency regularization and with a warm start using rotation prediction for pretraining.

In Table 1, we compare the accuracy of PAC with different recent approaches on DomainNet. Remarkably our simple approach outperforms other SSDA approaches by 3-5% on this benchmark with different backbones (with the exception of CDAC, with which it is competitive on most scenarios). In Table 2 holding the VisDA-17 results, besides our method, we report results of S+T and MME, that we replicated from the implementation of [85], on the 1-shot and 3-shot scenarios. In addition, for comparison with LIRR on equal footing, we evaluated PAC on scenarios with 1% and 5% of all the target examples are labelled. We see that PAC shows strong performance, with close to 10% improvement in accuracy over MME in the 3-shot scenario, and 4% improvement over LIRR.

Ablative analysis. In Table 3, we see what rotation prediction pretraining and consistency regularization do for final target classification performance separately. The two components provide boosts to the final performance individually, with the combination of both performing best. We see that in most cases consistency regularization helps performance significantly, especially in the 3-shot scenarios.

Rot ⁿ	CR	Target Accuracy			
		Alexnet		Resnet-34	
		1-shot	3-shot	1-shot	3-shot
		29.1	33.3	46.3	50.1
✓		35.1	37.9	54.1	56.1
	✓	32.5	45.9	64.3	68.9
✓	✓	38.7	48.2	68.7	70.2

Table 3: Ablation study for pretraining predicting rotations (Rotⁿ) and consistency regularization (CR) on the *real to sketch* scenario of Domainnet using both Alexnet and Resnet-34 backbones.

Feature space analysis. In Fig 3 we plot the 2-D TSNE [26] embeddings for features generated by 5 differently trained Alexnet backbones. The embeddings are plotted for all points from 5 randomly picked classes. The source domain points which are light colored circles, come from *real* images of Office-Home and the target domain points which are dark colored markers come from *clipart* images. The labelled target examples are marked with X’s. The two plots on the left compare differently pretrained backbones and the three on the right use backbones at the end of different SSDA training processes. In the plots we can see that pretraining the backbone for rotation prediction starts to align and cluster points according

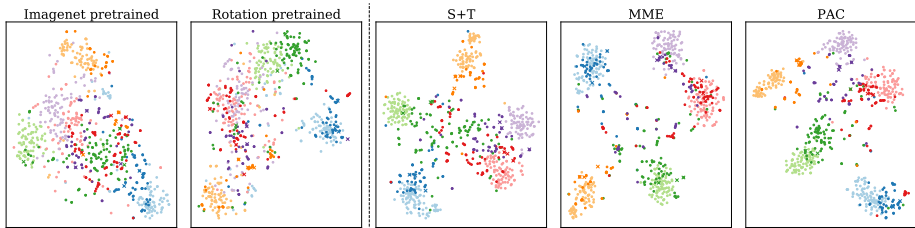


Figure 3: 2-D TSNE embeddings of features from 5 randomly chosen classes. Lighter colors represent source domain points and darker ones represent corresponding target domain points. The dark colored X’s are labelled target domain points (3 per class). Pretraining with rotation prediction begins to cluster points of the same class a little better over a backbone pretrained only on Imagenet. Our method on the right, separates classes just as well as MME. Refer to Section 4.1 for discussion. (Best viewed with color and under zoom)

to their classes a little better than what a backbone pretrained just on Imagenet can do. Out of the final classifiers on the right, we see that both PAC and MME create well separated classes in feature space allowing for the classifier to have decision boundaries in low-density regions. MME explicitly minimizes conditional entropy which may draw samples even further apart from the classifier boundaries, as compared to our method which simply tries to ensure that the classifier does not separate an example and its perturbed version.

In Table 4, we quantitatively analyze the features using different metrics : The \mathcal{A} -distance is a distance metric between the two domains in feature space computed using an SVM trained to classify domains as done in [2]. The higher the error of the SVM domain classifier, the lower is the \mathcal{A} -distance. The other two metrics are accuracies of distance based nearest neighbor (NN) classifiers in feature space. The first one, “NN Acc. (Target)” is the accuracy of a classifier that assigns any unlabelled target example, the class label of the target labelled examples closest to it on average in the feature space. “NN Acc. (Source)” similarly uses only the source examples, all of which are labelled, to compute the class label for an unlabelled target example. Finally BD is an empirical inter-class Bhattacharyya distance estimate (details in supp.). Comparing the pretrained backbones, we see that rotation pretraining improves the feature space both by bringing closer the features across the two domains (as indicated by the low \mathcal{A} -distance) and aligning them so that features within a class are closer (indicated by the higher NN accuracies). When it comes to final feature spaces of the SSDA methods, we see that MME, being a domain alignment method, reduces \mathcal{A} -distance more than PAC. However, PAC is able to better maintain the class-defined neighborhood of features, as indicated by the higher accuracies. Also, PAC has higher inter-class divergence (BD), which leads to lower error according to Eq. 1. This shows that in SSDA, for learning good target classification, low divergence between domains is not a necessity.

How does pretraining with rotation prediction compare to contrastive methods? Contrastive pretraining methods [9, 10, 13] have been shown to attain remarkable performance in learning features from unlabelled images that are useful for tasks like image recognition and object detection. We evaluate how momentum contrast (MoCo) [13] and SimSiam [10] perform for pretraining our feature extractor on both source and target images, compared with rotation prediction. Table 5 compares most of the same metrics as Table 4 with the addition of final model (training with labels and consistency regularization) performance on target classification. We see that, both MoCo and SimSiam improve the imagenet pretrained features to some extent helping the final classification performance. However, this improve-

Backbone	\mathcal{A} -dist	NN Acc. (Target)	NN Acc. (Source)	BD
Imagenet pt.	1.57	26.8	26.0	-
Rotation pt.	1.28	36.2	34.6	-
S+T	1.49	43.0	43.3	2.01
MME	1.24	51.2	51.5	4.03
PAC	1.45	56.4	56.8	13.43

Table 4: Feature space metrics for different Alexnet backbones on the *real to clipart* adaptation scenario of Office-Home. We see that rotation prediction helps improve the initial feature space. Also amongst the SSDA methods, PAC maintains a better class defined neighborhood both within and across domains, even though the two domains are not aligned as closely as in the case of MME.

Backbone pretraining	\mathcal{A} -dist	NN Acc. (Target)	NN Acc. (Source)	Final Acc.
Imagenet	1.57	26.8	26.0	54.1
MoCo	1.31	31.4	34.8	56.3
SimSiam	1.53	19.3	17.1	56.6
Rotation	1.28	36.2	34.6	58.8

Table 5: Feature space metrics and final method performance for differently pretrained Alexnet backbones on the *real to clipart* adaptation scenario of Office-Home. MoCo and SimSiam help over Imagenet pretraining, but not as much as rotation prediction.

ment is not as high as in the case of rotation prediction pretraining. We see that MoCo has marginally better class-defined structure across domains, but a poorer structure in the target domain indicated by the accuracies of the distance based NN classifiers. Interestingly, we see that nearest neighbors classifiers perform much poorer in the case of SimSiam pretraining as compared to other cases, while the pretraining still benefits final SSDA accuracy. This might indicate that this pretraining helps with optimization in some manner while not clustering the initial feature space as much.

5 Related Work

Cluster Assumption. As mentioned in the introduction, related to our approach is the cluster assumption that has been defined in prior semi-supervised learning work [8], suggesting that data points in a cluster in the input space should be classified similarly. Conditional entropy minimization [16] and consistency regularization [21, 28, 37] have been used as ways of achieving this.

Semi-supervised Domain Adaptation. As mentioned above, conditional entropy minimization has been used to enforce cluster assumption. Saito *et al.* [55] cleverly built this into a minimax optimization problem to propose an adversarial domain alignment method, called minimax entropy (MME) that also plays by the cluster assumption. They evaluated MME on SSDA, and found that it performed better than other domain alignment approaches without

conditional entropy minimization. Their approach uses entropy maximization with the classifier in an attempt to move the boundary close to and in between target unlabelled examples. Subsequent conditional entropy minimization using the feature extractor clusters the target unlabelled examples largely according to the “split” created by the classifier. Note that with this approach the neighborhood already gets defined by the classifier and if errors are made here, they are harder to fix. So, while PAC can fix errors like the ones in Fig. 1 (b), MME may find it harder to do so. We demonstrate this in supp. by comparing both methods trained from a randomly initialized backbone, and find that PAC performs better.

Saito *et al.* [15] also used a benchmark that has been used by subsequent approaches for evaluation of SSDA methods. Here we describe some of them. APE [20] uses different feature alignment objectives, within and across domains along with a perturbation consistency objective. BiAT [19] uses multiple adversarial perturbation strategies and consistency losses alongside MME. Li *et al.* [23] proposed an online meta-learning framework using target domain labelled data for meta-testing. They evaluated this approach with multiple domain alignment methods. We used their meta-MME model on SSDA for comparison. LIRR [22], optimizes for invariant classifier risk besides invariant representations across the two domains. We compared with two variants, the basic LIRR model and LIRR+CosC, which uses a cosine classifier. CDAC [24] uses a pseudo-labelling criterion based on ranks of pairwise-similarity between target examples, adversarially maximizing it using the classifier, while minimizing it using the feature representation.

Self Supervision and Domain Adaptation. In the absence of any labelled training data, different self-supervision objectives [6, 9, 15, 18, 27] have been proposed that can learn semantically meaningful image features for tasks like image classification and object detection. In domain adaptation some recent approaches [9, 38, 43] have used self-supervision tasks as an auxiliary objective to regularize their model. Saito *et al.* [36], used a self-supervised feature space clustering objective for universal domain adaptation. PAC differs from these approaches in that we use rotation prediction to pretrain our feature extractor (see supp. for a comparison with [38, 43]). This helps our initial features be more semantically meaningful and better class-wise clustered. In our experiments, we compared this with recent contrastive self-supervision approaches like MoCo [18] and SimSiam [10] and found features learnt using rotation prediction to have better properties for our task. This is also in line with the findings of Wallace *et al.* [42], where rotation prediction was found to be more semantically meaningful compared to other self-supervision objectives, for a range of classification tasks across different datasets.

6 Conclusion

We showed that consistency regularization and pretraining using rotation prediction are powerful techniques in SSDA. Our method, using simply a combination of these without requiring any domain alignment, could outperform recent state of the art on this task, most of which use adversarial alignment. With our approach we demonstrated that domain alignment is not a necessity for SSDA, and that achieving well-separated target clusters allows for low classifier error with a few labelled examples. We presented a thorough analysis of both the aforementioned techniques showing how they can improve target clustering and why they are better than other options for similar approaches. We hope our analysis can help inform future SSDA work.

References

- [1] Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, and Mario Marchand. Domain-adversarial neural networks. *arXiv preprint arXiv:1412.4446*, 2014.
- [2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19:137–144, 2006.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [4] Lukas Biewald. Experiment tracking with weights and biases, 2020.
- [5] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019.
- [6] Vittorio Castelli and Thomas M Cover. On the exponential value of labeled samples. *Pattern Recognition Letters*, 16(1):105–111, 1995.
- [7] Vittorio Castelli and Thomas M Cover. The relative value of labeled and unlabeled samples in pattern recognition with an unknown mixing parameter. *IEEE Transactions on information theory*, 42(6):2102–2117, 1996.
- [8] Olivier Chapelle and Alexander Zien. Semi-supervised classification by low density separation. In *AISTATS*, volume 2005, pages 57–64. Citeseer, 2005.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [10] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021.
- [11] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. 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, 2020.
- [12] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [13] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015.
- [14] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.

- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- [16] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [19] Pin Jiang, Aming Wu, Yahong Han, Yunfeng Shao, Meiyu Qi, and Bingshuai Li. Bidirectional adversarial training for semi-supervised domain adaptation.
- [20] Taekyung Kim and Changick Kim. Attract, perturb, and explore: Learning a feature alignment network for semi-supervised domain adaptation. *arXiv preprint arXiv:2007.09375*, 2020.
- [21] Alex Kurakin, Chun-Liang Li, Colin Raffel, David Berthelot, Ekin Dogus Cubuk, Han Zhang, Kihyuk Sohn, Nicholas Carlini, and Zizhao Zhang. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020.
- [22] Bo Li, Yezhen Wang, Shanghang Zhang, Dongsheng Li, Kurt Keutzer, Trevor Darrell, and Han Zhao. Learning invariant representations and risks for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1104–1113, 2021.
- [23] Da Li and Timothy Hospedales. Online meta-learning for multi-source and semi-supervised domain adaptation. *arXiv preprint arXiv:2004.04398*, 2020.
- [24] Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu. Cross-domain adaptive clustering for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2505–2514, 2021.
- [25] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1640–1650, 2018.
- [26] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [27] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020.
- [28] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.

- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [30] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- [31] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019.
- [32] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pages 213–226. Springer, 2010.
- [33] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *International Conference on Learning Representations*, 2018.
- [34] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018.
- [35] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8050–8058, 2019.
- [36] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. *arXiv preprint arXiv:2002.07953*, 2020.
- [37] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in neural information processing systems*, pages 1163–1171, 2016.
- [38] Yu Sun, Eric Tzeng, Trevor Darrell, and Alexei A Efros. Unsupervised domain adaptation through self-supervision. *arXiv preprint arXiv:1909.11825*, 2019.
- [39] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017.
- [40] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *(IEEE) Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

- [41] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 3635–3641. AAAI Press, 2019.
- [42] Bram Wallace and Bharath Hariharan. Extending and analyzing self-supervised learning across domains. *arXiv preprint arXiv:2004.11992*, 2020.
- [43] Jiaolong Xu, Liang Xiao, and Antonio M López. Self-supervised domain adaptation for computer vision tasks. *IEEE Access*, 7:156694–156706, 2019.
- [44] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael I Jordan. Bridging theory and algorithm for domain adaptation. *arXiv preprint arXiv:1904.05801*, 2019.