

UBR²S: Uncertainty-Based Resampling and Reweighting Strategy for Unsupervised Domain Adaptation

Tobias Ringwald
tobias.ringwald@kit.edu

Rainer Stiefelhagen
rainer.stiefelhagen@kit.edu

Institute for Anthropomatics and
Robotics (CV:HCI Lab)
Karlsruhe Institute of Technology
Karlsruhe, Germany

Abstract

Unsupervised domain adaptation (UDA) deals with the adaptation process of a model to an unlabeled target domain while annotated data is only available for a given source domain. This poses a challenging task, as the domain shift between source and target instances deteriorates a model's performance when not addressed. In this paper, we propose UBR²S – the Uncertainty-Based Resampling and Reweighting Strategy – to tackle this problem. UBR²S employs a Monte Carlo dropout-based uncertainty estimate to obtain per-class probability distributions, which are then used for dynamic resampling of pseudo-labels and reweighting based on their sample likelihood and the accompanying decision error. Our proposed method achieves state-of-the-art results on multiple UDA datasets with single and multi-source adaptation tasks and can be applied to any off-the-shelf network architecture. Code for our method is available at <https://gitlab.com/tringwald/UBR2S>.

1 Introduction

Modern convolutional neural networks (CNNs) require the optimization of millions of parameters by learning from a vast amount of training examples [1]. While this training data might be readily available for a given source domain, annotated data for the actual domain of interest – the target domain – might be nonexistent or very hard to obtain. For example, synthetic images could be generated en masse and used for training, while classifying unannotated real world data (e.g. medical images) is the actual objective. Unfortunately, the domain shift between the source and target data results in a severely degraded performance when utilizing current classification approaches.

Unsupervised domain adaptation (UDA) seeks to address the domain shift problem under the assumption that no annotated data for the target domain is available. Prior work in this area approached the problem from several different angles: Image and pixel-level methods were proposed for learning a direct image-to-image mapping between the different domains, thereby enabling the transfer of target data into the source domain (or vice versa), where straightforward training is then feasible [2, 3]. At the feature-level, UDA methods often rely on minimizing common divergence measures between the source and

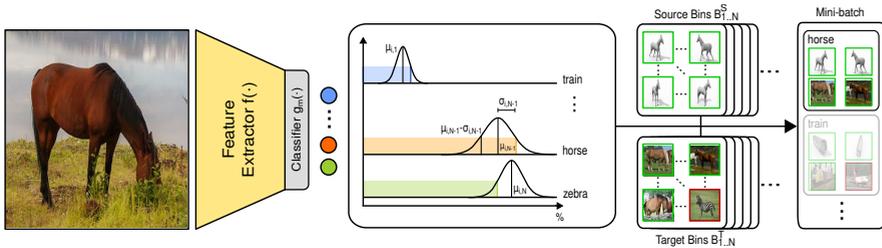


Figure 1: Overview of our proposed resampling strategy. Before each adaptation cycle, the current state of the network is frozen and used to extract class-wise uncertainty distributions quantified by Monte Carlo dropout. Based on these uncertainty distributions, class scores are resampled and converted into a pseudo-label. The instances are then grouped together with similar instances into bins, which are later used for sampling mixed mini-batches containing both source and target examples.

target distributions, such as the maximum mean discrepancy (MMD) [16], Kullback-Leibler divergence [25] or by enforcing features from both domains to be indistinguishable with the help of adversarial training [14, 28]. However, these approaches often require complicated training setups and additional stages such as domain discriminators [28], gradient reversal layers [9] or other domain-specific building blocks [2]. This adds both millions of parameters to be optimized and also further hyperparameters that have to be tuned. In this paper, we instead rely on a model’s inherent prediction uncertainty for the unsupervised domain adaptation task, which we quantify by Monte Carlo dropout [8]. Our proposed method leverages the extracted uncertainty for dynamic resampling, assignment and reweighting of pseudo-labels and can be directly applied to any off-the-shelf neural network architecture without any modification. Furthermore, we propose domain specific smoothing (DSS) – a label smoothing based improvement to the training pipeline for UDA setups.

To summarize, our contributions are as follows: (i) We propose UBR²S – the uncertainty-based resampling and reweighting strategy utilizing a model’s prediction uncertainty under Monte Carlo dropout. (ii) We introduce DSS – the domain specific smoothing operation for UDA tasks. (iii) We evaluate our method on multiple common UDA benchmarks and achieve state-of-the-art results on both single and multi-source adaptation tasks. (iv) We show that UBR²S works with a plethora of network architectures and outperforms recent methods while using only a fraction of their parameters.

2 Related Work

Unsupervised domain adaptation has been the subject of many prior works in recent years and was addressed in multiple different ways. The authors of [10] leverage GAN-based style transfer in order to translate synthetic data into the real domain and achieve domain adaptation for a monocular depth estimation task. Similarly, Deng *et al.* [5] apply unsupervised image-to-image translation for their person re-identification task and consider the self-similarity and domain-dissimilarity of source, target and translated images. Instead of aligning domains at the image-level, prior methods have also considered alignment at the feature-level: One of the first works in this direction was the gradient reversal layer (*RevGrad*)

proposed by Ganin *et al.* [9]. They achieve domain-invariant feature representations by forcing the feature distributions from the source and target domain to be as indistinguishable as possible with a domain classifier. Related to this, Pinheiro [28] proposes a similarity-based classifier that combines categorical prototypes with domain-adversarial learning. Park *et al.* [26] show that training with their proposed adversarial dropout can also help to improve generalization. In place of adversarial training, distribution divergence measures have also been applied successfully in order to align the source and target domain at the feature-level: Long *et al.* [17] use the maximum mean discrepancy (MMD) measure for alignment, which recently was extended by Kang *et al.* [16] for their proposed CDD loss and also used in the regularizer proposed by Gholami *et al.* [10]. In a similar way, Meng *et al.* [25] utilize the Kullback-Leibler divergence as another distribution difference measure. Recently, Hoffman *et al.* [12] have combined both feature- and image-level approaches in their proposed CyCADA framework, which adapts feature representations by enforcing local and global structural consistency and also employs cycle-consistent pixel transformations. In terms of pseudo-labeling target instances, Saito *et al.* [33] propose an asymmetric training method that consists of three separate networks where two networks act as pseudo-labelers and the third network is trained on said labels to obtain discriminative representations for the target domain samples. Related to this, Zhang *et al.* [39] propose an iterative pseudo-labeling and sample selection approach based on an image and domain classifier. This idea is also picked up by Chen *et al.* [8], who employ their progressive feature alignment network and an easy-to-hard transfer strategy for iterative training. Another direction was pursued by Chang *et al.* [7], who propose the use of domain-specific batch-normalization layers in order to deal with the distribution shift between domains. This concept was also employed in the contrastive adaptation network (CAN) proposed by [18] and proved to capture the domain specific image distributions.

Most related to our work in this paper, Long *et al.* [22] have explored the use of uncertainty for their proposed CDAN architecture and achieve domain adaptation by controlling the classifier uncertainty to guarantee transferability between domains. Han *et al.* [13] quantify model uncertainty under a general Rényi entropy regularization framework and utilize it for calibration of the prediction uncertainties between the source and target domain. Gholami *et al.* [10] consider the Shannon entropy of probability vectors in order to minimize a classifier’s uncertainty on unlabeled target domain instances. Similar to the approaches above, Manders *et al.* [24] propose an adversarial training setup forcing prediction uncertainties to be indistinguishable between domains. In this work, we explore the usage of prediction uncertainties quantified under the Monte Carlo dropout [8] approximation of Bayesian inference. Unlike prior work, our proposed UBR²S method leverages a model’s prediction uncertainty for dynamic resampling, assignment and reweighting of pseudo-labels. Our method does not require any image- or feature-level adjustments and can thus be applied to any off-the-shelf neural network. Nevertheless, it still achieves state-of-the-art results and is also competitive when using smaller feature extractors with a fraction of the usual parameters.

3 Methodology

Unsupervised domain adaptation (UDA) seeks to address the domain shift between a source and target domain in order to maximize a model’s generalization performance on the target domain while only given annotated data for the source domain. Formally, the annotated

source dataset \mathcal{D}_S consists of input-label pairs $\{x_i^s, y_i^s\} \in \mathcal{D}_S$ while the target dataset \mathcal{D}_T only contains unlabeled inputs $\{x_i^t\}$. Labels y_i^s are elements of class set $\mathcal{C} = \{1, 2, \dots, N\}$ with N classes. Given this definition, the objective of UDA tasks is to produce accurate predictions y_i^t for every input x_i^t of the target domain dataset.

The method discussed in this paper is presented in the context of deep neural networks that consist of a convolutional neural network (CNN) feature extractor $f(\cdot)$ followed by a classifier $g(\cdot)$ that projects f 's output into a probability distribution over the class set \mathcal{C} . A key part of our proposal is the concept of uncertainty quantified by Monte Carlo dropout (MCD) [8], which we will now formally introduce. Let \mathcal{M} be a set of size $|\mathcal{M}|$ containing binary masks $m_{1..|\mathcal{M}|}$ sampled from a Bernoulli distribution according to the Monte Carlo dropout rate, where $|\mathcal{M}|$ represents the number of MCD iterations. We then evaluate all dropout-masked classifiers $g_{m \in \mathcal{M}}(x_i^t)$ for a given target domain sample x_i^t and quantify its class-wise uncertainties by the mean μ and standard deviation σ for every class $c \in \mathcal{C}$ as follows:

$$\mu_c(x_i^t) = \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} [g_m(f(x_i^t))]_c, \quad \sigma_c(x_i^t) = \sqrt{\frac{1}{|\mathcal{M}| - 1} \sum_{m \in \mathcal{M}} [g_m(f(x_i^t))]_c - \mu_c(x_i^t)]^2} \quad (1)$$

For the sake of clarity, $\mu_c(x_i^t)$ will be shortened as $\mu_{i,c}$ in the following paragraphs ($\sigma_{i,c}$ likewise). With these prerequisites, we will now introduce our proposed uncertainty-based resampling and reweighting strategy (UBR²S).

3.1 Resampling Strategy

For unsupervised domain adaptation tasks, annotations are only available for the source domain. These are oftentimes used for model initialization via supervised pretraining and allow for initial pseudo-label estimates on the unannotated target domain. Due to the domain shift between source and target instances, these initial estimates are inherently noisy. Despite their noisy nature, recent research in this area [0, B9] often relies on the class predicted with maximum probability score in order to generate a pseudo-label, thereby neglecting the possibility of other classes. For UDA tasks, however, a model's maximum probability prediction for target domain data often does not correspond with the ground truth class after the *source-only* pretraining stage. Our resampling strategy will thus consider predictions other than the maximum for assignment of a pseudo-label.

Given f and g after supervised pretraining on the source domain dataset, we start by extracting uncertainty measures $\mu_{i,c}$ and $\sigma_{i,c}$ for the c -th class of the i -th target domain sample (see Equation 1). As this is a continuous distribution, we first resample the i -th target instance's probability scores as $\tilde{p}_{i,c} \sim \mathcal{N}(\mu_{i,c}, \sigma_{i,c})$ in order to obtain discrete values. Subsequently, \tilde{p}_i is re-normalized and then used for the assignment of a pseudo-label $\psi(\frac{\tilde{p}_i}{\sum_j \tilde{p}_{i,j}}) = \tilde{y}_i$ where $\psi: \mathbb{R}^{|\mathcal{C}|} \rightarrow \mathcal{C}$ is the weighted random sample function. This thus enables the usage of classes with non-maximum prediction scores for pseudo-labels based on the model's own predictive uncertainty. The resampling step ends by assigning the i -th target sample to the bin \mathcal{B}_v^T based on $v_i = \operatorname{argmax}_{c \in \mathcal{C}} \mu_{i,c}$. Here, bins are groups of similar samples that are later used for construction of mini-batches (see Section 3.4).

The above resampling process of our proposed method is also visualized in Figure 1.

3.2 Reweighting Strategy

Resampling a pseudo-label as described above allows for the consideration of non-maximum predictions. However, this is accompanied by the inherent risk of sampling the wrong class. After all, the maximum prediction is a reasonable estimate for some instances. We thus need to compensate for the potential decision error when choosing one class over another and also consider the likelihood of the current resampled value. We calculate the sample likelihood (SL) as $\lambda_{\text{SL}}^i = 1 - \frac{|\tilde{p}_{i,\tilde{y}_i} - \mu_{i,\tilde{y}_i}|}{2\sigma_{i,\tilde{y}_i}} \Big|_0^1$ where $\cdot \Big|_0^1$ clamps the value into range $[0, 1]$. Intuitively, this reflects the likelihood of the resampling step by measuring the deviation from the mean. However, it does not consider the risk involved with choosing the wrong class in the first place. For this reason, we determine this decision error based on the classes' uncertainty distributions by calculating the inverse of the cumulative probability Φ w.r.t. $\tilde{p}_{i,\tilde{y}_i}$ as per Equation 2:

$$\varphi(i, c) = 1 - \Phi(\tilde{p}_{i,\tilde{y}_i}, \mu_{i,c}, \sigma_{i,c}), \text{ with } \Phi(x, \mu, \sigma) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{x - \mu}{\sigma\sqrt{2}} \right) \right] \quad (2)$$

Here, erf denotes the Gauss error function. The final decision error λ_{DE} is then given by Equation 3 and calculated w.r.t. every class besides the current label estimation \tilde{y}_i . The graphical interpretation of this procedure is visualized in Figure 2a.

$$\lambda_{\text{DE}}^i = 1 - \max(\{\varphi(i, c) \mid \forall c \in \mathcal{C} \setminus \tilde{y}_i\}) \quad (3)$$

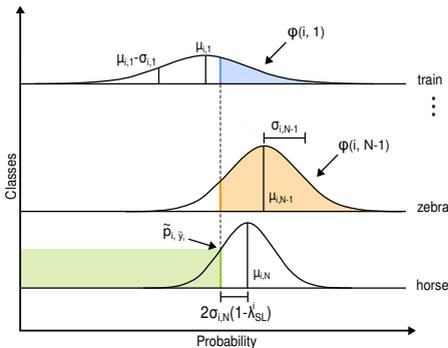
Finally, we normalize the product of λ_{SL} and λ_{DE} to a distribution with its center point at 1 and use it to dynamically reweigh the element-wise loss while training on target domain samples. Therefore, the loss contribution of a given sample depends on the certainty of the currently chosen pseudo-label. An in-depth description of this procedure will be given in Section 3.4.

3.3 Domain Specific Smoothing

This section will now introduce our proposed Domain Specific Smoothing (DSS) for UDA training. DSS is based on label smoothing, which was first proposed by Szegedy *et al.* [51] and is commonly used to curb overfitting and overconfident predictions. In a normal N class training setup, a label encoding function $v : \mathcal{C} \rightarrow \mathbb{R}^N$ would construct a discrete one-hot probability distribution so that one class is assigned 100% with all other $N - 1$ classes being at 0%. With label smoothing, v constructs a smoothed label vector by mapping a ground truth label (or estimated pseudo-label) $c \in \mathcal{C}$ into probability space according to Equation 4.

$$v(c)_i = \begin{cases} 1 - \varepsilon, & c = i \\ \frac{\varepsilon}{|\mathcal{C}| - 1}, & c \neq i \end{cases} \quad (4)$$

Here, ε is the smoothing factor. When training with a cross entropy loss, probabilities are needed for the loss calculation. However, neural networks usually output unnormalized logits. Softmax normalization is thus applied to convert the logits into probabilities: $\vartheta(\ell)_i = \frac{e^{\ell_i}}{\sum_j e^{\ell_j}}$ where ℓ is the logit vector. Cross entropy loss is then given as $-\sum_i y_i \log \vartheta(\ell)_i$.



(a) Graphical interpretation

Input. Source-trained f and g with weights θ ;

```

for  $j = 1; j \leq T_{\text{cycles}}$  do
   $\mu, \sigma \leftarrow \text{extractUncertainty}(\mathcal{D}_{\mathcal{T}}, f, g, \theta)$ ;
  for  $step = 0; step < T_{\text{steps}}$  do
    if  $step \bmod 10 = 0$  then
       $\bar{p} \leftarrow \text{resample}(\mu, \sigma)$ ;
       $\bar{y} \leftarrow \Psi\left(\frac{\bar{p}_i}{\sum_n \bar{p}_{i,n}}\right)$ ;
       $\forall i: v_i \leftarrow \text{argmax}_{c \in \mathcal{C}} \mu_{i,c}$ ;
       $\forall i: \mathcal{B}_{v_i}^T \leftarrow \mathcal{B}_{v_i}^T \cup \{x_i^c\} \in \mathcal{D}_{\mathcal{T}}$ ;
       $c \leftarrow \text{sampleClasses}(\mathcal{C}, \beta)$ ;
       $b \leftarrow \text{sampleBatch}(\mathcal{B}^S, \mathcal{B}^T, c)$ ;
       $\lambda_{\text{SL}}, \lambda_{\text{DE}} \leftarrow \text{calcError}(\mu, \sigma, \bar{p}, \bar{y})$ ;
       $\theta \leftarrow \text{train}(\theta, b, \lambda_{\text{SL}}, \lambda_{\text{DE}})$ ;

```

(b) Training loop

Figure 2: (a) Graphical interpretation of our proposed reweighting step. Using resampling provides the opportunity to consider a non-maximum prediction at the cost of sampling a wrong pseudo-label. Our loss reweighting step assesses this risk based on the class-wise uncertainty distributions by calculating the decision error based on $\varphi(i, c)$ and sample likelihood λ_{SL}^i for the i -th target instance in a N -class classification task. (b) Training loop of UBR²S.

Let output logit vector $\ell = [\ell_0, \ell_1]$ and training target $y = [1.0, 0.0]$ be subject of a training step with 2 classes. The objective of training with cross entropy loss is the assignment of target class c_0 to 100% and c_1 to 0%. Because of softmax normalization, this can only be the case when logit $\ell_0 \rightarrow \infty$ or $\ell_1 \rightarrow -\infty$. Due to the lack of floating point accuracy, this happens before approaching infinity in reality: $\ell = [19, 0]$ already results in a $\vartheta(\ell) \approx [1.0, 0.0]$ assignment. However, when using label smoothed target $\bar{y} = [0.8, 0.2]$, logits $\ell = [2 \log 2, 0.0]$ are already enough to match the target probabilities of \bar{y} (with $\varepsilon = 0.2$). As softmax is invariant to constant addition, $\ell = [q + 2 \log 2, q]$ leads to the same result for any choice of $q \in \mathbb{R}$. The absolute difference needed between ℓ_0 and ℓ_1 is therefore multiple times smaller for the smoothed target ($\frac{2 \log 2}{19} \approx 0.07$). For training with target y , this has the consequence of rapidly growing weights prior to the output layer in order to boost the logit values and thereby minimizing the cross entropy loss. This, however, is a prime example of overfitting as stated by Krogh *et al.* [18] and is also in conflict with Lawrence *et al.* [19] who found that smaller weights tend to generalize better. Thus, label smoothing can be seen as a regularization method that diminishes this adverse influence on training.

Similar to prior UDA setups [16, 39], our method constructs mini-batches using instances from both the source and target domain. While the presence of ground truth source instances can diminish the effect of wrong target pseudo-labels, constant training on source data will make the model overly focus on this domain even though the adaptation to the target domain is the actual objective. This is in conflict with prior research, which indicates that good transferability and generalization requires a non-saturated source classifier [3]. We extend this idea to domain adaptation tasks and propose domain specific smoothing (DSS): With DSS, label smoothing is only applied to source instances, even when training with a mixed batch. This leads to a mixture of one-hot pseudo-labels for the target domain instances and smoothed ground truth labels for the source instances. As later shown in our experiments, this helps to improve the generalization performance after the pretraining on source domain



Figure 3: From left to right: VisDA 2017 [27] with domains synthetic (train set) and real (validation and test set), Office-Caltech [11] with domains Amazon, Caltech, DSLR and Webcam, Office-31 [52] with domains Amazon, DSLR and Webcam and Office-Home [58] with domains Art, Clipart, Product and Real World.

data and also improves the domain adaptation capabilities of the final model.

3.4 Training Loop

With the major parts of our training pipeline described above, we provide an overview for the UBR²S training process in Figure 2b. After the uncertainty extraction and resampling process, β classes are sampled from class set \mathcal{C} . For every class c , $\frac{|b|}{2\beta}$ samples are randomly drawn from the source and target bins \mathcal{B}_c^S and \mathcal{B}_c^T where $|b|$ is the batch size. The source bins are constructed based on the available ground truth labels while the target bins are reconstructed based on the label estimation described in Section 3.1.

Finally, we calculate the sample likelihood and decision error for all target instances in a mini-batch and use it to reweigh their element-wise loss during training. Given the k -th target example x_k in a mini-batch with a total of K target instances we compute our proposed reweighted cross entropy loss as Equation 5 with weight ω , where $v(\cdot)$ represents the label smoothing function from Section 3.3 when using DSS and the one-hot encoding function otherwise.

$$\mathcal{L}(x_k, \tilde{y}_k) = -\omega_k \sum_{c \in \mathcal{C}} v(\tilde{y}_k)_c \log [g(f(x_k))_c], \text{ with } \omega_k = \frac{\lambda_{\text{DE}}^k \lambda_{\text{SL}}^k}{\frac{1}{K} \sum_j \lambda_{\text{DE}}^j \lambda_{\text{SL}}^j} \quad (5)$$

For source domain instances, weight ω is set to 1. Parameters θ of f and g are then updated by backpropagation according to this loss. Subsequent iterations use the updated weights θ for the uncertainty extraction and training process. Therefore, only a single neural network is needed during the complete training and adaptation phase.

4 Experiments

Datasets. We evaluate our proposed UBR²S method on four public benchmark datasets: *VisDA 2017* (also known as Syn2Real-C) [27] is a large scale dataset for the synthetic to real UDA task. It contains 12 classes in three domains: train (152,397 synthetic 3D renderings), validation (55,388 real world images from MS COCO [20]) and test (72,372 real world images from YouTube Bounding-Boxes [60]). For comparison to state-of-the-art methods, we

DSS _{Pre}	DSS _{Ada}	Reweigh	S→R _{test} Pre	S→R _{test}	Ar→Cl Pre	Ar→Cl	Pr→Ar Pre	Pr→Ar
×	×	×	49.4	78.7	44.0	52.5	51.1	58.1
S	×	×	51.3	82.5	46.0	54.1	54.1	58.8
S	S	×	51.3	83.8	46.0	54.6	54.1	61.8
S	T	×	51.3	67.0	46.0	38.6	54.1	47.5
S	S, T	×	51.3	67.8	46.0	47.4	54.1	53.6
S	S	SL	51.3	85.4	46.0	56.7	54.1	64.1
S	S	DE	51.3	89.5	46.0	57.4	54.1	65.1
S	S	DE+SL	51.3	89.8	46.0	58.3	54.1	67.0

Table 1: Ablation study on the VisDA 2017 S→R_{test} (mean class accuracy) and Office-Home Ar→Cl, Pr→Ar tasks (accuracy). Transfer tasks marked with $\overset{\rightarrow}{\text{Pre}}$ indicate results after the source only pretraining and before the adaptation step. The last table row represents our full UBR²S method.

calculate the mean class accuracy w.r.t. to the challenge evaluation protocol unless otherwise noted. *Office-31* [62] is one of the most used UDA datasets and contains 4,110 images of 31 classes in a generic office setting. The images come from the three domains Amazon (product images), DSLR and Webcam. The *Office-Caltech* [14] dataset is constructed from the 10 overlapping classes between the Caltech-256 [14] and Office-31 [62] datasets for a total of four domains: Amazon (958), Caltech (1,123), DSLR (157) and Webcam (295). The *Office-Home* [68] dataset offers 65 challenging classes from everyday office life. Its four domains Art, Clipart, Product and Real World contain a total of 15,588 images. Example images from all datasets are shown in Figure 3.

Setup. For our experiments, we follow the standard unsupervised domain adaptation setup (see [4, 16]) and use all labeled source domain and all unlabeled target domain images for training. For a detailed description of the employed setup and training procedure, please refer to the supplementary material.

4.1 Results

Ablation Study. We first conduct an ablation study w.r.t. to every part of our proposed UBR²S method. For this, we use ResNet-101 on the VisDA 2017 S→R (test set) task as well as ResNet-50 on the Ar→Cl and Pr→Ar transfer tasks from Office-Home. Results are reported in Table 1. First, we examine our proposed domain specific smoothing (DSS) method. Our baseline does not use DSS during pretraining (DSS_{Pre}[×]) or during the adaptation phase (DSS_{Ada}[×]). Expectedly, this baseline performs the worst for all three transfer tasks due to overfitting on the source domain. Applying DSS to the source samples during the pretraining phase (DSS_{Pre}^S) already improves pretrained results by almost 2% for VisDA (49.4% to 51.3%) and final results by almost 4% (78.7% to 82.5%). Similar trends can be observed for the Office-Home transfer tasks. Concerning DSS_{Ada}, we evaluate all possible combinations {×, S, T, S + T}. We find that applying label smoothing to the target domain (DSS_{Pre}^T and DSS_{Pre}^{S,T}) has a negative impact on the model’s accuracy. This is consistent for all three transfer tasks and can reduce accuracy by up to 16.8%. Conversely, adding label smoothing to the source domain always improves performance, considering both the DSS_{Ada}[×] to DSS_{Ada}^S and DSS_{Ada}^T to DSS_{Ada}^{S,T} transitions. Overall, DSS_{Pre}^S with DSS_{Ada}^S consistently achieved the best results. This confirms our hypothesis from Section 3.3 and implies that a non-saturated source classifier is needed for good transferability and generalization to new domains in UDA tasks. Further ablation studies are shown in the supplementary material.

We continue by examining the remaining parts of our UBR²S method in Table 1 using

Method	A			C			D			W			Avg.
	C	D	W	A	D	W	A	C	W	A	C	D	
CORAL [14, 15]	89.2	92.2	91.9	94.1	92.0	92.1	94.3	87.7	98.0	92.8	86.7	100.0	92.6
GTDA+LR [14]	91.5	98.7	94.2	95.4	98.7	89.8	95.2	89.0	99.3	95.2	90.4	100.0	94.8
RWA [15]	93.8	98.9	97.8	95.3	<u>99.4</u>	95.9	95.8	93.1	98.4	95.3	92.4	<u>99.2</u>	96.3
PrDA [15]	92.1	99.0	<u>99.3</u>	97.2	<u>99.4</u>	98.3	94.7	91.0	<u>99.7</u>	95.6	93.4	100.0	96.6
Rakshit <i>et al.</i> [14]*	92.8	98.9	97.0	96.0	99.0	97.0	<u>96.5</u>	97.0	99.5	95.5	91.5	100.0	96.8
ACDA [15]	<u>93.9</u>	100.0	100.0	96.2	100.0	100.0	96.7	<u>93.9</u>	100.0	96.6	93.9	100.0	97.6
UBR²S (ours)	95.5	<u>99.4</u>	<u>99.3</u>	<u>96.6</u>	94.9	<u>99.7</u>	96.2	<u>95.3</u>	100.0	<u>96.2</u>	95.4	100.0	<u>97.4</u>

Table 2: Classification accuracy (in %) for different methods using ResNet-50 on the **Office-Caltech** dataset with domains Amazon, Caltech, DSLR and Webcam. The method marked with * uses an ensemble setup with multiple classifiers.

Method	aero	bicyc	bus	car	horse	knife	motor	person	plant	skate	train	truck	Avg.
Source only	52.8	13.8	66.9	96.3	58.4	14.0	63.4	34.5	86.0	24.5	87.3	17.9	51.3
BUPT [17]	<u>95.7</u>	67.0	93.4	97.2	90.6	86.9	92.0	74.2	<u>96.3</u>	66.9	95.2	69.2	85.4
CAN [15]	—	—	—	—	—	—	—	—	—	—	—	—	87.4
SDAN [15]	94.3	86.5	86.9	95.1	91.1	90.0	<u>82.1</u>	77.9	96.4	77.2	86.6	88.0	87.7
UFAL [15]	94.9	<u>87.0</u>	<u>87.0</u>	<u>96.5</u>	<u>91.8</u>	95.1	76.8	78.9	96.5	<u>80.7</u>	<u>93.6</u>	86.5	<u>88.8</u>
UBR²S (ours)	96.6	90.8	87.9	94.6	92.2	<u>92.8</u>	77.8	<u>78.8</u>	95.3	89.2	92.6	88.9	89.8

Table 3: Per class accuracy (in %) for different methods on the **VisDA 2017 test set** as per challenge evaluation protocol. Results are obtained using the common ResNet-101 backbone.

the best performing DSS_{Pre}^S with DSS_{Ada}^S setup as baseline. As mentioned in Section 3.2, the resampling process in UBR²S comes with an inherent risk and the possibility for errors. This risk can be partially measured with the help of our proposed sample likelihood (SL) and used for reweighting, which already improves results by 1.6% for VisDA and up to 2.3% for the Office-Home tasks. It is also important to assess the current label estimation and how the chosen pseudo-label compares to other potential candidates. This is covered by our proposed decision error (DE), which – when solely used for reweighting – can improve results by 5.7% for VisDA and up to 3.3% for the Office-Home tasks. Collectively, the combination of SL and DE can further improve results by 6.0% for VisDA and up to 5.2% for Office-Home over the respective baselines and constitutes our full UBR²S method.

Comparison to state-of-the-art. We continue by comparing UBR²S to other recently proposed approaches. Results for the Office-Caltech dataset are shown in Table 2. Evidently, UBR²S can achieve domain adaptation even for Office-Caltech’s 12 diverse transfer tasks. Our proposed method achieves the best or second best result in 10 out of 12 transfer tasks (such as A→C) and is also on par with ACDA [15] for the overall average. Notably, UBR²S even manages to surpass the ensemble-based setup of Rakshit *et al.* [14] by 0.6%. Additionally, we report results for the VisDA 2017 test set in Table 3 and calculate the class accuracies as per challenge evaluation protocol. Our results indicate that UBR²S can also achieve domain adaptation for VisDA’s difficult synthetic to real transfer task and outperforms recently proposed methods. With 89.8% mean class accuracy, UBR²S also outperforms the VisDA challenge submissions SDAN [15] and BUPT [17]. Given the current VisDA 2017 challenge leaderboard [17], UBR²S would rank second place – only behind SE [9], a 5×ResNet-152 ensemble with results averaged over 16 test time augmentation runs. This, however, is not a fair comparison to our single ResNet-101 model, but still demonstrates UBR²S’ competitive UDA capabilities.

Visualization Finally, we visualize the embeddings learned by UBR²S in Figure 4. For this, we extract target domain features for the VisDA 2017 test set (real domain) before

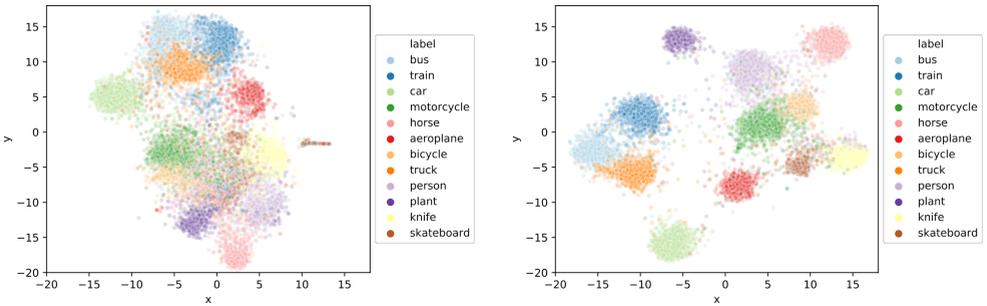


Figure 4: Visualizations of the VisDA 2017 test set features using t-SNE. Left: After the source-only pretraining phase. Right: After the adaptation phase using UBR²S.

and after the adaptation phase and project them into 2D space via t-SNE [23]. After the source-only pretraining phase, the model clearly has not learned discriminative representations for target domain samples. Features of all classes are accumulated in one big cluster with no clear separation. After UBR²S’ unsupervised domain adaptation phase, visually distinct clusters for each of VisDA’s 12 classes can be observed. This indicates that UBR²S is also able to learn discriminative target domain representations even in the absence of target domain annotations.

Due to space limitations, we provide further results in the supplementary material. This includes additional ablation studies, full results on the Office-Home [58] dataset, multi-source UDA results on the Office-31 [57] dataset, an evaluation with different network backbones and an analysis of training stability.

5 Conclusion

In this paper, we propose UBR²S, the uncertainty-based resampling and reweighting strategy. UBR²S’ resampling phase is based on a model’s prediction uncertainty quantified by Monte Carlo dropout. As resampling introduces the possibility of sampling a wrong pseudo-label, a dynamic reweighting stage is added to assess and incorporate this risk in the loss calculation. The efficacy of UBR²S is shown on multiple UDA benchmark datasets such as VisDA 2017, Office-Caltech, Office-31 and Office-Home in single and multi-source domain adaptation setups in which UBR²S outperforms recently proposed methods and achieves state-of-the-art results. Furthermore, we show that UBR²S can be applied to any off-the-shelf CNN and works even with very small networks (such as MobileNetV2) with extremely low parameter counts. Our code is made available on the project website for reproduction of our results and to encourage further research in the area of unsupervised domain adaptation.

6 Acknowledgement

This work was partially supported by the Competence Center Karlsruhe for AI Systems Engineering (CC-KING, <https://www.ai-engineering.eu>) sponsored by the Ministry of Economic Affairs, Labour and Housing Baden-Württemberg.

References

- [1] Amir Atapour-Abarghouei and Toby P Breckon. Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2800–2810, 2018.
- [2] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7354–7362, 2019.
- [3] Chaoqi Chen, Weiping Xie, Huang, et al. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 627–636, 2019.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [5] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 994–1003, 2018.
- [6] Geoff French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rkpoTaxA->.
- [7] B. Chidlovskii G. Csurka and S. Clinchant. VisDA Classification Challenge:Runner-Up Talk. https://ai.bu.edu/visda-2017/assets/attachments/VisDA_NaverLabs.pdf, 2020. Accessed: 2020-03-31.
- [8] Yarín Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML 15*, pages 1180–1189. JMLR.org, 2015. URL <http://dl.acm.org/citation.cfm?id=3045118.3045244>.
- [10] Behnam Gholami, Vladimir Pavlovic, et al. Punda: Probabilistic unsupervised domain adaptation for knowledge transfer across visual categories. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3581–3590, 2017.
- [11] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2066–2073. IEEE, 2012.
- [12] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007.

- [13] Ligong Han, Yang Zou, Ruijiang Gao, Lezi Wang, and Dimitris Metaxas. Unsupervised domain adaptation via calibrating uncertainties. In *CVPR Workshops*, volume 9, 2019.
- [14] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle consistent adversarial domain adaptation. In *International Conference on Machine Learning (ICML)*, 2018.
- [15] Kevin Hua and Yuhong Guo. Unsupervised domain adaptation with progressive domain augmentation. *arXiv preprint arXiv:2004.01735*, 2020.
- [16] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019.
- [17] Ben Usman Kate Saenko and Xingchao Peng. 2017 Visual Domain Adaptation (VisDA2017) Classification Challenge . <https://competitions.codalab.org/competitions/17052>, 2017. Accessed: 2021-04-06.
- [18] Anders Krogh and John A Hertz. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, pages 950–957, 1992.
- [19] Steve Lawrence and C Lee Giles. Overfitting and neural networks: conjugate gradient and backpropagation. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 1, pages 114–119. IEEE, 2000.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [21] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jianguang Sun, and Philip S Yu. Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE international conference on computer vision*, pages 2200–2207, 2013.
- [22] 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.
- [23] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [24] Jeroen Manders, Twan van Laarhoven, and Elena Marchiori. Adversarial alignment of class prediction uncertainties for domain adaptation. *arXiv preprint arXiv:1804.04448*, 2018.
- [25] Zhong Meng, Jinyu Li, Yifan Gong, and Biing-Hwang Juang. Adversarial teacher-student learning for unsupervised domain adaptation. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5949–5953. IEEE, 2018.
- [26] Sungrae Park, JunKeon Park, Su-Jin Shin, and Il-Chul Moon. Adversarial dropout for supervised and semi-supervised learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

- [27] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. *Visda: The visual domain adaptation challenge*, 2017.
- [28] Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8004–8013, 2018.
- [29] Sayan Rakshit, Ushasi Chaudhuri, Biplab Banerjee, and Subhasis Chaudhuri. Class consistency driven unsupervised deep adversarial domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [30] Esteban Real, Jonathon Shlens, Stefano Mazzocchi, Xin Pan, and Vincent Vanhoucke. Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5296–5305, 2017.
- [31] Tobias Ringwald and Rainer Stiefelagen. Unsupervised Domain Adaptation by Uncertain Feature Alignment. In *The British Machine Vision Conference (BMVC)*, 2020.
- [32] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*. Springer, 2010.
- [33] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2988–2997. JMLR. org, 2017.
- [34] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [35] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [36] Twan van Laarhoven and Elena Marchiori. Unsupervised domain adaptation with random walks on target labelings, 2017.
- [37] Sebastiano Vascon, Sinem Aslan, Alessandro Torcinovich, Twan van Laarhoven, Elena Marchiori, and Marcello Pelillo. Unsupervised domain adaptation using graph transduction games. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019.
- [38] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [39] Weichen Zhang, Wanli Ouyang, Wen Li, and Dong Xu. Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3801–3809, 2018.
- [40] Youshan Zhang and Brian D Davison. Adversarial continuous learning in unsupervised domain adaptation. In *International Conference on Pattern Recognition*, pages 672–687. Springer, 2021.