

Uncertainty-guided Source-free Domain Adaptation

Subhankar Roy^{1,2}, Martin Trapp³, Andrea Pilzer⁴, Juho Kannala³,
Nicu Sebe¹, Elisa Ricci^{1,2}, and Arno Solin³

¹ University of Trento, Trento, Italy

² Fondazione Bruno Kessler, Trento, Italy

³ Aalto University, Espoo, Finland

⁴ NVIDIA

subhankar.roy@unitn.it

Abstract. Source-free domain adaptation (SFDA) aims to adapt a classifier to an unlabelled target data set by only using a pre-trained source model. However, the absence of the source data and the domain shift makes the predictions on the target data unreliable. We propose quantifying the uncertainty in the source model predictions and utilizing it to guide the target adaptation. For this, we construct a probabilistic source model by incorporating priors on the network parameters inducing a distribution over the model predictions. Uncertainties are estimated by employing a Laplace approximation and incorporated to identify target data points that do not lie in the source manifold and to down-weight them when maximizing the mutual information on the target data. Unlike recent works, our probabilistic treatment is computationally lightweight, decouples source training and target adaptation, and requires no specialized source training or changes of the model architecture. We show the advantages of uncertainty-guided SFDA over traditional SFDA in the closed-set and open-set settings and provide empirical evidence that our approach is more robust to strong domain shifts even without tuning.

Keywords: Source-free domain adaptation, Uncertainty quantification

1 Introduction

Deep neural networks have proven to be very successful in a myriad of computer vision tasks such as categorization, detection, and retrieval. However, much of the success has come at the price of excessive human effort put into the manual data-labelling process. Since collecting annotated data can be prohibitive and impossible at times, domain adaptation (DA, see [8] for an overview) methods have gained increasing attention. They enable training on unlabelled target data by conjointly leveraging a previously labelled yet related source data set while mitigating *domain-shift* [60] between the two. Such methods predominantly comprise of minimizing statistical moments between distributions [62,37,57,53], using adversarial objectives to maximize domain confusion [14,61], or reconstructing data with generative methods [23].




Fig. 1. Illustrative sketch of source-free domain adaptation (SFDA) on a labelled source domain (🐦, 🚗) and an *unlabelled* target domain (🐦, 🚗) potentially containing additional classes (✈️). The **top-row** shows conventional methods which ignore model uncertainties; the **bottom-row** shows our method which incorporates uncertainties about the predictive model, enabling uncertainty-guided SFDA that is more robust to distribution shifts

Albeit successful, the preceding methods mandate access to the source data set during the target adaptation phase as they require an estimate of the source distribution for the alignment. With the emergence of regulations on data privacy and bottleneck in data transmission for large data sets, access to the source data can not always be guaranteed. Thus, paving the way to a relatively new and more realistic DA setting, called *source-free* DA (SFDA, [8]), where the task is to adapt to the target data set when the only source of supervision is a source-trained model. SFDA facilitates maintaining data anonymity in privacy-sensitive applications (*e.g.*, surveillance or medical applications) and at the same time reduces data transmission and storage overhead. Towards this goal, recently, several SFDA methods have been proposed that utilize the hypotheses learned from the source data [33,58,27]. Notably, SHOT [33] – an information maximization (IM) [17] based SFDA method – has demonstrated to work reasonably well on DA benchmarks, sometimes outperforming traditional DA methods. While promising, these conventional SFDA techniques do not account for the uncertainty in the predictions of the source model on the target data. As a by-product, solely maximizing mutual information [17] on the target data can lead to erroneous decision surfaces (see Fig. 1 top).

This work argues that quantification of the uncertainty in predictions is essential in SFDA. Depending on the inductive biases of the model, the source model may predict incorrect target pseudo-labels with high confidence, *e.g.*, due to the extrapolation property in ReLU networks [22] (see Fig. 2b left). In the literature, uncertainty-guided methods have been proposed in the context of traditional UDA and SFDA settings, employing Monte Carlo (MC) dropout to estimate the uncertainties in the model predictions [51,72]. However, MC

dropout requires specialized training and specialized model architecture, suffers from manual hyperparameter tuning [13], and is known to provide a poor approximation even for simple (*e.g.*, linear) models [46,47,10].

In this work, we propose to construct a probabilistic source model by incorporating priors on the network parameters, inducing a distribution over the model predictions, on the last layer of the source model. This enables us to perform an efficient local approximation to the posterior using a *Laplace approximation* (LA, [59,41]), see Fig. 2a. This principled Bayesian treatment leads to more robust predictions, especially when the target data set contains out-of-distribution (OOD) classes (see Fig. 1 bottom) or in case of strong domain shifts. Once the uncertainty in predictions is estimated, we selectively guide the target model to maximize the mutual information [17] in the target predictions. This alleviates the alignment of the target features with the wrong source hypothesis, resulting in a domain adaptation scheme that is robust to mild and strong domain shifts without tuning. We call our proposed method **Uncertainty-guided Source-Free Adaptation** (U-SFAN). Our approach requires no specialized source training or specialized architecture, opposed to exiting works (*e.g.*, [30,72]), introduces little computational overhead, and decouples source training and target adaptation.

We summarize our contributions as follows. *(i)* We emphasize the need to quantify uncertainty in the predictions for SFDA and propose to account for uncertainties by placing priors on the parameters of the source model. Our approach is computationally efficient by employing a last-layer Laplace approximation and greatly decouples the training of the source and target. *(ii)* We demonstrate that our proposed U-SFAN successfully guides the target adaptation without specialized loss functions or a specialised architecture. *(iii)* We empirical show the advantage of our method over SHOT [33] in the closed-set and the open-set setting for several benchmarks tasks and provide evidence for the improved robustness against mild and strong domain shifts.

2 Related Work

Closed-set Domain Adaptation, often abbreviated as UDA, refers to the family of DA methods that aim to learn a classifier for an unlabelled target data set while simultaneously using the labelled source data set, which differ in their underlying data distributions. In the literature [64] mainly three categories of UDA methods can be found. First, discrepancy-based UDA methods aim to diminish the domain-shift between the two domains with maximum mean discrepancy (MMD, [62,37,39]), or with correlation alignment [57,43,53]. The second category of UDA methods exploits the adversarial objective [18] to promote domain confusion between the two data distributions by using domain discriminator [14,61,38]. Finally, the third category comprises reconstruction-based UDA methods [4,16,23] that casts data reconstruction as an auxiliary objective in order to ensure invariance in the feature space. However, these methods can only work in the presence of the source data set during the adaptation stage, which might be limited in practice due to data privacy or storage concerns.

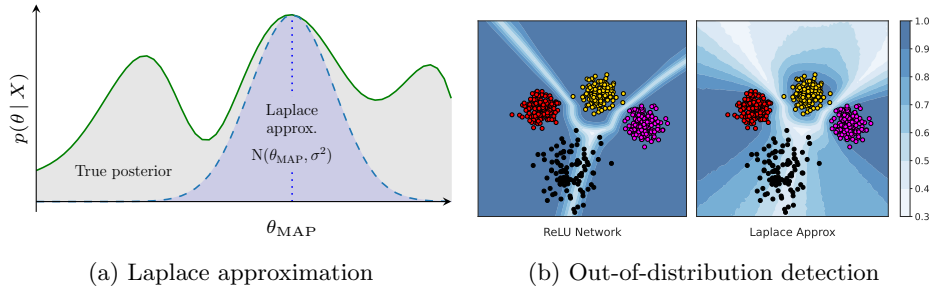


Fig. 2. (a) The Laplace approximation is mode-seeking and adapts to the local curvature around the mode θ_{MAP} . It does not necessarily capture the (intractable) full posterior, but gives a proxy for it, is principled, and efficient to evaluate. (b) Example of predictive uncertainty (un- certain) captured by a ReLU network vs. a Laplace approximation that assigns higher uncertainty to inputs (\bullet) of an unseen class

Open-set Domain Adaptation (OSDA), originally proposed in [48], refers to the DA setting where both the domains have some shared and private classes, with explicit knowledge about the shared classes. However, such a setting was deemed impractical, and later Saito *et al.* [56] proposed the open-set setting where the source labels are a subset of the target labels. Thereon, several OSDA methods have been proposed which use image-to-image translations [71], progressive filtering [36], ensemble of multiple classifiers [11] and one-vs-all classifiers [55] to detect OOD samples. Similar to the UDA, the OSDA methods also require support from the source data to detect target private classes, which make them unsuitable for source-free DA.

Source-free Domain Adaptation (SFDA) aims to adapt a model to the unlabelled target domain when only the source model is available and the source data set is absent during target adaptation. Existing SFDA methods use pseudo-label refinement [33,1,5], latent source feature generation using variational inference [70], or disparity among an ensemble of classifiers [30]. Certain SFDA methods resort to *ad hoc* source training protocols to enable the source model to be adapted on the target data. For instance, [30] requires an ensemble of classifiers to be trained during source training so that the disparity among them could be utilized for target adaptation. Similarly, USFDA [27] requires artificially generated negative samples in the source training stage for the model to detect OOD samples. Such coupled source and target training procedures make these SFDA methods less viable for practical applications. On the other hand, our proposed U-SFAN does not require specialized source training except a computationally lightweight approximate inference, which can be done with a single pass of the source data during the source training. Moreover, unlike [30,1], our U-SFAN works well on both closed-set and open-set SFDA without *ad hoc* modifications.

Uncertainty Quantification in the form of Bayesian deep learning (*e.g.*, [45,24]) is concerned with formalizing prior knowledge and representing uncertainty in model predictions, especially under domain-shift or out-of-distribution

samples. Even though the Bayesian methodology gives an explicit way of formalizing uncertainty, computation is often intractable. Thus, approximate inference methods such as Monte Carlo (MC) dropout [12], deep ensembles [29,66], other stochastic methods (*e.g.*, [42]), variational methods [3], or the Laplace approximation [52] are typically employed in practice. Prior works in semantic segmentation [72] and UDA [51,30,65,20,28] applied MC dropout or deep ensembles, respectively, for uncertainty quantification in DA. However, none of those above approaches can be considered practical for the more challenging source-free DA scenario as MC dropout, ensembles, and other stochastic methods do not lend themselves well to the source-free case. In particular, they either require retraining several models on the source, changing the model architecture or requiring a tailored learning procedure on the source data. Thus we take a Laplace approach which allows re-using the source model by linearizing around a point-estimate (see Fig. 2), which is *post hoc*, yet grounded in classical statistics [15].

3 Methods

Problem Definition and Notation We are given a labelled source data set, having $n^{[S]}$ instances, $\mathcal{D}^{[S]} = \{(\mathbf{x}_i^{[S]}, \mathbf{y}_i^{[S]})\}_{i=1}^{n^{[S]}}$, where $\mathbf{x}^{[S]} \in \mathcal{X}^{[S]}$ are D -dimensional inputs and $\mathbf{y}^{[S]} \in \mathcal{Y}^{[S]}$ where we assume K -dimensional one-hot encoded class labels, *i.e.*, $\mathcal{Y}^{[S]} = \mathbb{B}^K$. Moreover, we have $n^{[T]}$ unlabelled target observations $\mathcal{D}^{[T]} = \{\mathbf{x}_j^{[T]}\}_{j=1}^{n^{[T]}}$, where $\mathbf{x}^{[T]} \in \mathcal{X}^{[T]}$ are D -dimensional unlabelled inputs. As in any DA scenario, the assumption made is that the marginal distributions of the source and the target are different, but the semantic concept represented through class labels does not change. Formally, we assume that $p(\mathbf{y}^{[S]} | \mathbf{x}^{[S]}) \approx p(\mathbf{y}^{[S]} | \mathbf{x}^{[T]})$ and $p(\mathbf{x}^{[S]}) \neq p(\mathbf{x}^{[T]})$. In the SFDA scenario we further assume that the source data set is only available while learning the source function $f: \mathcal{X}^{[S]} \rightarrow \mathcal{Y}^{[S]}$ and becomes unavailable while adapting on the unlabelled data. The goal of SFDA is to adapt the source function f to the target domain solely by using the data in $\mathcal{D}^{[T]}$. The resulting target function, denoted as $f': \mathcal{X}^{[T]} \rightarrow \mathcal{Y}^{[T]}$, can then be used to infer the class assignment for $\mathbf{x}^{[T]} \in \mathcal{X}^{[T]}$. In this work we have considered two settings of the SFDA: i) *vanilla closed-set* SFDA where the label space of the source \mathcal{S} and the target \mathcal{T} is the same, $L^{[S]} = L^{[T]}$; and ii) *open-set* SFDA where the label space of the \mathcal{S} is a subset of the \mathcal{T} , *i.e.*, $L^{[S]} \subset L^{[T]}$, and $L^{[T]} \setminus L^{[S]}$ are denoted as *target-private* or OOD classes.

We model the source and target functions f with a neural network that is composed of two sub-networks: feature extractor g and hypothesis function h , such that $f = h \circ g$. The feature extractor g and the hypothesis function h are parameterized by parameters β and θ , respectively. During target adaptation, the model is initialized with parameters learned on $\mathcal{D}^{[S]}$ and subsequently the feature extractor parameters are updated using backpropagation, *i.e.*, the hypothesis function is kept frozen.

Overall Idea Our proposed method for SFDA operates in two stages. We begin the first stage (see Fig. 3a) by training a source model on the data set

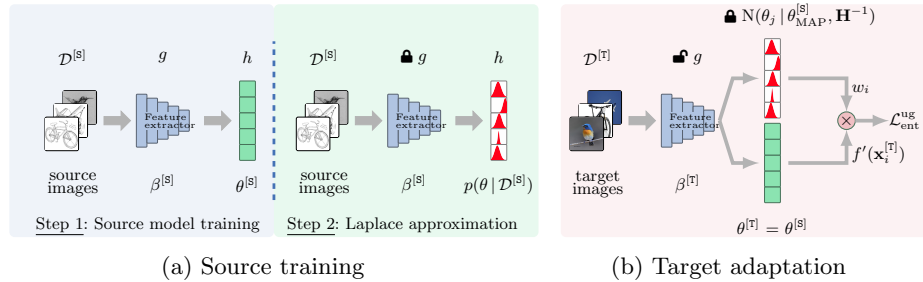


Fig. 3. The pipeline for U-SFAN: (a) Initial source model training (1) and the additional step (2) of composing a Laplace approximation for assessing the posterior over model parameters, $p(\theta | \mathcal{D}^{[S]})$. (b) At target adaptation, we keep the posterior over the parameters fixed (🔒) and train g under a uncertainty-aware composite loss that weights samples according to predictive uncertainty

$\mathcal{D}^{[S]}$, which gives us the maximum-a-posteriori (MAP) estimate of the source network parameters ($\{\beta_{\text{MAP}}^{[S]}, \theta_{\text{MAP}}^{[S]}\}$). The second stage (see Fig. 3b) comprises of maximization of mutual information [17] in the predictions for the target inputs $\mathcal{D}^{[T]}$. However, due to the overconfidence of ReLU networks [22], maximizing mutual information for all inputs equally, including those that are far away from the source data, could be detrimental. To overcome this pathology, we derive a per-sample weight using the model’s uncertainty and use it to modulate the mutual information objective in SHOT. To estimate the uncertainty in the predictions on the target data, we perform approximate posterior inference over the parameters of the hypothesis function, *i.e.*, $p(\theta^{[S]} | \mathcal{D}^{[S]})$. Inspired by recent works on approximate inference in Bayesian neural networks [40, 59, 25], we propose to estimate the posterior predictive distribution $p(\mathbf{y} | \mathbf{x}, \mathcal{D})$ using a Laplace approximation, introducing little computational overhead and without the need for specialized source training. We briefly describe the preliminaries to our approach in the following section.

3.1 Preliminaries

Liang *et al.* [33] proposed SHOT (Source **H**ypothesis **T**ransfer) for the task of SFDA, where the goal is to find a parameterization $\beta^{[T]}$ of the feature extractor g such that the distribution of latent features $\mathbf{z}^{[T]} = g_{\beta^{[T]}}(\mathbf{x}^{[T]})$ matches the distribution of the latent source features. This enables that the target data can be accurately classified by the hypothesis function parameterized by $\theta^{[S]}$. To this end, the authors address the SFDA task in two stages where the first and second stage comprise of source model training and maximizing the mutual information [17] between the latent representations and the classifier output, respectively.

The source model $f: \mathcal{X}^{[S]} \rightarrow \mathcal{Y}^{[T]}$ for a K -way classification task is learned using a label-smoothed cross-entropy objective [44], *i.e.*,

$$\mathcal{L}_{\text{src}} = -\mathbb{E}_{p(\mathbf{x}^{[S]}, \mathbf{y}^{[S]})} \sum_{k=1}^K \tilde{y}_k^{[S]} \log \phi_k(f(\mathbf{x}^{[S]})), \quad (1)$$

where $\phi_k(\mathbf{a}) = \exp(a_k) / \sum_j \exp(a_j)$ denotes the likelihood for the k^{th} component of the model output and $\tilde{y}_{i,k}^{[S]} = y_{i,k}^{[S]}(1 - \alpha) + \alpha/K$ the class label for the i^{th} label smoothed datum.

After the source training, the $\mathcal{D}^{[S]}$ is discarded and the target adaptation is conducted on $\mathcal{D}^{[T]}$ only. To adapt on the target domain, the target function f' is initialized based on the learned source function f and learned with the information maximization (IM) loss [17]. The IM loss ensures that the function mapping will produce one-hot predictions while at the same time enforcing diverse assignments, *i.e.*,

$$\mathcal{L}_{\text{ent}} = -\mathbb{E}_{p(\mathbf{x}^{[T]})} \sum_{k=1}^K \phi_k(f'(\mathbf{x}^{[T]})) \log \phi_k(f'(\mathbf{x}^{[T]})), \quad (2)$$

$$\mathcal{L}_{\text{div}} = \text{D}_{\text{KL}}(\hat{\mathbf{p}} \| K^{-1} \mathbf{1}_K) - \log K, \quad (3)$$

where $\mathbf{1}_K$ is a vector of all ones, and $\hat{p}_k = \mathbb{E}_{p(\mathbf{x}^{[T]})}[\phi_k(f'(\mathbf{x}^{[T]}))]$ is the expected network output for the k^{th} class. Intuitively, \mathcal{L}_{ent} is in charge of making the network output one-hot, while \mathcal{L}_{div} is responsible for equally partitioning the network prediction into K classes. In practice \mathcal{L}_{div} operates on a mini-batch level. In this work we start from SHOT-IM to adapt to the target domain.

3.2 Uncertainty-guided Source-free DA

Distributional shift between source and target data sets causes the network outputs to differ, even for the same underlying semantic concept [8]. In a standard UDA scenario, where the source data is available during target adaptation, it is still possible to align the marginal distributions by using a quantifiable discrepancy metric. The task becomes more challenging in the SFDA scenario because it is not possible to align the target feature distribution to a reference (or source) distribution. Moreover, standard ReLU networks are known to yield overconfident predictions for data points which lie far away from the training (source) data [22]. In other words, the MAP estimates of a neural network has no notion of uncertainty over the learned weights. Thus, blindly trusting the source model predictions for $\mathbf{x}^{[T]} \in \mathcal{D}^{[T]}$ while performing information maximization [33] or entropy minimization [19] can potentially lead to misalignment of clusters between the source and target.

In this work we propose to incorporate the uncertainty of the neural network's weights into the predictions. This mandates a Bayesian treatment of the networks parameters (θ), which gives a posterior distribution over the model parameters by conditioning onto observed data (\mathcal{D}), *i.e.*, $p(\theta | \mathcal{D}) = \frac{p(\theta)p(\mathcal{D}|\theta)}{p(\mathcal{D})} \propto p(\theta)p(\mathcal{D}|\theta)$. The prediction of the network h_θ for an observation \mathbf{x} is given by the predictive posterior distribution, *i.e.*,

$$p(y_k | \mathbf{x}, \mathcal{D}) = \int_{\theta} \phi_k(h_\theta(\mathbf{x})) p(\theta | \mathcal{D}) d\theta. \quad (4)$$

Note that the posterior $p(\theta | \mathcal{D})$ in Eq. (4) does not have an analytical solution in general and need to be approximated. For this, we employ a local approximation to the posterior using a Laplace approximation (LA, [59]). The LA locally

approximates the true posterior using a multivariate Gaussian distribution centred at a local maximum and with covariance matrix given by the inverse of the Hessian \mathbf{H} of the negative log-posterior, *i.e.*, $p(\theta | \mathcal{D}) \approx \mathcal{N}(\theta | \theta_{\text{MAP}}, \mathbf{H}^{-1})$ with $\mathbf{H} := -\nabla_{\theta}^2 \log p(\theta | \mathcal{D}) |_{\theta_{\text{MAP}}}$. Details can be found in the appendix. Note that the LA is a principled and simple, yet effective, approach to approximate posterior inference stemming from a second-order Taylor expansion of the true posterior around θ_{MAP} . Next we will discuss LA in the context of SFDA.

Bayesian Source Model Generation In the source training stage (see Fig. 3a), by optimizing Eq. (1), we obtain a MAP estimate of the weights for our source model, comprising β_{MAP} and θ_{MAP} for g and h , respectively. Since f is often modelled by a very deep neural network (*e.g.*, ResNet-50), computing the Hessian can be computationally infeasible owing to the large number of parameters. So we make another simplification by applying a Bayesian treatment only to hypothesis function h , known as the *last-layer* Laplace approximation [25]. This gives us a probabilistic source hypothesis with posterior distribution $p(\theta | \mathcal{D}^{[S]})$ for the parameters. The feature extractor g remains deterministic. Formally, let $\mathbf{z} = g_{\beta^{[S]}}(\mathbf{x})$ be the latent feature representation from the feature extractor. Following Eq. (4), the predictive posterior distribution is given as:

$$p(y_k | \mathbf{z}, \mathcal{D}^{[S]}) \approx \int_{\theta} \phi_k(h_{\theta}(\mathbf{z})) \mathcal{N}(\theta | \theta_{\text{MAP}}, \mathbf{H}^{-1}) d\theta. \quad (5)$$

While the *last-layer* LA greatly simplifies the computational overhead for large networks, the Hessian can still be difficult to compute in the case the number of classes is large. To simplify computations, we assume that \mathbf{H} can be Kronecker-factored $\mathbf{H} := \mathbf{V} \otimes \mathbf{U}$ and the resulting approximation is referred to as Kronecker-factored Laplace approximation (KFLA, [52]). Such probabilistic treatment allows us to quantify uncertainty in the predictions for data points from the target with little computational overhead. Also, the LA can be readily computed using a single forward pass of the source data through the network. Next, we describe how to use the uncertainty estimates during target adaptation.

Uncertainty-guided Information Maximization Upon completion of the source model generation stage, we exploit the probabilistic source hypothesis to guide the information maximization in the target adaptation stage. SHOT puts equal confidence on all the target predictions and do not make any distinction for the target feature that lies outside of the source manifold. We emphasize that in case of strong domain-shift naïvely maximizing the IM loss could lead to cluster misalignment. For that reason, we propose to weigh the entropy minimization objective (Eq. (2)) with a weight which is proportional to the certainty in the target predictions (see Fig. 3b). To get the per-sample weight for a $\mathbf{x}^{[T]}$ we need to compute the predictive posterior distribution, as outlined in Eq. (5). However, exactly solving the integration is intractable in many cases and we, therefore, resort to Monte Carlo (MC) integration. Let $\mathbf{z}^{[T]} = g_{\beta^{[T]}}(\mathbf{x}^{[T]})$, the approximate

predictive posterior distributions is:

$$p(y_k | \mathbf{z}^{[T]}, \mathcal{D}^{[S]}) \approx \frac{1}{M} \sum_{j=1}^M \phi_k \left(h_{\theta_j}(\mathbf{z}^{[T]}) \right), \quad (6)$$

where $\theta_j \sim \mathcal{N}(\theta_j | \theta_{\text{MAP}}, \mathbf{H}^{-1})$ and M denotes the number of MC steps. To encourage low entropy predictions we additionally scale the outputs of the hypothesis by $1/\tau$, where $0 < \tau \leq 1$. The final weight of each observation $\mathbf{x}_i^{[T]}$ is then computed as $w_i = \exp(-H)$ where H denotes the entropy of the predictive mean. The *uncertainty-guided entropy loss* is then given as:

$$\mathcal{L}_{\text{ent}}^{\text{ug}} = -\mathbb{E}_{p(\mathbf{x}^{[T]})} \sum_{k=1}^K w \sigma_k(f'(\mathbf{x}^{[T]})) \log \sigma_k(f'(\mathbf{x}^{[T]})). \quad (7)$$

The final training objective is then given as: $\mathcal{L}_{\text{U-SFAN}} = (1 - \gamma)\mathcal{L}_{\text{ent}}^{\text{ug}} + \gamma \mathcal{L}_{\text{div}}$. Pseudocode for our U-SFAN can be found in the appendix.

How does this differ from conventional uncertainty estimation? The importance and advantages of adopting a Laplace approximation (LA) over Monte Carlo (MC) dropout to estimate uncertainty in SFDA can be summarized as follows: (i) LA does not require specialized network architecture (*e.g.*, dropout layers), loss function, or re-training (as in MC dropout) to estimate predictive uncertainties. This greatly decouples the source training from target adaptation, which is essential to be applicable in SFDA; (ii) To have well-calibrated uncertainties, MC dropout requires a grid search over the dropout probabilities [13], a prohibitive operation in deep neural networks, especially as the future target data is not available at source training. LA is a more principled approach that does not require a grid search, making it better suited for SFDA. (iii) LA is computationally lightweight since it requires just a single forward pass of the source data through the network after the source training to estimate the posterior over the parameters of the sub-network. (iv) LA does not impact the training time during target adaptation because, unlike MC dropout, only a single forward pass is needed to quantify the predictive uncertainties. Because LA employs a Gaussian approximation to the posterior, MC integration is cheap and efficient to compute. (v) As used in our work, LA estimates the full posterior over the weights and biases, while MC dropout can only account for the uncertainties over the weights [12] and is known to be a poor approximation to the posterior [46,47,10]. (vi) LA preserves the decision boundary induced by the MAP estimate, which is not the case for MC dropout [25]. In summary, our contribution goes beyond the uncertainty re-weighting scheme [35,32] commonly used in UDA, while carrying many advantages over existing works.

4 Experiments

We conduct experiments on four standard DA benchmarks: OFFICE31 [54], OFFICE-HOME [63], VISDA-C [50], and the large-scale DOMAINNET [49] (**0.6**

million images). The details of the benchmarks are summarized in appendix. For the experiments in the open-set DA setting we follow the split of [33] for shared and target-private classes.

Evaluation protocol We report the classification accuracy for every possible pair of *source* \mapsto *target* directions, except for the VISDA-C where we are only concerned with the transfer from *synthetic* \mapsto *real* domain. For the open-set experiments, following the evaluation protocol in [33], we report the OS accuracy which includes the per-class accuracy of the known and the unknown class and is computed as $OS = \frac{1}{K+1} \sum_{k=1}^{K+1} acc_k$, where $k = \{1, 2, \dots, K\}$ denote the shared classes and $(K + 1)^{th}$ is the target-private or OOD classes. This metric is preferred over the known class accuracy, $OS^* = \frac{1}{K} \sum_{k=1}^K acc_k$, as it does not take into account the OOD classes.

Implementation details We adopted the network architectures used in the SFDA literature, which are ResNet-50 or ResNet-101 [21]. Following [33], we added a bottleneck layer containing 256 neurons which is then followed by a batch normalization layer. The network finally ends with a weight normalized linear classifier that is kept frozen during the target adaptation. Details about hyperparameters can be found in the appendix. For computing the KFLA we use the PyTorch package of Dangel *et al.* [9]. Our code is available at <https://github.com/roysubhankar/uncertainty-sfda>

4.1 Ablation Studies

As discussed in Sec. 3.2, conventional SFDA methods that rely on optimizing the IM loss on the unlabelled target data (*e.g.*, SHOT) are prone to misalignment of the target data with the source hypothesis under strong domain shift. To visually demonstrate this phenomenon, we design an experiment of a 3-way classification task on toy data (see Fig. 4). Given a set of source data points, belonging to three classes, we simulate two kinds of domain-shift: *mild* shift (Fig. 4a) and *strong* shift (Fig. 4b). In the case of mild shift, the target data points stay very close to the source manifold, and the conventional approach (only using the MAP estimate) can classify a majority of target data points without the need of adaptation. Whereas, in the case of strong shift, the target data points for the blue class, in particular, shift drastically away from the source points. The source model based on the MAP estimate misclassifies most of the target data points with high confidence. On the other hand, our uncertainty-guided source model remains certain only for those target points which lie within the source support and assigns low certainty otherwise (proportional to the strength of colours depicting the decision surface in Fig. 4), robustifying the adaptation on the target data in case of domain shift.

Given such a set-up, we optimize the IM loss (*i.e.*, SHOT-IM) for both the conventional and the uncertainty-guided source models. In the case of mild shift, both can reliably partition the target data points under the right decision surfaces (see Fig. 4a (right)). This is intuitive because the decision boundary of the

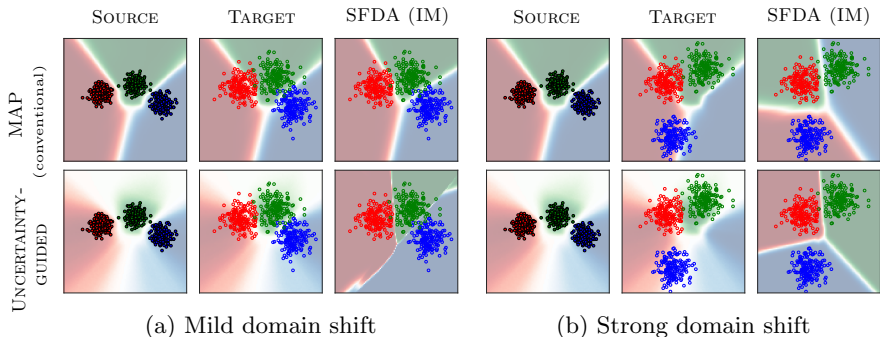


Fig. 4. Comparison of conventional IM (MAP) with our uncertainty-guided IM on target data under mild and strong domain-shift. The solid \bullet vs. hollow \circ circles represent the source and the target data, respectively. Each class is colour coded and the decision boundaries are shaded with the corresponding colours. Under strong domain-shift, IM, when used with a MAP estimate, finds a **completely flipped** decision boundary. U-SFAN finds the decision boundary by down-weighting the *far away* target data

target model already passes through the low-density regions. Hence, the optimization of the IM loss leads to correct target classification with both methods. However, when the domain shift is more substantial, the conventional approach results in *completely flipped* decision boundaries. This happens because most blue target points fall under the red decision surface, and thus, the IM loss assigns them to class ‘red’. On the contrary, our uncertainty-guided approach down-weights the blue points, and *safely* optimizes the IM loss as the model is uncertain about the class assignment for those points (Fig. 4b (right)). This protects from major changes in the decision boundaries and allows the optimization to find the correct decision boundaries for the target data. Therefore, highlighting the importance of having a notion of uncertainty in the model predictions during adaptation. We show later that this intuition also holds well for real-world data sets where our U-SFAN offers more robustness when the domain-shift in a data set becomes challenging.

To gain further insights, we visualize the entropy density plots of the source model predictions before and after adaptation with conventional (MAP estimate) and uncertainty-guided models on an image data set (CIFAR [26] as source data set and STL [7] as target). As shown in Fig. 5a, the MAP estimate has lower entropy predictions for both the correct and incorrect predictions, when compared to our uncertainty-guided model. Reduced over-confidence for our approach is expected before the adaptation phase, however, it is non-trivial that this behavior also bears in the post-adaptation phase. The reduced over-confident allows our U-SFAN to down-weight the incorrect predictions during target adaptation, resulting in improved target accuracy over SHOT-IM (77.04% for U-SFAN vs 75.69% for SHOT-IM). This effect can be noticed in Fig. 5b where U-SFAN has overall higher entropy incorrect predictions, which is desirable in SFDA.

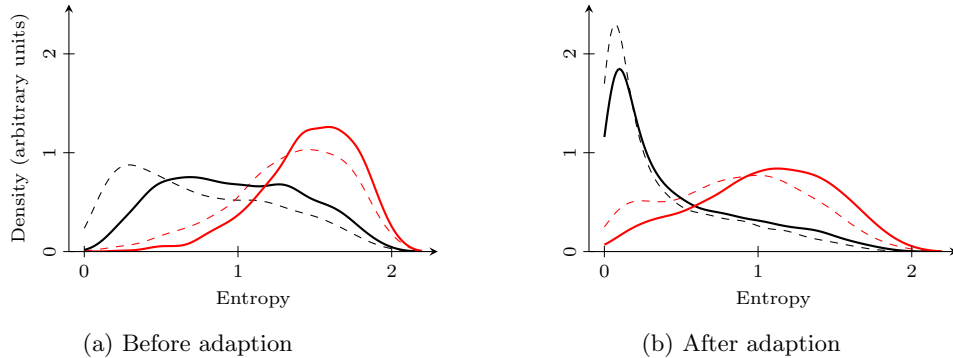


Fig. 5. Entropy density plots for CIFAR9 \rightarrow STL9 in the closed-set SFDA setting using the MAP estimate (--- correct, --- incorrect) or our approach (— correct, — incorrect). Our uncertainty-guided SFDA approach places less mass on low-entropy incorrect samples before and after adaptation

Table 1. Comparison of model performance using entropy weighting during target adaptation on the OFFICE-HOME data set. The weights computed using the LA is more beneficial than the weights computed with a MAP network

METHOD	SOURCE-ONLY	SHOT-IM [33]	SHOT-IM + ENT. WEIGHTING	U-SFAN (OURS)
Avg. Acc.	60.3	70.5	71.2	71.8

To further understand the contribution of our uncertainty-guided re-weighting, we run an ablation where the approximate posterior distribution of our method (Eq. (7)) is replaced by a weight computed from a point estimate from a MAP source model. This model is denoted as SHOT-IM + ENT. WEIGHTING in Table 1. We observe that such weighting scheme indeed improves the performance over SHOT-IM. However, it still lacks behind our proposed U-SFAN which uses weights computed from the uncertainty-guided model. This clearly shows that the improvement in performance with U-SFAN is not simply caused by the re-weighting but also due to better identification of target samples that are not well explained under the source model.

4.2 State-of-the-art Comparison

We compare our U-SFAN with UDA and SFDA methods on multiple data sets for closed-set and open-set settings. First, we compare U-SFAN with the baselines on the most common benchmark of OFFICE-HOME for both closed-set and open-set settings. As can be seen from Table 2 and Table 3 we improve the performance over majority of the baselines. Especially, we consistently improve over SHOT-IM with our method. We also combine the nearest centroid pseudo-labelling, used in SHOT [33], with U-SFAN (indicated as U-SFAN+ in Table 2 and Table 4a), and we find that it further helps improving the performance. Notably, the recently

Table 2. Comparison of the classification accuracy on the OFFICE-HOME for the closed-set setting using ResNet-50. High overall performance signifies *milder* distributional shift between domains. The improvement of U-SFAN upon SHOT is moderate, but competitive w.r.t. A²Net[67] or SHOT++ [34], which require complex training objectives

METHOD	A→C	A→P	A→R	C→A	C→P	C→R	P→A	P→C	P→R	R→A	R→C	R→P	AVG.
ResNet-50	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
DANN [14]	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
DWT [53]	50.3	72.1	77.0	59.6	69.3	70.2	58.3	48.1	77.3	69.3	53.6	82.0	65.6
CDAN [38]	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
SAFN [68]	52.0	71.7	76.3	64.2	69.9	71.9	63.7	51.4	77.1	70.9	57.1	81.5	67.3
SHOT-IM [33]	55.4	76.6	80.4	66.9	74.3	75.4	65.6	54.8	80.7	73.7	58.4	83.4	70.5
LSC [69]	57.9	78.6	81.0	66.7	77.2	77.2	65.6	56.0	82.2	72.0	57.8	83.4	71.3
U-SFAN (Ours)	58.5	78.6	81.1	66.6	75.2	77.9	66.3	57.9	80.6	73.6	61.4	84.1	71.8
A ² Net[67]	58.4	79.0	82.4	67.5	79.3	78.9	68.0	56.2	82.9	74.1	60.5	85.0	72.8
SHOT++ [34]	57.9	79.7	82.5	68.5	79.6	79.3	68.5	57.0	83.0	73.7	60.7	84.9	73.0
SHOT [33]	57.1	78.1	81.5	68.0	78.2	78.1	67.4	54.9	82.2	73.3	58.8	84.3	71.8
U-SFAN+ (Ours)	57.8	77.8	81.6	67.9	77.3	79.2	67.2	54.7	81.2	73.3	60.3	83.9	71.9

Table 3. Comparison of the OS classification accuracy on the OFFICE-HOME for the open-set setting using ResNet-50. U-SFAN improves over SHOT without the need for nearest-centroid pseudo-labelling in the case of open-set SFDA

METHOD	A→C	A→P	A→R	C→A	C→P	C→R	P→A	P→C	P→R	R→A	R→C	R→P	AVG.
ResNet-50	53.4	52.7	51.9	69.3	61.8	74.1	61.4	64.0	70.0	78.7	71.0	74.9	65.3
ATI-λ[48]	55.2	52.6	53.5	69.1	63.5	74.1	61.7	64.5	70.7	79.2	72.9	75.8	66.1
OpenMax [2]	56.5	52.9	53.7	69.1	64.8	74.5	64.1	64.0	71.2	80.3	73.0	76.9	66.7
STA [36]	58.1	53.1	54.4	71.6	69.3	81.9	63.4	65.2	74.9	85.0	75.8	80.8	69.5
SHOT-IM [33]	62.5	77.8	83.9	60.9	73.4	79.4	64.7	58.7	83.1	69.1	62.0	82.1	71.5
SHOT [33]	64.5	80.4	84.7	63.1	75.4	81.2	65.3	59.3	83.3	69.6	64.6	82.3	72.8
U-SFAN (Ours)	62.9	77.9	84.0	67.9	74.6	79.6	68.8	61.3	83.3	76.0	63.9	82.3	73.5

Table 4. (a) Comparison of the classification accuracy on the VISDA-C for the closed-set DA, pertaining to the *Synthetic* → *Real* direction, using ResNet-101. † indicates the numbers of [33] that are obtained using the official code from the authors. Note that several SFDA methods perform equally well for VISDA-C, hinting at saturating performance. (b) Comparison of the average accuracy on the DOMAINNET for the closed-set SFDA using ResNet-50. The SOURCE column indicates the domain where the source model has been trained. The data set being challenging (exhibiting *strong* domain-shift), the improvement with our U-SFAN over [33] is substantial

(a) VISDA-C		(b) DOMAINNET		
METHOD	ACC.	SOURCE	SHOT-IM [33]	U-SFAN
ResNet-101	52.4	CLIPART	25.04	30.88
CDAN+BSP [6]	75.9	INFOGRAPH	21.58	26.44
SAFN [68]	76.1	PAINTING	23.89	29.91
SHOT-IM† [33]	80.3	QUICKDRAW	10.76	10.44
U-SFAN (Ours)	81.2	REAL	21.74	29.32
3C-GAN [31]	81.6	SKETCH	28.87	29.99
A ² Net[67]	84.3	AVG.	21.98	26.13
SHOT† [33]	82.4			
U-SFAN+ (Ours)	82.7			

proposed A²Net [67] (which just addresses closed-set SFDA) outperforms our U-SFAN in a couple of data sets, but uses a combination of several loss functions. Interplay of multiple losses can be hard to tune in practice. On the other hand, our method is simpler, more versatile and works for both the SFDA settings. Due to lack of space, we report the numbers for OFFICE31 in the appendix. Given the performance of the SFDA baseline methods in OFFICE-HOME and VISDA-C are relatively high and closer to each other, the domain shift can be considered milder with respect to more challenging data set like DOMAIN-NET.

When we compare U-SFAN with SHOT-IM on the challenging SFDA benchmark DOMAIN-NET the advantage of our U-SFAN over SHOT-IM becomes imminent (*cf.* Table 4b), which is in line with the ablation study in Sec. 4.1. Different from the previous data sets, the difficulty in mitigating domain-shift for DOMAIN-NET is evident from the low overall performance of both SHOT-IM and U-SFAN. This data set can be seen as a real-world example of strong domain-shift. The improvement in the performance of U-SFAN over SHOT-IM for DOMAIN-NET demonstrates that incorporating the uncertainty in the model’s predictions plays a crucial role in SFDA. The conventional approach may overfit to noisy model predictions, leading to poor performance. Whereas, U-SFAN can capture the uncertainty in predictions and down-weight the impact of noisy predictions.

5 Discussion and Conclusion

In this work, we demonstrated the need for uncertainty quantification in SFDA and proposed U-SFAN for leveraging it during target adaptation. Our uncertainty-guided approach employs a Laplace approximation to the posterior, does not require specialized source training, and allows for efficient computation of predictive uncertainties. Our experiments showed that down-weighting *distant* target data points in our novel uncertainty-weighted IM loss alleviates the misalignment of target data with the source hypothesis. We ran experiments on closed and open-set SFDA settings and show that U-SFAN consistently improves upon the existing methods. Moreover, U-SFAN has shown to be robust under mild distribution shifts and shows promising results under severe distribution shifts.

While we mainly focused on the popular IM-based SFDA methods, our proposed uncertainty-guided adaptation is also applicable to other SFDA frameworks, *e.g.*, neighbourhood clustering [69] or extensions to the multi-source SFDA problem. Moreover, the principles we build upon are general, interpretable, and have strong backing in classical statistics. We believe that uncertainty-guided SFDA will become a backbone tool for future methods in DA that generalize over different problem domains, are less sensitive to the training setup, and will provide good results without extensive *ad hoc* tuning to each problem.

Acknowledgements. We acknowledge funding from EU H2020 projects SPRING (No. 871245) and AI4Media (No. 951911); the EUREGIO project OLIVER; Academy of Finland (No. 339730, 308640), and the Finnish Center for Artificial Intelligence (FCAI). We acknowledge the computational resources by the Aalto Science-IT project and CSC – IT Center for Science, Finland.

References

1. Ahmed, W., Morerio, P., Murino, V.: Adaptive pseudo-label refinement by negative ensemble learning for source-free unsupervised domain adaptation. In: Proceedings of the Winter Conference on Applications of Computer Vision (WACV) (2022)
2. Bendale, A., Boulton, T.E.: Towards open set deep networks. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR). pp. 1563–1572 (2016)
3. Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: A review for statisticians. *Journal of the American statistical Association* **112**(518), 859–877 (2017)
4. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 343–351 (2016)
5. Chen, W., Lin, L., Yang, S., Xie, D., Pu, S., Zhuang, Y., Ren, W.: Self-supervised noisy label learning for source-free unsupervised domain adaptation. arXiv preprint arXiv:2102.11614 (2021)
6. Chen, X., Wang, S., Long, M., Wang, J.: Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 1081–1090 (2019)
7. Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS). pp. 215–223. *Journal of Machine Learning Research Workshop and Conference Proceedings* (2011)
8. Csurka, G.: A comprehensive survey on domain adaptation for visual applications. *Domain Adaptation in Computer Vision Applications* pp. 1–35 (2017)
9. Dangel, F., Kunstner, F., Hennig, P.: Backpack: Packing more into backprop. In: Proceedings of the International Conference on Learning Representations (ICLR) (2020)
10. Foong, A., Burt, D., Li, Y., Turner, R.: On the expressiveness of approximate inference in bayesian neural networks. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 15897–15908 (2020)
11. Fu, B., Cao, Z., Long, M., Wang, J.: Learning to detect open classes for universal domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 567–583 (2020)
12. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 1050–1059 (2016)
13. Gal, Y., Hron, J., Kendall, A.: Concrete dropout. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 3581–3590 (2017)
14. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of Machine Learning Research* **17**(59), 1–35 (2016)
15. Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B.: *Bayesian Data Analysis*. Chapman and Hall/CRC, Boca Raton, FL, third edn. (2013)
16. Ghifary, M., Kleijn, W.B., Zhang, M., Balduzzi, D., Li, W.: Deep reconstruction-classification networks for unsupervised domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 597–613 (2016)

17. Gomes, R., Krause, A., Perona, P.: Discriminative clustering by regularized information maximization. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2010)
18. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2014)
19. Grandvalet, Y., Bengio, Y., et al.: Semi-supervised learning by entropy minimization. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 281–296 (2005)
20. Han, L., Zou, Y., Gao, R., Wang, L., Metaxas, D.: Unsupervised domain adaptation via calibrating uncertainties. In: *CVPR Workshops* (2019)
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
22. Hein, M., Andriushchenko, M., Bitterwolf, J.: Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 41–50 (2019)
23. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 1989–1998 (2018)
24. Kendall, A., Gal, Y.: What uncertainties do we need in Bayesian deep learning for computer vision? In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 5574–5584 (2017)
25. Kristiadi, A., Hein, M., Hennig, P.: Being Bayesian, even just a bit, fixes overconfidence in relu networks. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 5436–5446 (2020)
26. Krizhevsky, A.: *Learning Multiple Layers of Features from Tiny Images*. Master’s thesis, University of Tronto, Toronto, Canada (2009)
27. Kundu, J.N., Venkat, N., Babu, R.V., et al.: Universal source-free domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4544–4553 (2020)
28. Kurmi, V.K., Kumar, S., Nambodiri, V.P.: Attending to discriminative certainty for domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 491–500 (2019)
29. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
30. Lao, Q., Jiang, X., Havaei, M.: Hypothesis disparity regularized mutual information maximization. In: *Proceedings of the AAAI Conference on Artificial Intelligence* (2021)
31. Li, R., Jiao, Q., Cao, W., Wong, H.S., Wu, S.: Model adaptation: Unsupervised domain adaptation without source data. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9641–9650 (2020)
32. Liang, J., He, R., Sun, Z., Tan, T.: Exploring uncertainty in pseudo-label guided unsupervised domain adaptation. *Pattern Recognition* **96**, 106996 (2019)
33. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 6028–6039 (2020)

34. Liang, J., Hu, D., Wang, Y., He, R., Feng, J.: Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
35. Liang, J., Wang, Y., Hu, D., He, R., Feng, J.: A balanced and uncertainty-aware approach for partial domain adaptation. In: *Proceedings of European Conference on Computer Vision (ECCV)*. pp. 123–140. Springer (2020)
36. Liu, H., Cao, Z., Long, M., Wang, J., Yang, Q.: Separate to adapt: Open set domain adaptation via progressive separation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2927–2936 (2019)
37. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: *Proceedings of the International conference on machine learning (ICML)*. pp. 97–105 (2015)
38. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 1647–1657 (2018)
39. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 2208–2217 (2017)
40. MacKay, D.J.: A practical Bayesian framework for backpropagation networks. *Neural Computation* 4(3), 448–472 (1992)
41. MacKay, D.J.: *Information Theory, Inference and Learning Algorithms*. Cambridge University Press (2003)
42. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for Bayesian uncertainty in deep learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 13132–13143 (2019)
43. Morerio, P., Cavazza, J., Murino, V.: Minimal-entropy correlation alignment for unsupervised deep domain adaptation. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2018)
44. Müller, R., Kornblith, S., Hinton, G.: When does label smoothing help? In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 4694–4703 (2019)
45. Neal, R.M.: *Bayesian Learning for Neural Networks*. Springer Science & Business Media (2012)
46. Osband, I.: Risk versus uncertainty in deep learning: Bayes, bootstrap and the dangers of dropout. In: *NeurIPS workshop on Bayesian deep learning* (2016)
47. Osband, I., Aslanides, J., Cassirer, A.: Randomized prior functions for deep reinforcement learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 8617–8629 (2018)
48. Panareda Busto, P., Gall, J.: Open set domain adaptation. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 754–763 (2017)
49. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 1406–1415 (2019)
50. Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., Saenko, K.: Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924* (2017)
51. Ringwald, T., Stiefelhagen, R.: Unsupervised domain adaptation by uncertain feature alignment. *The British Machine Vision Conference (BMVC)* (2020)
52. Ritter, H., Botev, A., Barber, D.: A scalable laplace approximation for neural networks. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2018)

53. Roy, S., Siarohin, A., Sangineto, E., Bulo, S.R., Sebe, N., Ricci, E.: Unsupervised domain adaptation using feature-whitening and consensus loss. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9471–9480 (2019)
54. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 213–226 (2010)
55. Saito, K., Saenko, K.: Ovanet: One-vs-all network for universal domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 9000–9009 (2021)
56. Saito, K., Yamamoto, S., Ushiku, Y., Harada, T.: Open set domain adaptation by backpropagation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 153–168 (2018)
57. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 443–450 (2016)
58. Tian, J., Zhang, J., Li, W., Xu, D.: VDM-DA: Virtual domain modeling for source data-free domain adaptation. *IEEE Transactions on Circuits and Systems for Video Technology* (2021)
59. Tierney, L., Kadane, J.B.: Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* **81**(393), 82–86 (1986)
60. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1521–1528 (2011)
61. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Adversarial discriminative domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7167–7176 (2017)
62. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474* (2014)
63. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5018–5027 (2017)
64. Wang, M., Deng, W.: Deep visual domain adaptation: A survey. *Neurocomputing* **312**, 135–153 (2018)
65. Wen, J., Zheng, N., Yuan, J., Gong, Z., Chen, C.: Bayesian uncertainty matching for unsupervised domain adaptation. In: International Joint Conference on Artificial Intelligence (IJCAI) (2019)
66. Wilson, A.G.: The case for Bayesian deep learning. Tech. rep., New York University (2019)
67. Xia, H., Zhao, H., Ding, Z.: Adaptive adversarial network for source-free domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 9010–9019 (2021)
68. Xu, R., Li, G., Yang, J., Lin, L.: Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 1426–1435 (2019)
69. Yang, S., Wang, Y., van de Weijer, J., Herranz, L., Jui, S.: Generalized source-free domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 8978–8987 (2021)

70. Yeh, H.W., Yang, B., Yuen, P.C., Harada, T.: Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 474–483 (2021)
71. Zhang, H., Li, A., Han, X., Chen, Z., Zhang, Y., Guo, Y.: Improving open set domain adaptation using image-to-image translation. In: IEEE International Conference on Multimedia and Expo (ICME). pp. 1258–1263. IEEE (2019)
72. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision (IJCV)* (2021)