

Balancing Discriminability and Transferability for Source-Free Domain Adaptation

Jogendra Nath Kundu^{*1} Akshay Kulkarni^{*1} Suvaansh Bhambri^{*1} Deepesh Mehta¹ Shreyas Kulkarni¹
Varun Jampani² R. Venkatesh Babu¹

Abstract

Conventional domain adaptation (DA) techniques aim to improve domain transferability by learning domain-invariant representations; while concurrently preserving the task-discriminability knowledge gathered from the labeled source data. However, the requirement of simultaneous access to labeled source and unlabeled target renders them unsuitable for the challenging source-free DA setting. The trivial solution of realizing an effective original to generic domain mapping improves transferability but degrades task discriminability. Upon analyzing the hurdles from both theoretical and empirical standpoints, we derive novel insights to show that a mixup between original and corresponding translated generic samples enhances the discriminability-transferability trade-off while duly respecting the privacy-oriented source-free setting. A simple but effective realization³ of the proposed insights on top of the existing source-free DA approaches yields state-of-the-art performance with faster convergence. Beyond single-source, we also outperform multi-source prior-arts across both classification and semantic segmentation benchmarks.

1. Introduction

Generally, in machine learning, it is assumed that the test samples are drawn from the same distribution as the training samples. However, in practice, a model often encounters a shift in the input distribution (*i.e.* domain-shift), resulting in poor deployment performance. Unsupervised domain adaptation (DA) techniques seek to address this problem by transferring the task-discriminative knowledge from a

^{*}Equal contribution ¹Indian Institute of Science ²Google Research. Correspondence to: J. N. Kundu <jogendrak@iisc.ac.in>.

³<https://sites.google.com/view/mixup-sfda>

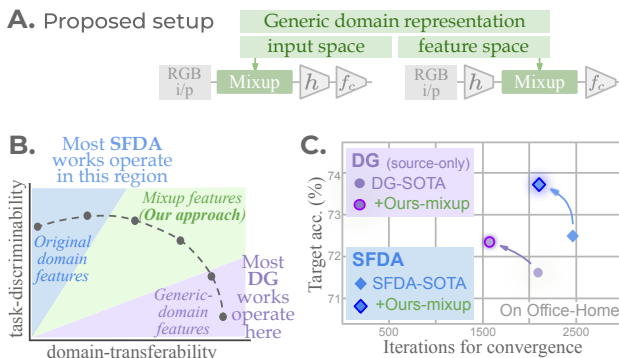


Figure 1. **A.** We propose an instance-level mixup between original and generic-domain samples. **B.** Compared to previous SFDA (blue) and DG works (purple), we operate in an intermediate region (green) with an improved discriminability-transferability trade-off. **C.** As a result, our approach complements prior DG and SFDA SOTA works and produces better and faster convergence.

labeled source domain to an unlabeled target domain.

From a theoretical viewpoint (Ben-David et al., 2006), the target error is upper bounded by three terms; **a)** the source error, **b)** the worst-case source-target domain discrepancy, and **c)** the error of a joint ideal hypothesis on source and target. With this line of thought, the widely adopted adversarial DA works (Ganin et al., 2016) aim to simultaneously minimize the first two terms using a feature extractor, task classifier, and domain classifier. The feature extractor is trained with two objectives, task classification on source and domain classifier fooling for source and target. While the first objective aims to preserve task-discriminability, the second objective aims to improve domain-transferability by encouraging domain-invariant representations. However, prior works (Chen et al., 2019a) show that the two objectives are at odds with each other *i.e.* improving transferability leads to degraded discriminability and vice-versa. Certain works (Yang et al., 2020a) explicitly address this by proposing ways to strike a suitable balance for better adaptation. But these works require joint access to both (labeled) source and (unlabeled) target data during training, making them unsuitable for source-free DA (Kundu et al., 2020b), where source and target data are not concurrently accessible. These DA scenarios with restricted data-sharing are privacy-oriented and thus, hold immense practical value.

Recognizing this, we seek answers to following questions,

1. What are the key hurdles in extending the available techniques for the challenging source-free setting?
2. What is the key design aspect to overcome these hurdles, thereby enabling effective source-free DA?

In prior non-source-free works, the transferability is assessed as the inability of a domain classifier to segregate samples from the source and target domain. Clearly, such a realization is incompatible with the source-free constraint. Though task-discriminability can be gathered on the source-side (via supervised training on labeled source), it cannot be preserved while adapting to the unlabeled target in the absence of labeled source. In other words, the discriminability-transferability trade-off becomes a more severe problem that remains unaddressed even in available source-free techniques. As shown in Fig. 1B, source-free DA works aim to preserve the task-discriminability while domain generalization (DG) works (Li et al., 2017) aim to achieve optimal transferability to generalize to unknown targets. Conversely, we believe an explicit effort to improve the trade-off (green area in Fig. 1B) could yield considerable improvements in source-free adaptation performance.

For the second question, we consider the idea of domain generic representations. A naive approach would be to develop *source-to-generic* and *target-to-generic* mappings which ideally leads to optimal transferability. In theory, the aim is to attain a generic representation that perfectly satisfies the following two criteria. First, the representation should be completely domain invariant. Second, the representation should not lose any task discriminative information. By definition, the discriminability-transferability trade-off would be well addressed. However, the key question remains, *is such a representation realizable in practice?*

To analyze this, we consider a synthetic to real-world road scene segmentation DA problem. The image samples for each domain can be considered an entangled mapping of two latent factors, *i.e.*, **a)** Shape (edge-related features) and **b)** Texture (other than edge-related features). Such a disentanglement is easily achievable by segregating a simple edge representation from the original image through traditional contour estimation techniques (see Fig. 2). Now, we analyze whether the edge can be considered as a generic representation. We empirically observe (plot on the right of Fig. 2) that though edge satisfies the first criteria of domain invariance (higher transferability), it fails on the second criteria with degraded task discriminability (lower adaptation accuracy). Thus, it cannot be considered an ideal generic domain. This is because, while the texture varies across domains, it also holds crucial task discriminative features, *e.g.*, differentiating between road, sidewalk, and hedges requires color/texture information. While we would ideally like to disentangle the purely task discriminative factors from the

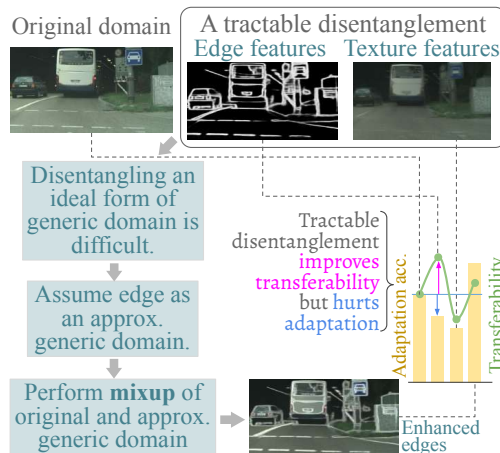


Figure 2. While an ideal generic domain is intractable, edges are a tractable disentanglement with higher transferability (pink arrow). The loss of task-discriminative information (blue arrow) is rectified by mixup between original and approximate generic domain.

domain-related ones, the entanglement is highly non-linear. Implying it is challenging to realize ideal source-to-generic and target-to-generic mappings in practice.

Now, the key question that arises is, *what could be a realizable approximation of the aforementioned ideal solution?* Following from the previous example, let us consider the edge representation as an approximate generic domain. Here, edge exhibits high domain-transferability with low task-discriminability (see Fig. 2). In comparison, the input representations of the original domains exhibit lower domain-transferability but higher task-discriminability. Aiming to attain an improved discriminability-transferability trade-off, we seek to develop an intermediate domain with a higher transferability than the original data domains and also a higher discriminability than the approximate generic domain. Following this overarching idea, we propose to realize the intermediate domain as an instance-level mixup of samples from the original and the approximate generic domains. This simple addition to existing source-free DA methods improves the implicit domain alignment, which in turn leads to better performance. The primary contributions of this work are:

- We are the first to analyze the source-free DA problem from the perspective of discriminability-transferability trade-off. As a remedy to the low task discriminability of realizable generic domain, we operate on an intermediate mixup domain between the original and the realizable generic domain. We theoretically show that the mixup domains achieve a tighter bound on target error leading to improved adaptation performance.
- Based on our insights, we propose novel ways to realize approx. generic domains for mixup. We find that simple input-space edge representation better suits dense

segmentation while feature-space mean-subdomain representation better suits non-dense classification.

- This simple modification, on top of existing DA approaches, yields state-of-the-art performance across source-free benchmarks for single-source and multi-source DA on both classification and segmentation.

2. Related Work

Domain translation for DA. Adversarial alignment (Long et al., 2018; Rangwani et al., 2022) methods translate input domains to a domain-invariant representation to improve feature-space transferability. Input space adaptation methods (Murez et al., 2018; Russo et al., 2018) explicitly translate *source-to-target* or *target-to-source* to improve transferability. Conversely, we propose separate source-to-generic and target-to-generic translations, realizable in both input and feature-space, to enable source-free DA.

Transferability-discriminability trade-off in DA. Chen et al. (2019a) use spectral analysis and find that transferability resides in few top eigenvectors, while discriminability is spread across eigenvectors, which creates the trade-off. Chen et al. (2020) study this problem for object detection DA and propose hierarchical calibration of transferability. Yang et al. (2020a) propose autoencoder-based adversarial adaptation to avoid loss of discriminability.

Source-free DA. Liang et al. (2020; 2021) use pseudo-labeling and information maximization to match target features with a fixed source classifier while Morerio et al. (2020) generate target-style samples from a GAN. Li et al. (2020); Yang et al. (2021a) focus on clustering-based regularization for source-free adaptation. Kundu et al. (2020c) focuses on class-incremental SFDA. Apart from these, Liu et al. (2021); Sivaprasad et al. (2021); Kundu et al. (2021) particularly target segmentation-specific source-free DA.

Mixup. Existing works interpolate different instances from the same class (Kim et al., 2021b), different classes (Chou et al., 2020), or even different domains (Na et al., 2021) to better separate the class or domain clusters. They use mixup between distinct images where the label changes to a convex combination of labels and mixup images look unnatural. Conversely, our *within-instance* mixup preserves the label (even for segmentation) with natural-looking mixup images.

3. Approach

Problem setup. Under closed set DA, consider a labeled source dataset $\mathcal{D}_s = \{(x_s, y_s) : x_s \in \mathcal{X}, y_s \in \mathcal{C}\}$ where \mathcal{X} is the input space and \mathcal{C} denotes the label set. x_s is drawn from the marginal distribution p_s . Also consider an unlabeled target dataset $\mathcal{D}_t = \{x_t : x_t \in \mathcal{X}\}$ where x_t is drawn from the marginal distribution p_t . The task is to assign a label for each target image x_t from the label set \mathcal{C} . Following Ganin et al. (2016); Long et al. (2015), we use a backbone




Figure 3. **A.** $\overrightarrow{p_1(x)}$ and $\overrightarrow{p_2(x)}$ are some particular directions of variance in the affine space of marginal probability distributions (Krueger et al., 2021). Mixup of original domains with realizable generic domains yields mixup-domains with lower $d_{\mathcal{H}}$ (in green). **B.** Prior works lower bound $\gamma_D + \gamma_T$ by τ while our proposed mixup enables a better trade-off with a higher lower-bound τ_m .

feature extractor $h : \mathcal{X} \rightarrow \mathcal{Z}$, where \mathcal{Z} is an intermediate representation space, and a task classifier $f_c : \mathcal{Z} \rightarrow \mathcal{C}$. We operate under the source-free constraint (Li et al., 2020) in a vendor-client paradigm (Kundu et al., 2020a). The vendor has the source dataset and can share a source-trained model with the client, without sharing the source data. The client adapts to the unlabeled target using the vendor-side model.

Theoretical background. The expected source risk of the classifier f_c with backbone h , with optimal labeling function $f_S : \mathcal{X} \rightarrow \mathcal{C}$, is $\epsilon_s(h) = \mathbb{E}_{x \sim p_s} [\mathbb{1}(f_c \circ h(x) \neq f_S(x))]$, where $\mathbb{1}(\cdot)$ is the indicator function. Similarly, $\epsilon_t(h)$ is the expected target risk with optimal labeling function $f_T : \mathcal{X} \rightarrow \mathcal{C}$. Ben-David et al. (2006) present a theoretical upper bound on the expected target risk $\epsilon_t(h)$. For any backbone hypothesis $h \in \mathcal{H}$ with \mathcal{H} as the hypothesis space and a domain classifier $f_d : \mathcal{Z} \rightarrow \{0, 1\}$ (0 for source, 1 for target),

$$\begin{aligned} \epsilon_t(h) &\leq \epsilon_s(h) + d_{\mathcal{H}}(p_s, p_t) + \kappa(p_s, p_t), \text{ where} \\ d_{\mathcal{H}}(p_s, p_t) &= \sup_{h' \in \mathcal{H}} \left| \mathbb{E}_{p_s} [\mathbb{1}(f_d \circ h'(x) = 1)] - \mathbb{E}_{p_t} [\mathbb{1}(f_d \circ h'(x) = 1)] \right| \\ \text{and } \kappa(p_s, p_t) &= \min_{h' \in \mathcal{H}} \epsilon_s(h') + \epsilon_t(h') \end{aligned} \quad (1)$$

Here, $d_{\mathcal{H}}(p_s, p_t)$ denotes the \mathcal{H} -divergence that indicates the distribution shift between the source and target domains. $\kappa(p_s, p_t)$ represents the joint optimal error *i.e.* the error of an ideal hypothesis on both source and target domains.

Discriminability and transferability. *Discriminability* refers to the ease of separating different categories by a supervised classifier trained on the features (Chen et al., 2019a). *Transferability* refers to the invariance of feature representations across domains. Exclusively improving the transferability leads to a drop in discriminability and vice versa (Chen et al., 2020). This is because the domain-specific information, that has to be removed to improve transferability, may contain entangled task-specific information important for discriminability. Based on the definitions,

Chen et al. (2019a) observe that the \mathcal{H} -divergence varies inversely with the transferability while the joint optimal error $\kappa(p_s, p_t)$ varies inversely with the discriminability of the backbone features. To quantify these two, we introduce a transferability metric γ_T and a discriminability metric γ_D :

$$\gamma_T = 1 - d_{\mathcal{H}}(p_s, p_t); \quad \gamma_D = 1 - \frac{1}{2}\kappa(p_s, p_t) \quad (2)$$

For empirical computation of γ_D and γ_T , we compute the expectation over the available domain samples. Note that $0 \leq d_{\mathcal{H}}(\cdot, \cdot) \leq 1$ while $0 \leq \kappa(\cdot, \cdot) \leq 2$ (from Eq. 1).

We are the first to study the problem of handling transferability and discriminability in source-free DA. To analyze this problem, we provide an insight towards the question, *what are the key hurdles in extending available discriminability-transferability based techniques for source-free DA?*

Insight 1. (Transferability-Discriminability in SFDA)

The vendor model needs to achieve a good tradeoff between transferability and discriminability in order to be transferable to multiple clients, each posing a diverse target domain, while reasonably preserving the task discriminability.

Remarks. In a vendor-client setting, the client desires a model with good discriminability in the target domain *i.e.* good task performance on target data. However, the client cannot significantly influence the discriminability of a vendor-side model without labeled data. Thus, the vendor needs to preserve the discriminability for the clients while improving the transferability to serve multiple future clients *i.e.* a good tradeoff benefits both vendor and client.

Domain translation without concurrent access. While existing domain translation methods are limited by the concurrent access constraint, a domain translation method can be devised to be employed separately on source and target without data sharing. As per Insight 1, vendor needs to improve transferability to serve multiple clients. So, a naive solution would be to translate source or target to a *generic-domain* with marginal distribution p_{s_g} or p_{t_g} , where features possess both high transferability and high discriminability. However, in practice, these *source-to-generic* and *target-to-generic* translations would result in a loss of discriminability as domain-specific and task-specific information is entangled in the original domains (see Fig. 2). Thus, to balance domain-specificity and task-specificity, we propose mixup (Zhang et al., 2018) between original domain samples and corresponding generic-domain translated samples. Formally, *source-to-mixup* translation in the input-space is,

$$x_{s_m} = \lambda x_{s_g} + (1 - \lambda)x_s \quad (3)$$

where x_{s_g} are the generic-domain samples (effectively drawn from the hypothetical p_{s_g}) corresponding to the original domain samples x_s . And x_{s_m} is the source-mixup sample (effectively drawn from the hypothetical p_{s_m}), while the mixup ratio λ is a scalar constant. The equivalent equation for *target-to-mixup* translation is $x_{t_m} = \lambda x_{t_g} + (1 - \lambda)x_t$.

3.1. Theoretical insights

Now, we theoretically analyze the impact of using mixup distributions instead of original distributions w.r.t. Eq. 1.

Insight 2. (Effect of mixup on $d_{\mathcal{H}}$ and κ) *We postulate that, keeping the same hypothesis space \mathcal{H} , the overall discriminability and transferability for the mixup distributions will be upper bounded by that for the original distributions,*

$$d_{\mathcal{H}}(p_{s_m}, p_{t_m}) + \kappa(p_{s_m}, p_{t_m}) \leq d_{\mathcal{H}}(p_s, p_t) + \kappa(p_s, p_t) \quad (4)$$

Remarks. Most prior works (Saito et al., 2018) attempt to minimize these two terms by optimizing on the \mathcal{H} -space. Conversely, we highlight the perspective of manipulating the input distributions. Prior works lower bound $\gamma_D + \gamma_T$ by some threshold τ (solid green line in Fig. 3B). Whereas by using the mixup distributions, the lower bound can be increased to τ_m (dotted green line in Fig. 3B). This leads to a better tradeoff between transferability and discriminability. Consequently, mixup achieves a tighter upper bound (Eq. 1) for $\epsilon_t(h)$ since the remaining term $\epsilon_s(h)$ can be minimized easily in both cases with the supervised source data.

a) Transferability for mixup vs. original distributions.

To support Insight 2, we first analyze the relation between the \mathcal{H} -divergences for the mixup and original distributions. Intuitively, the samples of mixup distributions p_{s_m} or p_{t_m} contain less domain-specific information compared to the original domain samples from p_s or p_t . Thus, the cross-domain feature transferability of p_{s_m} or p_{t_m} should improve. Now, we present a theoretical result to support our intuition.

Theorem 1. (Mixup \mathcal{H} -divergence) *Assume that original source p_s and target p_t are easily separable *i.e.* perfect accuracy for domain classifier f_d . Also assume that source-generic p_{s_g} and target-generic p_{t_g} domains are impossible to separate *i.e.* accuracy of f_d imitates that of a random classifier. For a linear domain classifier f_d :*

$$d_{\mathcal{H}}(p_{s_m}, p_{t_m}) \leq d_{\mathcal{H}}(p_s, p_t) \quad (5)$$

Remarks. We make three assumptions in this proof. First, original domains should be easily separable, giving perfect accuracy for the domain classifier f_d used in $d_{\mathcal{H}}$ (see Eq. 1). This is reasonable since realistic domain shifts are found to be easily separable for empirical domain classifiers (even for linear classifiers *i.e.* our second assumption). Finally, we assume that the source-generic and target-generic domains are impossible to separate. This is also a fair assumption since a generic-domain has high transferability by definition. Proof is provided in Appendix B.1. We illustrate the same result of lower mixup \mathcal{H} -divergence in Fig. 3A.

b) Discriminability for mixup vs. original distributions.

As discussed earlier, the discriminability varies inversely with the joint optimal error $\kappa(p_s, p_t)$. Since transferability and discriminability are at odds with each other, the lower $d_{\mathcal{H}}(p_{s_m}, p_{t_m})$ should increase $\kappa(p_{s_m}, p_{t_m})$ w.r.t. $\kappa(p_s, p_t)$.

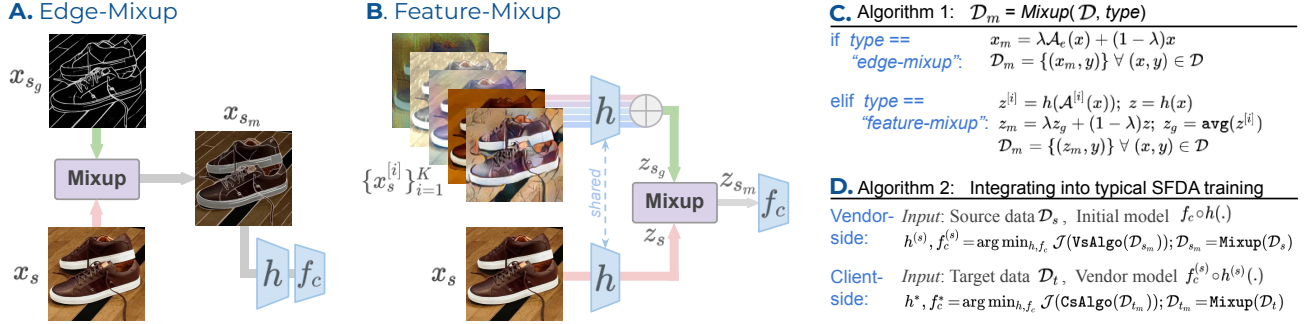


Figure 4. **A.** Edge-mixup involves mixup of an input sample x_s and the corresponding edge representation x_{s_g} (Sec. 3.2.1). **B.** Feature-mixup involves mixup of the input features $z_s = h(x_s)$ and corresponding generic domain features z_{s_g} , obtained as the feature mean of augmented sub-domain samples $\{x_s^{[i]}\}_{i=1}^K$ (Sec. 3.2.2). **C.** Algorithm for edge-mixup and feature-mixup to convert an original dataset \mathcal{D} into a mixup dataset \mathcal{D}_m . **D.** Algorithm to integrate our proposed mixup into a typical wrapper for source-free DA training.

However, mixup with corresponding generic-domain samples implies that the task discriminative features will be preserved to a good extent since the common characteristics across the original and generic-domain samples are preserved. Unlike Theorem 1, it is not possible to theoretically arrive at a result relating $\kappa(p_{s_m}, p_{t_m})$ and $\kappa(p_s, p_t)$, as the assumptions on f_c become unrealistic (see Appendix B.2). However, as long as task discriminative features are preserved (this is observed for low λ values empirically in Fig. 5A), mixup discriminability $\kappa(p_{s_m}, p_{t_m})$ cannot drop w.r.t. the original $\kappa(p_s, p_t)$ and Eq. 4 holds. Thus, mixup provides an improved transferability-discriminability tradeoff.

3.1.1. Criteria to realize generic-domain mixup

There are two aspects involved in realizing the generic-domain mixup. First, *how to obtain generic-domain samples corresponding to original samples?* The ideal generic-domain p_{s_g} or p_{t_g} has features with high domain-transferability. Thus, the criteria for choosing an original-to-generic domain mapping would be the effectiveness of the translation technique in disentangling the domain-generic characteristics. This can be measured through the transferability metric γ_T for source and target generic-domain samples. We empirically observe (Fig. 5A) that our proposed generic domains ($\lambda = 1$) exhibit higher γ_T .

Second, *how to perform the mixup?* In the proposed method, we perform mixup of generic-domain samples with corresponding original samples as a convex combination with a fixed mixup ratio λ (see Eq. 3). While we show the performance to be insensitive over a wide range of λ values (see Fig. 5C), there may be other ways to perform mixup. For instance, we may employ a learnable combination via meta-learning (Wei et al., 2021). However, we choose the simple convex combination, which validates our theoretical insights. Other options can be explored in future work.

3.2. Training algorithm and realizable generic-domains

We aim to demonstrate that our proposed mixup is complementary to existing DA methods (both source-free and

non-source-free). To this end, we simply replace the original training datasets \mathcal{D}_s and \mathcal{D}_t in existing methods with our devised mixup datasets \mathcal{D}_{s_m} and \mathcal{D}_{t_m} respectively. Note the use of only the devised datasets, different from *data augmentation* which would also use the original datasets.

Vendor-side training (Fig. 4C). Consider an existing vendor-side training algorithm $\text{VsAlgo}(\mathcal{D}_s)$ that trains the backbone h and classifier f_c on the original source dataset \mathcal{D}_s . We minimize the objective of the same training algorithm but on the source-mixup dataset \mathcal{D}_{s_m} . Formally,

$$\min_{h, f_c} \mathcal{J}(\text{VsAlgo}(\mathcal{D}_{s_m})) \quad (6)$$

where $\mathcal{J}(\cdot)$ represents the objective or loss function.

Client-side training (Fig. 4D). Consider an existing client-side training algorithm $\text{CsAlgo}(\mathcal{D}_t)$. Similar to the vendor-side training, we train h and f_c on the target-mixup dataset \mathcal{D}_{t_m} using the objective of CsAlgo . Formally,

$$\min_{h, f_c} \mathcal{J}(\text{CsAlgo}(\mathcal{D}_{t_m})) \quad (7)$$

Next, we discuss an input-space generic-domain representation to obtain the mixup datasets \mathcal{D}_{s_m} and \mathcal{D}_{t_m} . Following this, we devise a more flexible feature-space generic-domain representation that emulates the expected properties.

3.2.1. Edge representation as a generic-domain

Motivation. General image classification or semantic segmentation tasks are based primarily on recognizing object shapes. Thus, any candidate representation for a generic-domain needs to possess at least the shape information. Intuitively, edge representation satisfies this criteria as it preserves the shape information while removing domain-variant information like color or texture.

Datasets. Consider the labeled source-edge-mixup dataset $\mathcal{D}_{s_m} = \{(x_{s_m}, y_s) : x_{s_m} = \lambda x_{s_g} + (1 - \lambda)x_s, (x_s, y_s) \in \mathcal{D}_s\}$ where \mathcal{D}_s is the original source dataset and $x_{s_g} = \mathcal{A}_e(x_s)$ is the edge representation of x_s . We use Soria et al. (2020) for the edge detector $\mathcal{A}_e(\cdot)$. Similarly, the unlabeled target-edge-mixup dataset is $\mathcal{D}_{t_m} = \{x_{t_m}\}$.

Table 1. Single-Source Domain Adaptation (SSDA) on Office-31 and VisDA benchmarks. SF indicates *source-free* adaptation and (+x.x) indicates absolute improvements over the corresponding baseline methods NRC and SHOT++ (current source-free SOTA).

Method	SF	Office-31							VisDA
		A→D	A→W	D→W	W→D	D→A	W→A	Avg	S → R
CDAN+RADA (Jin et al., 2021)	✗	96.1	96.2	99.3	100.0	77.5	77.4	91.1	76.3
FAA (Huang et al., 2021a)	✗	94.4	92.3	99.2	99.7	80.5	78.7	90.8	82.7
FixBi (Na et al., 2021)	✗	95.0	96.1	99.3	100.0	78.7	79.4	91.4	87.2
HCL (Huang et al., 2021b)	✓	90.8	91.3	98.2	100.0	72.7	72.7	87.6	83.5
CPGA (Qiu et al., 2021)	✓	94.4	94.1	98.4	99.8	76.0	76.6	89.9	84.1
A ² Net (Xia et al., 2021)	✓	94.5	94.0	99.2	100.0	76.7	76.1	90.1	84.3
VDM-DA (Tian et al., 2021)	✓	93.2	94.1	98.0	100.0	75.8	77.1	89.7	85.1
NRC (Yang et al., 2021a)	✓	96.0	90.8	99.0	100.0	75.3	75.0	89.4	85.9
<i>Ours</i> (edge-mixup) + NRC	✓	96.1	92.4	99.2	100.0	76.9	77.1	90.3 (+0.9)	86.4 (+0.5)
<i>Ours</i> (feat-mixup) + NRC	✓	96.3	92.8	99.2	100.0	77.4	77.5	90.5 (+1.1)	87.3 (+1.4)
SHOT++ (Liang et al., 2021)	✓	94.3	90.4	98.7	99.9	76.2	75.8	89.2	87.3
<i>Ours</i> (edge-mixup) + SHOT++	✓	94.4	92.0	98.9	100.0	77.8	77.9	90.2 (+1.0)	87.5 (+0.2)
<i>Ours</i> (feat-mixup) + SHOT++	✓	94.6	93.2	98.9	100.0	78.3	78.9	90.7 (+1.5)	87.8 (+0.5)

3.2.2. Feature-space generic-domain representation

Motivation. While edge representation is intuitive as a generic-domain, it poses some limitations. First, edges may not represent a generic-domain for all possible tasks. For instance, in texture classification, shape is not the primary discriminative characteristic. Here, edge-mixup would hinder the discriminability more than it assists the transferability. Second, while an input-space generic-domain (like edges) seems intuitive for some tasks, it may be very challenging to devise the same for other tasks. To address these limitations, we introduce a more flexible feature-space generic-domain.

Augmented sub-domains. We perform the mixup in the feature-space of the backbone h . Consider the features $z_s = h(x_s)$ for a source sample x_s . We choose a set of task-preserving image augmentations $\{\mathcal{A}^{[i]}(\cdot)\}_{i=1}^K$ from a pool of commonly used augmentations. The augmentation should not distort task discriminative factors while significantly modifying domain variant factors. For example, a strong stylization technique like AdaIN (Huang et al., 2017) satisfies this criteria. Whereas weaker augmentations like Gaussian blurring are not chosen, as they fail to significantly alter the domain variant factors. Intuitively, manipulating the domain-variant factors simulates a sample from an *augmented sub-domain*. See Appendix C.2 for more details.

Concretely, we extract a set of K features $\{z_s^{[i]} : z_s^{[i]} = h(x_s^{[i]})\}_{i=1}^K$ where $x_s^{[i]} = \mathcal{A}^{[i]}(x_s)$ is the i^{th} augmentation of x_s . Since each of the augmented sub-domain features represents the same task discriminative features, but vary in domain-variant features, we use the feature mean as the generic-domain features, *i.e.* $z_{s_g} = \frac{1}{K} \sum_{i=1}^K z_s^{[i]}$.

Usually, DA works (Saito et al., 2018; Rangwani et al., 2021) apply domain alignment or other losses at feature-level as deep features exhibit high domain discrepancy (Stephenson et al., 2021; Kundu et al., 2022). Thus, different augmentations of the same instance exhibit high feature-level domain-variance. Then, feature mean (generic-domain) effectively diffuses this variance while preserving task-related factors.

Datasets. Consider labeled source-feature-mixup dataset $\mathcal{D}_{s_m} = \{(z_{s_m}, y_s) : z_{s_m} = \lambda z_{s_g} + (1 - \lambda)z_s, (x_s, y_s) \in \mathcal{D}_s\}$ where $z_s = h(x_s)$. Similarly, unlabeled target-feature-mixup dataset is $\mathcal{D}_{t_m} = \{z_{t_m}\}$. Note that using the same augmentations for both source and target does not require concurrent source-target access *i.e.* does not violate the source-free constraints. We re-use the dataset notations from edge-mixup to avoid introducing more notations.

During training, the backbone h is only updated through the computation graph of z_s , *i.e.* the gradients do not flow through the computation graph of z_{s_g} . This is because the gradients corresponding to each augmentation may be at odds with each other and should not be aggregated directly. Also note that the final adapted model is finetuned on the original target data using the client-side algorithm *i.e.* $\text{CsAlg}_0(\mathcal{D}_t)$, for a small number of iterations, to remove the mixup requirement at inference-time. This ensures fair comparisons with prior arts using the same architecture.

4. Experiments

We thoroughly assess our technique against numerous state-of-the-art methods in different DA scenarios.

Datasets. We use four object classification DA benchmarks. Office-31 (Saenko et al., 2010) has three domains, 31 classes each: Amazon (**A**), DSLR (**D**), and Webcam (**W**). Office-Home (Venkateswara et al., 2017) contains four domains, 65 classes each: Artistic (**Ar**), Clipart (**Cl**), Product (**Pr**), and Real-world (**Rw**). VisDA (Peng et al., 2018) is a large-scale dataset with synthetic source and real target domains. DomainNet (Peng et al., 2019) is the most challenging with six domains, 345 classes each: Clipart (**C**), Real (**R**), Infograph (**I**), Painting (**P**), Sketch (**S**), and Quickdraw (**Q**). For semantic segmentation DA, we use synthetic GTA5 (**G**) (Richter et al., 2016), SYNTHIA (**Y**) (Ros et al., 2016), Synscapes (**S**) (Wrenninge et al., 2018) as source datasets and real-world Cityscapes (Cordts et al., 2016) as the target data. See Appendix C.1 for more details.

Balancing Discriminability and Transferability for Source-Free Domain Adaptation

Table 2. Multi-Source Domain Adaptation (MSDA) on DomainNet and Office-Home. * indicates results from released code. We outperform *source-free* (SF) prior arts despite not using domain labels. (+x.x) indicates improvements over the source-free SOTA NRC.

Method	SF	w/o Domain Labels	DomainNet						Office-Home					
			→C	→I	→P	→Q	→R	→S	Avg	→Ar	→Cl	→Pr	→Rw	Avg
SImpAl ₅₀ (Venkat et al., 2020)	✗	✗	66.4	26.5	56.6	18.9	68.0	55.5	48.6	70.8	56.3	80.2	81.5	72.2
CMSDA (Scalbert et al., 2021)	✗	✗	70.9	26.5	57.5	21.3	68.1	59.4	50.4	71.5	67.7	84.1	82.9	76.6
DRT (Li et al., 2021b)	✗	✗	71.0	31.6	61.0	12.3	71.4	60.7	51.3	-	-	-	-	-
STEM (Nguyen et al., 2021)	✗	✗	72.0	28.2	61.5	25.7	72.6	60.2	53.4	-	-	-	-	-
Source-combine	✗	✓	57.0	23.4	54.1	14.6	67.2	50.3	44.4	58.0	57.3	74.2	77.9	66.9
SHOT (Liang et al., 2020)-Ens	✓	✗	58.6	25.2	55.3	15.3	70.5	52.4	46.2	72.2	59.3	82.8	82.9	74.3
DECISION (Ahmed et al., 2021)	✓	✗	61.5	21.6	54.6	18.9	67.5	51.0	45.9	74.5	59.4	84.4	83.6	75.5
CAiDA (Dong et al., 2021)	✓	✗	-	-	-	-	-	-	-	75.2	60.5	84.7	84.2	76.2
NRC (Yang et al., 2021a)*	✓	✓	65.8	24.1	56.0	16.0	69.2	53.4	47.4	70.6	60.0	84.6	83.5	74.7
Ours (edge-mixup) + NRC	✓	✓	74.8	24.1	56.8	19.6	66.9	55.5	49.6 (+2.2)	72.1	62.9	86.4	84.8	76.6 (+1.9)
Ours (feature-mixup) + NRC	✓	✓	75.4	24.6	57.8	23.6	65.8	58.5	51.0 (+3.6)	72.6	67.4	85.9	83.6	77.4 (+2.7)

Table 3. Single-Source DA (SSDA) on Office-Home. SF indicates *source-free*, (+x.x) indicates gains over NRC and SHOT++ respectively.

Method	SF	Office-Home												
		Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg
RSDA-MSTN (Gu et al., 2020)	✗	53.2	77.7	81.3	66.4	74.0	76.5	67.9	53.0	82.0	75.8	57.8	85.4	70.9
SENTRY (Prabhu et al., 2021)	✗	61.8	77.4	80.1	66.3	71.6	74.7	66.8	63.0	80.9	74.0	66.3	84.1	72.2
FixBi (Na et al., 2021)	✗	58.1	77.3	80.4	67.7	79.5	78.1	65.8	57.9	81.7	76.4	62.9	86.7	72.7
SCDA (Li et al., 2021a)	✗	60.7	76.4	82.8	69.8	77.5	78.4	68.9	59.0	82.7	74.9	61.8	84.5	73.1
SHOT (Liang et al., 2020)	✓	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
A ² Net (Xia et al., 2021)	✓	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
GSFDA (Yang et al., 2021b)	✓	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
CPGA (Qiu et al., 2021)	✓	59.3	78.1	79.8	65.4	75.5	76.4	65.7	58.0	81.0	72.0	64.4	83.3	71.6
NRC (Yang et al., 2021a)	✓	57.7	80.3	82.0	68.1	79.8	78.6	65.3	56.4	83.0	71.0	58.6	85.6	72.2
Ours (edge-mixup) + NRC	✓	60.8	80.1	81.6	67.2	79.3	78.5	65.4	61.0	83.8	70.2	63.1	85.3	73.0 (+0.8)
Ours (feat-mixup) + NRC	✓	61.6	80.9	82.5	68.1	80.1	79.1	66.0	61.8	84.5	71.2	63.7	86.1	73.8 (+1.6)
SHOT++ (Liang et al., 2021)	✓	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
Ours (edge-mixup) + SHOT++	✓	61.0	80.4	82.1	67.6	79.8	78.8	67.1	60.7	84.3	73.0	63.5	85.7	73.7 (+0.7)
Ours (feat-mixup) + SHOT++	✓	61.8	81.2	83.0	68.5	80.6	79.4	67.8	61.5	85.1	73.7	64.1	86.5	74.5 (+1.5)

Implementation details. For object classification DA, we primarily use the source-free NRC (Yang et al., 2021a) for $VsAlgo(\cdot)$ and $CsAlgo(\cdot)$. Unless otherwise specified, *Ours* implies NRC as $VsAlgo(\cdot)$ and $CsAlgo(\cdot)$. We follow NRC, with a ResNet-50 (He et al., 2016) backbone for Office-Home, Office-31, and DomainNet, and a ResNet-101 for VisDA. For semantic segmentation DA, we follow Kundu et al. (2021) and use the standard DeepLabv2 (Chen et al., 2017) with a ResNet-101 backbone. We empirically set $\lambda = 0.1$ for both edge and feature-mixup during both vendor-side and client-side training. We find that $\lambda = 0.1$ works well across all settings and tasks. See Appendix C.2 for extensive implementation details.

4.1. Comparison with prior arts

a) Single Source Domain Adaptation (SSDA). Table 3 reports the results for object classification SSDA on Office-Home. Adding our proposed techniques to NRC (Yang et al., 2021a) improves NRC by 0.8% for edge-mixup and 1.6% for feature-mixup. Similarly, SHOT++ (Liang et al., 2021) with our edge-mixup improves by 0.7% while feature-mixup gives an improvement of 1.5%. Further, we outperform even the non-source-free works on Office-Home. Table 1 shows the results on Office-31 and the VisDA dataset. Similar

to Office-Home, we observe consistent improvements over both NRC and SHOT++ after adding feature-mixup (1.1% and 1.5%) on Office-31. We also outperform the non-source-free works on the large-scale VisDA dataset.

b) Multi Source Domain Adaptation (MSDA). Table 2 shows the results for object classification MSDA on Office-Home and the large-scale DomainNet benchmark. We use our proposed mixup techniques with NRC, without using domain labels. We observe improvements of 1.9% for edge-mixup and 2.7% for feature-mixup on Office-Home. With the same settings, we observe consistent gains of 2.2% for edge-mixup and 3.6% for feature-mixup on DomainNet. In Table 5, we compare our method to source-free and non-source-free prior works to assess our performance on Office-31 dataset. Even compared to non-source-free works, our technique obtains *state-of-the-art* performance on the Office-31 benchmark despite not using domain labels.

c) Multi-Target Domain Adaptation (MTDA). We adhere to the experimental protocols employed in Chen et al. (2019b); Mitsuzumi et al. (2021). We evaluate all possible combinations of one source domain and three target domains for MTDA for Office-31 benchmark in Table 4. Even without domain labels, our technique outperforms all others.

Table 4. Multi-Target Domain Adaptation (MTDA) on Office-31. * indicates taken from Roy et al. (2021)

Method	SF	w/o Domain Labels	Office-31			
			Amazon→	DSLR→	Webcam→	Avg.
Source train	✗	✓	68.6	70.0	66.5	68.4
MT-MTDA (Nguyen-Meidine et al., 2021)	✗	✗	87.9	83.7	84.0	85.2
HGAN (Yang et al., 2020b)	✗	✗	88.0	84.4	84.9	85.8
D-CGCT (Roy et al., 2021)	✗	✗	93.4	86.0	87.1	88.8
JAN (Long et al., 2017)*	✗	✓	84.2	74.4	72.0	76.9
CDAN (Long et al., 2018)*	✗	✓	93.6	80.5	81.3	85.1
AMEAN (Chen et al., 2019b)	✗	✓	90.1	77.0	73.4	80.2
GDA (Mitsuzumi et al., 2021)	✗	✓	88.8	74.5	73.2	87.9
CGCT (Roy et al., 2021)	✗	✓	93.9	85.1	85.6	88.2
Ours (edge-mixup)	✓	✓	90.3	87.1	86.9	88.1
Ours (feature-mixup)	✓	✓	92.5	88.4	86.5	89.1

Table 5. Multi-Source DA (MSDA) on Office-31.

Method	SF	Office-31			
		→A	→W	→D	Avg.
PFSA (Fu et al., 2021)	✗	57.0	97.4	99.7	84.7
DCTN (Xu et al., 2018)	✗	64.2	98.2	99.3	87.2
SimpAI (Venkat et al., 2020)	✗	70.6	97.4	99.2	89.0
WAMDA (Aggarwal et al., 2020)	✗	72.0	98.6	99.6	90.0
MFSAN (Zhu et al., 2019)	✗	72.7	98.5	99.5	90.2
MIAN (Park & Lee, 2021)	✗	76.2	98.4	99.2	91.3
MLAN (Xu et al., 2022)	✗	75.7	98.8	99.6	91.4
Source-combine	✗	65.2	94.6	98.4	86.1
SHOT(Liang et al., 2020)-Ens	✓	75.0	94.9	97.8	89.3
DECISION (Ahmed et al., 2021)	✓	75.4	98.4	99.6	91.1
CAiDA (Dong et al., 2021)	✓	75.8	98.9	99.8	91.6
Ours (edge-mixup)	✓	76.3	99.1	99.4	91.6
Ours (feature-mixup)	✓	76.9	99.1	99.3	91.8

d) DA for Semantic Segmentation. We use GtA (Kundu et al., 2021) with our proposed mixup techniques for both SSDA and MSDA (Table 6). For edge-mixup, we obtain consistent gains of 1% and 1.2% over GtA on SSDA for GTA5→Cityscapes (G) and SYNTHIA→Cityscapes (Y) respectively. For MSDA with edge-mixup, we obtain consistent gains (average 2.2%) over GtA on 4 settings (combinations of different source datasets). We also outperform the non-source-free SOTA MSDA-CL (He et al., 2021) by an average 0.4%. In contrast to object classification DA, the gains of feature-mixup over GtA (average 1.1%) are lower than those for edge-mixup. This is because feature-mixup better suits non-dense vectorized prediction tasks like object classification (where mixup is performed after the global average pooling layer). Whereas feature-mixup in semantic segmentation DA is performed on spatial conv. features.

e) Faster and better convergence. Fig. 5B shows the improved and faster convergence for our approach on both SSDA and MSDA on Office-Home. Since we effectively improve the transferability *i.e.* reduce the domain gap with the mixup domains, the adaptation becomes faster. Further, as per Insight 2, the lowered target error upper bound leads to improved performance *i.e.* better convergence.

Table 6. Domain adaptation for semantic segmentation. G, S, Y indicate GTA5, Synscapes, SYNTHIA datasets as source, SF indicates source-free, mIoU is 19-class for G column and 13-class for all others, (+x.x) indicates gains over prior source-free SOTA.

Method (→ Cityscapes)	SF	SSDA		MSDA				
		G	Y	G+S	S+Y	G+Y	G+S+Y	Avg
FDA (Yang et al. 2020)	✗	50.5	52.5	-	-	-	-	-
ProDA (Zhang et al., 2021)	✗	57.5	62.0	-	-	-	-	-
Src-combine (He et al., 2021)	✗	-	-	57.2	53.6	55.5	58.0	56.1
MSDA-CL (He et al., 2021)	✗	-	-	65.8	63.1	59.4	67.1	63.8
SFDA (Liu et al., 2021)	✓	43.1	45.9	-	-	-	-	-
URMA (Teja et al., 2021)	✓	45.1	45.0	-	-	-	-	-
SFUDA (Ye et al., 2021)	✓	49.4	51.9	-	-	-	-	-
GtA (Kundu et al., 2021)	✓	51.6	55.5	63.5	62.8	58.3	63.4	62.0
Ours (feature-mixup)	✓	51.9	55.6	63.6	63.2	61.4	64.3	63.1
		+0.3	+0.1	+0.1	+0.4	+3.1	+0.9	+1.1
Ours (edge-mixup)	✓	52.6	56.7	64.6	65.4	61.8	64.9	64.2
		+1.0	+1.2	+1.1	+2.6	+3.5	+1.5	+2.2

4.2. Analysis

a) Effect of sub-domain augmentations. We performed an ablation study for SSDA on Office-Home (see Table 8) to disentangle the gains from the use of sub-domain augmentations in feature-mixup. While sub-domain augmentations improve NRC and SHOT++ by 0.6% and 0.5% respectively, we observe that feature-mixup provides further improvements of ~1% in both cases. Thus, the gains from feature-mixup can be attributed more to the proposed mixup than the sub-domain augmentations.

b) Empirical transferability vs. discriminability. Fig. 5A illustrates the empirical curve between discriminability and transferability for different λ in edge-mixup and feature-mixup. Transferability γ_T is evaluated on mixup domain data with Eq. 1, 2 using a domain-classifier trained on original source-target data. Discriminability γ_D is evaluated with Eq. 1, 2 by computing the accuracy of a source-target joint supervised task classifier. Both domain classifier and task classifier are trained on the features of a frozen backbone h trained with $\text{VsAlgo}(\mathcal{D}_{s_m})$ with a particular λ . The empirical plots are similar to the conceptual plot of Fig. 3B.




Figure 5. **A.** Empirical discriminability vs. transferability for CI→Rw (Office-Home) at different mixup ratios λ . The plots are similar to the conceptual plot of Fig. 3B and $\lambda=0.1$ presents a good trade-off. **B.** Faster and better convergence w.r.t. existing source-free works on both SSSDA and MSDA for Office-Home, as mixup reduces the domain gap. **C.** Sensitivity to mixup ratio λ for MSDA on Office-Home.

Table 7. **Compatibility with non-source-free DA** works on Office-Home. SSSDA and MSDA indicate single-source and multi-source DA. (+x.x) indicates gains over corresponding baseline.

Method	Office-Home	
	SSDA	MSDA
DANN (Ganin et al., 2016)	57.6	64.6
DANN + feature-mixup	67.2 (+9.6)	72.9 (+8.3)
CDAN+E (Long et al., 2018)	65.8	69.4
CDAN+E + feature-mixup	72.2 (+6.4)	74.1 (+4.7)
SRDC (Tang et al., 2020)	71.3	73.1
SRDC + feature-mixup	72.3 (+1.0)	75.1 (+2.0)

Table 8. **Ablation study on sub-domain augmentations** for SSSDA on Office-Home benchmark. (+x.x) indicates improvements over NRC and SHOT++ respectively.

Method	Average Acc.
NRC (Yang et al., 2021a)	72.2
NRC + sub-domain augs.	72.8 (+0.6)
Ours (feat-mixup) + NRC	73.8 (+1.6)
SHOT++ (Liang et al., 2021)	73.0
SHOT++ + sub-domain augs.	73.5 (+0.5)
Ours (feat-mixup) + SHOT++	74.5 (+1.5)

c) Sensitivity to mixup ratio λ . We conduct a sensitivity analysis for mixup ratio λ (Fig. 5C) for MSDA on Office-Home. For edge-mixup, we observe gains over the baseline ($\lambda=0$) until $\lambda=0.1$, followed by a drop w.r.t. the baseline for $\lambda>0.2$. For higher λ , the proportion of edges increases. The significant loss of texture and color causes a sizeable drop in discriminability (Fig. 5A), leading to poor adaptation performance. In contrast, the flexibility of feature-space mixup enables it to outperform the baseline consistently upto $\lambda=0.8$. Thus, a relatively low λ is preferable for edge-mixup while feature-mixup is more invariant to λ .

d) Qualitative analysis for semantic segmentation DA. Figure 6 illustrates qualitative results on Cityscapes (target dataset) for different SSSDA (G) and MSDA settings (G+S, S+Y, G+Y, G+S+Y) for semantic segmentation. We compare with the vendor-side source-only baseline and the prior source-free SOTA (Kundu et al., 2021) and highlight the improvement regions with white circles.

4.3. Compatibility with non-source-free DA

Table 7 examines the complementary nature of the proposed mixup to prior non-source-free SSSDA works on Office-Home. For MSDA comparisons, multiple sources are merged into a single source. For SSSDA, adding feature-mixup enhances DANN by 9.8%, CDAN by 6.4% and SRDC by 1%. Similarly, for MSDA, adding feature-mixup improves DANN by 8.3%, CDAN by 4.7% and SRDC by 2%. These consistent improvements demonstrate our general compatibility with non-source-free DA works.

5. Conclusion

We study the perspective of discriminability-transferability trade-off in source-free DA. Identifying the key hurdles to extending available trade-off-based techniques, we investigate the idea of generic domain representations without concurrent source-target access. Based on our insights, we propose novel ways to realize generic domains, but observe degraded task-discriminability. As a remedy, we operate on intermediate mixup domains and theoretically demonstrate that the mixup domains achieve a tighter bound on target error, thereby improving the DA performance. This simple modification, added to prior DA approaches, yields state-of-the-art performance across source-free benchmarks for single-source and multi-source DA on both classification and segmentation. Since the procedure for realizing a generic-domain is somewhat task-dependent, future work can focus on learnable realizations of generic-domains.

Acknowledgements. This work was supported by MeitY (Ministry of Electronics and Information Technology) project (No. 4(16)2019-ITEA), Govt. of India.

References

- Aggarwal, S., Kundu, J. N., Radhakrishnan, V. B., and Chakraborty, A. WAMDA: Weighted alignment of sources for multi-source domain adaptation. In *BMVC*, 2020.
- Ahmed, S. M., Raychaudhuri, D. S., Paul, S., Oymak, S., and Roy-Chowdhury, A. K. Unsupervised multi-source domain adaptation without access to source data. In *CVPR*, 2021.
- Ahmed, W., Morerio, P., and Murino, V. Cleaning noisy labels

- by negative ensemble learning for source-free unsupervised domain adaptation. In *WACV*, 2022.
- Awais, M., Zhou, F., Xu, H., Hong, L., Luo, P., Bae, S.-H., and Li, Z. Adversarial robustness for unsupervised domain adaptation. In *ICCV*, 2021.
- Ben-David, S., Blitzer, J., Crammer, K., and Pereira, F. Analysis of representations for domain adaptation. In *NeurIPS*, 2006.
- Blitzer, J., Dredze, M., and Pereira, F. Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*, 2007.
- Chen, C., Zheng, Z., Ding, X., Huang, Y., and Dou, Q. Harmonizing transferability and discriminability for adapting object detectors. In *CVPR*, 2020.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- Chen, X., Wang, S., Long, M., and Wang, J. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In *ICML*, 2019a.
- Chen, Z., Zhuang, J., Liang, X., and Lin, L. Blending-target domain adaptation by adversarial meta-adaptation networks. In *CVPR*, 2019b.
- Chou, H.-P., Chang, S.-C., Pan, J.-Y., Wei, W., and Juan, D.-C. Remix: Rebalanced mixup. In *ECCV*, 2020.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- Dong, J., Fang, Z., Liu, A., Sun, G., and Liu, T. Confident anchor-induced multi-source free domain adaptation. In *NeurIPS*, 2021.
- Fu, Y., Zhang, M., Xu, X., Cao, Z., Ma, C., Ji, Y., Zuo, K., and Lu, H. Partial feature selection and alignment for multi-source domain adaptation. In *CVPR*, 2021.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Gu, X., Sun, J., and Xu, Z. Spherical space domain adaptation with robust pseudo-label loss. In *CVPR*, 2020.
- He, J., Jia, X., Chen, S., and Liu, J. Multi-source domain adaptation with collaborative learning for semantic segmentation. In *CVPR*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, 2016.
- Hou, W., Wang, J., Tan, X., Qin, T., and Shinozaki, T. Cross-domain speech recognition with unsupervised character-level distribution matching. In *Interspeech*, 2021.
- Huang, J., Guan, D., Xiao, A., and Lu, S. RDA: Robust domain adaptation via fourier adversarial attacking. In *ICCV*, 2021a.
- Huang, J., Guan, D., Xiao, A., and Lu, S. Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. In *NeurIPS*, 2021b.
- Huang, X. and Belongie, S. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- Jin, X., Lan, C., Zeng, W., and Chen, Z. Re-energizing domain discriminator with sample relabeling for adversarial domain adaptation. In *ICCV*, 2021.
- Jung, A. B., Wada, K., Crall, J., Tanaka, S., Graving, J., Reinders, C., Yadav, S., Banerjee, J., Vecsei, G., Kraft, A., Rui, Z., Borovec, J., Vallentin, C., Zhydenko, S., Pfeiffer, K., Cook, B., Fernández, I., De Rainville, F.-M., Weng, C.-H., Ayala-Acevedo, A., Meudec, R., Laporte, M., et al. [imgaug](https://github.com/aleju/imgaug). <https://github.com/aleju/imgaug>, 2020. Online; accessed 01-Feb-2020.
- Karouzos, C. F., Paraskevopoulos, G., and Potamianos, A. UDALM: Unsupervised domain adaptation through language modeling. In *NAACL*, 2021.
- Kim, D., Saito, K., Oh, T.-H., Plummer, B. A., Sclaroff, S., and Saenko, K. CDS: Cross-domain self-supervised pre-training. In *ICCV*, 2021a.
- Kim, J.-H., Choo, W., Jeong, H., and Song, H. O. Co-mixup: Saliency guided joint mixup with supermodular diversity. In *ICLR*, 2021b.
- Kim, Y., Cho, D., Han, K., Panda, P., and Hong, S. Domain adaptation without source data. *IEEE Transactions on Artificial Intelligence*, 2(6):508–518, 2021c.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Krueger, D., Caballero, E., Jacobsen, J.-H., Zhang, A., Binas, J., Zhang, D., Priol, R. L., and Courville, A. Out-of-distribution generalization via risk extrapolation (rex). In *ICML*, 2021.
- Kundu, J. N., Venkat, N., and Babu, R. V. Universal source-free domain adaptation. In *CVPR*, 2020a.
- Kundu, J. N., Venkat, N., Revanur, A., V. R. M., and Babu, R. V. Towards inheritable models for open-set domain adaptation. In *CVPR*, 2020b.
- Kundu, J. N., Venkatesh, R. M., Venkat, N., Revanur, A., and Babu, R. V. Class-incremental domain adaptation. In *ECCV*, 2020c.
- Kundu, J. N., Kulkarni, A., Singh, A., Jampani, V., and Babu, R. V. Generalize then adapt: Source-free domain adaptive semantic segmentation. In *ICCV*, 2021.
- Kundu, J. N., Kulkarni, A., Bhambrri, S., Jampani, V., and Babu, R. V. Amplitude spectrum transformation for open compound domain adaptive semantic segmentation. In *AAAI*, 2022.
- Li, D., Yang, Y., Song, Y.-Z., and Hospedales, T. M. Deeper, broader and artier domain generalization. In *ICCV*, 2017.

- Li, R., Jiao, Q., Cao, W., Wong, H.-S., and Wu, S. Model adaptation: Unsupervised domain adaptation without source data. In *CVPR*, 2020.
- Li, S., Xie, M., Lv, F., Liu, C. H., Liang, J., Qin, C., and Li, W. Semantic concentration for domain adaptation. In *ICCV*, 2021a.
- Li, Y., Yuan, L., Chen, Y., Wang, P., and Vasconcelos, N. Dynamic transfer for multi-source domain adaptation. In *CVPR*, 2021b.
- Liang, J., Hu, D., and Feng, J. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020.
- Liang, J., Hu, D., Wang, Y., He, R., and Feng, J. Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Liu, Y., Zhang, W., and Wang, J. Source-free domain adaptation for semantic segmentation. In *CVPR*, 2021.
- Long, M., Cao, Y., Wang, J., and Jordan, M. Learning transferable features with deep adaptation networks. In *ICML*, 2015.
- Long, M., Zhu, H., Wang, J., and Jordan, M. I. Deep transfer learning with joint adaptation networks. In *International conference on machine learning*, pp. 2208–2217. PMLR, 2017.
- Long, M., Cao, Z., Wang, J., and Jordan, M. I. Conditional adversarial domain adaptation. In *NeurIPS*, 2018.
- Mathur, A., Kawsar, F., Berthouze, N., and Lane, N. D. Libri-adapt: a new speech dataset for unsupervised domain adaptation. In *ICASSP*, 2020.
- Mitsuzumi, Y., Irie, G., Ikami, D., and Shibata, T. Generalized domain adaptation. In *CVPR*, 2021.
- Morerio, P., Volpi, R., Ragonesi, R., and Murino, V. Generative pseudo-label refinement for unsupervised domain adaptation. In *WACV*, 2020.
- Murez, Z., Kolouri, S., Kriegman, D., Ramamoorthi, R., and Kim, K. Image to image translation for domain adaptation. In *CVPR*, 2018.
- Na, J., Jung, H., Chang, H. J., and Hwang, W. FixBi: Bridging domain spaces for unsupervised domain adaptation. In *CVPR*, 2021.
- Nguyen, V.-A., Nguyen, T., Le, T., Tran, Q. H., and Phung, D. STEM: An approach to multi-source domain adaptation with guarantees. In *ICCV*, 2021.
- Nguyen-Meidine, L. T., Belal, A., Kiran, M., Dolz, J., Blais-Morin, L.-A., and Granger, E. Unsupervised multi-target domain adaptation through knowledge distillation. In *WACV*, 2021.
- Park, G. Y. and Lee, S. W. Information-theoretic regularization for multi-source domain adaptation. In *ICCV*, 2021.
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. VisDA: The visual domain adaptation challenge. In *CVPRW*, 2018.
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., and Wang, B. Moment matching for multi-source domain adaptation. In *ICCV*, 2019.
- Prabhu, V., Khare, S., Kartik, D., and Hoffman, J. SENTRY: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *ICCV*, 2021.
- Qiu, Z., Zhang, Y., Lin, H., Niu, S., Liu, Y., Du, Q., and Tan, M. Source-free domain adaptation via avatar prototype generation and adaptation. In *IJCAI*, 2021.
- Rangwani, H., Jain, A., Aithal, S. K., and Babu, R. V. S3VAADA: Submodular subset selection for virtual adversarial active domain adaptation. In *ICCV*, 2021.
- Rangwani, H., Aithal, S. K., Mishra, M., Jain, A., and Babu, R. V. A closer look at smoothness in domain adversarial training. In *ICML*, 2022.
- Richter, S. R., Vineet, V., Roth, S., and Koltun, V. Playing for data: Ground truth from computer games. In *ECCV*, 2016.
- Ros, G., Sellart, L., Materzynska, J., Vazquez, D., and Lopez, A. M. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, 2016.
- Roy, S., Krivosheev, E., Zhong, Z., Sebe, N., and Ricci, E. Curriculum graph co-teaching for multi-target domain adaptation. In *CVPR*, 2021.
- Russo, P., Carlucci, F. M., Tommasi, T., and Caputo, B. From source to target and back: symmetric bi-directional adaptive GAN. In *CVPR*, 2018.
- Saenko, K., Kulis, B., Fritz, M., and Darrell, T. Adapting visual category models to new domains. In *ECCV*, 2010.
- Saito, K., Watanabe, K., Ushiku, Y., and Harada, T. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018.
- Salimans, T. and Kingma, D. P. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *NeurIPS*, 2016.
- Scalbert, M., Vakalopoulou, M., and Couzini'e-Devy, F. Multi-source domain adaptation via supervised contrastive learning and confident consistency regularization. In *BMVC*, 2021.
- Sivaprasad, P. T. and Fleuret, F. Uncertainty reduction for model adaptation in semantic segmentation. In *CVPR*, 2021.
- Soria, X., Riba, E., and Sappa, A. Dense extreme inception network: Towards a robust cnn model for edge detection. In *WACV*, 2020.
- Stephenson, C., Padhy, S., Ganesh, A., Hui, Y., Tang, H., and Chung, S. On the geometry of generalization and memorization in deep neural networks. In *ICLR*, 2021.
- Tang, H., Chen, K., and Jia, K. Unsupervised domain adaptation via structurally regularized deep clustering. In *CVPR*, 2020.
- Tian, J., Zhang, J., Li, W., and Xu, D. VDM-DA: Virtual domain modeling for source data-free domain adaptation. *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *CVPR*, 2017.

- Venkat, N., Kundu, J. N., Singh, D. K., Revanur, A., and Babu, R. V. Your classifier can secretly suffice multi-source domain adaptation. In *NeurIPS*, 2020.
- Venkateswara, H., Eusebio, J., Chakraborty, S., and Panchanathan, S. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017.
- Wei, G., Lan, C., Zeng, W., and Chen, Z. MetaAlign: Coordinating domain alignment and classification for unsupervised domain adaptation. In *CVPR*, 2021.
- Wen, J., Greiner, R., and Schuurmans, D. Domain aggregation networks for multi-source domain adaptation. In *ICML*, 2020.
- Wrenninge, M. and Unger, J. Synscapes: A photorealistic synthetic dataset for street scene parsing, 2018.
- Xia, H., Zhao, H., and Ding, Z. Adaptive adversarial network for source-free domain adaptation. In *ICCV*, 2021.
- Xu, R., Chen, Z., Zuo, W., Yan, J., and Lin, L. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, 2018.
- Xu, Y., Kan, M., Shan, S., and Chen, X. Mutual learning of joint and separate domain alignments for multi-source domain adaptation. In *WACV*, 2022.
- Yang, J., Zou, H., Zhou, Y., Zeng, Z., and Xie, L. Mind the discriminability: Asymmetric adversarial domain adaptation. In *ECCV*, 2020a.
- Yang, S., van de Weijer, J., Herranz, L., Jui, S., et al. Exploiting the intrinsic neighborhood structure for source-free domain adaptation. In *NeurIPS*, 2021a.
- Yang, S., Wang, Y., van de Weijer, J., Herranz, L., and Jui, S. Generalized source-free domain adaptation. In *ICCV*, 2021b.
- Yang, X., Deng, C., Liu, T., and Tao, D. Heterogeneous graph attention network for unsupervised multiple-target domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020b.
- Yang, Y. and Soatto, S. FDA: Fourier domain adaptation for semantic segmentation. In *CVPR*, 2020.
- Ye, M., Zhang, J., Ouyang, J., and Yuan, D. Source data-free unsupervised domain adaptation for semantic segmentation. In *ACMMM*, 2021.
- Yue, Z., Sun, Q., Hua, X.-S., and Zhang, H. Transporting causal mechanisms for unsupervised domain adaptation. In *ICCV*, 2021.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In *ICLR*, 2018.
- Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., and Wen, F. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *CVPR*, 2021.
- Zhao, H., Zhang, S., Wu, G., Moura, J. M., Costeira, J. P., and Gordon, G. J. Adversarial multiple source domain adaptation. In *NeurIPS*, 2018.
- Zhu, Y., Zhuang, F., and Wang, D. Aligning domain-specific distribution and classifier for cross-domain classification from multiple sources. In *AAAI*, 2019.

Appendix

In this appendix, we provide more details of our approach and theoretical insights, extensive implementation details, additional quantitative and qualitative performance analysis and ablation studies. Towards reproducible research, we will publicly release our complete codebase and trained network weights. The appendix is organized as follows:

- Section A: Notations (Table 9)
- Section B: Discussions related to theory
- Section C: Experiments
 - Implementation details (Sec. C.1)
 - Experimental settings (Sec. C.2)
 - Additional results (Sec. C.3, Table 10, Fig. 6)
 - Extended comparisons (Sec. C.4, Table 11, 12, 13)

A. Notations

We summarize the notations used throughout the paper in Table 9. The notations are listed under 7 groups *i.e.* models, distributions, theory-related, datasets, samples, spaces and miscellaneous.

Table 9. Notation Table

	Symbol	Description
Models	h	Backbone feature extractor
	f_c	Task classifier
	f_d	Domain classifier
Distributions	p_s	Source marginal distribution
	p_t	Target marginal distribution
	p_{s_g}	Source-generic marginal distr.
	p_{t_g}	Target-generic marginal distr.
	p_{s_m}	Source-mixup marginal distr.
	p_{t_m}	Target-mixup marginal distr.
Theory-related	ϵ_s	Expected source risk
	ϵ_t	Expected target risk
	$d_{\mathcal{H}}$	\mathcal{H} -divergence
	\mathcal{H}	Backbone hypothesis space
Datasets	\mathcal{D}_s	Labeled source dataset
	\mathcal{D}_{s_m}	Labeled source-mixup dataset
	\mathcal{D}_t	Unlabeled target dataset
	\mathcal{D}_{t_m}	Unlabeled target-mixup dataset
Samples	(x_s, y_s)	Labeled source sample
	x_{s_g}	Generic domain sample of x_s
	(x_{s_m}, y_s)	Labeled source-mixup sample
	x_t	Unlabeled target sample
	x_{t_m}	Unlabeled target-mixup sample
	z_s	Features of sample x_s
	$z_s^{[i]}$	Features of sample with i^{th} aug.
	z_{s_g}	Generic domain features of z_s
	z_{s_m}	Source-mixup features
Spaces	\mathcal{X}	Input space
	\mathcal{Z}	Backbone feature space
	\mathcal{C}	Label set for goal task
Miscellaneous	\mathcal{A}_e	Edge estimation method
	$\mathcal{A}^{[i]}$	i^{th} augmentation function

B. Discussions related to theory

B.1. Theorem 1 and Proof

Theorem 1. (Mixup \mathcal{H} -divergence) Assume that original source p_s and target p_t are easily separable i.e. perfect accuracy for domain classifier f_d . Also assume that source-generic p_{s_g} and target-generic p_{t_g} domains are impossible to separate i.e. accuracy of f_d imitates that of a random classifier. For a linear domain classifier f_d :

$$d_{\mathcal{H}}(p_{s_m}, p_{t_m}) \leq d_{\mathcal{H}}(p_s, p_t) \quad (8)$$

Proof. From Eq. 1,

$$\begin{aligned} d_{\mathcal{H}}(p_s, p_t) &= \sup_{h' \in \mathcal{H}} \left| \mathbb{E}_{x \sim p_s} [\mathbb{1}(f_d(z)=1)] - \mathbb{E}_{x \sim p_t} [\mathbb{1}(f_d(z)=1)] \right| \\ &= \sup_{h' \in \mathcal{H}} |\phi(p_s) - \phi(p_t)|; \text{ where } z = h'(x) \end{aligned} \quad (9)$$

Note that domain classifier produces 0 for source and 1 for target. Thus, $\phi(p_s)$ is the source-domain-classification *error* while $\phi(p_t)$ is the target-domain-classification *accuracy* since the conditions in both indicator functions are for 1 (i.e. target). Recall that p_{s_g} and p_{t_g} are the source and target generic distributions respectively.

First, we focus on $\phi(p_{s_m})$ i.e. source-domain-classification error for the source-mixup distribution p_{s_m} . We replace the condition of $f_d(z) = 1$ to $f_d(z) > 0$ since f_d would practically be a binary domain classifier relying on sigmoid activation, but we require f_d to simply be a linear function without any non-linear activation function. Then,

$$\phi(p_{s_m}) = \mathbb{E}_{x \sim p_s} [\mathbb{1}(f_d(\lambda z_{s_g} + (1-\lambda)z) > 0)]; \text{ where } z = h(x)$$

We assume that f_d is a linear function i.e. $f_d(ka+b) = kf_d(a) + f_d(b)$ for a scalar constant k . And we use the property that expectation over indicator function is a probability,

$$\phi(p_{s_m}) = \Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) + (1-\lambda)f_d(z) > 0]$$

Now we consider the possible separate cases where $\lambda f_d(z_{s_g}) + (1-\lambda)f_d(z) > 0$ is satisfied (as shown in the Table below).

	$\lambda f_d(z_{s_g})$	$(1-\lambda)f_d(z)$	Comments
1	> 0	> 0	> 0 always
2.1	> 0	< 0	> 0 if $\lambda f_d(z_{s_g}) > (1-\lambda) f_d(z) $
2.2	> 0	< 0	< 0 otherwise
3.1	< 0	> 0	> 0 if $\lambda f_d(z_{s_g}) < (1-\lambda) f_d(z) $
3.2	< 0	> 0	< 0 otherwise
4	< 0	< 0	< 0 always

Using 1, 2.1, 3.1 from the table above, we split the probability term into a sum of product of probabilities of the three cases,

$$\begin{aligned} \phi(p_{s_m}) &= \Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) > 0] \Pr_{x \sim p_s} [(1-\lambda)f_d(z) > 0] + \\ &\Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) > 0] \Pr_{x \sim p_s} [(1-\lambda)f_d(z) < 0] \Pr_{x \sim p_s} [\lambda|f_d(z_{s_g})| > (1-\lambda)|f_d(z)|] + \\ &\Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) < 0] \Pr_{x \sim p_s} [(1-\lambda)f_d(z) > 0] \Pr_{x \sim p_s} [\lambda|f_d(z_{s_g})| < (1-\lambda)|f_d(z)|] \end{aligned}$$

Note that λ and $(1-\lambda)$ terms do not matter when the product of λ or $(1-\lambda)$ with $f_d(\cdot)$ is compared with 0, as both are positive scalar constants. Further, using $\phi(p_s) = \Pr_{x \sim p_s} [f_d(z) > 0]$ and $\phi(p_{s_g}) = \Pr_{x \sim p_s} [f_d(z_{s_g}) > 0]$, we get,

$$\begin{aligned} \phi(p_{s_m}) &= \phi(p_{s_g})\phi(p_s) + \phi(p_{s_g})(1-\phi(p_s)) \Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) > (1-\lambda)f_d(z)] + \\ &(1-\phi(p_{s_g}))\phi(p_s) \Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) < (1-\lambda)f_d(z)] \end{aligned}$$

To simplify, we define $\zeta_s = \Pr_{x \sim p_s} [\lambda f_d(z_{s_g}) > (1 - \lambda) f_d(z)]$ where $\zeta_s \in [0, 1]$. Note that the other probability term is $(1 - \zeta_s)$ due to the complement rule of probabilities.

$$\phi(p_{s_m}) = \phi(p_{s_g})\phi(p_s) + \phi(p_{s_g})(1 - \phi(p_s))\zeta_s + (1 - \phi(p_{s_g}))\phi(p_s)(1 - \zeta_s)$$

Now, we use the assumption that the source-generic domain p_{s_g} and target-generic domain p_{t_g} are impossible to separate *i.e.* accuracy of f_d imitates that of a random classifier. Formally, $\phi(p_{s_g})$ represents the source-domain-classification error for the source-generic domain. Thus, $\phi(p_{s_g}) \rightarrow \frac{1}{2}$ (since binary classification). Using this and simplifying,

$$\phi(p_{s_m}) = \phi(p_s)(1 - \zeta_s) + \frac{\zeta_s}{2} \quad (10)$$

Similarly, we can derive the same expression for the target-mixup domain *i.e.* $\phi(p_{t_m})$ except that the assumption of target-domain-classification accuracy, imitating the accuracy of a random classifier, for target-generic domain has to be used *i.e.* $\phi(p_{t_g}) \rightarrow \frac{1}{2}$. We get,

$$\phi(p_{t_m}) = \phi(p_t)(1 - \zeta_t) + \frac{\zeta_t}{2} \quad (11)$$

We know that $0 \leq \zeta_s, \zeta_t \leq 1$ since both are probabilities. Further, we had assumed that original source and target domains are easily separable *i.e.* perfect accuracy (or zero error) for domain classifier. Thus, $\phi(p_s) \rightarrow 0$ and $\phi(p_t) \rightarrow 1$ as per their interpretations of error and accuracy respectively. With these constraints on $\phi(p_s), \phi(p_t), \zeta_s, \zeta_t$, it can be easily shown that,

$$\phi(p_{s_m}) \geq \phi(p_s) \quad \forall \zeta_s \in [0, 1]; \quad \phi(p_{t_m}) \leq \phi(p_t) \quad \forall \zeta_t \in [0, 1] \quad (12)$$

The interpretation is that source-domain-classification error increases while target-domain-classification accuracy decreases after mixup. In other words, the mixup domain features are more domain-invariant compared to the original domain features. Note that while λ does not directly appear in Eq. 12, both ζ_s and ζ_t vary with λ . Thus, λ will control both the decrease in accuracy and the increase in error *i.e.* the \mathcal{H} -divergence relation between mixup domains and original domains is not independent of λ .

Now, we add the two inequalities of Eq. 12 and rearrange the terms,

$$\begin{aligned} \phi(p_s) + \phi(p_{t_m}) &\leq \phi(p_{s_m}) + \phi(p_t) \\ \phi(p_{t_m}) - \phi(p_{s_m}) &\leq \phi(p_t) - \phi(p_s) \end{aligned}$$

Finally, applying absolute value function on both sides (without sign change as both sides are positive) and taking supremum on both sides over $h \in \mathcal{H}$,

$$\sup_{h \in \mathcal{H}} |\phi(p_{t_m}) - \phi(p_{s_m})| \leq \sup_{h \in \mathcal{H}} |\phi(p_t) - \phi(p_s)|$$

Applying the definition of \mathcal{H} -divergence from Eq. 9, we arrive at our result *i.e.* Eq. 8. \square

B.2. Regarding result relating $\kappa(p_{s_m}, p_{t_m})$ and $\kappa(p_s, p_t)$

As noted in the main paper, the assumptions on f_c (involved in the computation of $\kappa(\cdot, \cdot)$ in Eq. 1) become unrealistic if we use a similar strategy as in the proof of Theorem 1. Concretely, we may assume f_c to be a linear classifier in order to separate the mixup terms over f_c . However, we cannot make realistic assumptions on the source-generic error $\epsilon_{s_g}(h)$ or target-generic error $\epsilon_{t_g}(h)$ in the way that assumptions were made for the domain classifier f_d . This is because the realizable generic domains have high transferability by definition, but lack a strict condition on their discriminability (*i.e.* on $\epsilon_{s_g}(h)$ or $\epsilon_{t_g}(h)$). However, we observe in practice that the discriminability is preserved over a reasonable range of λ (see Fig. 5A).

C. Experiments

C.1. Implementation details

Datasets. We assess the efficacy of our approach on four object classification DA benchmarks. The **Office-31** (Saenko et al., 2010) benchmark contains three domains from office settings: Amazon (**A**), DSLR (**D**), and Webcam (**W**), each with 31 object categories. **Office-Home** (Venkateswara et al., 2017), a more complex dataset, consists of images of everyday objects from four domains, each having 65 classes: Artistic (**Ar**), Clipart (**Cl**), Product (**Pr**), and Real-world (**Rw**). **VisDA** (Peng et al., 2018) is a large-scale dataset with 152,397 synthetic images in the source domain and 55,388 real-world images in the target domain. Finally, **DomainNet** (Peng et al., 2019) is the most challenging due to its extremely diversified domains, high class imbalance, and 6 domains with 345 classes each: Clipart (**C**), Real (**R**), Infograph (**I**), Painting (**P**), Sketch (**S**), and Quickdraw (**Q**).

C.2. Experimental settings

For object classification DA, we primarily use the source-free NRC (Yang et al., 2021a) for $VsAlgo(\cdot)$ and $CsAlgo(\cdot)$. We follow NRC, with a ResNet-50 (He et al., 2016) backbone for Office-Home, Office-31, and DomainNet, and a ResNet-101 backbone for VisDA. For edge-mixup, we choose a CNN-based edge detector (Soria et al., 2020), pretrained on a dataset with images irrelevant to the usual DA benchmark datasets. We use the Adam optimizer (Kingma & Ba, 2014) with a learning rate of $1e-3$, momentum of 0.9, and batch size of 64 for training with label smoothing following (Yang et al., 2021a; Liang et al., 2020). For semantic segmentation DA, we follow Kundu et al. (2021) and use the standard DeepLabv2 (Chen et al., 2017) with a ResNet-101 backbone. We omit the cPAE as it presents an extra model that complicates the training process and is orthogonal to our contributions.

Sub-domain augmentations. We selected to depict the sub-domains using the following four strong augmentations (total 5 sub-domains) on the original data:

- a) FDA: In this augmentation (Yang & Soatto, 2020), the image is stylized based on a reference image by exchanging the low frequency FFT spectrum. We use images from a style transfer dataset (Huang et al., 2017) as reference images.
- b) Weather augmentation: We employ Jung et al. (2020)’s frost (weather condition) augmentation. There are five severity levels, and we randomly choose one from the lowest three to govern the augmentation strength.
- c) AdaIN: By altering the convolutional feature statistics in an instance normalization layer (Ulyanov et al., 2017), this approach (Huang et al., 2017) stylizes the image depending on a reference image. To guarantee adequate content retention, we set the stylization strength to 0.3 (on a scale of 0 to 1, *i.e.* original image to severely stylized). We use images from a style transfer dataset (Huang et al., 2017) as reference images.
- d) Cartoonization: We employ the Jung et al. (2020) cartoonization augmentation. The strength of the augmentation has no controllable parameter.

Feature-mixup training. We use the same network architecture as NRC (Yang et al., 2021a), substituting the classifier with a fully connected layer with batch normalization (Ioffe et al., 2015) and a fully connected layer with weight normalization (Salimans et al., 2016). We divide the network into a backbone h (upto ResLayer-4 followed by Global Average Pooling) and the classifier f_c , as defined in Sec. 3. First, the network is initialized from a pre-trained model on the source dataset \mathcal{D}_s with random sub-domain augmentations. During vendor-side and client-side training, the mixup operation is performed between the random augmentation feature z extracted by the backbone and the corresponding domain generic feature z_g obtained by averaging over the sub-domain augmentation features as demonstrated in Fig. 4B. The obtained mixup feature z_m is then passed through the classifier f_c and the loss computation is done as per the predefined $VsAlgo(\cdot)$ and $CsAlgo(\cdot)$.

Edge-mixup training. In contrast, the edge-mixup optimization algorithm is much simpler. Here, the RGB image x is directly mixed-up with the edge representation of the image *i.e.* $\mathcal{A}_e(x)$ extracted using an edge estimation model (Soria et al., 2020), with the mixup ratio λ . The mixup image with enhanced edge features is directly used as input for the vendor-side and client-side training.

C.3. Additional results

Evaluation on other modalities like text and speech. Our key theoretical insights build on data-type-agnostic DA theory (Ben-David et al., 2006) and are useful for any ML data. While *contours* are vision-specific, similar generic-domains can be formed using *task-specific* knowledge (e.g.

via FFT for speech). Moreover, our feature-mixup relies on **a**) deeper layer feature space and **b**) augmentations which are prevalent in DL for any data-type and task. We demonstrate the gains of feature-mixup on audio DA and text DA benchmarks (Table 10) using open source audio¹ and speech² augmentations.

Table 10. Evaluation on Libri-Adapt benchmark (Mathur et al., 2020) and Amazon reviews benchmark (Blitzer et al., 2007).

DA for Speech Recog. (audio) on Libri-Adapt dataset		DA for Sentiment Clsf. (text) on Amazon Reviews dataset	
CMatch (Hou et al., 2021)	71.8	UDALM (Karouzos et al., 2021)	91.7
+ feat-mixup	73.5	+ feat-mixup	93.1

C.4. Extended comparisons

We show extended comparisons with more prior arts due to lack of space in the main paper, but the main results are the same. Table 11 shows SSDA for Office-Home (Table 3 in the main paper). Table 12 shows SSDA for Office-31 and VisDA (Table 1 in the main paper). Finally, Table 13 shows MSDA on DomainNet and Office-Home (Table 2 in the main paper).

Table 11. Single-Source Domain Adaptation (SSDA) on Office-Home benchmarks. SF indicates *source-free* adaptation.

Method	SF	Office-Home												Avg
		Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	
SRDC (Tang et al., 2020)	✗	52.3	76.3	81.0	69.5	76.2	78.0	68.7	53.8	81.7	76.3	57.1	85.0	71.3
SENTRY (Prabhu et al., 2021)	✗	61.8	77.4	80.1	66.3	71.6	74.7	66.8	63.0	80.9	74.0	66.3	84.1	72.2
TCM (Yue et al., 2021)	✗	58.6	74.4	79.6	64.5	74.0	75.1	64.6	56.2	80.9	74.6	60.7	84.7	70.7
CDS (Kim et al., 2021a)	✗	55.9	74.6	77.7	65.5	73.3	75.0	67.8	54.5	79.5	73.7	59.3	82.4	69.9
SCDA (Li et al., 2021a)	✗	60.7	76.4	82.8	69.8	77.5	78.4	68.9	59.0	82.7	74.9	61.8	84.5	73.1
SHOT (Liang et al., 2020)	✓	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
A ² Net (Xia et al., 2021)	✓	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
GSFDA (Yang et al., 2021b)	✓	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
CPGA (Qiu et al., 2021)	✓	59.3	78.1	79.8	65.4	75.5	76.4	65.7	58.0	81.0	72.0	64.4	83.3	71.6
SFDA (Kim et al., 2021c)	✓	48.4	73.4	76.9	64.3	69.8	71.7	62.7	45.3	76.6	69.8	50.5	79.0	65.7
NRC (Yang et al., 2021a)	✓	57.7	80.3	82.0	68.1	79.8	78.6	65.3	56.4	83.0	71.0	58.6	85.6	72.2
Ours (edge-mixup) + NRC	✓	60.8	80.1	81.6	67.2	79.3	78.5	65.4	61.0	83.8	70.2	63.1	85.3	73.0
Ours (feat-mixup) + NRC	✓	61.6	80.9	82.5	68.1	80.1	79.1	66.0	61.8	84.5	71.2	63.7	86.1	73.8
SHOT++ (Liang et al., 2021)	✓	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
Ours (edge-mixup) + SHOT++	✓	61.0	80.4	82.1	67.6	79.8	78.8	67.1	60.7	84.3	73.0	63.5	85.7	73.7
Ours (feat-mixup) + SHOT++	✓	61.8	81.2	83.0	68.5	80.6	79.4	67.8	61.5	85.1	73.7	64.1	86.5	74.5

¹<https://github.com/iver56/audiomentations>

²<https://github.com/makcedward/nlpaug>

Balancing Discriminability and Transferability for Source-Free Domain Adaptation

Table 12. Single-Source Domain Adaptation (SSDA) on Office-31 and VisDA benchmarks. SF indicates *source-free* adaptation.

Method	SF	Office-31						Avg	VisDA
		A→D	A→W	D→W	W→D	D→A	W→A		S → R
FixBi (Na et al., 2021)	✗	95.0	96.1	99.3	100.0	78.7	79.4	91.4	87.2
FAA (Huang et al., 2021a)	✗	94.4	92.3	99.2	99.7	80.5	78.7	90.8	82.7
CDAN+RADA (Jin et al., 2021)	✗	96.1	96.2	99.3	100.0	77.5	77.4	91.1	76.3
RFA (Awais et al., 2021)	✗	93.0	92.8	99.1	100.0	78.0	77.7	90.2	79.4
SCDA (Li et al., 2021a)	✗	95.4	95.3	99.0	100.0	77.2	75.9	90.5	-
SHOT (Liang et al., 2020)	✓	94.0	90.1	98.4	99.9	74.7	74.3	88.6	82.9
3C-GAN (Li et al., 2020)	✓	92.7	93.7	98.5	99.8	75.3	77.8	89.6	81.6
CPGA (Qiu et al., 2021)	✓	94.4	94.1	98.4	99.8	76.0	76.6	89.9	84.1
HCL (Huang et al., 2021b)	✓	90.8	91.3	98.2	100.0	72.7	72.7	87.6	83.5
SFDA (Kim et al., 2021c)	✓	92.2	91.1	98.2	99.5	71.0	71.2	87.2	76.7
VDM-DA (Tian et al., 2021)	✓	93.2	94.1	98.0	100.0	75.8	77.1	89.7	85.1
A ² Net (Xia et al., 2021)	✓	94.5	94.0	99.2	100.0	76.7	76.1	90.1	84.3
NRC (Yang et al., 2021a)	✓	96.0	90.8	99.0	100.0	75.3	75.0	89.4	85.9
<i>Ours (edge-mixup) + NRC</i>	✓	96.1	92.4	99.2	100.0	76.9	77.1	90.3	86.4
<i>Ours (feat-mixup) + NRC</i>	✓	96.3	92.8	99.2	100.0	77.4	77.5	90.5	87.3
SHOT++ (Liang et al., 2021)	✓	94.3	90.4	98.7	99.9	76.2	75.8	89.2	87.3
<i>Ours (edge-mixup) + SHOT++</i>	✓	94.4	92.0	98.9	100.0	77.8	77.9	90.2	87.5
<i>Ours (feat-mixup) + SHOT++</i>	✓	94.6	93.2	98.9	100.0	78.3	78.9	90.7	87.8

Table 13. Multi-Source Domain Adaptation (MSDA) on DomainNet and Office-Home. SF indicates *source-free* adaptation. The middle section compares SSDA works by adapting each source-target pair and reporting the best and worst results for each target.

Method	SF	w/o Domain Labels	DomainNet							Office-Home				
			→C	→I	→P	→Q	→R	→S	Avg	→Ar	→Cl	→Pr	→Rw	Avg
MDAN (Zhao et al., 2018)	✗	✗	52.4	21.3	46.9	8.6	54.9	46.5	38.4	68.1	67.0	81.0	82.8	74.7
SImpAl ₅₀ (Venkat et al., 2020)	✗	✗	66.4	26.5	56.6	18.9	68.0	55.5	48.6	70.8	56.3	80.2	81.5	72.2
CMSDA (Scalbert et al., 2021)	✗	✗	70.9	26.5	57.5	21.3	68.1	59.4	50.4	71.5	67.7	84.1	82.9	76.6
DARN (Wen et al., 2020)	✗	✗	-	-	-	-	-	-	-	70.0	68.4	82.7	83.9	76.3
DRT (Li et al., 2021b)	✗	✗	71.0	31.6	61.0	12.3	71.4	60.7	51.3	-	-	-	-	-
STEM (Nguyen et al., 2021)	✗	✗	72.0	28.2	61.5	25.7	72.6	60.2	53.4	-	-	-	-	-
SHOT(Liang et al., 2020)-worst	✓	✗	14.8	1.0	3.5	3.8	6.6	11.9	7.0	66.6	53.8	77.9	80.8	69.8
SHOT(Liang et al., 2020)-best	✓	✗	58.3	22.7	53.0	18.7	65.9	48.4	44.5	72.1	57.2	83.4	81.3	73.5
<i>Ours-worst</i>	✓	✗	17.0	2.4	8.7	4.9	12.1	18.1	10.5	66.0	61.6	80.1	79.1	71.7
<i>Ours-best</i>	✓	✗	68.6	21.9	53.8	17.4	65.4	58.1	47.5	71.2	63.7	86.1	84.5	76.4
Source-combine	✗	✓	57.0	23.4	54.1	14.6	67.2	50.3	44.4	58.0	57.3	74.2	77.9	66.9
SHOT (Liang et al., 2020)-Ens	✓	✗	58.6	25.2	55.3	15.3	70.5	52.4	46.2	72.2	59.3	82.8	82.9	74.3
DECISION (Ahmed et al., 2021)	✓	✗	61.5	21.6	54.6	18.9	67.5	51.0	45.9	74.5	59.4	84.4	83.6	75.5
NEL (Ahmed et al., 2022)	✓	✗	68.3	22.1	54.7	22.8	67.3	57.1	48.7	-	-	-	-	-
CAiDA (Dong et al., 2021)	✓	✗	-	-	-	-	-	-	-	75.2	60.5	84.7	84.2	76.2
<i>Ours (edge-mixup) + NRC</i>	✓	✓	74.8	24.1	56.8	19.6	66.9	55.5	49.6	72.1	62.9	86.4	84.8	76.6
<i>Ours (feature-mixup) + NRC</i>	✓	✓	75.4	24.6	57.8	23.6	65.8	58.5	51.0	72.6	67.4	85.9	83.6	77.4

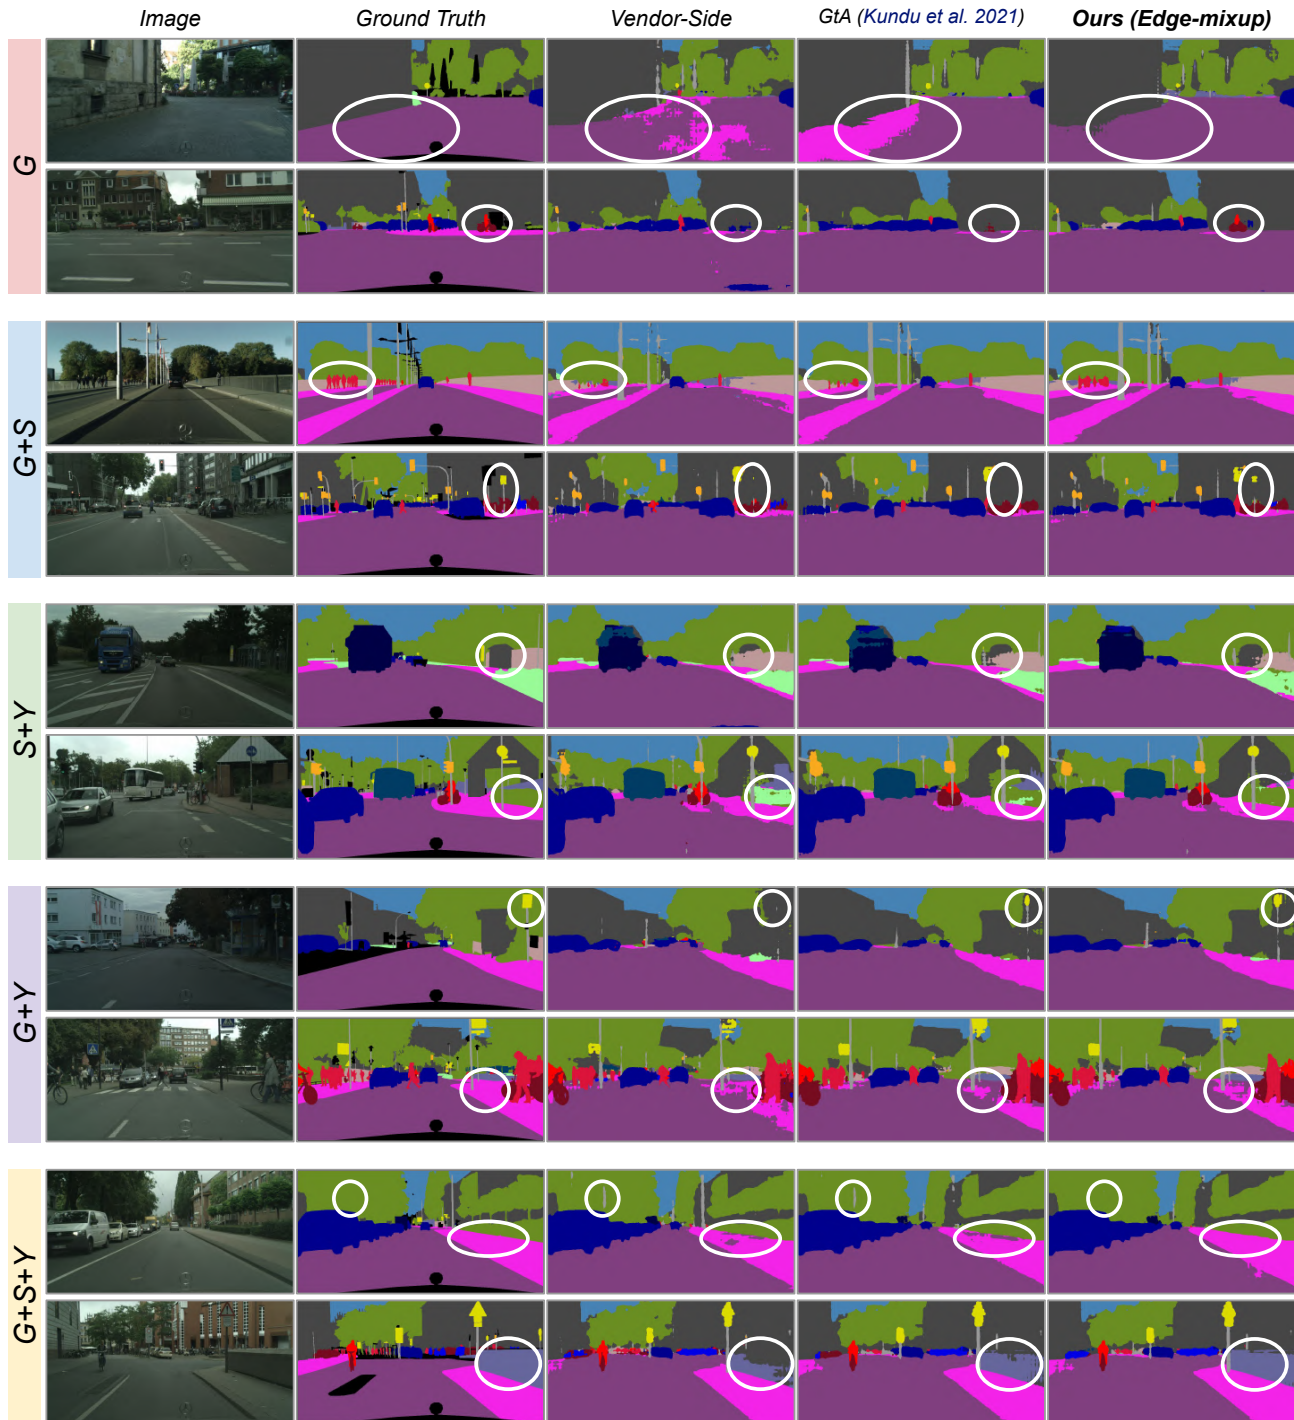


Figure 6. Qualitative analysis on target Cityscapes validation set for semantic segmentation SSDA (G) and MSDA (G+S, S+Y, G+Y, G+S+Y) settings. White circles indicate areas of improvement w.r.t. prior art and vendor-side baseline. *Best viewed in color.*