

# OPTIMAL TRANSPORT MEETS NOISY LABEL ROBUST LOSS AND MIXUP REGULARIZATION FOR DOMAIN ADAPTATION

**Kilian Fatras** \*

Mila - Quebec AI Institute, McGill University  
Montréal, Québec, Canada

**Hiroki Naganuma & Ioannis Mitliagkas**

Mila - Quebec AI Institute, Université de Montréal  
Montréal, Québec, Canada

## ABSTRACT

It is common in computer vision to be confronted with domain shift: images which have the same class but different acquisition conditions. In domain adaptation (DA), one wants to classify unlabeled target images using source labeled images. Unfortunately, deep neural networks trained on a source training set perform poorly on target images which do not belong to the training domain. One strategy to improve these performances is to align the source and target image distributions in an embedded space using optimal transport (OT). However OT can cause negative transfer, *i.e.* aligning samples with different labels, which leads to overfitting especially in the presence of label shift between domains. In this work, we mitigate negative alignment by explaining it as a noisy label assignment to target images. We then mitigate its effect by appropriate regularization. We propose to couple the MixUp regularization (Zhang et al., 2018a) with a loss that is robust to noisy labels in order to improve domain adaptation performance. We show in an extensive ablation study that a combination of the two techniques is critical to achieve improved performance. Finally, we evaluate our method, called MIXUNBOT, on several benchmarks and real-world DA problems.

## 1 INTRODUCTION

Deep neural networks have reached state-of-the-art performance on classification problems due to their ability to fit complex dataset while also generalizing within a specific domain (He et al., 2016). However, in applications like computer vision, it is standard to have same-class samples coming from different domains, for instance when they have different backgrounds or colorspace. Unfortunately, the generalization of deep neural networks across different domains is poor, and the object of intense research. Diversifying the domains with new samples in the training dataset is also a challenging task. For instance in medicine, collecting and annotating data is time-consuming and prone to errors. The problem of *domain adaptation*, tackles the case where we have access to two domains sharing the same classes where one has labeled data, called the source domain, and the other has unlabeled data, called the target domain. The purpose of domain adaptation problems is to classify the unlabeled target samples using the labeled source samples (Pan & Yang, 2010; Patel et al., 2015). A popular and efficient method to solve this problem is to use the *alignment strategy*, where same-class samples from different domains are aligned in an embedded space. To align the embedded source and target samples, several techniques exist such as adversarial training (Ganin et al., 2016; Long et al., 2018; Chen et al., 2020a) or optimal transport (Courty et al., 2017; Courty et al., 2017; Redko et al., 2017).

Optimal Transport (Peyré & Cuturi, 2019) has become one of the most used methods to compare probability distributions in machine learning. It has been popular for tasks such as generative models (Arjovsky et al., 2017; Genevay et al., 2018; Salimans et al., 2018; Burnel et al., 2021) or supervised learning problems (Frogner et al., 2015; Fatras et al., 2021a). In domain adaptation, it has been used to transport the source domain to the target domains (Courty et al., 2017), or to align the domains in a joint space of data and labels (Courty et al., 2017; Bhushan Damodaran et al., 2018). To reduce the OT cost, Bhushan Damodaran et al. (2018) used a minibatch approximation of OT (Fatras et al., 2020), which led to non-optimal connections between domains due to the minibatch samplings and the marginal constraints of exact OT. This phenomenon corresponds to negative transfer in domain adaptation problems. To mitigate the non-optimal connections of minibatch OT, DEEPJDOT (Bhushan Damodaran et al., 2018) required large batch sizes—otherwise small batch sizes lead deep neural networks to overfit. Another workaround to mitigate the non-optimal connections was to use Unbalanced Optimal Transport (UOT), an OT variant with relaxed marginals, as proposed in Fatras et al. (2021b).

In this paper, we explain non-optimal connections as assigning noisy labels to target samples. To avoid overfitting from noisy labels, we propose to regularize the neural networks and to use a noisy label robust loss in the transfer

\* corresponding author: kilian.fatras@mila.quebec

term. We propose to couple two techniques: *i*) regularize the neural networks using the MixUp regularization (Zhang et al., 2018a) on both source and target domains, MixUp interpolates samples from the same distribution uniformly at random as well as their label when available; *ii*) Use the symmetric cross-entropy loss in the transfer term, which has been proven to be robust to label noise (Wang et al., 2019). Our findings show that it is the combination of the MixUp and symmetric cross-entropy which leads to an increase in the performances of models as shown in an extensive ablation study, while the use of the techniques separately does not lead to any increase. It can be used for different optimal transport loss as we show in our experiments.

This paper is structured as follows. In Section 2, we review the different methods to solve domain adaptation and we define optimal transport as well as its use in domain adaptation problems. In Section 3, we present our method MIXUNBOT as well as the components it is built upon such as DEEPIEDOT, MixUp, and the symmetric cross-entropy (SCE) loss. In Section 4, we present extensive domain adaptation and partial domain adaptation experiments.

## 2 RELATED WORK ON DOMAIN ADAPTATION AND OPTIMAL TRANSPORT

In this section, we review the basic strategies for solving domain adaptation problems. Then we define formally optimal transport and we discuss its empirical implementation in deep learning applications, its current limitations and its use in domain adaptation.

**Unsupervised domain adaptation and notations** We recall that in unsupervised domain adaptation, we want to classify unlabeled target data using labeled source data (Pan & Yang, 2010). Formally, let  $\mathcal{D}_s$  (resp.  $\mathcal{D}_t$ ) be the labeled source (resp. target) dataset which is composed of  $n$  *i.i.d* random labeled (resp. unlabeled) vectors in  $\mathbb{R}^d$  drawn from a distribution  $\mu$  (resp.  $\nu$ ), *i.e.*,  $\mathcal{D}_s = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^n, \mathbf{x}_i^s \in \mathbb{R}^d$  (resp.  $\mathcal{D}_t = \{(\mathbf{x}_j^t)\}_{j=1}^n, \mathbf{x}_j^t \in \mathbb{R}^d$ ). We denote the tuple of source (resp. target) vectors as  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  (resp.  $\mathbf{Z} = (z_1, \dots, z_n)$ ) and the product measure  $\mu^{\otimes m}$  denotes a sample of  $m$  independent random variables following the distribution  $\mu$ . When a map  $g$  is applied to a tuple  $\mathbf{X}$ , we have  $g(\mathbf{X}) = \{g(\mathbf{x}_1), \dots, g(\mathbf{x}_n)\}$ . In vanilla domain adaptation, both domains share the same label space  $\mathcal{Y}_s = \mathcal{Y}_t$ . Other variants can be considered such as partial domain adaptation (Cao et al., 2018b), where some source classes are not in the target domain  $\mathcal{Y}_t \subset \mathcal{Y}_s$  and open set domain adaptation (Busto & Gall, 2017), where some target classes are not in the source domain  $\mathcal{Y}_s \subset \mathcal{Y}_t$ . In this work, we focus on vanilla and partial domain adaptation.

Several strategies exist to classify the target samples using deep neural networks, and they can be divided into three groups. The first group introduces a domain discriminator, which predicts if a sample is from the source or the target domain, and it is trained in an adversarial training manner as invented in Goodfellow et al. (2014). The adversarial training encourages domain confusion in order to fool the discriminator so that it can not decide which domain the samples come from Ganin et al. (2016). Several variants of the initial idea were published, for instance Long et al. (2018); Chen et al. (2020a). Another group of methods is the reconstruction of both source and target samples. In Ghifary et al. (2016), authors learn an embedding where source samples are classified, and target samples are encoded, then the learned representation not only preserves the ability to discriminate but also encodes useful information from the target domain. The final group of methods aims at minimizing the domain discrepancy by minimizing different distances between probability distributions like maximum mean discrepancy (MMD) (Long et al., 2015; 2017) or optimal transport (Courty et al., 2017; Courty et al., 2017; Bhushan Damodaran et al., 2018; Fatras et al., 2021b). Our methods are built upon the works on optimal transport, and its definition is the topic of the next paragraph.

**Optimal transport** The optimal transport cost measures the minimal displacement cost of moving a probability distribution  $\mu$  to another probability distribution  $\nu$  with respect to a ground metric  $c$  on the data space  $\mathcal{X}$  and  $\mathcal{Z}$  (Peyré & Cuturi, 2019). Formally, let  $(\mu, \nu) \in \mathcal{M}_+^1(\mathcal{X}) \times \mathcal{M}_+^1(\mathcal{Z})$ , where  $\mathcal{M}_+^1(\mathcal{X})$  (resp.  $\mathcal{M}_+^1(\mathcal{Z})$ ) is the set of probability measures on  $\mathcal{X}$  (resp.  $\mathcal{Z}$ ). For a ground cost  $c : \mathcal{X} \times \mathcal{Z} \mapsto \mathbb{R}_+$ , the optimal transport cost between the two distributions  $\mu, \nu$  is

$$W_c(\mu, \nu) = \min_{\pi \in \mathcal{U}(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Z}} c(\mathbf{x}, \mathbf{z}) d\pi(\mathbf{x}, \mathbf{z}), \quad (1)$$

where  $\mathcal{U}(\mu, \nu)$  is the set of joint probability distributions with marginals  $\mu$  and  $\nu$  such that  $\mathcal{U}(\mu, \nu) = \{\pi \in \mathcal{M}_+^1(\mathcal{X}, \mathcal{Z}) : \mathbf{P}_{\mathcal{X}} \# \pi = \mu, \mathbf{P}_{\mathcal{Z}} \# \pi = \nu\}$ .  $\mathbf{P}_{\mathcal{X}} \# \pi$  (resp.  $\mathbf{P}_{\mathcal{Z}} \# \pi$ ) is the marginalization of  $\pi$  over  $\mathcal{X}$  (resp.  $\mathcal{Z}$ ). These constraints enforce that all the mass from  $\mu$  is transported to  $\nu$  and vice-versa. This optimization problem is called the Kantorovitch formulation of OT and the optimal  $\pi$  is called an optimal transport plan. The ground cost  $c$  is usually chosen as the Euclidean or squared Euclidean distance on  $\mathbb{R}^d$ . In the case where  $c$  is a distance,  $W_c$  is a distance between probability distributions called the Wasserstein distance. When the distributions are discrete, the problem becomes a discrete linear program that can be solved with a cubical complexity in the size of the distributions support, which can be prohibitive for applications with large datasets.

That is why the optimal transport cost has been computed in practice between random minibatches of input distributions (Genevay et al., 2018; Bhushan Damodaran et al., 2018). The minibatch approximation of optimal transport has been theoretically investigated as a loss function in (Fratras et al., 2020; 2021c) and its deviation from exact OT in (Sommerfeld et al., 2019). Instead of solving the optimal transport cost between the full input distributions, the minibatch approximation computes the expectation of optimal transport cost between minibatches drawn uniformly at random. While it defines a transport problem in the sense that all the mass of input distributions are transported, Fratras et al. (2020) nonetheless showed that minibatch OT acts like a regularization on the transportation plan, leading to a denser plan. Those are due to the non-optimal connections from the random sampling of minibatches and the marginal constraints in the OT definition. They also showed that the number of connections increases when the minibatch size decreases. The consequences are that minibatch OT is not a distance but has some appealing statistical and optimization properties. Despite these weaknesses, optimal transport and its minibatch approximation have found numerous applications to compare probability distributions from different domains as in domain adaptation.

**Optimal transport for domain adaptation** In domain adaptation, optimal transport has been used to transport the source domain to the target domain with theoretical guarantees (Redko et al., 2017). In Courty et al. (2017), source samples were transported with their labels to the target domain and a classifier was then trained on the transported labeled samples. It was achieved by using the optimal transport plan  $\pi$  as a barycentric mapping between the source and the target distributions. A more recent strategy was to align the joint distributions of features and labels of the different domains (Courty et al., 2017; Bhushan Damodaran et al., 2018) instead of only considering the feature distributions. To learn the joint space, OT based techniques used a ground cost which incorporates a term on the (embedded) sample space as well as a term on the label space. This OT cost enables one to learn a mapping which tries to assign labels to samples which have close embeddings. However, OT can make connections between samples from different classes notably in the case of label shift, when there is a different proportion of labels between domains (Redko et al., 2019). This problem is exacerbated when one uses the minibatch OT approximation to reduce the OT complexity as done in Bhushan Damodaran et al. (2018), where the non-optimal connections between distributions from the minibatch approximation led to even more negative transfer in domain adaptation problems. Indeed by sampling samples uniformly at random, it is likely to draw minibatches which have a shift in the labels especially when there is a shift in the labels between the full distributions. For instance, drawing a source minibatch full of images from one class and a target minibatch full of samples from another class. Thus, the DEEPPDOT method requires large batch sizes in order to be performing and this is inconvenient as the purpose of using minibatches is to decrease the OT cubical complexity. In order to mitigate the effect of these connections, (Fratras et al., 2021b) proposed to rely on unbalanced optimal transport at the minibatch level that we formally define. Consider  $D_{KL}$  to be the Kullback-Leibler divergence, we define the unbalanced OT cost between two probability distributions  $\mu, \nu$  as:

$$\text{UOT}_c^\tau(\mu, \nu) = \min_{\pi \geq 0} \int_{\mathcal{X} \times \mathcal{Z}} c(\mathbf{x}, \mathbf{z}) d\pi(\mathbf{x}, \mathbf{z}) + \tau(D_{KL}(\pi|\mu) + D_{KL}(\pi|\nu)). \quad (2)$$

It is an OT variant with relaxed marginal constraints which does not transport the mass when it is costly to transport, at the cost of choosing well one extra hyper-parameter, thus decreasing the non-optimal connections from the shift in the labels between domains and the minibatch sampling. However, it does not avoid all non-optimal connections. In the next section, we present our method which aims at regularizing the neural networks in order to mitigate the effect of the non-optimal connections.

### 3 PROPOSED METHOD

We are now ready to present our proposed approach to solve the domain adaptation problem. We start by introducing the different methods which are at the heart of our approach.

#### 3.1 BACKGROUND

**DEEPPDOT** Our approach is built upon DEEPPDOT (Bhushan Damodaran et al., 2018) and JUMBOT (Fratras et al., 2021b). As described in the previous section, DEEPPDOT aims at finding a joint distribution map between the source and the target distributions. It does so thanks to the ground cost used in the optimal transport problem which incorporates a term on a neural network embedding space and a term on the label space between domains. Formally, let  $g_\theta$  be a feature extractor where the input is mapped from the space of images to a latent space and  $f_\phi$  which maps the latent space to the label space. The embedding space in our experiment is the penultimate layer of a neural network. For two

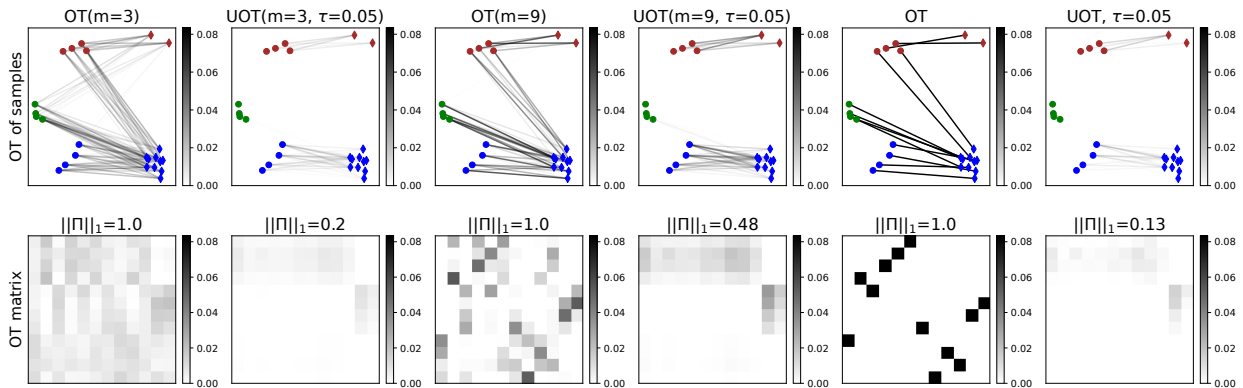


Figure 1. Overview of OT plans between 2D distributions with 12 samples. The distributions have a cluster-like structure, where the source distribution has three (3) balanced clusters and the target distribution has only two (2) imbalanced clusters which are in the source domain. The first row shows the transported mass as connections between samples for different batch size, the second row illustrates the minibatch OT and UOT plans  $\Pi$  as well as their mass.

probability distributions  $\mu, \nu$ , feature extractor  $g_\theta$  and classification map  $f_\phi$ , the transfer term is:

$$\mathbb{E}_{(\mathbf{X}^s, \mathbf{Y}^s) \sim \mu^{\otimes m}, \mathbf{X}^t \sim \nu^{\otimes m}} h_{C_{\theta, \phi}}((g_\theta(\mathbf{X}^s), \mathbf{Y}^s), (g_\theta(\mathbf{X}^t), f_\phi(g_\theta(\mathbf{X}^t)))), \quad (3)$$

with  $(C_{\theta, \phi})_{i, j} = \eta_1 \|g_\theta(\mathbf{x}_i^s) - g_\theta(\mathbf{x}_j^t)\|_2^2 + \eta_2 \mathcal{L}(\mathbf{y}_i^s, f_\phi(g_\theta(\mathbf{x}_j^t)))$ ,

where  $h$  is the considered OT variant,  $\mathcal{L}(\cdot, \cdot)$  is the cross-entropy loss and  $\eta_1, \eta_2$  are positive constants. Note that we make an abuse of notation between a tuple of samples  $\mathbf{X}$  and its empirical distribution. We also add a cross-entropy term on the source data. Hence our optimization problem is:

$$\min_{\theta, \phi} \sum_i \mathcal{L}(f_\phi(g_\theta(\mathbf{x}_i^s)), \mathbf{y}_i^s) + \eta_3 \mathbb{E}_{(\mathbf{X}^s, \mathbf{Y}^s) \sim \mu^{\otimes m}, \mathbf{X}^t \sim \nu^{\otimes m}} h_{C_{\theta, \phi}}((\mathbf{X}^s, \mathbf{Y}^s), (\mathbf{X}^t, f_\phi(g_\theta(\mathbf{X}^t)))). \quad (4)$$

Intuitively, the OT plan would send mass between source and target samples which have close embeddings as well as close classifications. Thus close embeddings should have the same classification and samples which have the same classification should have close embeddings. However, due to the hard marginal constraints of exact OT, the shift in the labels which occurs between the source and the target domains leads to transport mass from the source to the target samples which have a different label. This is especially the case in partial domain adaptation where there is a serious shift in the labels (classes which are in the source domain but not in the target domain). Furthermore, the minibatch approximation of OT strengthens the non-optimal connections because of the sampling effect. To address this issue, [Fratras et al. \(2021b\)](#) developed JUMBOT which relied on the use of unbalanced OT, which mitigates the transfer of mass when a sample is costly to transport. The only difference between the methods JUMBOT and DEEPIJDOT is the use of entropically regularized unbalanced OT instead of exact OT. Another practical advantage of unbalanced OT is that it can be used to solve the partial domain adaptation problem ([Fratras et al., 2021b](#)).

Unfortunately, as discussed in [Fratras et al. \(2021b\)](#), there is still a transport of mass between samples of different labels. To illustrate this phenomenon, in Figure 1 we consider the partial domain adaptation problem in a 2D scenario with 12 samples. The source distribution has three classes and the target distribution has only two classes which are in the source domain. The classes have a cluster-like structure and are balanced in the source domain and imbalanced in the target domain (10 elements for the blue class and only 2 for the brown class). The OT cost and minibatch OT cost create a lot of connections between samples from different classes due to marginal constraints and minibatch sampling for the latter. However, the unbalanced OT based variants suffer less from these connections at the price of choosing a regularization coefficient  $\tau$ . Furthermore, for small  $\tau$  values, the mass of the transport plan can be small as shown in Figure 1, thus decreasing the transfer between domains. To mitigate this problem, [Fratras et al. \(2021b\)](#) used an extra parameter before the OT term to increase the strength of the transfer loss. To help the transfer between domains, we propose to regularize neural networks to mitigate the influence of non-optimal connections.

**MixUp regularization** is a training strategy based on a data augmentation technique proposed in [Zhang et al. \(2018a\)](#). It artificially augments the training dataset by interpolating samples randomly between them, as well as their labels when they are available, in order to create new training samples. As we work on a joint distribution,

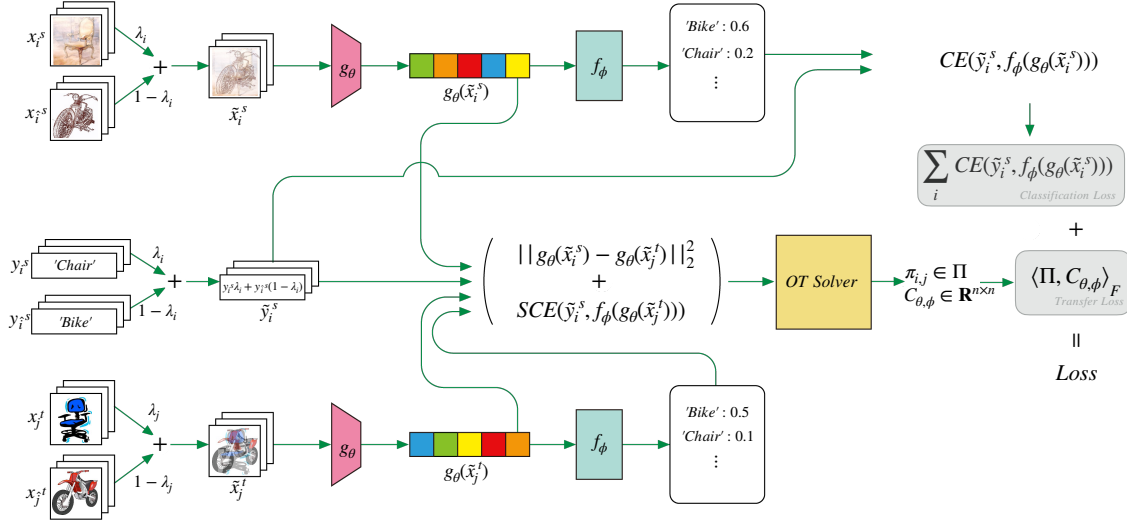


Figure 2. (Best viewed in color) Overview of the proposed MIXUNBOT method between full source and target distributions (we use a minibatch computation in practice). We denote the feature extractor  $g_\theta$  and the classifier  $f_\phi$ . The latent representations and labels are used to compute a ground cost matrix that is fed to an optimal transport solver. We then minimize a cross-entropy term on the source domain and an optimal transport alignment term between domains.

our new neighbouring source distribution considers the interpolated samples and their label while the neighbouring target distribution considers the interpolated sample and their prediction from the neural network. They are defined as follows:

$$\tilde{\mu} = \frac{1}{n} \sum_{i,j=0}^n \mathbb{E}_{\lambda \sim \beta(\alpha)} \delta_{(\lambda \cdot x_i^s + (1-\lambda) \cdot x_j^i, \lambda \cdot y_i^s + (1-\lambda) \cdot y_j^i)} = \mathbb{E}_{\lambda \sim \beta(\alpha)} \tilde{\mu}_\lambda, \quad (5)$$

$$\tilde{\nu} = \frac{1}{n} \sum_{i,j=0}^n \mathbb{E}_{\lambda \sim \beta(\alpha)} \delta_{(\lambda \cdot x_i^t + (1-\lambda) \cdot x_j^t, f(\lambda \cdot x_i^t + (1-\lambda) \cdot x_j^t))} = \mathbb{E}_{\lambda \sim \beta(\alpha)} \tilde{\nu}_\lambda. \quad (6)$$

Where  $\tilde{\mu}_\lambda$  (resp.  $\tilde{\nu}_\lambda$ ) is the measure associated to interpolated samples from  $\mu$  (resp.  $\nu$ ) for a fixed  $\lambda$  and  $n$  the number of samples in measures  $\mu$  and  $\nu$ . The parameter  $\lambda$  follows a beta distribution of parameter alpha  $\beta(\alpha)$ , where  $\alpha$  is set to 0.2 in our experiment. We study its impact on the performance in a sensitivity analysis (see Section 4.1). While it is a simple technique, it has allowed neural networks to reach state-of-the-art results on classification tasks while being more robust to adversarial examples (Zhang et al., 2021). Moreover, MixUp regularization acts like four different regularizations: perturbed data, temperature scaling, label smoothing, and Jacobian regularization (Carratino et al., 2020). MixUp also finds adaptively the good parameters to make the four regularization work together.

MixUp has been previously used in domain adaptation (Mao et al., 2019), however authors coupled it with the method DIRT-T (Shu et al., 2018) and the case where it was coupled with optimal transport has not been studied. MixUp data augmentation was also applied between samples from different domains in Xu et al. (2020). It was also used in variants of domain adaptation as semi-supervised DA (Yang et al., 2021).

In this work, we propose to compute the OT cost between the neighbouring probability distributions  $\tilde{\mu}$  and  $\tilde{\nu}$  to inherit the benefits from the MixUp regularization on both domains. However at each iteration, we compute the OT cost between minibatches from  $\tilde{\mu}_\lambda$  and  $\tilde{\nu}_\lambda$  for a fixed  $\lambda \sim \beta(\alpha)$ , as done in Zhang et al. (2018a). While it works well in practice, this implementation leads us to compute an upper bound of  $W(\tilde{\mu}, \tilde{\nu})$  as shown in the following Proposition:

**Proposition 1** *Let  $\mu, \nu$  be two probability distributions and  $\tilde{\mu}, \tilde{\nu}$  their neighbouring respective distributions. We have the following upper bound for the Wasserstein distance:*

$$W(\tilde{\mu}, \tilde{\nu}) \leq \mathbb{E}_{\lambda, \lambda'} W(\tilde{\mu}_\lambda, \tilde{\nu}_{\lambda'}). \quad (7)$$

The proof is straightforward thanks to the convexity of the Wasserstein distance (Peyré & Cuturi, 2019, Section 9.1) and Jensen inequality. In this paper, we empirically show that adding MixUp to DEEPJDOT only improves slightly the



performances of DEEPJDOT contrary to the findings from Mao et al. (2019), where MixUp really increased the results of DIRT-T. While the use of the MixUp data augmentation leads to the use of different regularizations, we also want to make the neural network robust to non-optimal connections which assign a noisy label to a target image.

**Label noise.** The cross-entropy loss has been used to get state-of-the-art classifiers on many classification tasks such as ImageNet (Krizhevsky et al., 2012). However, it also comes with downsides. Using cross-entropy loss gives a huge ability to neural networks to memorize the label which leads to poor generalization properties in the case of noisy labels (Zhang et al., 2017). In the case of domain adaptation, minibatch OT and label shift create connections between samples belonging to different classes. Going back to our example where a source minibatch is composed of one class and the target minibatch of another, DEEPJDOT would transport all images from the source to the target domain even if they have different class. The loss on the classifier would be then:

$$\sum_j \text{CE}(y_1, f(\mathbf{x}_j)),$$

where  $y_1$  is the label of images from the source minibatch. This forces the classifier to give noisy labels to the target samples and can be seen as learning with noisy labels. To overcome this issue, we replace the cross-entropy loss used in (Bhushan Damodaran et al., 2018) by the symmetric cross-entropy loss developed in (Wang et al., 2019). This has the effect to make the classifier more robust to label noise and to non-optimal connections between source and target samples. For two probability vectors  $q$  and  $q'$ , we define the symmetric cross-entropy loss as:

$$\text{SCE}(q, q') = \eta_4 \text{CE}(q, q') + \text{CE}(q', q). \quad (8)$$

In all our experiments we have  $\eta_4 = 0.01$  to reduce the tuning of hyper-parameters. After describing the DEEPJDOT method, MixUp regularization and the SCE loss, we have all the different ingredients to present our new methods.

### 3.2 PROPOSED APPROACH MIXUNBOT

We now present our method called MIXUNBOT which is based on the methods described in the previous section. We propose to use the neighbouring distributions  $\tilde{\mu}, \tilde{\nu}$  in the OT cost from DEEPJDOT like algorithms, with minibatches drawn from  $\tilde{\mu}_\lambda^{\otimes m}$  and  $\tilde{\nu}_\lambda^{\otimes m}$ , instead of considering the initial probability distributions  $\mu$  and  $\nu$ . By doing so, we would inherit the regularization effect of MixUp regularization on the neural network. In addition to this change, we also use the symmetric cross-entropy loss in the ground cost, instead of the standard cross-entropy loss, to make the model robust to noisy labels from negative transfer. Thus our full loss function is:

$$\begin{aligned} \min_{\phi, \theta} \sum_i \mathcal{L}(f_\phi(g_\theta(\mathbf{x}_i^s)), \mathbf{y}_i^s) + \eta_3 \mathbb{E}_{\lambda \sim \beta(\alpha), (\mathbf{X}^s, \mathbf{Y}^s) \sim \tilde{\mu}_\lambda^{\otimes m}, \mathbf{X}^t \sim \tilde{\nu}_\lambda^{\otimes m}} h_{C_{\theta, \phi}}((\mathbf{X}^s, \mathbf{Y}^s), (\mathbf{X}^t, f_\phi(g_\theta(\mathbf{X}^t)))), \\ \text{with } (C_{\theta, \phi})_{i, j} = \eta_1 \|g_\theta(\mathbf{x}_i^s) - g_\theta(\mathbf{x}_j^t)\|_2^2 + \eta_2 \text{SCE}(\mathbf{y}_i^s, f_\phi(g_\theta(\mathbf{x}_j^t))). \end{aligned} \quad (9)$$

In the case where  $h$  is the exact optimal transport cost as in DEEPJDOT, we call our method MIXOT. When  $h$  is the unbalanced OT cost, we call our method MIXUNBOT. When connections are made between samples with different labels due to OT, our proposed regularization would help the neural network not to learn the label. We show that it is the combination of the two methods which leads to an improvement in the performances in an ablation study on domain adaptation (see Section 4.1). Our method is fully described in Figure 2. After presenting our method and the different related methods, we now evaluate it on different domain adaptation scenarios.

## 4 EXPERIMENTS

In this section, we evaluate our methods on several standard, real-world problems of domain adaptation and partial domain adaptation. For our experiments, we relied on the PyTorch deep learning library (Paszke et al., 2017) as well as the POT library for the optimal transport solvers (Flamary et al., 2021).

### 4.1 VANILLA DOMAIN ADAPTATION

In this paragraph, we evaluate and discuss the practical advantage of our proposed methods on vanilla domain adaptation experiments. In vanilla domain adaptation, the source and target domains share the same classes with a likely shift in the labels.

**Datasets.** We start the experimental section with **digits** datasets. We follow the evaluation protocol of (Bhushan Damodaran et al., 2018) on three domain adaptation scenarios based on digits datasets: USPS to MNIST

	Methods/DA	U $\mapsto$ M	M $\mapsto$ MM	S $\mapsto$ M	Avg
Exp.	DANN	93.2 $\pm$ 0.2	91.7 $\pm$ 6.5	89.2 $\pm$ 0.5	91.4
	CDAN-E	<b>99.2 <math>\pm</math> 0.1</b>	96.0 $\pm$ 1.5	95.6 $\pm$ 1.8	96.9
	ALDA	97.9 $\pm$ 1.3	<b>96.7 <math>\pm</math> 0.6</b>	95.3 $\pm$ 1.5	96.6
	DEEPPDOT	95.9 $\pm$ 0.1	81.2 $\pm$ 2.0	94.2 $\pm$ 0.4	90.3
	JUMBOT	98.3 $\pm$ 0.4	96.6 $\pm$ 0.4	98.8 $\pm$ 0.2	97.9
	MIXOT	97.6 $\pm$ 0.1	86.2 $\pm$ 5.0	98.2 $\pm$ 0.1	94.0
	MIXUNBOT	98.8 $\pm$ 0.2	96.4 $\pm$ 0.6	<b>99.1 <math>\pm</math> 0.1</b>	<b>98.1</b>
Ablation	DEEPPDOT	95.9 $\pm$ 0.1	-	94.2 $\pm$ 0.4	95.1
	DEEPPDOT (RCE)	95.9 $\pm$ 0.1	-	94.6 $\pm$ 0.2	95.3
	MIXOT (CE)	96.2 $\pm$ 0.1	-	94.6 $\pm$ 0.2	95.4
	MIXOT	<b>97.6 <math>\pm</math> 0.1</b>	-	<b>98.2 <math>\pm</math> 0.1</b>	<b>97.9</b>

Methods	Accuracy (in %)
CDAN-E	70.1
ALDA	70.5
DEEPPDOT	68.0
ROBUST OT	66.3
JUMBOT	70.2
MIXUNBOT (Ours)	<b>71.3</b>

Table 1: DA results on digit and Visda-2017 datasets. (Left table) The Above part sums up experiments on DA digits. Each experiment was run three times with different seeds and the best score is reported in **bold**. Lower parts gather an ablation study of method over two DA tasks on digits datasets. (Right table) Summary table of DA results on Visda-2017 datasets (ResNet-50). All results are reported at the end of training.

	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
DA	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	44.3	59.8	69.8	48.0	58.3	63.0	49.7	42.7	70.6	64.0	51.7	78.3	58.3
	CDAN-E	52.5	71.4	76.1	59.7	69.9	71.5	58.7	50.3	77.5	70.5	57.9	83.5	66.6
	DEEPPDOT	50.7	68.6	74.4	59.9	65.8	68.1	55.2	46.3	73.8	66.0	54.9	78.3	63.5
	ALDA	52.2	69.3	76.4	58.7	68.2	71.1	57.4	49.6	76.8	70.6	57.3	82.5	65.8
	ROT	47.2	71.8	76.4	58.6	68.1	70.2	56.5	45.0	75.8	69.4	52.1	80.6	64.3
	JUMBOT	55.2	75.5	80.8	65.5	74.4	74.9	65.2	52.7	79.2	73.0	59.9	83.4	70.0
	MIXOT (ours)	<b>56.4</b>	<b>76.8</b>	<b>81.5</b>	<b>66.4</b>	<b>76.2</b>	<b>77.0</b>	<b>64.0</b>	<b>53.5</b>	<b>81.3</b>	<b>72.9</b>	<b>58.6</b>	<b>84.4</b>	<b>70.8</b>
	MIXUNBOT (ours)	<b>56.0</b>	<b>77.5</b>	<b>81.4</b>	<b>68.4</b>	<b>74.4</b>	<b>76.7</b>	<b>66.7</b>	<b>54.2</b>	<b>81.5</b>	<b>76.3</b>	<b>60.4</b>	<b>84.9</b>	<b>71.5</b>
	MIXOT (oracle)	56.5	76.9	81.7	66.8	76.2	77.0	65.3	53.9	81.5	73.5	58.9	84.4	71.1
	MIXUNBOT (oracle)	56.1	77.8	81.4	68.8	74.8	77.1	66.7	54.5	81.5	76.6	60.4	85.0	71.7
PDA	RESNET-50	46.3	67.5	75.9	59.1	59.9	62.7	58.2	41.8	74.9	67.4	48.2	74.2	61.4
	DEEPPDOT	48.2	66.2	76.6	56.1	57.8	64.5	58.3	42.7	73.5	65.7	48.2	73.7	60.9
	BA3US	56.7	76.0	<b>84.8</b>	73.9	67.8	<b>83.7</b>	72.7	56.5	84.9	77.8	64.5	83.8	73.6
	JUMBOT	<b>62.7</b>	<b>77.5</b>	84.4	<b>76.0</b>	<b>73.3</b>	80.5	<b>74.7</b>	<b>60.8</b>	<b>85.1</b>	<b>80.2</b>	<b>66.5</b>	<b>83.9</b>	<b>75.5</b>
	MIXUNBOT (ours)	<b>63.3</b>	<b>80.0</b>	<b>89.0</b>	<b>78.9</b>	<b>76.6</b>	<b>87.5</b>	<b>76.4</b>	<b>65.9</b>	<b>88.3</b>	<b>81.4</b>	<b>67.6</b>	<b>86.8</b>	<b>78.5</b>
PDA oracle	PADA	51.9	67.0	78.7	52.2	53.8	59.0	52.6	43.2	78.8	73.7	56.6	77.1	62.1
	IWAN	53.9	54.4	78.1	61.3	47.9	63.3	54.2	52.0	81.3	76.5	56.7	82.9	63.5
	SAN	44.4	68.7	74.6	67.5	65.0	77.8	59.8	44.7	80.1	72.2	50.2	78.7	65.3
	DRCN	54.0	76.4	83.0	62.1	64.5	71.0	70.8	49.8	80.5	77.5	59.1	79.9	69.0
	ETN	59.2	77.0	79.5	62.9	65.7	75.0	68.3	55.4	84.4	75.7	57.7	84.5	70.4
	SAFN	59.0	76.2	81.4	70.4	73.0	77.8	72.4	55.3	80.4	75.8	60.4	79.9	71.8
	RTNET	<b>63.2</b>	<b>80.1</b>	80.7	66.7	69.3	77.2	71.6	53.9	84.6	77.4	57.9	85.5	72.3
	JUMBOT	62.7	77.8	86.1	<b>76.1</b>	73.3	80.5	72.9	60.8	85.1	80.2	<b>66.5</b>	83.9	75.5
	BA3US	60.6	83.2	88.4	71.8	72.8	83.4	75.4	61.6	86.5	79.2	62.8	<b>86.0</b>	76.0
	AR	62.1	79.2	<b>89.1</b>	73.9	<b>75.6</b>	<b>84.4</b>	<b>78.4</b>	<b>61.9</b>	<b>87.9</b>	<b>82.2</b>	65.4	85.3	<b>77.1</b>
	MIXUNBOT (ours)	<b>63.3</b>	<b>80.0</b>	<b>89.1</b>	<b>78.9</b>	<b>76.7</b>	<b>87.5</b>	<b>76.4</b>	<b>65.9</b>	<b>88.3</b>	<b>81.5</b>	<b>67.8</b>	<b>86.8</b>	<b>78.6</b>

Table 2: (Best view in color) Summary table of domain adaptation and partial domain adaptation results on Office-Home dataset with a ResNet-50. Results are reported at the end of training except for PDA oracle, where best score along training is reported. **Best scores** are in **bold** and *second best scores* are in *blue*.

(U $\mapsto$ M), MNIST to MNIST-M (M $\mapsto$ MM), and SVHN to MNIST (S $\mapsto$ M). MNIST (LeCun & Cortes, 2010) has 60,000 images of handwritten digits, MNIST-M has the 60,000 MNIST images with color patches (Ganin et al., 2016) and USPS (Hull, 1994) has 7,291 digit images. Street View House Numbers (SVHN)(Netzer et al., 2011) is a dataset of 73, 257 images with digits and numbers in natural scenes. We report the evaluation results on the test target datasets at the end of training. **Office-Home** (Venkateswara et al., 2017) is a real world dataset for unsupervised domain adaptation (UDA) composed of 15,500 images from four different domains: Artistic (A), Clip Art (C), Product (P) and Real-World (R) images. For each domain, the dataset has images of 65 object categories that are common in office and home scenarios. We evaluate our methods in the 12 possible scenarios. Our last dataset is **VisDA-2017** (Peng et al., 2017). It is a large-scale dataset from simulation to real images. Visda-2017 contains 152,397 synthetic images as the source domain and 55,388 real-world images as the target domain. There are 12 different object categories in the two domains. Following (Long et al., 2018; Chen et al., 2020a), we evaluate all methods on Visda-2017 validation set.




Figure 3. (Best view in color) (Two left plots) Classification accuracy on classes along training for the SVHN to MNIST DA task of our method and DEEPPDOT. The performances of our method reach higher accuracies than DEEPPDOT and are more stable. (Two right plots) T-SNE embeddings of 1500 test samples from SVHN (source) and MNIST (target) for DEEPPDOT and our method MIXOT at test time. It shows the ability of the methods to discriminate classes (samples are colored w.r.t. their classes).

**Results.** We compare our MIXOT methods against state-of-the-art methods: DANN(Ganin et al., 2016), CDAN-E (Long et al., 2018), ALDA (Chen et al., 2020a), ROT (Balaji et al., 2020). We also consider the methods which are at the heart of MIXOT variants, namely DEEPPDOT (Bhushan Damodaran et al., 2018) and JUMBOT (Fatras et al., 2021b). We reported the accuracy on the testset at the end of training for all benchmarks and methods. We believe it is a more fair method than reporting the best results along training on the testset, because evaluating the method along training breaks the assumptions on labels on the target domain. However, on Office-Home, we reported the results with the highest prediction along training for our method to show that we are not prone to overfitting (see when oracle is mentioned). As our setting is identical to Fatras et al. (2021b), we reported their results for the different competitors on Office-Home and VisDA-2017. The full details of our training procedure and architectures can be found in appendix.

The results on digit datasets can be found in the left above part of Table 1. We conducted each experiment three times, with different seeds, and report their average results and variance. We only resize and normalize the image without data augmentation. We see that our method MIXOT performs 4% better than DEEPPDOT on average while MIXUNBOT performs slightly better than JUMBOT. Furthermore, we see an important 8% increase of the performance compared to DEEPPDOT. Regarding Office-Home experiments, the results are gathered in the above part of Table 2. For fair comparison with previous work, we used a similar data pre-processing, training procedure, architectures, hyperparameter selection on the AC task, and we used the ten-crop technique (Long et al., 2018; Chen et al., 2020a) for testing our methods. Our method MIXUNBOT achieves better performances on all considered tasks than all competitors and is on average 2% higher than the best competitor JUMBOT. It achieves the best performances on 10 tasks over 12 and is only beaten by our method MIXOT. Regarding MIXOT, it achieves a better accuracy than JUMBOT on average and on 10 of the 12 scenarios from the Office-Home dataset, furthermore, it is higher than DEEPPDOT by more than 7% showing the regularization effect of our strategy. When MIXOT does not achieve the best score, it is in second position behind MIXUNBOT. Finally, on Visda-2017 datasets, the results can be found in the right part of Table 1. MIXUNBOT performs better than all competitors, in particular, it is 1% above our reproduced JUMBOT score.

**Analysis and sensitivity study.** In this paragraph, we provide an empirical analysis of our method. We conducted an ablation study between DEEPPDOT and MIXOT as the difference between MIXUNBOT and MIXOT is the use of unbalanced OT instead of exact OT at the minibatch level. The ablation study can be found in the lower left part of Table 1. The difference between DEEPPDOT and MIXOT is first the use of the symmetric cross-entropy loss instead of cross-entropy loss between labels in the transfer loss. The second difference is the use of MixUp data augmentation on source and target samples. As we can see in the results, the use of the SCE loss in DEEPPDOT and the use of MixUp data with the standard cross-entropy loss do not improve results over DEEPPDOT. However, when the MixUp regularization is coupled with the SCE loss, we see a 1.7% increase on the USPS to MNIST task and a 3.5% increase on the SVHN to MNIST task, demonstrating the practical advantage of our approach. The ablation study shows that we can not take fully advantage of the MixUp regularization with noisy labels. We conjecture that the success of the combination is due to the SCE loss which prevents the neural network to overfit on noisy labels. It allows MixUp to regularize the neural network as it would have done in a clean label environment.

We also provide a classification accuracy along training as well as a T-SNE plot (van der Maaten & Hinton, 2008) on embedded source and target test samples between DEEPPDOT and MIXOT (see Figure 3). Regarding the classification accuracy along training, we see that our method reach higher accuracies on classes and that these accuracies are stable with a low variance. However for DEEPPDOT, we see that the training is not stable, with some variance and a small overfitting issue on some classes after 50 epochs. Moreover, the accuracy on some classes is lower than the ones






Figure 4. (Best view in color) Sensitivity analysis of our methods MIXOT and MIXUNBOT for several hyperparameter variations ( $\alpha$ ,  $\eta_4$ ,  $\eta_5$  and batch size  $m$ ). We report the average classification accuracy and variance on three runs with different seeds for DA tasks USPS  $\mapsto$  MNIST (U-M) and SVHN  $\mapsto$  MNIST (S-M). Networks have been pretrained for 2 epochs except on (U-M) batch size sensitivity analysis (10 epochs).

achieved by our method as justified in our experiments. Regarding the TSNE plot, we can see that our method produces well defined clusters while the DEEPIJDOT method has some overlaps between them.

Finally, we conducted a sensitivity analysis on the DA tasks USPS to MNIST and SVHN to MNIST for the following hyper-parameters: coefficients  $\alpha$ ,  $\eta_4$ ,  $\eta_5$  and batch size  $m$ .  $\eta_4$ ,  $\eta_5$  are constants in the SCE loss  $\text{SCE}(q, q') = \eta_4 \text{CE}(q, q') + \eta_5 \text{CE}(q', q)$  to evaluate the influence of the term  $\text{CE}(q', q)$ . All results are summarized in Figure 4. We have higher accuracies for values of  $\alpha$  greater than 0.1 and smaller than 0.5 as larger values lead to high perturbed samples. Regarding  $\eta_4$ , we can see that the accuracy remains stable for MIXUNBOT and decreases for MIXOT when  $\eta_4$  gets closer to 1. We also see a decrease in the accuracy for small values  $\eta_5$ , *i.e.*, when we recover the typical cross-entropy. It shows that MixUp leads to improvement when used with the SCE loss. Finally, we can see that small minibatch sizes lead to poor performances for MIXOT probably due to the high negative transfer between domains.

## 4.2 PARTIAL DOMAIN ADAPTATION

In this section, we study the performances of our proposed method MIXUNBOT on partial domain adaptation (PDA) experiments. In partial domain adaptation, some classes in the source domain do not exist in the target domain and thus, the target labels are a subset of the source labels, *i.e.*,  $\mathcal{Y}_t \subset \mathcal{Y}_s$ . The different samples which are in these classes can be considered as outliers at the level of the target domain as they would produce negative transfer. That is why methods based on exact Optimal Transport fail to solve this problem as they would transport these samples to the target domain, however JUMBOT has been shown to perform well on this problem thanks to Unbalanced OT. We thus want to evaluate the performances of MIXUNBOT in partial DA.

We evaluate our method on the dataset partial Office-Home. We follow the setting from Cao et al. (2018b) which selects the first 25 categories (in alphabetic order) as target classes to form the partial target domain. We compare our method against the state-of-the-art partial domain adaptation methods: PADA (Cao et al., 2018b), ETN (Cao et al., 2019), BA3US (Jian et al., 2020), SAFN (Xu et al., 2019), RTNET (Chen et al., 2020b), IWAN (Zhang et al., 2018b), SAN (Cao et al., 2018a), DRCN (Li et al., 2021), JUMBOT (FAtlas et al., 2021b) and AR (Gu et al., 2021). For fair comparison we followed the experimental setting of PADA, ETN, BA3US and JUMBOT. We selected the best hyper parameters on the A-C task as validation and used them on the remaining tasks. We report all the training and architecture details to the appendix. We evaluate our methods in two different ways. We first report the performances of our method at the end of training against BA3US, DEEPIJDOT and JUMBOT. These results are reported from (FAtlas et al., 2021b). Then we report the best score along training of our methods (oracle), where we evaluated the models every 1250 iterations, against all mentioned competitors where we reported the scores from their paper for our setting (neural network architecture and training procedure). All results can be found in the lower part of Table 2.

For results at the end of training, we can see that MIXUNBOT achieves the highest performances on all partial domain adaptation tasks and is on average 3% above the best competitor JUMBOT, showing the relevance to regularize the neural networks as we do. Regarding the oracle score, we see that MIXUNBOT achieves the highest performances on 10 partial domain adaptation tasks over 12 and that our method is 1.5% above the best competitor on average. The different results between our oracle score and the score at the end of training are similar on each task which show that our method does not suffer from an overfitting issue in the training.

## 5 CONCLUSION

In this paper we proposed new methods to tackle the domain adaptation problem. When label shift occurs between domains, optimal transport based methods create negative alignments. These negative alignments assign noisy labels to target samples which can decrease the performances of models. We proposed a new method to mitigate the influ-

ence of these noisy labels on our models. We first used the MixUp regularization on both source and target domains. We then used the symmetric cross-entropy loss, a noisy label robust loss function, when aligning the source and the target distributions. Afterwards, we evaluated our methods on several benchmark and real-world domain adaptation and partial domain adaptation scenarios. We showed that our methods are able to outperform recent state-of-the-art methods. In an ablation study, we showed that it is the combination of the MixUp regularization and the symmetric cross-entropy loss which increases performance, while individually they do not improve results. Future work will consider a mathematical analysis on the use of MixUp augmentation and symmetric cross-entropy to justify the empirical performances or the use of manifold MixUp regularization (Verma et al., 2019).

#### AUTHOR CONTRIBUTIONS

In this section, we state the contributions of each author on each part of the present manuscript.

- Formalism : KF
- Proof-of-concepts : KF
- Experiments : KF, HN
- Experiments review : KF, HN
- Writing—original draft preparation : KF
- Writing—review and editing : KF, HN, IM
- Supervision : KF, IM
- Funding acquisition : IM

#### ACKNOWLEDGMENTS

This work was partially supported by NSERC Discovery grant (RGPIN-2019-06512), a Samsung grant and a Canada CIFAR AI chair. Authors thank Nicolas Courty and Rémi Flamary for fruitful discussions.

#### REFERENCES

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 214–223, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/arjovsky17a.html>.
- Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Robust optimal transport with applications in generative modeling and domain adaptation. In *Advances in Neural Information Processing Systems*, 2020.
- Bharath Bhushan Damodaran, Benjamin Kellenberger, Remi Flamary, Devis Tuia, and Nicolas Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- Jean-Christophe Burnel, Kilian Fatras, Rémi Flamary, and Nicolas Courty. Generating natural adversarial remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 2021.
- Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 754–763, 2017. doi: 10.1109/ICCV.2017.88.
- Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I. Jordan. Partial transfer learning with selective adversarial networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2724–2732, 2018a.
- Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018b.
- Zhangjie Cao, Kaichao You, Mingsheng Long, Jianmin Wang, and Qiang Yang. Learning to transfer examples for partial domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

- Luigi Carratino, Moustapha Cisse, Rodolphe Jenatton, and Jean-Philippe Vert. On mixup regularization. *ArXiv*, abs/2006.06049, 2020.
- Minghao Chen, Shuai Zhao, Haifeng Liu, and Deng Cai. Adversarial-learned loss for domain adaptation. *arXiv*, abs/2001.01046, 2020a.
- Zhihong Chen, Chao Chen, Zhaowei Cheng, Ke Fang, and Xinyu Jin. Selective transfer with reinforced transfer network for partial domain adaptation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12703–12711, 2020b.
- Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Math. Comput.*, 87(314):2563–2609, 2018. doi: 10.1090/mcom/3303.
- N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, 2017.
- Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 3730–3739. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6963-joint-distribution-optimal-transportation-for-domain-adaptation.pdf>.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 26*, pp. 2292–2300. Curran Associates, Inc., 2013. URL <http://papers.nips.cc/paper/4927-sinkhorn-distances-lightspeed-computation-of-optimal-transport.pdf>.
- Kilian Fatras, Younes Zine, Rémi Flamary, Remi Gribonval, and Nicolas Courty. Learning with minibatch wasserstein : asymptotic and gradient properties. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2131–2141, Online, 26–28 Aug 2020. PMLR. URL <http://proceedings.mlr.press/v108/fatras20a.html>.
- Kilian Fatras, Bharath Bhushan Damodaran, Sylvain Lobry, Remi Flamary, Devis Tuia, and Nicolas Courty. Wasserstein adversarial regularization for learning with label noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021a. doi: 10.1109/TPAMI.2021.3094662.
- Kilian Fatras, Thibault Sejourne, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3186–3197. PMLR, 18–24 Jul 2021b. URL <http://proceedings.mlr.press/v139/fatras21a.html>.
- Kilian Fatras, Younes Zine, Szymon Majewski, Rémi Flamary, Rémi Gribonval, and Nicolas Courty. Minibatch optimal transport distances; analysis and applications, 2021c.
- Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.
- Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning with a wasserstein loss. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 28*, pp. 2053–2061. Curran Associates, Inc., 2015. URL <http://papers.nips.cc/paper/5679-learning-with-a-wasserstein-loss.pdf>.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.
- Aude Genevay, Gabriel Peyre, and Marco Cuturi. Learning generative models with sinkhorn divergences. In Amos Storkey and Fernando Perez-Cruz (eds.), *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1608–1617, Playa Blanca, Lanzarote, Canary Islands, 09–11 Apr 2018. PMLR. URL <http://proceedings.mlr.press/v84/genevay18a.html>.

- Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (eds.), *Computer Vision – ECCV 2016*, pp. 597–613, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46493-0.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- Xiang Gu, Xi Yu, Yan Yang, Jian Sun, and Zongben Xu. Adversarial reweighting for partial domain adaptation. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=f5liPryFRoA>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.
- Jonathan Hull. Database for handwritten text recognition research. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 16, 1994. doi: 10.1109/34.291440.
- Liang Jian, Wang Yunbo, Hu Dapeng, He Ran, and Feng Jiashi. A balanced and uncertainty-aware approach for partial domain adaptation. In *European Conference on Computer Vision (ECCV)*, August 2020.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>.
- Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010.
- Shuang Li, Chi Harold Liu, Qiuxia Lin, Qi Wen, Limin Su, Gao Huang, and Zhengming Ding. Deep residual correction network for partial domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(7): 2329–2344, Jul 2021. ISSN 1939-3539. doi: 10.1109/tpami.2020.2964173. URL <http://dx.doi.org/10.1109/TPAMI.2020.2964173>.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 97–105, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/long15.html>.
- Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Deep transfer learning with joint adaptation networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2208–2217. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/long17a.html>.
- Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pp. 1645–1655, 2018.
- Xudong Mao, Yun Ma, Zhenguo Yang, Yangbin Chen, and Qing Li. Virtual mixup training for unsupervised domain adaptation, 2019.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010. doi: 10.1109/TKDE.2009.191.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE Signal Processing Magazine*, 32(3):53–69, 2015. doi: 10.1109/MSP.2014.2347059.

- Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *CoRR*, abs/1710.06924, 2017.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 2019.
- Ievgen Redko, Amaury Habrard, and Marc Sebban. Theoretical Analysis of Domain Adaptation with Optimal Transport. In *ECML PKDD 2017*, Skopje, Macedonia, September 2017. URL <https://hal.archives-ouvertes.fr/hal-01613564>.
- Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. Optimal transport for multi-source domain adaptation under target shift. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 849–858. PMLR, 16–18 Apr 2019. URL <https://proceedings.mlr.press/v89/redko19a.html>.
- Tim Salimans, Han Zhang, Alec Radford, and Dimitris Metaxas. Improving gans using optimal transport, 2018.
- Rui Shu, Hung H. Bui, Hirokazu Narui, and Stefano Ermon. A dirt-t approach to unsupervised domain adaptation, 2018.
- Max Sommerfeld, Jörn Schrieber, Yoav Zemel, and Axel Munk. Optimal transport: Fast probabilistic approximation with exact solvers. *Journal of Machine Learning Research*, 2019.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9 (86):2579–2605, 2008. URL <http://jmlr.org/papers/v9/vandermaaten08a.html>.
- Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *(IEEE) Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 6438–6447. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/verma19a.html>.
- Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *IEEE International Conference on Computer Vision*, 2019.
- Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pp. 6502–6509. AAAI Press, 2020.
- Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- Luyu Yang, Yan Wang, Mingfei Gao, Abhinav Shrivastava, Kilian Q. Weinberger, Wei-Lun Chao, and Ser-Nam Lim. Deep co-training with task decomposition for semi-supervised domain adaptation, 2021.
- C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *The International Conference on Learning Representations*, 2017.
- Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018a. URL <https://openreview.net/forum?id=r1Ddp1-Rb>.
- Jing Zhang, Zewei Ding, W. Li, and Philip Ogunbona. Importance weighted adversarial nets for partial domain adaptation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8156–8164, 2018b.
- Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=8yKEo06dKNo>.






Figure 5. (Best view in color) Sensitivity analysis of our methods MIXOT and MIXUNBOT for several hyperparameter variations ( $\lambda$ ,  $\eta_4$ ,  $\eta_5$ ). We report the classification accuracy on the AC task of the Office-Home dataset.

## APPENDIX

### A ADDITIONAL EXPERIMENTS

In this section, we include a sensitivity analysis of our methods on the AC task of the Office-Home dataset (see Figure 5).

### B EXPERIMENTAL SETTINGS

#### B.1 IMPLEMENTATION AND ENVIRONMENT FOR EXPERIMENT

Experiments were conducted on two 1080-Ti NVIDIA GPUs. As a software environment, we use CentOS 7, gcc 9.3.0, Python 3.8.2, Pytorch 1.10.0, cuDNN 8.2.0, and CUDA 11.4. Our code can be found here <https://github.com/kilianFAtlas/MixOT>.

#### B.2 SYMMETRIC CROSS-ENTROPY LOSS

For two probability vectors  $q, q' \in \Sigma_k$ , where  $\Sigma_k$  is the simplex of dimension  $k$ , the symmetric cross-entropy loss is defined as:

$$\text{SCE}(q, q') = \text{CE}(q, q') + \text{CE}(q', q),$$

$$\text{where } \text{CE}(q, q') = \sum_{i=1}^K q_i \log(q'_i).$$

However when  $q$  is a one-hot vector, like a label for instance,  $\text{CE}(q', q)$  is not directly computable because of the logarithm. That’s why we clip the value of  $q$  with a small constant of  $1e^{-7}$  in our experiments. A similar strategy was used in Wang et al. (2019).

#### B.3 UNBALANCED OT

Our method MIXUNBOT is based on unbalanced optimal transport. In practice, we used the entropically-regularized variant of Unbalanced OT as done in Fatras et al. (2021b). It allows us to use a generalized formulation of the Sinkhorn algorithm (Cuturi, 2013; Chizat et al., 2018) which is simple to implement in practice. For digits experiments the entropic regularization coefficient was set to 0.1. For the VisDA and Office Home experiments, it was set to 0.01.

#### B.4 MODEL AND DATASETS

In the digits experiment, we used a 6 CNN layer neural network as feature extractor and 1 dense layer for classification as proposed in Bhushan Damodaran et al. (2018). The number of input channels is set to 3 for the experiments for SVHN to MNIST and MNIST to MNISTM, while the number of channels is set to 1 only for the experiments for USPS to MNIST. In the Office-Home and Visda-2017 experiments, we use the pre-trained model of ImageNet-1K on Resnet-50<sup>1</sup>.

<sup>1</sup>You can find the pre-trained model at <https://download.pytorch.org/models/resnet50-0676ba61.pth>

## B.5 HYPER-PARAMETERS AND DETAILED CONFIGURATION

Note that for all experiments we relied on a stratified source minibatches, meaning that we had the same number of samples per class in source minibatches, as it was done in (Bhushan Damodaran et al., 2018; Fatras et al., 2021b).

### B.5.1 DIGITS EXPERIMENTS

In the digits experiment, we basically followed the experimental setup from DEEPJDOT (Bhushan Damodaran et al., 2018) which uses Adam as optimizer with a learning rate set to  $2e^{-4}$  and without applying weight decay. We pre-train our neural network on the source domain for 2 epochs and then, we train our model for 100 epochs. The batch size is set to 120. The same experimental procedure is used for the ablation study the sensitivity analysis. The best hyper-parameters for MIXOT are  $\eta_1 = 0.1, \eta_2 = 0.1$  and for MIXUNBOT  $\eta_1 = 0.1, \eta_2 = 0.1, \eta_3 = 1., \tau = 1.$

Regarding competitors, we use the official implementations with the considered architecture and training procedure except that we use Adam with weight decay (set to  $1e^{-5}$ ) and a learning rate of  $2e^{-4}$ , since we observe that simply adaptation of Adam significantly degrades the performance of these methods, and their experiments in their paper are reported by using Momentum SGD and weight decay. The USPS to MNIST task served as validation of hyper-parameters.

### B.5.2 OFFICE-HOME EXPERIMENTS

In the Office-Home experiment, the number of steps was fixed at 10K, and the batch size was 65, the same number as the number of classes as done in the Fatras et al. (2021b) experiments. Nesterov’s acceleration method was used as an optimization method, and the momentum coefficient was 0.9. As a regularization, the weight decay is set to  $5e^{-4}$ , and learning rate decay is applied. We schedule the learning rate with the strategy in Ganin et al. (2016), it is adjusted by  $\chi_p = \frac{\chi_0}{(1+\zeta l)^\kappa}$ , where  $l$  is the training progress linearly changing from 0 to 1,  $\chi_0 = 0.01, \zeta = 10, \kappa = 0.75$ .

Since Office-Home has four domains, we used the AC task as validation to determine the best hyper-parameters. Evaluation was done using the test 10 crop technique as in Fatras et al. (2021b); Jian et al. (2020). The hyper-parameters search of the experiments shown in Table 2 of Section 4 are explained in Table 3. Hyperparameter  $\alpha$  for Beta distribution used in MixUp is set to 0.2. The best hyper-parameters for MIXOT on vanilla Office-Home are  $\eta_1 = 0.0025, \eta_2 = 0.05$ . and for MIXUNBOT, the best hyper-parameters are  $\eta_1 = 0.005, \eta_2 = 0.1, \tau = 0.5$ .

Method	$\eta_1$	$\eta_2$	$\tau$
MIXOT	{0.0005, 0.001, 0.0025, 0.005, 0.01, 0.05, 0.1}	{1.0, 0.5, 0.1, 0.05, 0.01}	-
MIXUNBOT	{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1}	{1.0, 0.5, 0.1, 0.05, 0.01}	{0.3, 0.5, 0.7, 1.0}

Table 3: Hyperparameter Search Range: Office-Home Dataset

### B.5.3 VISDA-2017 EXPERIMENTS

Visda-2017 dataset is composed of two domains: synthetic images and real images. The Visda-2017 experiment has almost the same configuration as the Office-Home experiment. The difference is that we set the batch size to 72, which is 6 times the number of Visda-2017 classes 12. They are set to  $\eta_1 = 0.001, \eta_2 = 0.5, \tau = 0.25$ . The hyper-parameters search range of the Visda-2017 experiments shown in Table 1 of Section 4 are explained in Table 4.

Method	$\eta_1$	$\eta_2$	$\tau$
MIXUNBOT	{0.0001, 0.005, 0.001, 0.05, 0.1}	{0.005, 0.01, 0.05, 0.1, 0.5}	{0.25, 0.5, 0.7}

Table 4: Hyperparameter Search Range: Visda-2017 Dataset

### B.5.4 PARTIAL OFFICE-HOME

For Partial Domain Adaptation, we considered the same neural network architecture and the same training procedure as in Fatras et al. (2021b); Jian et al. (2020); Cao et al. (2018b; 2019). We used a pre-trained Resnet-50 on imagenet as pre-trained neural network. Our hyper-parameters are chosen by using the AC task as validation. They are set to  $\eta_1 = 0.001, \eta_2 = 0.05, \tau = 0.05$ , and finally  $\eta_3$  was set to 15. We trained the network for 5000 iterations with a stratified batch size of 65 (one element of each class per batch) and for optimization procedure, we used the same as in

Jian et al. (2020). We do not use the ten crop technique to evaluate our model on the test set like Fatras et al. (2021b); Jian et al. (2020). Furthermore, we do not know if the official reported results of competitors were evaluated at the end of training or during training.