

How does the Combined Risk Affect the Performance of Unsupervised Domain Adaptation Approaches?

Li Zhong^{1,2,*}, Zhen Fang^{2,*}, Feng Liu^{2,*}, Jie Lu^{2,†}, Bo Yuan¹, Guangquan Zhang²

¹Shenzhen International Graduate School, Tsinghua University

²Centre for Artificial Intelligence, University of Technology Sydney

zhongl18@mails.tsinghua.edu.cn, feng.liu@uts.edu.au, zhen.fang@student.uts.edu.au,

yuanb@sz.tsinghua.edu.cn, {guangquan.zhang, jie.lu}@uts.edu.au

Abstract

Unsupervised domain adaptation (UDA) aims to train a target classifier with labeled samples from the source domain and unlabeled samples from the target domain. Classical UDA learning bounds show that target risk is upper bounded by three terms: source risk, distribution discrepancy, and combined risk. Based on the assumption that the combined risk is a small fixed value, methods based on this bound train a target classifier by only minimizing estimators of the source risk and the distribution discrepancy. However, the combined risk may increase when minimizing both estimators, which makes the target risk uncontrollable. Hence the target classifier cannot achieve ideal performance if we fail to control the combined risk. To control the combined risk, the key challenge takes root in the unavailability of the labeled samples in the target domain. To address this key challenge, we propose a method named E-MixNet. E-MixNet employs *enhanced mixup*, a generic vicinal distribution, on the labeled source samples and pseudo-labeled target samples to calculate a proxy of the combined risk. Experiments show that the proxy can effectively curb the increase of the combined risk when minimizing the source risk and distribution discrepancy. Furthermore, we show that if the proxy of the combined risk is added into loss functions of four representative UDA methods, their performance is also improved.

Introduction

Domain Adaptation (DA) aims to train a target-domain classifier with samples from source and target domains (Lu et al. 2015). When the labels of samples in the target domain are unavailable, DA is known as *unsupervised DA* (UDA) (Zhong et al. 2020; Fang et al. 2020), which has been applied to address diverse real-world problems, such as computer vision (Zhang et al. 2020b; Dong et al. 2019, 2020b), natural language processing (Lee and Jha 2019; Guo, Pasunuru, and Bansal 2020), and recommender system (Zhang et al. 2017; Yu, Wang, and Yuan 2019; Lu et al. 2020)

Significant theoretical advances have been achieved in UDA. Pioneering theoretical work was proposed by Ben-David et al. (2007). This work shows that the target risk is

*Equal Contribution. Work done at AAIL, UTS.

†Corresponding Author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The values of combined risk and accuracy on the task C \rightarrow P on Image-CLEF. The left figure shows the value of combined risk. The right figure shows the accuracy of the task. Blue line: ignore the optimization of combined risk. Green line: optimize the combined risk by source samples and target samples with high confidence. Orange line: optimize combined risk by the proxy formulated by mixup. Purple line: optimize the proxy formulated by e-mixup.

upper bounded by three terms: source risk, marginal distribution discrepancy, and *combined risk*. This earliest learning bound has been extended from many perspectives, such as considering more surrogate loss functions (Zhang et al. 2019a) or distributional discrepancies (Mohri and Medina 2012; Shen et al. 2018) (see (Redko et al. 2020) as a survey). Recently, Zhang et al. (2019a) proposed a new distributional discrepancy termed Margin Disparity Discrepancy and developed a tighter and more practical UDA learning bound.

The UDA learning bounds proposed by (Ben-David et al. 2007, 2010) and the recent UDA learning bounds proposed by (Shen et al. 2018; Xu et al. 2020; Zhang et al. 2020a) consist of three terms: source risk, marginal distribution discrepancy, and combined risk. Minimizing the source risk aims to obtain a source-domain classifier, and minimizing the distribution discrepancy aims to learn domain-invariant features so that the source-domain classifier can perform well on the target domain. The combined risk embodies the *adaptability* between the source and target domains (Ben-David et al. 2010). In particular, when the hypothesis space is fixed, the combined risk is a constant.

Based on the UDA learning bounds where the combined risk is assumed to a small constant, many existing UDA methods focus on learning domain-invariant features (Fang et al. 2019; Dong et al. 2020c,a; Liu et al. 2019) by minimiz-

ing the estimators of the source risk and the distribution discrepancy. In the learned feature space, the source and target distributions are similar while the source-domain classifier is required to achieve a small error. Furthermore, the generalization error of the source-domain classifier is expected to be small in the target domain.

However, the combined risk may increase when learning the domain-invariant features, and the increase of the combined risk may degrade the performance of the source-domain classifier in the target domain. As shown in Figure 1, we calculate the value of the combined risk and accuracy on a real-world UDA task (see the green line). The performance worsens with the increased combined risk. Zhao et al. (2019) also pointed out the increase of combined risk causes the failure of source-domain classifier on the target domain.

To investigate how the combined risk affect the performance on the domain-invariant features, we rethink and develop the UDA learning bounds by introducing feature transformations. In the new bound (see Eq. (5)), the combined risk is a function related to feature transformation but not a constant (compared to bounds in (Ben-David et al. 2010)). We also reveal that the combined risk is deeply related to the conditional distribution discrepancy (see Theorem 3). Theorem 3 shows that, the conditional distribution discrepancy will increase when the combined risk increases. Hence, it is hard to achieve satisfied target-domain accuracy if we only focus on learning domain-invariant features and omit to control the combined risk.

To estimate the combined risk, the key challenge takes root in the *unavailability* of the labeled samples in the target domain. A simple solution is to leverage the pseudo labels with high confidence in the target domain to estimate the combined risk. However, since samples with high confidence are insufficient, the value of the combined risk may still increase (see the green line in Figure 1). Inspired by semi-supervised learning methods, an advanced solution is to directly use the *mixup* technique to augment pseudo-labeled target samples, which can slightly help us estimate the combined risk better than the simple solution (see the orange line in Figure 1).

However, the target-domain pseudo labels provided by the source-domain classifier may be inaccurate due to the discrepancy between domains, which causes that *mixup* may not perform well with inaccurate labels. To mitigate the issue, we propose *enhanced mixup* (e-mixup) to substitute *mixup* to compute a proxy of the combined risk. The purple line in Figure 1 shows that the proxy based on e-mixup can significantly boost the performance. Details of the proxy is shown in section Motivation.

To the end, we design a novel UDA method referred to E-MixNet. E-MixNet learns the target-domain classifier by simultaneously minimizing the source risk, the marginal distribution discrepancy, and the proxy of combined risk. Via minimizing the proxy of combined risk, we control the increase of combined risk effectively, thus, control the conditional distribution discrepancy between two domains.

We conduct experiments on three public datasets (Office-31, Office-Home, and Image-CLEF) and compare E-MixNet with a series of existing state-of-the-art methods. Further-

more, we introduce the proxy of the combined risk into four representative UDA methods (*i.e.*, DAN (Long et al. 2015), DANN (Ganin et al. 2016), CDAN (Long et al. 2018), SymNets (Zhang et al. 2019b)). Experiments show that E-MixNet can outperform all baselines, and the four representative methods can achieve better performance if the proxy of the combined risk is added into their loss functions.

Problem Setting and Concepts

In this section, we introduce the definition of UDA, then introduce some important concepts used in this paper.

Let $\mathcal{X} \subset \mathbb{R}^d$ be a feature space and $\mathcal{Y} := \{\mathbf{y}_c\}_{c=1}^K$ be the label space, where the label $\mathbf{y}_c \in \mathbb{R}^K$ is a one-hot vector, whose c -th coordinate is 1 and the other coordinates are 0.

Definition 1 (Domains in UDA). *Given random variables $X_s, X_t \in \mathcal{X}$, $Y_s, Y_t \in \mathcal{Y}$, the source and target domains are joint distributions $P_{X_s Y_s}$ and $P_{X_t Y_t}$ with $P_{X_s Y_s} \neq P_{X_t Y_t}$.*

Then, we propose the UDA problem as follows.

Problem 1 (UDA). *Given independent and identically distributed (*i.i.d.*) labeled samples $D_s = \{(\mathbf{x}_s^i, \mathbf{y}_s^i)\}_{i=1}^{n_s}$ drawn from the source domain $P_{X_s Y_s}$ and *i.i.d.* unlabeled samples $D_t = \{\mathbf{x}_t^i\}_{i=1}^{n_t}$ drawn from the target marginal distribution P_{X_t} , the aim of UDA is to train a classifier $f: \mathcal{X} \rightarrow \mathcal{Y}$ with D_s and D_t such that f classifies accurately target data D_t .*

Given a loss function $\ell: \mathbb{R}^K \times \mathbb{R}^K \rightarrow \mathbb{R}_{\geq 0}$ and any scoring functions \mathbf{C}, \mathbf{C}' from function space $\{\mathbf{C}: \mathcal{X} \rightarrow \mathbb{R}^K\}$, source risk, target risk and classifier discrepancy are

$$\begin{aligned} R_s^\ell(\mathbf{C}) &:= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P_{X_s Y_s}} \ell(\mathbf{C}(\mathbf{x}), \mathbf{y}), \\ R_t^\ell(\mathbf{C}) &:= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P_{X_t Y_t}} \ell(\mathbf{C}(\mathbf{x}), \mathbf{y}), \\ R_s^\ell(\mathbf{C}', \mathbf{C}) &:= \mathbb{E}_{\mathbf{x} \sim P_{X_s}} \ell(\mathbf{C}'(\mathbf{x}), \mathbf{C}(\mathbf{x})), \\ R_t^\ell(\mathbf{C}', \mathbf{C}) &:= \mathbb{E}_{\mathbf{x} \sim P_{X_t}} \ell(\mathbf{C}'(\mathbf{x}), \mathbf{C}(\mathbf{x})). \end{aligned}$$

Lastly, we define the disparity discrepancy based on double losses, which will be used to design our method.

Definition 2 (Double Loss Disparity Discrepancy). *Given distributions P, Q over some feature space $\tilde{\mathcal{X}}$, two losses ℓ_s, ℓ_t , a hypothesis space $\mathcal{H} \subset \{\mathbf{C}: \tilde{\mathcal{X}} \rightarrow \mathbb{R}^K\}$ and any scoring function $\mathbf{C} \in \mathcal{H}$, then the double loss disparity discrepancy $d_{\mathbf{C}, \mathcal{H}}^{\ell_s, \ell_t}(P, Q)$ is*

$$\sup_{\mathbf{C}' \in \mathcal{H}} (R_P^{\ell_t}(\mathbf{C}', \mathbf{C}) - R_Q^{\ell_s}(\mathbf{C}', \mathbf{C})), \quad (1)$$

where

$$\begin{aligned} R_P^{\ell_t}(\mathbf{C}', \mathbf{C}) &:= \mathbb{E}_{\mathbf{x} \sim P} \ell_t(\mathbf{C}'(\mathbf{x}), \mathbf{C}(\mathbf{x})), \\ R_Q^{\ell_s}(\mathbf{C}', \mathbf{C}) &:= \mathbb{E}_{\mathbf{x} \sim Q} \ell_s(\mathbf{C}'(\mathbf{x}), \mathbf{C}(\mathbf{x})). \end{aligned}$$

When losses ℓ_s, ℓ_t are the margin loss (Zhang et al. 2019a), the double loss disparity discrepancy is known as the Margin Disparity Discrepancy (Zhang et al. 2019a).

Compared with the classical discrepancy distance (Mansour, Mohri, and Rostamizadeh 2009):

$$d_{\mathcal{H}}^\ell(P, Q) := \sup_{\mathbf{C}', \mathbf{C} \in \mathcal{H}} |R_P^\ell(\mathbf{C}', \mathbf{C}) - R_Q^\ell(\mathbf{C}', \mathbf{C})|, \quad (2)$$

double loss disparity discrepancy is tighter and more flexible.

Theorem 1 (DA Learning Bound). *Given a loss ℓ satisfying the triangle inequality and a hypothesis space $\mathcal{H} \subset \{\mathbf{C} : \mathcal{X} \rightarrow \mathbb{R}^K\}$, then for any $\mathbf{C} \in \mathcal{H}$, we have*

$$R_t^\ell(\mathbf{C}) \leq R_s^\ell(\mathbf{C}) + d_{\mathcal{H}}^\ell(P_{X_s}, P_{X_t}) + \lambda^\ell,$$

where $d_{\mathcal{H}}^\ell$ is the discrepancy distance defined in Eq. (2) and $\lambda^\ell := \min_{\mathbf{C}^* \in \mathcal{H}} (R_s^\ell(\mathbf{C}^*) + R_t^\ell(\mathbf{C}^*))$ known as the combined risk.

In Theorem 1, when the hypothesis space \mathcal{H} and the loss ℓ are fixed, the combined risk is a fixed constant.

Theoretical Analysis

Here we introduce our main theoretical results. All proofs can be found at <https://github.com/zhonglii/E-MixNet>.

Rethinking DA Learning Bound

Many existing UDA methods (Wang and Breckon 2020; Zou et al. 2019; Tang and Jia 2020) learn a suitable feature transformation \mathbf{G} such that the discrepancy between transformed distributions $P_{\mathbf{G}(X_s)}$ and $P_{\mathbf{G}(X_t)}$ is reduced. By introducing the transformation \mathbf{G} in the classical DA learning bound, we discover that the combined error λ^ℓ is not a fixed value.

Theorem 2. *Given a loss ℓ satisfying the triangle inequality, a transformation space $\mathcal{G} \subset \{\mathbf{G} : \mathcal{X} \rightarrow \mathcal{X}_{\text{new}}\}$ and a hypothesis space $\mathcal{H} \subset \{\mathbf{C} : \mathcal{X}_{\text{new}} \rightarrow \mathbb{R}^K\}$, then for any $\mathbf{G} \in \mathcal{G}$ and $\mathbf{C} \in \mathcal{H}$,*

$$R_t^\ell(\mathbf{C} \circ \mathbf{G}) \leq R_s^\ell(\mathbf{C} \circ \mathbf{G}) + d_{\mathcal{H}}^\ell(P_{\mathbf{G}(X_s)}, P_{\mathbf{G}(X_t)}) + \lambda^\ell(\mathbf{G}),$$

where $d_{\mathcal{H}}^\ell$ is the discrepancy distance defined in Eq. (2) and

$$\lambda^\ell(\mathbf{G}) := \min_{\mathbf{C}^* \in \mathcal{H}} (R_s^\ell(\mathbf{C}^* \circ \mathbf{G}) + R_t^\ell(\mathbf{C}^* \circ \mathbf{G})) \quad (3)$$

known as the combined risk.

According to Theorem 2, it is not enough to minimize the source risk and distribution discrepancy by seeking the optimal classifier \mathbf{C} and optimal transformation \mathbf{G} from spaces \mathcal{H} and \mathcal{G} . Because we cannot guarantee the value of combined risk $\lambda^\ell(\mathbf{G})$ is always small during the training process.

For convenience, we define

$$\Lambda^\ell(\mathbf{C}, \mathbf{G}) := R_s^\ell(\mathbf{C} \circ \mathbf{G}) + R_t^\ell(\mathbf{C} \circ \mathbf{G}), \quad (4)$$

hence, $\lambda^\ell(\mathbf{G}) = \min_{\mathbf{C}^* \in \mathcal{H}} \Lambda^\ell(\mathbf{C}^*, \mathbf{G})$.

Meaning of Combined Risk $\lambda^\ell(\mathbf{G})$

To future understand the meaning of the combined risk $\lambda^\ell(\mathbf{G})$, we prove the following Theorem.

Theorem 3. *Given a symmetric loss ℓ satisfying the triangle inequality, a feature transformation $\mathbf{G} \in \mathcal{G} \subset \{\mathbf{G} : \mathcal{X} \rightarrow \mathcal{X}_{\text{new}}\}$, a hypothesis space $\mathcal{H} \subset \{\mathbf{C} : \mathcal{X}_{\text{new}} \rightarrow \mathbb{R}^K\}$, and*

$$\mathbf{C}_s = \operatorname{argmin}_{\mathbf{C} \in \mathcal{H}} R_s^\ell(\mathbf{C} \circ \mathbf{G}), \quad \mathbf{C}_t = \operatorname{argmin}_{\mathbf{C} \in \mathcal{H}} R_t^\ell(\mathbf{C} \circ \mathbf{G}),$$

then

$$\begin{aligned} 2\lambda^\ell(\mathbf{G}) - \delta &\leq R_s^\ell(\mathbf{C}_t \circ \mathbf{G}) + R_t^\ell(\mathbf{C}_s \circ \mathbf{G}) \\ &\leq 2\lambda^\ell(\mathbf{G}) + d_{\mathcal{H}}^\ell(P_{\mathbf{G}(X_s)}, P_{\mathbf{G}(X_t)}) + \delta, \end{aligned}$$

where $\lambda^\ell(\mathbf{G})$ is defined in Eq. (3), $\delta := R_s^\ell(\mathbf{C}_s \circ \mathbf{G}) + R_t^\ell(\mathbf{C}_t \circ \mathbf{G})$ known as the approximation error and $d_{\mathcal{H}}^\ell$ is the discrepancy distance defined in Eq. (2).

Theorem 3 implies that the combined risk $\lambda^\ell(\mathbf{G})$ is deeply related to the optimal classifier discrepancy

$$R_s^\ell(\mathbf{C}_t \circ \mathbf{G}) + R_t^\ell(\mathbf{C}_s \circ \mathbf{G}),$$

which can be regarded as the conditional distribution discrepancy between $P_{Y_s|\mathbf{G}(X_s)}$ and $P_{Y_t|\mathbf{G}(X_t)}$. If $\lambda^\ell(\mathbf{G})$ increases, the conditional distribution discrepancy is larger.

Double Loss DA Learning Bound

Note that there exist methods, such as MDD (Zhang et al. 2019a), whose source and target losses are different. To understand these UDA methods and bridge the gap between theory and algorithms, we develop the classical DA learning bound to a more general scenario.

Theorem 4. *Given losses ℓ_s and ℓ_t satisfying the triangle inequality, a transformation space $\mathcal{G} \subset \{\mathbf{G} : \mathcal{X} \rightarrow \mathcal{X}_{\text{new}}\}$ and a hypothesis space $\mathcal{H} \subset \{\mathbf{C} : \mathcal{X}_{\text{new}} \rightarrow \mathbb{R}^K\}$, then for any $\mathbf{G} \in \mathcal{G}$ and $\mathbf{C} \in \mathcal{H}$, then $R_t^{\ell_t}(\mathbf{C} \circ \mathbf{G})$ is bounded by*

$$R_s^{\ell_s}(\mathbf{C} \circ \mathbf{G}) + d_{\mathcal{C}, \mathcal{H}}^{\ell_s, \ell_t}(P_{\mathbf{G}(X_s)}, P_{\mathbf{G}(X_t)}) + \lambda^{\ell_s, \ell_t}(\mathbf{G}), \quad (5)$$

where $d_{\mathcal{C}, \mathcal{H}}^{\ell_s, \ell_t}$ is the double loss disparity discrepancy defined in Eq. (1) and $\lambda^{\ell_s, \ell_t}(\mathbf{G})$ is the combined risk:

$$\lambda^{\ell_s, \ell_t}(\mathbf{G}) := \min_{\mathbf{C}^* \in \mathcal{H}} \Lambda^{\ell_1, \ell_2}(\mathbf{C}^*, \mathbf{G}), \quad \text{here} \quad (6)$$

$$\Lambda^{\ell_1, \ell_2}(\mathbf{C}^*, \mathbf{G}) := R_s^{\ell_s}(\mathbf{C}^* \circ \mathbf{G}) + R_t^{\ell_t}(\mathbf{C}^* \circ \mathbf{G}). \quad (7)$$

In Theorem 4, the condition that ℓ_s and ℓ_t satisfy the triangle inequality, can be replaced by a weaker condition:

$$\begin{aligned} R_t^{\ell_t}(\mathbf{C} \circ \mathbf{G}) &\leq R_t^{\ell_t}(\mathbf{C}' \circ \mathbf{G}, \mathbf{C} \circ \mathbf{G}) + R_t^{\ell_t}(\mathbf{C}' \circ \mathbf{G}), \\ R_s^{\ell_s}(\mathbf{C}' \circ \mathbf{G}, \mathbf{C} \circ \mathbf{G}) &\leq R_s^{\ell_s}(\mathbf{C}' \circ \mathbf{G}) + R_s^{\ell_s}(\mathbf{C} \circ \mathbf{G}). \end{aligned}$$

If we set ℓ_s, ℓ_t as the margin loss, ℓ_s, ℓ_t do not satisfy the triangle inequality but they satisfy above condition.

Proposed Method: E-MixNet

Here we introduce motivation and details of our method.

Motivation

Theorem 3 has shown that the combined risk is related to the conditional distribution discrepancy. As the increase of the combined risk, the conditional distribution discrepancy is increased. Hence, omitting the importance of the combined risk may make negative impacts on the target-domain accuracy. Figure 1 (blue line) verifies our observation.

To control the combined risk, we consider the following problem.

$$\min_{\mathbf{G} \in \mathcal{G}} \lambda^\ell(\mathbf{G}) = \min_{\mathbf{G} \in \mathcal{G}, \mathbf{C}^* \in \mathcal{H}} \Lambda^\ell(\mathbf{C}^*, \mathbf{G}), \quad (8)$$

where $\Lambda^\ell(\mathbf{C}^*, \mathbf{G})$ is defined in Eq. (4). Eq. (8) shows we can control the combined risk by minimizing $\Lambda^\ell(\mathbf{C}^*, \mathbf{G})$. However, it is prohibitive to directly optimize the combined risk, since the labeled target samples are indispensable to estimate $\Lambda^\ell(\mathbf{C}^*, \mathbf{G})$.

To alleviate the above issue, a simple method is to use the target pseudo labels with high confidence to estimate the

$\Lambda^\ell(\mathbf{C}^*, \mathbf{G})$. Given the source samples $\{(\mathbf{x}_s^i, \mathbf{y}_s^i)\}_{i=1}^{n_s}$ and the target samples with high confidence $\{(\mathbf{x}_T^i, \mathbf{y}_T^i)\}_{i=1}^{n_T}$, the empirical form of $\Lambda^\ell(\mathbf{C}^*, \mathbf{G})$ can be computed by

$$\frac{1}{n_s} \sum_{i=1}^{n_s} \ell(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_s^i), \mathbf{y}_s^i) + \frac{1}{n_T} \sum_{i=1}^{n_T} \ell(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_T^i), \mathbf{y}_T^i).$$

However, the combined risk may still increase as shown in Fig 1 (green line). The reason may be that the target samples, whose pseudo labels with high confidence, are insufficient.

Inspired by semi-supervised learning, an advanced solution is to use *mixup* technique (Zhang et al. 2018) to augment pseudo-labeled target samples. Mixup produces new samples by a convex combination: given any two samples $(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)$,

$$\mathbf{x} = \alpha \mathbf{x}_1 + (1 - \alpha) \mathbf{x}_2, \quad \mathbf{y} = \alpha \mathbf{y}_1 + (1 - \alpha) \mathbf{y}_2,$$

where α is a hyper-parameter. Zhang et al. (2018) has shown that mixup not only reduces the memorization to adversarial samples, but also performs better than *Empirical Risk Minimization* (ERM) (Vapnik and Chervonenkis 2015). By applying mixup on the target samples with high confidence, new samples $\{(\mathbf{x}_m^i, \mathbf{y}_m^i)\}_{i=1}^n$ are produced, then we propose a proxy $\tilde{\Lambda}^\ell(\mathbf{C}^*, \mathbf{G})$ of Λ^ℓ as follows:

$$\frac{1}{n_s} \sum_{i=1}^{n_s} \ell(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_s^i), \mathbf{y}_s^i) + \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_m^i), \mathbf{y}_m^i).$$

The aforementioned issue can be mitigated, since mixup can be regarded as a data augmentation (Zhang et al. 2018). However, the target-domain pseudo labels provided by the source-domain classifier may be inaccurate due to the discrepancy between domains, which causes that mixup may not perform well with inaccurate labels. We propose *enhanced mixup* (e-mixup) to substitute the mixup to compute the proxy. E-mixup introduces the pure true-labeled source-samples to mitigate the issue caused by bad pseudo labels.

Furthermore, to increase the diversity of new samples, e-mixup produces each new sample using two distant samples, where the distance of the two distant samples is expected to be large. Compared the ordinary mixup technique (i.e., producing new samples using randomly selected samples), e-mixup can produce new examples more effectively. We also verify that e-mixup can further boost the performance (see Table 5). Details of the e-mixup are shown in Algorithm 1. Corresponding to the double loss situation, denoted by samples $\{(\mathbf{x}_e^i, \mathbf{y}_e^i)\}_{i=1}^n$ produced by e-mixup, the proxy $\tilde{\Lambda}^{\ell_s \ell_t}(\mathbf{C}^*, \mathbf{G})$ of $\Lambda^{\ell_s \ell_t}$ (defined in Eq. (7)) as

$$\frac{1}{n_s} \sum_{i=1}^{n_s} \ell_s(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_s^i), \mathbf{y}_s^i) + \frac{1}{n} \sum_{i=1}^n \ell_t(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_e^i), \mathbf{y}_e^i). \quad (9)$$

The purple line in Figure 1 and ablation study show that e-mixup can further boost performance.

Algorithm

The optimization of the combined risk plays a crucial role in UDA. Accordingly, we propose a method based on the aforementioned analyses to solve UDA more deeply.

Algorithm 1: e-mixup

Input: samples $\{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n$.

Parameter: α , the number of class K .

Output: new samples $\{(\mathbf{x}_e^i, \mathbf{y}_e^i)\}_{i=1}^n$.

- 1: **for** $i = 1, 2, \dots, n$ **do**
 - 2: $y^i = \operatorname{argmin}_{c \in \{1, \dots, K\}} y_c^i$ % y^i is the c -th coordinate value of vector \mathbf{y}^i .
 - 3: Select one from the samples whose label is y^i and denoted it by $(\mathbf{x}_r^i, \mathbf{y}_r^i)$.
 - 4: $\mathbf{x}_e^i = \alpha \mathbf{x}^i + (1 - \alpha) \mathbf{x}_r^i$, $\mathbf{y}_e^i = \alpha \mathbf{y}^i + (1 - \alpha) \mathbf{y}_r^i$
 - 5: **end for**
-

Objective Function According to the theoretical bound in Eq. (5), we need to solve the following problem

$$\min_{\mathbf{C} \in \mathcal{H}, \mathbf{G} \in \mathcal{G}} \left(\widehat{R}_s^{\ell_s}(\mathbf{C} \circ \mathbf{G}) + d_{\mathbf{C}, \mathcal{H}}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}(X_s)}, \widehat{P}_{\mathbf{G}(X_t)}) + \min_{\mathbf{C}^* \in \mathcal{H}} \tilde{\Lambda}^{\ell_s \ell_t}(\mathbf{C}^*, \mathbf{G}) \right),$$

where $d_{\mathbf{C}, \mathcal{H}}^{\ell_s \ell_t}$ is the double losses disparity discrepancy defined in Eq. (1) and $\tilde{\Lambda}^{\ell_s \ell_t}(\mathbf{C}^*, \mathbf{G})$ is defined in Eq. (9).

Minimizing double loss disparity discrepancy is a *minimax game*, since the double losses disparity discrepancy is defined as the supremum over hypothesis space \mathcal{H} . Thus, we revise the above problem as follows:

$$\min_{\mathbf{C} \in \mathcal{H}, \mathbf{G} \in \mathcal{G}} \left(\widehat{R}_s^{\ell_s}(\mathbf{C} \circ \mathbf{G}) + \tilde{\Lambda}^{\ell_s \ell_t}(\tilde{\mathbf{C}}, \mathbf{G}) + \eta d_{\mathbf{C}, \mathbf{C}', \gamma}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}(X_s)}, \widehat{P}_{\mathbf{G}(X_t)}) \right), \quad (10)$$

where γ, η are parameters to make our model more flexible,

$$\begin{aligned} & d_{\mathbf{C}, \mathbf{D}, \gamma}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}(X_s)}, \widehat{P}_{\mathbf{G}(X_t)}) \\ &= \widehat{R}_t^{\ell_t}(\mathbf{D} \circ \mathbf{G}, \mathbf{C} \circ \mathbf{G}) - \gamma \widehat{R}_s^{\ell_s}(\mathbf{D} \circ \mathbf{G}, \mathbf{C} \circ \mathbf{G}), \\ \mathbf{C}' &= \operatorname{argmax}_{\mathbf{D} \in \mathcal{H}} d_{\mathbf{C}, \mathbf{D}, \gamma}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}(X_s)}, \widehat{P}_{\mathbf{G}(X_t)}), \\ \tilde{\mathbf{C}} &= \operatorname{argmin}_{\mathbf{C}^* \in \mathcal{H}} \tilde{\Lambda}^{\ell_s \ell_t}(\mathbf{C}^*, \mathbf{G}). \end{aligned} \quad (11)$$

To solve the problem (10), we construct a deep method. The network architecture is shown in Fig. 2(b), which consists of a generator \mathbf{G} , a discriminator \mathbf{D} , and two classifiers \mathbf{C}, \mathbf{C}' . Next, we introduce the details about our method.

We use standard cross-entropy as the source loss ℓ_s and use modified cross-entropy (Goodfellow et al. 2014; Zhang et al. 2019a) as the target loss ℓ_t .

For any scoring functions $\mathbf{F}, \mathbf{F}' : \mathcal{X} \rightarrow \mathbb{R}^K$,

$$\begin{aligned} \ell_s(\mathbf{F}(\mathbf{x}), \mathbf{F}'(\mathbf{x})) &:= -\log(\sigma_{h'(\mathbf{x})}(\mathbf{F}(\mathbf{x}))), \\ \ell_t(\mathbf{F}(\mathbf{x}), \mathbf{F}'(\mathbf{x})) &:= \log(1 - \sigma_{h'(\mathbf{x})}(\mathbf{F}(\mathbf{x}))), \end{aligned} \quad (12)$$

where σ is softmax function: for any $\mathbf{y} = [y_1, \dots, y_K] \in \mathbb{R}^K$,

$$\sigma_c(\mathbf{y}) = \frac{\exp(y_c)}{\sum_{k=1}^K \exp(y_k)}, \quad \text{for } c = 1, \dots, K,$$

and

$$h'(\mathbf{x}) = \operatorname{argmax}_{c \in \{1, \dots, K\}} F'_c(\mathbf{x}), \quad (13)$$

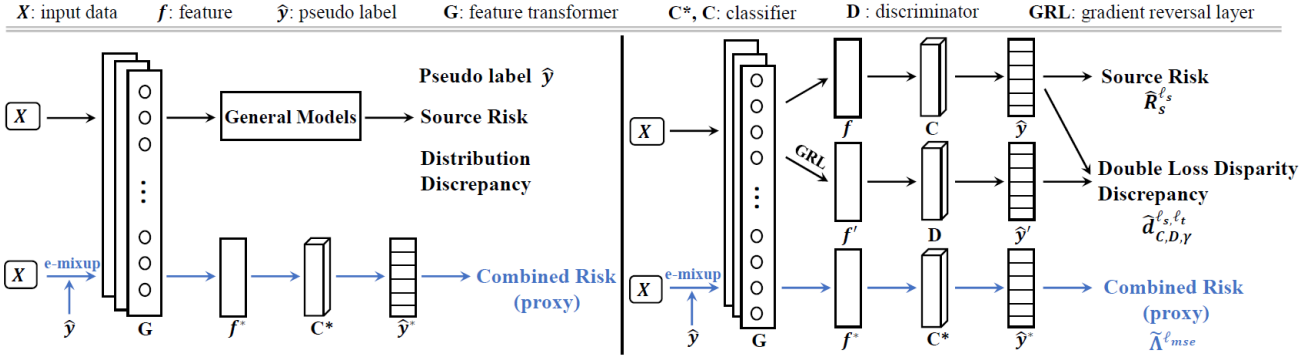


Figure 2: The network architecture of applying the proxy of the combined risk. The left figure is the general model for adding the proxy into existing UDA models. The right figure is a specific model based on double loss disparity discrepancy.

here F'_c is the c -th coordinate function of function \mathbf{F}' .

Source risk. Given the source samples $\{(\mathbf{x}_s^i, \mathbf{y}_s^i)\}_{i=1}^{n_s}$, then

$$\widehat{R}_s^{\ell_s}(\mathbf{C} \circ \mathbf{G}) = \frac{1}{n_s} \sum_{i=1}^{n_s} -\log(\sigma_{y_s^i}(\mathbf{C} \circ \mathbf{G}(\mathbf{x}_s^i))), \quad (14)$$

where y_s^i is the label corresponding to one-hot vector \mathbf{y}_s^i .

Double loss disparity discrepancy. Given the source and target samples $\{(\mathbf{x}_s^i, \mathbf{y}_s^i)\}_{i=1}^{n_s}, \{(\mathbf{x}_t^i, \mathbf{y}_t^i)\}_{i=1}^{n_t}$, then

$$\begin{aligned} & d_{\mathbf{C}, \mathbf{D}, \gamma}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}}(X_s), \widehat{P}_{\mathbf{G}}(X_t)) \\ &= -\frac{\gamma}{n_s} \sum_{i=1}^{n_s} \ell_s(\mathbf{D} \circ \mathbf{G}(\mathbf{x}_s^i), \mathbf{C} \circ \mathbf{G}(\mathbf{x}_s^i)) \\ & \quad + \frac{1}{n_t} \sum_{i=1}^{n_t} \ell_t(\mathbf{D} \circ \mathbf{G}(\mathbf{x}_t^i), \mathbf{C} \circ \mathbf{G}(\mathbf{x}_t^i)), \end{aligned} \quad (15)$$

where ℓ_s, ℓ_t are defined in Eq. (12).

Combined risk. As discussed in Motivation, the combined risk cannot be optimized directly. To mitigate this problem, we use the proxy $\widetilde{\Lambda}^{\ell_s \ell_t}$ in Eq. (9) to substitute it.

Further, motivated by (Berthelot et al. 2019), we use mean square error (ℓ_{mse}) to calculate the proxy of the combined risk, because, unlike the cross-entropy loss, ℓ_{mse} is bounded and less sensitive to incorrect predictions. Denoted by $\{(\mathbf{x}_e^i, \mathbf{y}_e^i)\}_{i=1}^n$ the output of the *e-mixup*. Then the proxy is calculated by

$$\begin{aligned} \widetilde{\Lambda}^{\ell_{mse}}(\mathbf{C}^*, \mathbf{G}) &= \frac{1}{n_s} \sum_{i=1}^{n_s} \ell_{mse}(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_s^i), \mathbf{y}_s^i) \\ & \quad + \frac{1}{n} \sum_{i=1}^n \ell_{mse}(\mathbf{C}^* \circ \mathbf{G}(\mathbf{x}_e^i), \mathbf{y}_e^i). \end{aligned} \quad (16)$$

Training Procedure Finally, the UDA problem can be solved by the following minimax game.

$$\begin{aligned} & \min_{\mathbf{C} \in \mathcal{H}, \mathbf{G} \in \mathcal{G}} (\widehat{R}_s^{\ell_s}(\mathbf{C} \circ \mathbf{G}) + \widetilde{\Lambda}^{\ell_{mse}}(\mathbf{C}^*, \mathbf{G})) \\ & \quad + \eta d_{\mathbf{C}, \mathbf{D}, \gamma}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}}(X_s), \widehat{P}_{\mathbf{G}}(X_t)), \\ & \max_{\mathbf{D} \in \mathcal{H}} d_{\mathbf{C}, \mathbf{D}, \gamma}^{\ell_s \ell_t}(\widehat{P}_{\mathbf{G}}(X_s), \widehat{P}_{\mathbf{G}}(X_t)), \\ & \min_{\mathbf{C}^* \in \mathcal{H}} \widetilde{\Lambda}^{\ell_{mse}}(\mathbf{C}^*, \mathbf{G}). \end{aligned} \quad (17)$$

Algorithm 2: The training procedure of E-MixNet

Input: source, target samples $\{(\mathbf{x}_s^i, \mathbf{y}_s^i)\}_{i=1}^{n_s}, \{(\mathbf{x}_t^i, \mathbf{y}_t^i)\}_{i=1}^{n_t}$.

Parameter: learning rate l , batch size n_b , the number of iteration T , network parameters Θ .

Output: the predicted target label $\widehat{\mathbf{y}}_t$.

- 1: Initialize Θ
- 2: **for** $j = 1, 2, \dots, T$ **do**
- 3: Fetch source minibatch D_s^m
- 4: Fetch target minibatch D_t^m
- 5: Calculate $\widehat{R}_s^{\ell_s}$ using D_s^m
- 6: Calculate $d_{\mathbf{C}, \mathbf{D}, \gamma}^{\ell_s \ell_t}$ using D_s^m, D_t^m
- 7: Obtain highly confident target samples $\{(\mathbf{x}_T^i, \mathbf{y}_T^i)\}_{i=1}^{n_T}$ predicted by $\mathbf{C} \circ \mathbf{G}$ on D_t^m
- 8: $\{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n = D_s^m \cup \{(\mathbf{x}_T^i, \mathbf{y}_T^i)\}_{i=1}^{n_T}$
- 9: $\{(\mathbf{x}_e^i, \mathbf{y}_e^i)\}_{i=1}^n = \text{e-mixup}(\{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n)$
- 10: Calculate $\widetilde{\Lambda}^{\ell_{mse}}$ using $D_s^m, \{(\mathbf{x}_e^i, \mathbf{y}_e^i)\}_{i=1}^n$
- 11: Update Θ according to Eq. (17)
- 12: **end for**

The training procedure is shown in Algorithm 2.

Experiments

We evaluate E-Mixnet on three public datasets, and compare it with several existing state-of-the-art methods. Codes will be available at <https://github.com/zhonglii/E-MixNet>. Furthermore, we further test the efficacy of the proxy of the combined risk by adding it into the loss function of four representative UDA methods.

Datasets

Three common UDA datasets are used to evaluate the efficacy of E-MixNet.

Office-31 (Saenko et al. 2010) contains three domains with a slight discrepancy: *amazon* (A), *dslr* (D) and *webcam* (W). Each domain contains 31 kinds of objects. So there are 6 domain adaptation tasks on *Office-31*: $A \rightarrow D, A \rightarrow W, D \rightarrow A, D \rightarrow W, W \rightarrow A, W \rightarrow D$.

Office-Home (Venkateswara et al. 2017) contains four domains with more obvious domain discrepancy than *Office-*

Method	A→W	D→W	W→D	A→D	D→A	W→A	Avg
ResNet-50 (He et al. 2016)	68.4±0.2	96.7±0.1	99.3±0.1	68.9±0.2	62.5±0.3	60.7±0.3	76.1
DAN (Long et al. 2015)	80.5±0.4	97.1±0.2	99.6±0.1	78.6±0.2	63.6±0.3	62.8±0.2	80.4
RTN (Long et al. 2016)	84.5±0.2	96.8±0.1	99.4±0.1	77.5±0.3	66.2±0.2	64.8±0.3	81.6
DANN (Ganin et al. 2016)	82.0±0.4	96.9±0.2	99.1±0.1	79.7±0.4	68.2±0.4	67.4±0.5	82.2
ADDA (Tzeng et al. 2017)	86.2±0.5	96.2±0.3	98.4±0.3	77.8±0.3	69.5±0.4	68.9±0.5	82.9
JAN (Long et al. 2013)	86.0±0.4	96.7±0.3	99.7±0.1	85.1±0.4	69.2±0.3	70.7±0.5	84.6
MADA (Pei et al. 2018)	90.0±0.1	97.4±0.1	99.6±0.1	87.8±0.2	70.3±0.3	66.4±0.3	85.2
SimNet (Pinheiro 2018)	88.6±0.5	98.2±0.2	99.7±0.2	85.3±0.3	73.4±0.8	71.8±0.6	86.2
MCD (Saito et al. 2018)	89.6±0.2	98.5±0.1	100.0±0.0	91.3±0.2	69.6±0.1	70.8±0.3	86.6
CDAN+E (Long et al. 2018)	94.1±0.1	98.6±0.1	100.0±0.0	92.9±0.2	71.0±0.3	69.3±0.3	87.7
SymNets (Zhang et al. 2019b)	90.8±0.1	98.8±0.3	100.0±0.0	93.9±0.5	74.6±0.6	72.5±0.5	88.4
MDD (Zhang et al. 2019a)	94.5±0.3	98.4±0.1	100.0±0.0	93.5±0.2	74.6±0.3	72.2±0.1	88.9
E-MixNet	93.0±0.3	99.0±0.1	100.0±0.0	95.6±0.2	78.9±0.5	74.7±0.7	90.2

Table 1: Results on Office-31 (ResNet-50)

Method	I→P	P→I	I→C	C→I	C→P	P→C	Avg
ResNet-50 (He et al. 2016)	74.8±0.3	83.9±0.1	91.5±0.3	78.0±0.2	65.5±0.3	91.2±0.3	80.7
DAN (Long et al. 2015)	74.5±0.4	82.2±0.2	92.8±0.2	86.3±0.4	69.2±0.4	89.8±0.4	82.5
DANN (Ganin et al. 2016)	75.0±0.6	86.0±0.3	96.2±0.4	87.0±0.5	74.3±0.5	91.5±0.6	85.0
JAN (Long et al. 2013)	76.8±0.4	88.0±0.2	94.7±0.2	89.5±0.3	74.2±0.3	91.7±0.3	85.8
MADA (Pei et al. 2018)	75.0±0.3	87.9±0.2	96.0±0.3	88.8±0.3	75.2±0.2	92.2±0.3	85.8
CDAN+E (Long et al. 2018)	77.7±0.3	90.7±0.2	97.7±0.3	91.3±0.3	74.2±0.2	94.3±0.3	87.7
SymNets (Zhang et al. 2019b)	80.2±0.3	93.6±0.2	97.0±0.3	93.4±0.3	78.7±0.3	96.4±0.1	89.9
E-MixNet	80.5±0.4	96.0±0.1	97.7±0.3	95.2±0.4	79.9±0.2	97.0±0.3	91.0

Table 2: Results on Image-CLEF (ResNet-50)

31. These domains are *Artistic* (A), *Clipart* (C), *Product* (P), *Real-World* (R). Each domain contains 65 kinds of objects. So there are 12 domain adaptation tasks on Office-Home: A → C, A → P, A → R, ..., R → P.

ImageCLEF-DA¹ is a dataset organized by selecting the 12 common classes shared by three public datasets (domains): *Caltech-256* (C), *ImageNet ILSVRC 2012* (I), and *Pascal VOC 2012* (P). There are six transfer tasks: I→P, P→I, I→C, C→I, C→P, P→C.

Experimental Setup

Following the standard protocol in (Long et al. 2018), all labeled source samples and unlabeled target samples are used in the training process and we report the average classification accuracy based on three random experiments. γ in Eq. (15) is selected from 2, 4, 8, and it is set to 2 for Office-Home, 4 for Office-31, and 8 for Image-CLEF.

ResNet-50 (He et al. 2016) pretrained on ImageNet is employed as the backbone network (G). C, D and C* are all two fully connected layers where the hidden unit is 1024. The mini-batch stochastic gradient descent with momentum 0.9 is employed as the optimizer, and The learning rate is adjusted by $l_i = l_0(1 + \delta i)^{-\beta}$, where i linearly increase from 0 to 1 during the training process, $\delta = 10$, $l_0 = 0.04$. We follow (Zhang et al. 2019a) to employ a progressive strategy for η : $\eta = \frac{2\eta_0}{1 + \exp(\delta * i) - \eta_0}$, η_0 is set to 0.1. The α in e-mixup is set to 0.6 in all experiments.

¹<http://imageclef.org/2014/adaptation/>

Results

The results on *Office-31* are reported in Tabel 1. E-MixNet achieves the best results and exceeds the baselines for 4 of 6 tasks. Compared to the competitive baseline MDD, E-MixNet surpasses it by 4.3% for the difficult task D → A.

The results on *Image-CLEF* are reported in Table 2. E-MixNet significantly outperforms the baselines for 5 of 6 tasks. For the hard task C → P, E-MixNet surpasses the competitive baseline SymNets by 2.7%.

The results on *Office-Home* are reported in Table 3. Despite *Office-Home* is a challenging dataset, E-MixNet still achieves better performance than all the baselines for 9 of 12 tasks. For the difficult tasks A → C, P → A, and R → C, E-MixNet has significant advantages.

Furthermore, we add the proxy into the loss functions of four representative UDA methods. As shown in Fig. 2(a), we add a new classifier that is the same as the classifier in the original method to formulate the proxy of the combined risk. The results are shown in Table 4. The four methods can achieve better performance after optimizing the proxy. The experiments adequately demonstrate the proxy can indeed curb the increase of the combined risk.

Ablation Study and Parameter Analysis

Ablation Study. To further verify the efficacy of the proxy of combined risk calculated by mixup and e-mixup respectively. Ablation experiments are shown in Tabel 5, where s indicates that the source samples are introduced to augment the target samples, t indicates augmenting the target

Method	A→C	A→P	A→R	C→A	C→P	C→R	P→A	P→C	P→R	R→A	R→C	R→P	Avg
ResNet-50 (He et al. 2016)	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
DAN (Long et al. 2015)	54.0	68.6	75.9	56.4	66.0	67.9	57.1	50.3	74.7	68.8	55.8	80.6	64.7
DANN (Ganin et al. 2016)	44.1	66.5	74.6	57.9	62.0	67.2	55.7	40.9	73.5	67.5	47.9	77.7	61.3
JAN (Long et al. 2013)	45.9	61.2	68.9	50.4	59.7	61.0	45.8	43.4	70.3	63.9	52.4	76.8	58.3
CDAN+E (Long et al. 2018)	47.0	69.4	75.8	61.0	68.8	70.8	60.2	47.1	77.9	70.8	51.4	81.7	65.2
SymNets (Zhang et al. 2019b)	46.0	73.8	78.2	64.1	69.7	74.2	63.2	48.9	80.0	74.0	51.6	82.9	67.2
MDD (Zhang et al. 2019a)	54.9	73.7	77.8	60.0	71.4	71.8	61.2	53.6	78.1	72.5	60.2	82.3	68.1
E-MixNet	57.7	76.6	79.8	63.6	74.1	75.0	63.4	56.4	79.7	72.8	62.4	85.5	70.6

Table 3: Results on Office-Home (ResNet-50)

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
DAN (Long et al. 2015)	54.0	68.6	75.9	56.4	66.0	67.9	57.1	50.3	74.7	68.8	55.8	80.6	64.7
DAN+ $\tilde{\Lambda}^{\ell_{mse}}$	57.0	71.0	77.9	59.9	72.6	70.1	58.1	57.1	77.3	72.7	64.7	84.6	68.6
DANN (Ganin et al. 2016)	44.1	66.5	74.6	57.9	62.0	67.2	55.7	40.9	73.5	67.5	47.9	77.7	61.3
DANN+ $\tilde{\Lambda}^{\ell_{mse}}$	50.9	69.6	77.8	61.9	70.7	71.6	60.0	49.5	78.4	71.8	55.7	83.7	66.8
CDAN+E (Long et al. 2018)	47.0	69.4	75.8	61.0	68.8	70.8	60.2	47.1	77.9	70.8	51.4	81.7	65.2
CDAN+E+ $\tilde{\Lambda}^{\ell_{mse}}$	49.5	70.1	77.8	64.3	71.3	74.2	61.6	50.6	80.0	73.5	56.6	84.1	67.8
SymNets (Zhang et al. 2019b)	46.0	73.8	78.2	64.1	69.7	74.2	63.2	48.9	80.0	74.0	51.6	82.9	67.2
SymNets+ $\tilde{\Lambda}^{\ell_{mse}}$	48.8	74.7	79.7	64.9	72.5	75.6	63.9	47.0	80.8	73.9	52.4	83.9	68.2

Table 4: Results on Office-Home (ResNet-50)

s	t	m	e	I→P	P→I	I→C	C→I	C→P	P→C	Avg
				80.2	94.2	96.7	94.7	79.2	95.5	90.1
	✓			79.9	92.2	97.7	93.8	79.4	96.5	89.9
	✓	✓		79.7	93.7	97.5	94.5	79.7	96.2	90.2
✓	✓	✓		79.4	95.0	97.8	94.8	81.4	96.5	90.8
✓	✓		✓	80.5	96.0	97.7	95.2	79.9	97.0	91.0

Table 5: Ablation experiments on Image-CLEF

samples, m denotes mixup, and e denotes e-mixup. Table 5 shows that E-MixNet achieves the best performance, which further shows that the combined risk can be controlled by the proxy $\tilde{\Lambda}^{\ell_{mse}}$.

To study how the parameter γ affects the performance and the efficiencies of *mean square error* (MSE) and cross-entropy for the proxy of combined risk. Firstly, as shown in Fig. 3(a), a relatively larger γ can obtain better performance and faster convergence. Secondly, when *mixup* behaves between two samples, the accuracy of the pseudo labels of the target samples are much important. To against the adversarial samples, MSE is employed to substitute cross-entropy. As shown in Fig. 3(b), MSE can obtain more stable and better performance. Furthermore, \mathcal{A} -distance is also an important indicator showing the distribution discrepancy, which is defined as $dis_{\mathcal{A}} = 2(1 - 2\epsilon)$ where ϵ is the test error. As shown in Fig. 3 (c), compared with ResNet (R), DANN (D), and CDAN (C), E-MixNet (E) achieves a better performance of adaptation, implying the efficiency of the proposed proxy.

Conclusion

Though numerous UDA methods have been proposed and have achieved significant success, the issue caused by com-




Figure 3: The impact of γ is shown in (a). The impact of the loss functions for the proxy of the combined risk is shown in (b). Comparison of \mathcal{A} -distance.

bin risk has not been brought to the forefront and none of the proposed methods solve the problem. This paper reveals that the combined risk is deeply related to the conditional distribution discrepancy and plays a crucial role for transfer performance. Furthermore, this paper propose a method termed E-MixNet, which employs *enhanced mixup* to calculate a proxy of the combined risk. Experiments show that our method achieves a comparable performance compared with existing state-of-the-art methods and the performance of the four representative methods can be boosted by adding the proxy into their loss functions.

Acknowledgments

The work presented in this paper was supported by the Australian Research Council (ARC) under DP170101632 and FL190100149. The first author particularly thanks the support of UTS-AAII during his visit.

References

- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning* 79(1-2): 151–175.
- Ben-David, S.; Blitzer, J.; Crammer, K.; and Pereira, F. 2007. Analysis of representations for domain adaptation. In *NeurIPS*, 137–144.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 5049–5059.
- Dong, J.; Cong, Y.; Sun, G.; and Hou, D. 2019. Semantic-Transferable Weakly-Supervised Endoscopic Lesions Segmentation. In *ICCV*, 10711–10720.
- Dong, J.; Cong, Y.; Sun, G.; Liu, Y.; and Xu, X. 2020a. CSCL: Critical Semantic-Consistent Learning for Unsupervised Domain Adaptation. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J.-M., eds., *ECCV*, 745–762. Cham: Springer International Publishing. ISBN 978-3-030-58598-3.
- Dong, J.; Cong, Y.; Sun, G.; Yang, Y.; Xu, X.; and Ding, Z. 2020b. Weakly-Supervised Cross-Domain Adaptation for Endoscopic Lesions Segmentation. *IEEE Transactions on Circuits and Systems for Video Technology* 1–1. doi: 10.1109/TCSVT.2020.3016058.
- Dong, J.; Cong, Y.; Sun, G.; Zhong, B.; and Xu, X. 2020c. What Can Be Transferred: Unsupervised Domain Adaptation for Endoscopic Lesions Segmentation. In *CVPR*, 4022–4031.
- Fang, Z.; Lu, J.; Liu, F.; Xuan, J.; and Zhang, G. 2020. Open set domain adaptation: Theoretical bound and algorithm. *IEEE Transactions on Neural Networks and Learning Systems*.
- Fang, Z.; Lu, J.; Liu, F.; and Zhang, G. 2019. Unsupervised domain adaptation with sphere retracting transformation. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17: 2096–2030.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *NeurIPS*, 2672–2680. Curran Associates, Inc.
- Guo, H.; Pasunuru, R.; and Bansal, M. 2020. Multi-Source Domain Adaptation for Text Classification via DistanceNet-Bandits. In *AAAI*, 7830–7838.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Lee, S.; and Jha, R. 2019. Zero-shot adaptive transfer for conversational language understanding. In *AAAI*, volume 33, 6642–6649.
- Liu, F.; Lu, J.; Han, B.; Niu, G.; Zhang, G.; and Sugiyama, M. 2019. Butterfly: A panacea for all difficulties in wildly unsupervised domain adaptation. *arXiv preprint arXiv:1905.07720*.
- Long, M.; Cao, Y.; Wang, J.; and Jordan, M. I. 2015. Learning transferable features with deep adaptation networks. In *ICML*, 97–105.
- Long, M.; Cao, Z.; Wang, J.; and Jordan, M. I. 2018. Conditional adversarial domain adaptation. In *NeurIPS*, 1640–1650.
- Long, M.; Wang, J.; Ding, G.; Sun, J.; and Yu, P. S. 2013. Transfer feature learning with joint distribution adaptation. In *ICCV*, 2200–2207.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2016. Unsupervised domain adaptation with residual transfer networks. In *NeurIPS*, 136–144.
- Lu, J.; Behbood, V.; Hao, P.; Zuo, H.; Xue, S.; and Zhang, G. 2015. Transfer learning using computational intelligence: A survey. *Knowledge-Based Systems* 80: 14–23.
- Lu, W.; Yu, Y.; Chang, Y.; Wang, Z.; Li, C.; and Yuan, B. 2020. A Dual Input-aware Factorization Machine for CTR Prediction. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009. Domain Adaptation: Learning Bounds and Algorithms. In *COLT*.
- Mohri, M.; and Medina, A. M. 2012. New analysis and algorithm for learning with drifting distributions. In *ALT*, 124–138. Springer.
- Pei, Z.; Cao, Z.; Long, M.; and Wang, J. 2018. Multi-adversarial domain adaptation. *arXiv preprint arXiv:1809.02176*.
- Pinheiro, P. O. 2018. Unsupervised domain adaptation with similarity learning. In *CVPR*, 8004–8013.
- Redko, I.; Morvant, E.; Habrard, A.; Sebban, M.; and Ben-nani, Y. 2020. A survey on domain adaptation theory. *arXiv preprint arXiv:2004.11829*.
- Saenko, K.; Kulis, B.; Fritz, M.; and Darrell, T. 2010. Adapting visual category models to new domains. In *ECCV*, 213–226. Springer.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 3723–3732.
- Shen, J.; Qu, Y.; Zhang, W.; and Yu, Y. 2018. Wasserstein distance guided representation learning for domain adaptation. In *AAAI*.
- Tang, H.; and Jia, K. 2020. Discriminative Adversarial Domain Adaptation. In *AAAI*, 5940–5947.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial discriminative domain adaptation. In *CVPR*, 7167–7176.
- Vapnik, V. N.; and Chervonenkis, A. Y. 2015. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, 11–30. Springer.

- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 5018–5027.
- Wang, Q.; and Breckon, T. P. 2020. Unsupervised Domain Adaptation via Structured Prediction Based Selective Pseudo-Labeling. In *AAAI*, 6243–6250. AAAI Press.
- Xu, M.; Zhang, J.; Ni, B.; Li, T.; Wang, C.; Tian, Q.; and Zhang, W. 2020. Adversarial Domain Adaptation with Domain Mixup. In *AAAI*, 6502–6509. AAAI Press.
- Yu, Y.; Wang, Z.; and Yuan, B. 2019. An Input-aware Factorization Machine for Sparse Prediction. In *IJCAI*, 1466–1472.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *ICLR*.
- Zhang, Q.; Wu, D.; Lu, J.; Liu, F.; and Zhang, G. 2017. A cross-domain recommender system with consistent information transfer. *Decision Support Systems* 104: 49–63.
- Zhang, Y.; Deng, B.; Tang, H.; Zhang, L.; and Jia, K. 2020a. Unsupervised multi-class domain adaptation: Theory, algorithms, and practice. *arXiv preprint arXiv:2002.08681*.
- Zhang, Y.; Liu, F.; Fang, Z.; Yuan, B.; Zhang, G.; and Lu, J. 2020b. Clarinet: A One-step Approach Towards Budget-friendly Unsupervised Domain Adaptation. *arXiv preprint arXiv:2007.14612*.
- Zhang, Y.; Liu, T.; Long, M.; and Jordan, M. 2019a. Bridging Theory and Algorithm for Domain Adaptation. In Chaudhuri, K.; and Salakhutdinov, R., eds., *ICML*, volume 97 of *PMLR*, 7404–7413. PMLR.
- Zhang, Y.; Tang, H.; Jia, K.; and Tan, M. 2019b. Domain-symmetric networks for adversarial domain adaptation. In *CVPR*, 5031–5040.
- Zhao, H.; des Combes, R. T.; Zhang, K.; and Gordon, G. 2019. On Learning Invariant Representation for Domain Adaptation. *ICML*.
- Zhong, L.; Fang, Z.; Liu, F.; Yuan, B.; Zhang, G.; and Lu, J. 2020. Bridging the Theoretical Bound and Deep Algorithms for Open Set Domain Adaptation. *arXiv preprint arXiv:2006.13022*.
- Zou, H.; Zhou, Y.; Yang, J.; Liu, H.; Das, H. P.; and Spanos, C. J. 2019. Consensus adversarial domain adaptation. In *AAAI*, volume 33, 5997–6004.