

SimT: Handling Open-set Noise for Domain Adaptive Semantic Segmentation

Xiaoqing Guo¹Jie Liu¹Tongliang Liu²Yixuan Yuan^{1*}¹City University of Hong Kong²University of Sydney

{xqguo.ee, jliu.ee}@my.cityu.edu.hk tongliang.liu@sydney.edu.au yxyuan.ee@cityu.edu.hk

Abstract

This paper studies a practical domain adaptive (DA) semantic segmentation problem where only pseudo-labeled target data is accessible through a black-box model. Due to the domain gap and label shift between two domains, pseudo-labeled target data contains mixed closed-set and open-set label noises. In this paper, we propose a simplex noise transition matrix (SimT) to model the mixed noise distributions in DA semantic segmentation and formulate the problem as estimation of SimT. By exploiting computational geometry analysis and properties of segmentation, we design three complementary regularizers, i.e. volume regularization, anchor guidance, convex guarantee, to approximate the true SimT. Specifically, volume regularization minimizes the volume of simplex formed by rows of the non-square SimT, which ensures outputs of segmentation model to fit into the ground truth label distribution. To compensate for the lack of open-set knowledge, anchor guidance and convex guarantee are devised to facilitate the modeling of open-set noise distribution and enhance the discriminative feature learning among closed-set and open-set classes. The estimated SimT is further utilized to correct noise issues in pseudo labels and promote the generalization ability of segmentation model on target domain data. Extensive experimental results demonstrate that the proposed SimT can be flexibly plugged into existing DA methods to boost the performance. The source code is available at <https://github.com/CityU-AIM-Group/SimT>.

1. Introduction

DA semantic segmentation aims to adapt a segmentation model for the target domain through transferring knowledge from a source domain [22]. Considering the inefficient transmission and privacy issues in source domain knowledge, effective adaptation with only pseudo-labeled target data obtained through a black-box model is practical in real-world deployments [19]. This paper mainly focuses on this realistic scenario, and the challenge of this DA

semantic segmentation problem lies in noisy pseudo labels generated for target domain data. To mitigate noisy adaptation, some methods leverage confidence score [13, 25, 50] or uncertainty [17, 46] to select reliable pseudo labels. For example, prior work [50] defines a threshold to eliminate low-confidence pseudo-labeled pixels. However, simply dropping *confusing pixels* with highly potential noise labels might lead to learn from a biased data distribution [48]. In fact, confusing pixels are more prone to be mislabeled and critical for accurate predictions, hence need to be carefully addressed. Instead of discarding unreliable pseudo-labeled target data, other methods denoise pseudo labels based on prototypes [42] or correct supervision signal through noise transition matrix (NTM) [11]. Nevertheless, these approaches heavily rely on the assumption of a shared label set between two domains and only consider the closed-set noisy pseudo labels. In real-world scenarios, the label space of target domain may not be consistent with that of source domain, and *open-set* class instances are commonly encountered in target domain. Hence, enhancing the segmentation performance on target domain needs to carefully address *confusing pixels* and *open-set label noises*.

With the aforementioned aims, we model the closed-set and open-set noise distributions in target pseudo labels through a SimT, which is a class-dependent and instance-independent transition matrix^a. Then the modeled noise distribution is leveraged to rectify supervision signals (*i.e.* loss) derived from noisy labels, thus all target samples are fully utilized. Therefore, *the problem of alleviating closed-set and open-set noise issues can be treated as the problem of estimating SimT*. To give an intuitive explanation of SimT, we interpret SimT through a geometric analysis on a toy example in Figure 1. In this toy example, SimT is defined as a matrix $\mathbf{T} \in [0, 1]^{5 \times 3}$, as illustrated in Figure 1 (a). Each row, *i.e.* $\mathbf{T}_{c,:}$, indicates the probability of label c flipping to three classes, which can be regarded as a *three-dimensional point* within the *blue triangle* of Figure 1 (b) since coordinates of this point add up to 1. The first three rows $[\mathbf{T}_{1,:}, \mathbf{T}_{2,:}, \mathbf{T}_{3,:}]$ of SimT model closed-set noise dis-

*Yixuan Yuan is the corresponding author.

This work was supported by Hong Kong Research Grants Council (RGC) General Research Fund 11211221(CityU 9043152).

^aThe complex instance-dependent label noise in real-world can be well approximated by class-dependent label noise when the noise rate is low [4].




Figure 1. Geometric illustration of SimT with the setting of 3 pseudo label classes and 5 ground truth label classes. Note that coordinates (x, y, z) of any points in **blue triangle** satisfy conditions of $x + y + z = 1$ and $0 < x, y, z < 1$.

tribution, and the rest rows $[T_{4,:}, T_{5,:}]$ model open-set noise distribution. Given an input instance \mathbf{x} , the noisy class posterior probability $p(\tilde{\mathbf{y}} | \mathbf{x}) = [p(\tilde{y} = 1 | \mathbf{x}), p(\tilde{y} = 2 | \mathbf{x}), p(\tilde{y} = 3 | \mathbf{x})]$ can be regarded a **red point** scattered in the **blue triangle**. Since SimT bridges noisy class posterior and clean class posterior $p(\mathbf{y} | \mathbf{x}) = [p(y = 1 | \mathbf{x}), p(y = 2 | \mathbf{x}), \dots, p(y = 5 | \mathbf{x})]$ through $p(\tilde{\mathbf{y}} | \mathbf{x}) = p(\mathbf{y} | \mathbf{x}) \mathbf{T}$, the noisy class posterior probability $p(\tilde{\mathbf{y}} | \mathbf{x})$ is interpreted as a convex combination among rows of \mathbf{T} with the factor of $p(\mathbf{y} | \mathbf{x})$. In the *three-dimensional* space, the closer $p(\tilde{\mathbf{y}} | \mathbf{x})$ is to the vertex $T_{c,:}$ of **green pentagon** $Sim\{\mathbf{T}\}$, the more likely the input \mathbf{x} belongs to class c . Hence, the simplex formed by rows of \mathbf{T} should enclose $p(\tilde{\mathbf{y}} | \mathbf{x})$ for any \mathbf{x} . However, since there exists an infinite number of simplices enclosing $p(\tilde{\mathbf{y}} | \mathbf{x})$, SimT is not identifiable if no assumption has been made.

A feasible solution to estimate SimT is volume minimization based on the sufficiently scattered assumption where class posterior distribution is far from uniform [16]. As demonstrated in [16], finding the minimum-volume simplex enclosing $p(\tilde{\mathbf{y}} | \mathbf{x})$ recovers the ground truth NTM and $p(\mathbf{y} | \mathbf{x})$. In semantic segmentation, interior pixels away from object edge usually exhibit high confidence for classification. Those confident instances show discriminative distribution and thus have class posteriors far from uniform. In this regards, the sufficiently scattered assumption is easy to satisfy in segmentation, hence we follow the idea of volume minimization to estimate SimT. Nevertheless, NTM in [16] only models closed-set noisy labels, and directly applying volume minimization that constrains a square NTM to our SimT is problematic with two main reasons. Firstly, SimT is a non-square matrix, whose volume cannot be calculated through its determination. Secondly, the success of volume minimization depends on the diagonally dominant guarantee of the square matrix, while the diagonally dominant assumption is not available for the open-set part of SimT. Hence, by minimizing volume of SimT, open-set part of SimT will collapse to a trivial solution (Figure 1 (a) and $Sim\{\mathbf{T}''''\}$ in (c)), leading to indiscriminating open-set feature distributions and non-convex problem. These motivate us to design a computational geometry based algorithm

which can detect open-set noises from pseudo-labeled target data and approximate the true SimT.

To this end, we propose three regularizers, *i.e.* *volume regularization*, *anchor guidance* and *convex guarantee*, for the estimation of SimT. Specifically, *volume regularization* for non-square SimT estimation is designed to approximate the true SimT, which is the one with minimum volume ($Sim\{\mathbf{T}\}$) among all simplices enclosing $p(\tilde{\mathbf{y}} | \mathbf{x})$ (e.g., $Sim\{\mathbf{T}'\}$) in Figure 1 (c). The minimum volume of SimT ensures outputs of segmentation model to fit into clean class posteriors. To prevent the trivial solution of open-set part in SimT, we devise *anchor guidance* regularization. In particular, we first define **anchor point** as the most representative pixel of each class with the largest clean class posterior probability as in Figure 1 (c). $Sim\{\mathbf{T}''\}$ formed by all anchor points is an approximation to the true SimT, hence we leverage anchor points to guide the estimation of SimT. To further facilitate the detection of anchor points in open-set classes, an auxiliary loss enhancing the discriminative feature learning among closed-set and open-set data is designed. Moreover, the trivial solution of open-set part may lead to non-convex $Sim\{\mathbf{T}''''\}$. Hence, we advance *convex guarantee* to preserve the convex property of SimT, and the convex SimT can implicitly push the learned feature distribution of open-set classes away from closed-set classes.

We summarize our contributions in three aspects.

- We present a general and practical DA semantic segmentation framework to robustly learn from noisy pseudo-labeled target data, including closed-set and open-set label noises. To our best knowledge, it represents the first attempt to address the problem of open-set noisy label in DA.
- We, for the first time, model the closed-set and open-set noise distributions of target pseudo labels through SimT, and utilize computational geometry analysis to estimate it. The optimized SimT is used to benefit loss correction for noisy pseudo-labeled target data.
- Extensive experiments are conducted to verify the effect of the proposed SimT and demonstrate that it can be plugged into existing DA methods, such as UDA or SFDA, to boost the performance.

2. Related Work

2.1. Domain Adaptive Semantic Segmentation

DA semantic segmentation aims to adapt a model for the target domain through transferring knowledge from a source domain. Unsupervised domain adaptation (UDA) methods leverage adversarial learning to align data distributions of two domains [5, 14, 34]. Source-free domain adaptation (SFDA) methods [15, 23, 28, 45] mine useful knowledge embedded in source model to perform DA without source data. These DA works require the coexistence of source data (or the detailed specification of source models) and tar-

get data since concurrent access can help to characterize the domain gap [15]. However, it is not always practical in real-world deployment scenarios, where the source domain data transmission may be inefficient when it is extremely large [45], data sharing may be restrained due to privacy [28] or the specific knowledge of source models may be not accessible [19]. Different from UDA and SFDA, we study a more general and practical DA problem where only pseudo-labeled target data is accessible through a black-box model.

Without specific requirements of source domain data or model, self-training strategy is already a common choice to adapt the model for unlabeled target data [11, 13, 23, 25, 42, 46, 50]. It optimizes the segmentation model with pseudo-labeled target data, which enhances the discriminative feature representation of target data [11]. To tackle the target pseudo label noise issue, some methods filter out noisy pixels with respect to confidence or uncertainty [13, 23, 25, 46]. Kundu *et al.* [23] train the source model with virtually extended multi-source data and select top 33% confident target pixels for self-training. Zheng *et al.* [46] utilize uncertainty estimation to enable dynamic threshold. Instead of filtering out potentially noisy pseudo-labeled target data, other methods denoise pseudo labels [42] or correct supervision signal [11]. Zhang *et al.* [42] learn segmentation net from denoised pseudo labels according to class prototypes. Guo *et al.* [11] model the noise distribution of pseudo labels by NTM and solve it in a meta-learning strategy.

However, almost all self-training based methods ignore the open-set pseudo label noises in target domain, which are commonly encountered in practice [20]. Although confidence or uncertainty based self-training methods omit potentially noisy data and may reject some open-set pseudo label noises, simply dropping pixels will lead to a biased data distribution learning especially when pixels lie around the decision boundary. In this paper, we model the closed-set and open-set noise distributions of pseudo labels in target domain with a SimT, and optimize the segmentation model with corrected supervision signals on all target data.

2.2. Deep Learning with Noisy Labels

Types of noisy labels can be divided into two categories: **Closed-set noisy labels.** ‘Closed-set’ means that true class labels of noisy labeled samples are within known classes of noisy data. Many efforts have been devoted to tackling the closed-set noisy label issue, such as *sample reweighting* [33], *label correction* [44], *loss correction* [11, 38]. Shu *et al.* [33] leverage multiple layer perceptions to automatically suppress noisy samples. Zhang *et al.* [44] conduct a dynamic linear combination of noisy label and prediction to refurbish labels. Wang *et al.* [38] utilize a set of trusted clean-labeled samples to estimate NTM for loss correction.

Open-set noisy labels. ‘Open-set’ indicates that true class labels of noisy labeled samples are outside known classes of noisy data. Contrast to the widely explored

closed-set noisy label problem, it has been seldom invested, which is our focus in this study. Prior works concentrate on *resampling* approach [20, 31, 36]. This solution is intuitive but suboptimal as the discarded data may contain useful information for shaping the decision boundary. *Label correction* and *loss correction* are vulnerable with open-set label noises since these methods are based on the assumption that training samples belong to known classes in the dataset. Hence, Yu *et al.* [41] transfer the ‘open-set’ problem to ‘closed-set’ one and then correct noisy labels. Xia *et al.* [39] extend the traditional NTM with an open-set class, so as to model the open-set label noises. Our work represents the first effort to combat open-set noisy labels in DA.

3. Methodology

3.1. Self-training with SimT

We focus on the problem of domain adaptive semantic segmentation, wherein we only have access to target domain images $\mathcal{X}_{\mathcal{T}} = \{X_t \in \mathbb{R}^{H \times W \times 3}\}_{t \in \mathcal{T}}$ with spatial dimension of $H \times W$ and the corresponding pseudo labels $\mathcal{Y}_{\mathcal{T}} = \{\tilde{Y}_t \in \{0, 1\}^{H \times W \times C}\}_{t \in \mathcal{T}}$. Note that Y_t is the pixel-wise label map with C classes, and pseudo labels are derived from a given black-box model $f_b(\cdot)$. We aim to learn a segmentation model $f(\cdot)_{\mathbf{w}}$ parameterized by \mathbf{w} that can correctly categorize pixels for target data $\mathcal{X}_{\mathcal{T}}$. Self-training based methods [11, 42, 46, 50] utilize pseudo labels as approximated ground truth labels for model training, and the cross-entropy loss over the target dataset for self-training is $\mathcal{L}_{ST} = -\sum_{t \in \mathcal{T}} \tilde{Y}_t \log f(X_t)_{\mathbf{w}}$. By minimizing \mathcal{L}_{ST} on target data, the optimized model can be discriminatively adapted to the target domain. However, noises in pseudo-labeled target domain data, including closed-set label noises and open-set label noises, inevitably deteriorate the performance of self-training based DA methods.

To enhance the noise tolerance property of \mathcal{L}_{ST} , we propose SimT to benefit loss correction, as shown in Figure 2. SimT is formulated as $\mathbf{T} \in [0, 1]^{(C+n) \times C}$, where $\mathbf{T}_{1:C,:}$: models closed-set noise distribution and $\mathbf{T}_{C+1:C+n,:}$: models open-set noise distribution in target domain. n is a hyper-parameter that indicates the potential open-set class number, and we set $n > 1$ to implicitly model the diverse semantics within open-set classes. The defined SimT specifies the probability of ground truth label ($Y_t = j$) flipping to noisy label ($\tilde{Y}_t = k$) by $\mathbf{T}_{jk} = p(\tilde{Y}_t = k | Y_t = j)$, and we have $\sum_{k=1}^C \mathbf{T}_{jk} = 1$. Assuming the class posterior probability for noisy pseudo label is $p(\tilde{Y}_t = k | X_t; \mathbf{w}) \in [0, 1]^C$ and for ground truth label is $p(Y_t = j | X_t; \mathbf{w}) \in [0, 1]^{C+n}$, SimT bridges the generated pseudo label \tilde{Y}_t and the ground truth label Y_t through:

$$p(\tilde{Y}_t = k | X_t; \mathbf{w}) = \sum_{j=1}^{C+n} \mathbf{T}_{jk} \cdot p(Y_t = j | X_t; \mathbf{w}), \quad (1)$$

$$\Rightarrow p(\tilde{Y}_t | X_t; \mathbf{w}) = p(Y_t | X_t; \mathbf{w}) \mathbf{T}.$$

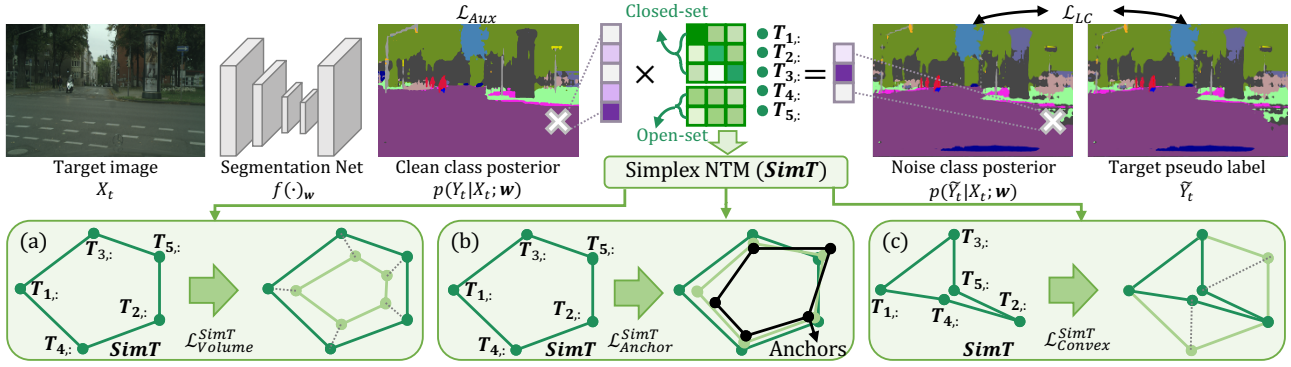


Figure 2. The proposed domain adaption framework contains a segmentation net and a SimT. Target images are passed through the segmentation net $f(\cdot)_w$ to perform semantic segmentation, and supervision signals are corrected by the learnable SimT. SimT is estimated through the proposed (a) volume regularization, (b) anchor guidance and (c) convex guarantee.

With the formulated noise distribution (\mathbf{T}), we correct the self-training loss over noisy pseudo-labeled target data as

$$\mathcal{L}_{LC} = - \sum_{t \in \mathcal{T}} \tilde{Y}_t \log[f(X_t)_w \mathbf{T}]. \quad (2)$$

This corrected loss encourages the similarity between noisy class posterior probabilities $p(\tilde{Y}_t | X_t; \mathbf{w}) = f(X_t)_w \mathbf{T}$ and noisy pseudo labels \tilde{Y}_t . It is obvious that once the true SimT is obtained, we can recover the desired estimation of clean class posterior probability $p(Y_t | X_t; \mathbf{w})$ by the softmax of first C classes output $f(X_t)_w$ even training the segmentation model with noisy data. Considering no clean-labeled target domain data is accessible, previous work estimates NTM through a meta-learning strategy [11]. Different from prior art, we formulate NTM to make up a simplex and solve it through computational geometry analysis without the requirement of a clean labeled target data set.

3.2. Volume Regularization of SimT

Semantic segmentation can be regarded as pixel-wise classification, thus we make geometry analysis at pixel level. Given a pixel \mathbf{x} within target domain image X_t , the noisy class posterior probability $p(\tilde{\mathbf{y}} | \mathbf{x}) = [p(\tilde{y} = 1 | \mathbf{x}), p(\tilde{y} = 2 | \mathbf{x}), \dots, p(\tilde{y} = C | \mathbf{x})]$ can be regarded as a point in C -dimensional space, and the clean class posterior probability $p(\mathbf{y} | \mathbf{x}) = [p(y = 1 | \mathbf{x}), p(y = 2 | \mathbf{x}), \dots, p(y = C + n | \mathbf{x})]$ is a $(C + n)$ dimensional vector. These two class posteriors can be bridged by the proposed SimT via $p(\tilde{\mathbf{y}} | \mathbf{x}) = p(\mathbf{y} | \mathbf{x}) \mathbf{T}$ as in Eq. (1). In the meanwhile, conditions of $\sum_{j=1}^{C+n} p(y = j | \mathbf{x}) = 1$ and $p(\mathbf{y} | \mathbf{x}) > 0$ are satisfied, thus, the noisy class posterior probability $p(\tilde{\mathbf{y}} | \mathbf{x})$ can be interpreted as a convex combination of \mathbf{T} rows with factor of $p(\mathbf{y} | \mathbf{x})$. In other words, the simplex formed by rows of \mathbf{T} encloses $p(\tilde{\mathbf{y}} | \mathbf{x})$ for any input pixel \mathbf{x} [16]. However, SimT is not identifiable if no assumption is made since there exists an infinite number of simplices enclosing $p(\tilde{\mathbf{y}} | \mathbf{x})$.

To approximate the true SimT, we follow the sufficiently scattered assumption [16] and devise volume regularization.

Among all simplices enclosing $p(\tilde{\mathbf{y}} | \mathbf{x})$, the true SimT should be the one with minimum volume under the sufficiently scattered assumption. In this circumstance, the volume consisted by the target clean class posterior will converge the maximum one, and outputs of optimized segmentation model thus approximate to the ground truth label distribution. This motivates us to propose the following volume regularization for the optimization of SimT:

$$\mathcal{L}_{Volume}^{SimT} = \log[\text{Vol}(\mathbf{T})] = \log \sqrt{\det(\mathbf{T}^T \mathbf{T})} / (N!) \quad (3)$$

where $\sqrt{\det(\mathbf{T}^T \mathbf{T})} / (N!)$ denotes the volume of a r -parallelogram defined by the non-square matrix \mathbf{T} in geometry [8, 10]. Since $N!$ is a constant number, we discard this denominator in experiments. $\log(\cdot)$ function is to stabilize the optimization for computed determinant, following [8].

To make \mathbf{T} differentiable and satisfy the condition of $\mathbf{T} \in [0, 1]^{(C+n) \times C}$, $\sum_{k=1}^C \mathbf{T}_{jk} = 1$, we utilize the reparameterization method. Specifically, we first randomly initialize a matrix $\mathbf{U} \in \mathbb{R}^{(C+n) \times C}$. To preserve the diagonally dominant property of closed-set part of SimT ($\mathbf{T}_{1:C, :}$), the diagonal prior \mathbf{I} is introduced, which has been widely used in the literature of NTM estimation [16, 26]. Considering segmentation model tends to classify samples into majority categories, the class distribution of pseudo labels \mathbf{C} is involved. Incorporating these prior informations, we obtain $\mathbf{V} = \mathbf{C} \cdot \sigma(\mathbf{U}) + \mathbf{I}$, where $\sigma(\cdot)$ is the sigmoid function to avoid negative value in SimT. Then we do the normalization of $\mathbf{T}_{jk} = \frac{\mathbf{V}_{jk}}{\sum_{k=1}^C \mathbf{V}_{jk}}$ to derive SimT (\mathbf{T}). Since both sigmoid function and normalization operation are differentiable, SimT can be updated through gradient descent on \mathbf{U} .

3.3. Anchor Guidance of SimT

Considering the diagonally dominant assumption is not available for the open-set part $\mathbf{T}_{C:C+n, :}$ of SimT, the open-set part will eventually collapse to a trivial solution by minimizing the volume regularization $\mathcal{L}_{Volume}^{SimT}$. To prevent the trivial solution of open-set part in SimT, we devise anchor guidance, where anchor points are detected to guide

the estimation of SimT. Formally, anchor point of class c is defined as the pixel whose clean class posterior are $p(y = c | \mathbf{x}^c) = 1$ and $p(y \neq c | \mathbf{x}^c) = 0$. Thus we have $p(\tilde{\mathbf{y}} = k | \mathbf{x}^c) = \sum_{j=1}^{C+n} \mathbf{T}_{jk} \cdot p(y = j | \mathbf{x}^c) = \mathbf{T}_{ck}$, indicating that the noisy class posterior probability of anchor point $p(\tilde{\mathbf{y}} | \mathbf{x}^c)$ lies at one vertex of SimT ($\mathbf{T}_{c,\cdot}$). Hence, the simplex formed by noisy class posteriors of anchor points can recover the true SimT in theory. However, under mixed open-set and closed-set noisy labels, anchor points are hard to detect. Therefore, we define anchor point as the most representative pixel within each closed-set and open-set class with the largest clean class posterior probability. The anchor point in class c is thus derived via

$$\mathbf{x}^c = \underset{\mathbf{x}}{\operatorname{argmax}} p(y = c | \mathbf{x}; \mathbf{w}). \quad (4)$$

In geometrical perspective, the estimated \mathbf{x}^c is the point, whose noisy class posterior $p(\tilde{\mathbf{y}} | \mathbf{x}^c)$ is closest to the vertex $\mathbf{T}_{c,\cdot}$ in C -dimensional space. Then we leverage computed anchor points to guide the approximation of the true SimT, and the anchor guidance is formulated as

$$\mathcal{L}_{Anchor}^{SimT} = \sum_{c=1}^{C+n} \sum_{k=1}^C \left\| \mathbf{T}_{ck} - p(\tilde{\mathbf{y}} = k | \mathbf{x}^c; \mathbf{w}^{fixed}) \right\|^2, \quad (5)$$

where the noisy class posterior of anchor point $p(\tilde{\mathbf{y}} = k | \mathbf{x}^c; \mathbf{w}^{fixed})$ is computed from the segmentation model warmed up with pseudo labels, denoted as $f(\cdot)_{\mathbf{w}^{fixed}}$ with parameter of \mathbf{w}^{fixed} . For each target image, only occurred classes would derive the corresponding anchor points. Then those detected anchor points are leveraged to calculate the regularization of anchor guidance. The accurate detection of anchor points in open-set classes can prevent the trivial solution and guide to approximate the true SimT.

Auxiliary loss for anchor point detection. To facilitate the detection of anchor points in open-set classes, an auxiliary loss enhancing the discriminative feature learning among closed-set and open-set data is designed. Firstly, we introduce a threshold to select confident pixels as non-noisy samples, and class posterior probabilities of those selected pixels are directly supervised to optimize the segmentation model. This direct supervision explicitly enables the model to encode the diverse semantics for both closed-set and open-set classes. Specifically, the confidence score is derived from the segmentation model $f(\cdot)_{\mathbf{w}^{fixed}}$. We select confident closed-set pixels $X_{\mathcal{K}} = \{\mathbf{x}_k\}_{k \in \mathcal{K}}$ through the condition of $\max f(\mathbf{x}_k)_{\mathbf{w}^{fixed}} > \tau_{high}$, and the corresponding labels $\tilde{\mathbf{y}}_k$ are one-hot vectors of $\tilde{\mathbf{y}}_k = \operatorname{argmax} f(\mathbf{x}_k)_{\mathbf{w}^{fixed}}$. Confident open-set pixels $X_{\mathcal{U}} = \{\mathbf{x}_u\}_{u \in \mathcal{U}}$ are those with $\max f(\mathbf{x}_u)_{\mathbf{w}^{fixed}} < \tau_{low}$ & $\operatorname{argmax} f(\mathbf{x}_u)_{\mathbf{w}} > C$, and their one-hot ground truth labels $\tilde{\mathbf{y}}_u$ are derived from $\tilde{\mathbf{y}}_u = \operatorname{argmax} f(\mathbf{x}_u)_{\mathbf{w}}$. Note that τ_{high}, τ_{low} are hyper-parameters and empirically set to be 0.8, 0.2. To further detect open-set class samples, we make the open-set classifier to output the second-largest

probability for closed-set pixels $X_{\mathcal{K}}$, which is constrained by the second term of following auxiliary loss:

$$\begin{aligned} \mathcal{L}_{Aux} = & - \sum_{i \in \mathcal{K} \& \mathcal{U}} \tilde{\mathbf{y}}_i \log f(\mathbf{x}_i)_{\mathbf{w}} \\ & - \lambda \sum_{k \in \mathcal{K}} \tilde{\mathbf{y}}^o \log[f(\mathbf{x}_k)_{\mathbf{w}} \setminus \tilde{\mathbf{y}}_k], \end{aligned} \quad (6)$$

where $\tilde{\mathbf{y}}^o$ is the one-hot form of $\tilde{\mathbf{y}}^o = \operatorname{argmax}_{c \in [C, C+n]} f(\mathbf{x}_k)_{\mathbf{w}}$,

and λ is a hyper-parameter to control the contribution of the second term in auxiliary loss. The first item maintains performance in the closed-set classes and detects open-set class pixels. The second term matches the masked-probability with the most possible open-set class, which prevents the open-set classifier far away from closed-set classifier and facilitates the discovery of open-set class pixels. The discriminative feature learning induced by auxiliary loss simplifies the detection of anchor points, hence enhancing the performance of anchor guidance.

3.4. Convex Guarantee of SimT

As analyzed in §1, the trivial solution of open-set part will lead to non-convex SimT. In other word, one of vertices in SimT may be the convex combination of the rest vertices. This means different ground truth classes may show the similar noise transition probability, and results in multiple clean class posteriors $p(\mathbf{y} | \mathbf{x})$ correspond to the same noisy class posterior $p(\tilde{\mathbf{y}} | \mathbf{x})$. Considering similar noise transition probabilities reveal close feature representations [39], the non-convex SimT may lead the segmentation model to produce undiscriminating features. Hence, we advance convex guarantee to prevent SimT from becoming non-convex. Inspired by the insight that each vertex of the convex hull cannot be represented by any convex combinations of rest vertices, we first introduce a weighting matrix $\mathbf{u} \in [0, 1]^{(C+n) \times (C+n)}$ with constraints of $\mathbf{u}_{j,k=j} = -1$ and $\sum_k \mathbf{u}_{j,k \neq j} = 1$ to denote convex combination coefficients. The reparameterization method is introduced to make the weighting matrix \mathbf{u} differentiable and satisfy its constraints. Then the convex guarantee is formulated as a minimax problem, where the optimization objective is $\max_{\mathbf{T}} \min_{\mathbf{u}} \|\mathbf{u}\mathbf{T}\|^2$. By updating \mathbf{u} to minimize the optimization objective, the weighting matrix learns to represent each vertex with the convex combination of rest vertices. In turn, by updating \mathbf{T} to maximize the optimization objective, each vertex is pushed away from the convex combination of rest vertices, and SimT is encouraged to be convex. For simplification, convex guarantee regularization is formulated as

$$\mathcal{L}_{Convex}^{SimT} = - \|\mathbf{u}\mathbf{T}\|^2, \text{ where } \mathbf{u} = \operatorname{argmin}_{\mathbf{u}} \|\mathbf{u}\mathbf{T}\|^2. \quad (7)$$

The convex guaranteed SimT enforces the class posterior distribution to be scattered, thereby the segmentation model is encouraged to extract discriminative features among closed-set and open-set classes.

Table 1. Results of UDA and SFDA on adapting GTA5 to Cityscapes.

Methods	road	sidewalk	building	wall	fence	pole	t-light	t-sign	veg.	terrain	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
UDA																				
AdaptSegNet (18') [34]	86.5	36.0	79.9	23.4	23.3	23.9	35.2	14.8	83.4	33.3	75.6	58.5	27.6	73.7	32.5	35.4	3.9	30.1	28.1	42.4
LTIR (20') [14]	92.9	55.0	85.3	34.2	31.1	34.9	40.7	34.0	85.2	40.1	87.1	61.0	31.1	82.5	32.3	42.9	0.3	36.4	46.1	50.2
MLSL (20') [13]	89.0	45.2	78.2	22.9	27.3	37.4	46.1	43.8	82.9	18.6	61.2	60.4	26.7	85.4	35.9	44.9	36.4	37.2	49.3	49.0
IntraDA (20') [25]	90.6	37.1	82.6	30.1	19.1	29.5	32.4	20.6	85.7	40.5	79.7	58.7	31.1	86.3	31.5	48.3	0.0	30.2	35.8	46.3
RPLL (21') [46]	90.4	31.2	85.1	36.9	25.6	37.5	48.8	48.5	85.3	34.8	81.1	64.4	36.8	86.3	34.9	52.2	1.7	29.0	44.6	50.3
MetaCorrection (21') [11]	92.8	58.1	86.2	39.7	33.1	36.3	42.0	38.6	85.5	37.8	87.6	62.8	31.7	84.8	35.7	50.3	2.0	36.8	48.0	52.1
DPL-Dual (21') [5]	92.8	54.4	86.2	<u>41.6</u>	32.7	36.4	<u>49.0</u>	34.0	85.8	41.3	86.0	63.2	34.2	87.2	39.3	44.5	<u>18.7</u>	42.6	43.1	53.3
UncerDA (21') [37]	90.5	38.7	86.5	41.1	32.9	40.5	48.2	42.1	86.5	36.8	84.2	64.5	38.1	87.2	34.8	50.4	0.2	41.8	54.6	52.6
DSP (21') [9]	92.4	48.0	87.4	33.4	35.1	36.4	41.6	46.0	87.7	43.2	89.8	66.6	32.1	89.9	<u>57.0</u>	56.1	0.0	44.1	<u>57.8</u>	55.0
ProDA (21') [42]	87.8	56.0	79.7	46.3	44.8	45.6	53.5	53.5	88.6	45.2	82.1	70.7	39.2	88.8	45.5	<u>59.4</u>	1.0	<u>48.9</u>	56.4	57.5
BAPA-Net (21') [22]	94.4	61.0	<u>88.0</u>	26.8	<u>39.9</u>	38.3	46.1	<u>55.3</u>	<u>87.8</u>	<u>46.1</u>	<u>89.4</u>	<u>68.8</u>	<u>40.0</u>	<u>90.2</u>	60.4	<u>59.0</u>	0.0	45.1	54.2	57.4
Ours (SimT)	<u>94.2</u>	<u>60.0</u>	88.5	30.3	39.7	<u>41.2</u>	47.8	60.8	88.6	47.3	89.3	71.5	45.0	90.7	54.2	60.2	0.0	51.8	58.4	58.9
SFDA																				
Source only	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6
Test-time BN (20') [24]	79.5	79.5	81.3	27.2	29.8	74.1	28.0	58.8	25.3	22.2	19.4	11.0	22.6	17.0	24.5	19.9	<u>34.3</u>	2.4	14.7	37.7
SHOT (20') [18]	88.7	34.9	82.1	27.6	22.5	30.9	31.4	24.7	83.6	37.7	76.8	58.1	24.9	83.9	34.1	39.5	6.0	26.8	24.6	44.1
TENT (21') [35]	87.3	79.8	83.8	85.0	39.0	77.7	21.2	57.9	34.7	19.6	24.3	4.5	16.6	20.8	24.9	17.8	25.1	2.0	16.6	38.9
SFDA (21') [23]	84.2	<u>82.7</u>	82.4	80.0	<u>39.2</u>	85.3	25.9	<u>58.7</u>	30.5	22.1	27.5	30.6	21.9	31.5	33.1	22.1	31.1	3.6	27.8	43.2
S4T (21') [28]	89.7	84.4	86.8	85.0	46.7	<u>79.3</u>	<u>39.5</u>	61.2	41.8	29.0	25.7	9.3	36.8	28.2	19.3	26.7	45.1	5.3	11.8	44.8
URMA (21') [30]	92.3	55.2	81.6	30.8	18.8	37.1	17.7	12.1	84.2	35.9	83.8	57.7	24.1	81.7	27.5	44.3	6.9	24.1	40.4	45.1
SFDAseg (21') [15]	<u>91.7</u>	53.4	86.1	37.6	32.1	37.4	38.2	35.6	<u>86.7</u>	48.5	89.9	<u>62.6</u>	34.3	<u>87.2</u>	51.0	<u>50.8</u>	4.2	<u>42.7</u>	<u>53.9</u>	<u>53.4</u>
Ours (SimT)	92.3	55.8	<u>86.3</u>	34.4	31.7	37.8	39.9	41.4	87.1	<u>47.8</u>	<u>88.5</u>	64.7	<u>36.3</u>	87.3	<u>41.7</u>	55.2	0.0	47.4	57.6	54.4

3.5. Optimization summary for SimT

With the proposed regularizers, SimT is optimized to estimate the mixed closed-set and open-set noise distributions, thereby benefiting the loss correction of pseudo-labeled target data. In summary, the overall objective function for our DA semantic segmentation framework is formulated as

$$\mathcal{L}(\mathbf{w}, \mathbf{T}) = \mathcal{L}_{LC}(\mathbf{w}, \mathbf{T}) + \mathcal{L}_{Aux}(\mathbf{w}) + \alpha \mathcal{L}_{Volume}^{SimT}(\mathbf{T}) + \beta \mathcal{L}_{Anchor}^{SimT}(\mathbf{T}) + \gamma \mathcal{L}_{Convex}^{SimT}(\mathbf{T}). \quad (8)$$

where corrected loss $\mathcal{L}_{LC}(\mathbf{w}, \mathbf{T})$ (Eq. (2)) and $\mathcal{L}_{Aux}(\mathbf{w})$ (Eq. (6)) optimize the segmentation model. $\mathcal{L}_{LC}(\mathbf{w}, \mathbf{T})$ and three regularizers $\mathcal{L}_{Volume}^{SimT}(\mathbf{T})$ (Eq. (3)), $\mathcal{L}_{Anchor}^{SimT}(\mathbf{T})$ (Eq. (5)), $\mathcal{L}_{Convex}^{SimT}(\mathbf{T})$ (Eq. (7)) are minimized to approximate the true SimT. α, β, γ are regularization coefficients to balance contributions of three regularizers.

4. Experiments

4.1. Training Details

Datasets. We evaluate the performance of our method on two DA semantic segmentation tasks, including a synthetic-to-real scenario (GTA5 [29]→Cityscapes [7]) and a surgical instrument type segmentation task (Endovis17 [2]→Endovis18 [1]). *GTA5* contains 24,966 images captured from a video game, and 19 classes in pixel-wise annotations are compatible with Cityscapes. *Cityscapes* is a real-world semantic segmentation dataset obtained in driving scenarios, containing 34 classes according to the manual annotations. Specifically, 2,975 unlabeled images are regarded as the target domain training data, and evaluations are performed on 500 validation images with manual annotations. *Endovis17* contains 1800 source images captured from ab-

dominal porcine procedure, and 3 instrument type classes are compatible with Endovis18. *Endovis18* is an abdominal procedure dataset, consisting of 2384 target images with 6 instrument types. Following [21], we divide Endovis18 into 1639 unlabeled training images and 596 validation images.

Implementation Details. We adopt DeepLab-v2 [3] backbone with encoder of ResNet-101 [12] as our segmentation model. Given the pseudo-labeled target domain data derived from a black-box model $f_b(\cdot)$, we warm up the segmentation model to obtain $f(\cdot)_{\mathbf{w}^{fixed}}$ with C -way output probabilities. The classifier of segmentation model is then extended to output $(C+n)$ -way clean class posterior probabilities. The extended model is denoted as $f(\cdot)_{\mathbf{w}}$ we aim to optimize, and we obtain final predictions from first C -way of $f(\cdot)_{\mathbf{w}}$ during inference phase. Our method is implemented with the PyTorch library on Nvidia 2080Ti. We adopt polynomial learning rate scheduling to optimize the feature extractor with the initial learning rate of $6e-4$, while it is set as $6e-3$ for the optimization of classifier and SimT. The batch size is set as 1, and the maximum iteration number is 40000. Hyper-parameters of $n, \lambda, \alpha, \beta, \gamma$ are set as 15, 0.1, 1.0, 1.0, 0.1 in our implementation. Performances are evaluated by the widely utilized metrics, intersection-over-union (IoU) of each class and the mean IoU (mIoU).

4.2. Results on GTA5→Cityscapes

UDA. We first verify the effectiveness of our approach in the GTA5→Cityscapes scenario, and the corresponding comparison results are listed in Table 1 with the first and second best results highlighted in bold and underline. Under the setting of UDA, our black-box model $f_b(\cdot)$ can be any existing UDA models, and we leverage model of [22] to generate pseudo labels for unlabeled target domain data.

Table 2. Results on adapting Endovis17 to Endovis18.

Methods	scissor	needle driver	forceps	mIoU
UDA				
I2I (19') [27]	60.78	15.34	52.55	42.89
FDA (20') [40]	68.25	18.29	50.29	45.61
S2RC (20') [32]	64.48	22.79	49.34	45.54
IntraDA (20') [25]	68.07	25.88	52.63	48.86
ASANet (20') [47]	70.19	11.34	50.70	44.08
EI-MUNIT(21') [6]	62.73	29.83	48.32	46.96
RPLL (21') [46]	65.11	13.17	51.75	43.34
IGNet (21') [21]	<u>70.65</u>	<u>37.60</u>	<u>52.97</u>	<u>53.74</u>
Ours (SimT)	76.24	39.83	58.94	58.34
SFDA				
Source only	45.97	22.08	38.92	35.66
SHOT (20') [18]	<u>68.44</u>	22.90	47.95	46.43
SFDA (21') [23]	66.05	24.90	45.79	45.58
SFDASeg (21') [15]	67.90	<u>35.31</u>	<u>52.46</u>	<u>51.89</u>
Ours (SimT)	75.56	36.46	53.66	55.23

Overall, the proposed SimT method arrives at the state-of-the-art mIoU score of 58.9%, surpassing all other UDA methods by a large margin. In terms of the per class IoU score, the proposed SimT method shows the superior performance. To be specific, we arrive at the best on 10 out of 19 categories and the second best on 3 out of 19 categories. Moreover, SimT outperforming other self-training based methods [9, 11, 13, 22, 25, 42, 46, 49, 50] demonstrates the effectiveness of the proposed method in mitigating the closed-set and open-set label noise problems.

SFDA. In the GTA5→Cityscapes scenario, we then verify the proposed SimT method under the SFDA setting, where the black-box model $f_b(\cdot)$ is borrowed from a SFDA model in [15]. The corresponding comparison results of our method and baseline models are listed in Table 1. It is observed that the proposed SimT surpasses all other SFDA methods with the state-of-the-art mIoU of 54.4%, outperforming the model learned only from the source domain data by a large margin of 18.8% in mIoU.

4.3. Results on Endovis17→Endovis18

UDA. We further verify the effectiveness of SimT in a medical scenario, *i.e.* Endovis17→Endovis18, and the quantitative comparison results are listed in Table 2. We utilize the UDA model of [21] to generate pseudo labels for unlabeled target domain data. Obviously, the proposed SimT method obtains the state-of-the-art performance with 58.12% mIoU, overpassing all other UDA methods by a large margin. This observation proves the impact of the proposed method in medical image analysis.

SFDA. In the SFDA setting of Endovis17→Endovis18, performance of the proposed SimT is evaluate based on that the black-box model $f_b(\cdot)$ is borrowed from [15]. The corresponding comparison results of our method and baseline models in Table 2 show that the proposed SimT is superior to all other SFDA methods with mIoU of 55.23%.

Table 3. Ablation study. ‘Pseudo Label’ denotes we employ pseudo labels generated by the corresponding UDA model.

Method	Pseudo Label	GTA5 → Cityscapes	Δ
AdaptSegNet (CVPR18') [34]	—	42.4	—
Self-training (MRENT, ICCV 19' [49])	AdaptSegNet	45.1	2.7
Self-training (Threshold, ECCV18' [50])	AdaptSegNet	44.4	2.0
Self-training (Uncertainty, IJCV21' [46])	AdaptSegNet	46.1	3.7
Self-training (MetaCorrection, CVPR21' [11])	AdaptSegNet	47.3	4.9
Ours (w/o $\mathcal{L}_{Volume}^{SimT}$)	AdaptSegNet	46.8	4.4
Ours (w/o $\mathcal{L}_{Anchor}^{SimT}$)	AdaptSegNet	46.9	4.5
Ours (w/o \mathcal{L}_{Aux}^{SimT})	AdaptSegNet	46.7	4.3
Ours (w/o $\mathcal{L}_{Convex}^{SimT}$)	AdaptSegNet	47.3	4.9
Ours	AdaptSegNet	48.0	5.6
LTIR (CVPR20') [14]	—	50.2	—
Self-training (MetaCorrection, CVPR21' [11])	LTIR	52.1	1.9
Ours	LTIR	52.3	2.1
DSP (AAAI21'') [9]	—	55.0	—
Self-training (MetaCorrection, CVPR21' [11])	DSP	56.4	1.4
Ours	DSP	57.3	2.3
BAPA-Net (ICCV21') [22]	—	57.4	—
Self-training (MetaCorrection, CVPR21' [11])	BAPA-Net	57.7	0.3
Ours	BAPA-Net	58.9	1.5




Figure 3. Ablation study for volume regularization, anchor guidance and convex guarantee of SimT.

4.4. Ablation Study

Ablation Experiments. To investigate effects of individual components within SimT, we design ablation experiments under UDA setting (GTA5→Cityscapes), as shown in Table 3. In particular, we validate three proposed regularizers, *i.e.* volume regularization, anchor guidance and convex guarantee, and the auxiliary loss devised for anchor guidance is also ablated to validate its effectiveness. We adopt AdaptSegNet [34] as the black-box model to generate pseudo labels for target domain data. It is observed that incorporating SimT (*row* 11) to model mixed closed-set and open-set noise distributions can correct supervision signals on target data and obtain the performance gain with 5.6% increment in mIoU. Ablating each proposed component (*rows* 7-10) induces the performance degradation. It is because that the designed regularizers are complementary and combining all regularizers with properly chosen coefficients (*i.e.* $\alpha = 1.0$, $\beta = 1.0$, $\gamma = 0.1$) can reach a trade-off to better approximate the true SimT, as shown in Figure 3.

Effectiveness of SimT. We then compare SimT with four typical self-training based UDA models [11, 46, 49, 50]. For a fair comparison, all competed methods use AdaptSegNet to derive pseudo labels. As listed in Table 3, the proposed SimT (*row* 11) is superior to other self-training methods, including entropy minimization [49] (*row* 3), threshold [50] (*row* 4) and uncertainty [46] (*row* 5) based sample selection, MetaCorrection based rectification [11] (*row* 6), yielding increments of 2.9%, 3.6%, 1.9%, 0.7% mIoU.




Figure 4. Qualitative results of UDA in GTA5→Cityscapes scenario. (a) Target image, (b) Ground truth, Predictions from (c) source only model, (d) BAPA-Net [22], (e) ours (SimT).




Figure 5. Qualitative results of UDA in Endovis17→Endovis18 scenario. (a) Target image, (b) Ground truth, Predictions from (c) source only model, (d) I2Net [21], (e) ours (SimT).

With entropy minimization, considerable noisy labels inevitably lead to unsatisfactory adaptation. Although self-training methods in [46, 50] discard noisy samples and may filter out open-set label noises, but lose useful information in those discarded pixels, especially when they are around the decision boundary. Both MetaCorrection [11] and SimT distill effective information from all samples. Comparing with MetaCorrection that only models closed-set noise distribution, the proposed SimT additionally ameliorates open-set noise issues and shows the superior performance.

Robustness to Various Types of Noise. To explore the robustness of our SimT, we further evaluate it under different types of label noises. To be specific, we adopt four UDA models, including AdaptSegNet [34], LTIR [14], DSP [9] and BAPA-Net [22], as black-box models to generate noisy pseudo labels. Then SimT is applied to mitigate the closed-set and open-set noise problem within pseudo labels. As illustrated in Table 3, these four UDA models have significantly improved performance with the proposed SimT.

4.5. Visualization Results

Segmentation Visualization. As illustrated in Figure 4, we provide some typical qualitative segmentation results of target data in GTA5→Cityscapes scenario. Compare with the baseline BAPA-Net [22], the proposed SimT shows superior performance in identifying small-scale objectives (e.g., ‘traffic sign’, rows 2, 3, 5) and differentiating confused class samples (e.g., ‘rider’, ‘person’ and ‘bike’, rows

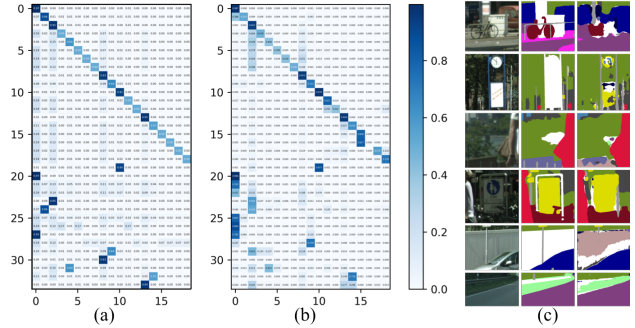


Figure 6. Visualization results. (a) The learned SimT, (b) confusion matrix of pseudo labels, (c) open-set class visualization.

1, 4). In practice, we observe that label noises in pseudo-labeled target data tend to appear in the minor categories and ambiguous categories, leading to the unsatisfactory performance of baselines. On the contrary, the proposed SimT corrects supervision signals for target data and alleviates the noise issues. We also illustrate qualitative results of target Endovis18 data under UDA setting, as shown in Figure 5.

SimT Visualization. Inspired by the fact that confusion matrix is a crude estimator of the true NTM [43], we visualize the learned SimT and the confusion matrix of pseudo labels for reference, as illustrated in Figure 6 (a, b). It is obvious that the optimized SimT is similar to confusion matrix, demonstrating the proposed SimT can model the closed-set and open-set noise distributions and generate corrected supervision signals for pseudo-labeled target data.

Open-set Class Visualization. We further visualize detected open-set class regions, which are denoted by ‘white’ areas in Figure 6 (c). From left to right, they are input image, ground truth and clean class posterior prediction, respectively. We observe the proposed method is able to detect open-set class pixels and inherently has the capability of alleviating open-set label noise issue in target data.

5. Conclusion

In this paper, we study a general and practical DA semantic segmentation setting where only pseudo-labeled target data is accessible through a black-box model. SimT is advanced to model mixed close-set and open-set noise distributions within pseudo-labeled target data. We formulate columns of SimT to make up a simplex and interpret it through the computational geometry analysis. Based on the geometry analysis, we devise three regularizers, *i.e.* volume regularization, anchor guidance, convex guarantee, to approximate the true SimT, which is further utilized to correct noise issues in pseudo labels and enhance the generalization ability of segmentation model on the target domain data. Extensive experimental results demonstrate that the proposed SimT can be flexibly plugged into existing UDA and SFDA methods to boost the performance. In the future, we plan to introduce open-set class prior to enable the semantic separation inner detected open-set regions.

References

- [1] Max Allan, Satoshi Kondo, Sebastian Bodenstedt, Stefan Leger, Rahim Kadkhodamohammadi, Imanol Luengo, Felix Fuentes, Evangello Flouty, Ahmed Mohammed, Marius Pedersen, et al. 2018 robotic scene segmentation challenge. *arXiv preprint arXiv:2001.11190*, 2020. **6**
- [2] Max Allan, Alex Shvets, Thomas Kurmann, Zichen Zhang, Rahul Duggal, Yun-Hsuan Su, Nicola Rieke, Iro Laina, Niveditha Kalavakonda, Sebastian Bodenstedt, et al. 2017 robotic instrument segmentation challenge. *arXiv preprint arXiv:1902.06426*, 2019. **6**
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. **6**
- [4] Jiacheng Cheng, Tongliang Liu, Kotagiri Ramamohanarao, and Dacheng Tao. Learning with bounded instance and label-dependent label noise. In *ICML*, pages 1789–1799. PMLR, 2020. **1**
- [5] Yiting Cheng, Fangyun Wei, Jianmin Bao, Dong Chen, Fang Wen, and Wenqiang Zhang. Dual path learning for domain adaptation of semantic segmentation. In *ICCV*, pages 9082–9091, 2021. **2, 6**
- [6] Emanuele Colleoni and Danail Stoyanov. Robotic instrument segmentation with image-to-image translation. *IEEE robot. autom. lett.*, 6(2):935–942, 2021. **7**
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016. **6**
- [8] Xiao Fu, Kejun Huang, Bo Yang, Wing-Kin Ma, and Nicholas D Sidiropoulos. Robust volume minimization-based matrix factorization for remote sensing and document clustering. *IEEE Trans. Signal Process.*, 64(23):6254–6268, 2016. **4**
- [9] Li Gao, Jing Zhang, Lefei Zhang, and Dacheng Tao. Dsp: Dual soft-paste for unsupervised domain adaptive semantic segmentation. In *ACM MM*, pages 2825–2833, 2021. **6, 7, 8**
- [10] Peter Gritzmann, Victor Klee, and David Larman. Largest j-simplices in n-polytopes. *Discrete Comput. Geom.*, 13(3):477–515, 1995. **4**
- [11] Xiaoqing Guo, Chen Yang, Baopu Li, and Yixuan Yuan. Metacorrection: Domain-aware meta loss correction for unsupervised domain adaptation in semantic segmentation. In *CVPR*, pages 3927–3936, 2021. **1, 3, 4, 6, 7, 8**
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. **6**
- [13] Javed Iqbal and Mohsen Ali. Msl: Multi-level self-supervised learning for domain adaptation with spatially independent and semantically consistent labeling. In *WACV*, pages 1864–1873, 2020. **1, 3, 6, 7**
- [14] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. In *CVPR*, pages 12975–12984, 2020. **2, 6, 7, 8**
- [15] Jogendra Nath Kundu, Akshay Kulkarni, Amit Singh, Varun Jampani, and R Venkatesh Babu. Generalize then adapt: Source-free domain adaptive semantic segmentation. In *ICCV*, pages 7046–7056, 2021. **2, 3, 6, 7**
- [16] Xuefeng Li, Tongliang Liu, Bo Han, Gang Niu, and Masashi Sugiyama. Provably end-to-end label-noise learning without anchor points. In *ICML*, pages 6403–6413, 2021. **2, 4**
- [17] Jian Liang, Ran He, Zhenan Sun, and Tieniu Tan. Exploring uncertainty in pseudo-label guided unsupervised domain adaptation. *Pattern Recognit.*, 96:106996, 2019. **1**
- [18] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, pages 6028–6039, 2020. **6, 7**
- [19] Jian Liang, Dapeng Hu, Ran He, and Jiashi Feng. Distill and fine-tune: Effective adaptation from a black-box source model. *arXiv preprint arXiv:2104.01539*, 2021. **1, 3**
- [20] Huafeng Liu, Chuanyi Zhang, Yazhou Yao, Xiushen Wei, Fumin Shen, Zhenmin Tang, and Jian Zhang. Exploiting web images for fine-grained visual recognition by eliminating open-set noise and utilizing hard examples. *IEEE Trans. Multimed.*, 2021. **3**
- [21] Jie Liu, Xiaoqing Guo, and Yixuan Yuan. Graph-based surgical instrument adaptive segmentation via domain-common knowledge. *IEEE Trans. Med. Imag.*, 2021. **6, 7, 8**
- [22] Yahao Liu, Jinhong Deng, Xincheng Gao, Wen Li, and Lixin Duan. Bapa-net: Boundary adaptation and prototype alignment for cross-domain semantic segmentation. In *ICCV*, pages 8801–8811, 2021. **1, 6, 7, 8**
- [23] Yuang Liu, Wei Zhang, and Jun Wang. Source-free domain adaptation for semantic segmentation. In *CVPR*, pages 1215–1224, 2021. **2, 3, 6, 7**
- [24] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020. **6**
- [25] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *CVPR*, pages 3764–3773, 2020. **1, 3, 6, 7**
- [26] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, pages 1944–1952, 2017. **4**
- [27] Micha Pfeiffer, Isabel Funke, Maria R Robu, Sebastian Bodenstedt, Leon Strenger, Sandy Engelhardt, Tobias Roß, Matthew J Clarkson, Kurinchi Gurusamy, Brian R Davidson, et al. Generating large labeled data sets for laparoscopic image processing tasks using unpaired image-to-image translation. In *MICCAI*, pages 119–127. Springer, 2019. **7**
- [28] Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. S4t: Source-free domain adaptation for semantic segmentation via self-supervised selective self-training. *arXiv preprint arXiv:2107.10140*, 2021. **2, 3, 6**

- [29] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, pages 102–118. Springer, 2016. [6](#)
- [30] Prabhu Teja S and Francois Fleuret. Uncertainty reduction for model adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9613–9623, 2021. [6](#)
- [31] Ragav Sachdeva, Filipe R Cordeiro, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. Evidentialmix: Learning with combined open-set and closed-set noisy labels. In *WACV*, pages 3607–3615, 2021. [3](#)
- [32] Manish Sahu, Ronja Strömsdörfer, Anirban Mukhopadhyay, and Stefan Zachow. Endo-sim2real: Consistency learning-based domain adaptation for instrument segmentation. In *MICCAI*, pages 784–794. Springer, 2020. [7](#)
- [33] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, pages 1919–1930, 2019. [3](#)
- [34] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, pages 7472–7481, 2018. [2](#), [6](#), [7](#), [8](#)
- [35] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*, 2021. [6](#)
- [36] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *CVPR*, pages 8688–8696, 2018. [3](#)
- [37] Yuxi Wang, Junran Peng, and ZhaoXiang Zhang. Uncertainty-aware pseudo label refinery for domain adaptive semantic segmentation. In *ICCV*, pages 9092–9101, 2021. [6](#)
- [38] Zhen Wang, Guosheng Hu, and Qinghua Hu. Training noise-robust deep neural networks via meta-learning. In *CVPR*, pages 4524–4533, 2020. [3](#)
- [39] Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Jiankang Deng, Jiatong Li, and Yinian Mao. Extended t: Learning with mixed closed-set and open-set noisy labels. *arXiv preprint arXiv:2012.00932*, 2020. [3](#), [5](#)
- [40] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *CVPR*, pages 4085–4095, 2020. [7](#)
- [41] Qing Yu and Kiyoharu Aizawa. Unknown class label cleaning for learning with open-set noisy labels. In *ICIP*, pages 1731–1735. IEEE, 2020. [3](#)
- [42] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *CVPR*, 2021. [1](#), [3](#), [6](#), [7](#)
- [43] Yivan Zhang, Gang Niu, and Masashi Sugiyama. Learning noise transition matrix from only noisy labels via total variation regularization. In *ICML*, volume 139, pages 12501–12512, 2021. [8](#)
- [44] Zizhao Zhang, Han Zhang, Serkan O Arik, Honglak Lee, and Tomas Pfister. Distilling effective supervision from severe label noise. In *CVPR*, pages 9294–9303, 2020. [3](#)
- [45] Yuyang Zhao, Zhun Zhong, Zhiming Luo, Gim Hee Lee, and Nicu Sebe. Source-free open compound domain adaptation in semantic segmentation. *arXiv preprint arXiv:2106.03422*, 2021. [2](#), [3](#)
- [46] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *Int. J. Comput. Vis.*, pages 1–15, 2021. [1](#), [3](#), [6](#), [7](#), [8](#)
- [47] Wei Zhou, Yukang Wang, Jiajia Chu, Jiehua Yang, Xiang Bai, and Yongchao Xu. Affinity space adaptation for semantic segmentation across domains. *IEEE Trans. Image Process.*, 30:2549–2561, 2020. [7](#)
- [48] Zhaowei Zhu, Tongliang Liu, and Yang Liu. A second-order approach to learning with instance-dependent label noise. In *CVPR*, pages 10113–10123, 2021. [1](#)
- [49] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, pages 5982–5991, 2019. [7](#)
- [50] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, pages 289–305, 2018. [1](#), [3](#), [7](#), [8](#)