# CoPT: Unsupervised Domain Adaptive Segmentation using Domain-Agnostic Text Embeddings

Cristina Mata, Kanchana Ranasinghe, and Michael S. Ryoo

Stony Brook University, Stony Brook, NY, USA {cfmata, kranasinghe, mryoo}@cs.stonybrook.edu

Abstract. Unsupervised domain adaptation (UDA) involves learning class semantics from labeled data within a source domain that generalize to an unseen target domain. UDA methods are particularly impactful for semantic segmentation, where annotations are more difficult to collect than in image classification. Despite recent advances in large-scale visionlanguage representation learning, UDA methods for segmentation have not taken advantage of the domain-agnostic properties of text. To address this, we present a novel **Covariance-based Pixel-Text** loss, CoPT, that uses domain-agnostic text embeddings to learn domain-invariant features in an image segmentation encoder. The text embeddings are generated through our LLM Domain Template process, where an LLM is used to generate source and target domain descriptions that are fed to a frozen CLIP model and combined. In experiments on four benchmarks we show that a model trained using CoPT achieves the new state of the art performance on UDA for segmentation. The code can be found at https://github.com/cfmata/CoPT.

**Keywords:** Unsupervised Domain Adaptation · Semantic Segmentation · Vision Language Models

# 1 Introduction

Semantic segmentation models are used in safety-critical applications, such as surgical robotics and autonomous vehicles, where it is of the utmost importance to properly segment small, rare and moving objects. The pixel-wise classification task requires a copious amount of labeled training data to perform well, especially on small objects, but human-collected annotations are difficult to obtain at scale. In addition, annotations collected automatically using large internet-scale pre-trained models such as Segment Anything Model (SAM) [13] are often noisy or lack semantic class associations, resulting in downstream errors. To avoid issues caused by imperfect labels, practitioners still train on synthetic data that emulates the target domain, or on real labeled data in a different source domain, before adapting the model to the target domain. The challenge remains to develop an unsupervised domain adaptation (UDA) method that learns strong semantic representations from available labels in a source domain while generalizing to the unlabeled target domain.



**Fig. 1:** CoPT encourages the distance between class centroids in pixel feature space to be the same as in domain-agnostic text embedding space.

The dominant approach taken by UDA methods is self-training in a teacherstudent paradigm [6–8], where an EMA-updated image encoder teacher produces pseudo-labels on target images for training the student, essentially bootstrapping unlabeled data with noisy labels to learn target domain features gradually. Selftraining is an attractive approach due to its straightforward implementation and impressive results, but remains susceptible to noisy predictions from the teacher and dependent on the strength of the pretrained image encoder upon initialization. An orthogonal line of work is data augmentation [6], which aims to make the model robust to changes in an class' appearance by augmenting the source samples directly with transformations. Data augmentation cannot simulate the true structure of objects in target domains, which inherently brings limited improvements to performance. Previous approaches do not target domain-specific features in latent representations learned by the image encoder.

We propose a novel **Co**variance-based **P**ixel-**T**ext loss, CoPT, which learns domain-agnostic representations for objects in images using domain-agnostic representations from a different modality: text. Text embeddings of class names are domain-agnostic [1] because they do not contain the visual domain-specific bias that prevents source-supervised segmentation networks from performing well on the target domain. The word for "cat" remains the same regardless of the animal's appearance in a real or synthetic image. It follows that text can be used as a regularization tool to encourage learning domain-invariant features in visual embeddings. This idea was developed and validated by RISE [9] for image classification, but in that work both text and visual embeddings came from the same CLIP [21] model. CoPT claims that this type of regularization works even when text and visual embeddings lie in separate latent spaces and backs this claim with strong empirical results (Section 5).

To guide learning domain-invariant pixel features in a segmentation encoder, CoPT enforces the distance between pixel features from different classes to be the same in a segmentation encoder's latent space as it is in a text encoder's latent space. To circumvent the hindrance that pixel and text features lie in different



**Fig. 2:** Class pixel features are generated by first extracting  $\frac{H}{r} \times \frac{W}{r} \times 512$  resolution feature maps from source images using  $\mathcal{E}_{\psi}$ , where *r* is the encoder's spatial downsampling factor. Each class' binary ground truth mask is downsampled to the same resolution. The class mask is applied to the low resolution feature, then spatial averaging is used to generate a  $1 \times 1 \times 512$  feature for each class.

latent spaces, CoPT minimizes the distance between the covariance matrix of pixel and text features. Figure 1 conceptually shows how CoPT is guided by domain-agnostic text embeddings to find domain-agnostic pixel features. Rather than relying on hand-crafted text templates to generate domain-agnostic text embeddings, we propose LLM Domain Template, a process through which an LLM (Large Language Model) is used to describe domain properties.

We validate CoPT on four UDA segmentation benchmarks and show that our method surpasses the current state of the art. We perform ablations on  $\text{GTA}\rightarrow\text{Cityscapes}$  to show that CoPT can boost the performance of multiple UDA methods and to show the benefits of domain-agnostic text embeddings generated from LLM Domain Template over hand-crafted ones. Our main contributions are,

- 1. The first method to leverage vision-language representations in UDA for semantic segmentation.
- 2. LLM Domain Template, a framework for using LLMs to generate domainagnostic text embeddings.
- 3. CoPT, a novel covariance-based pixel-text loss that bridges separate image and text latent spaces in order to learn domain-agnostic features.

### 2 Related Work

### 2.1 Unsupervised Domain Adaptation for Semantic Segmentation

Unsupervised domain adaptation (UDA) for semantic segmentation methods can be categorized through their use of four strategies: adversarial learning, data augmentation, self-supervision and prior learning. Self-training in particular has been shown to give large boosts in performance and many variations of it have been explored [4,7,8,14,17,18,27,30,31]. While this approach strengthens representations of common classes, the predictions can be noisy and the model suffers on small and rare classes. Data augmentation strategies in the form of strongly augmented self-training [6,7] and source-target image/label mixing [3] has brought strong improvements.

Our work presents a major departure from previous approaches. To the best of our knowledge we present the first work to incorporate vision-language representations into UDA for semantic segmentation. While this idea has been studied for image classification [9], it is not directly transferable to segmentation representations because of the need to conserve spatial resolution in the encoder. A comprehensive study of the best ways to use vision-language representation learning for UDA in segmentation is needed. We hope our work will encourage further methods to take advantage of recent advances in vision-language representation learning and the domain-agnostic properties of text.

### 2.2 Domain Adaptation using Vision-Language Embeddings

Since the introduction of CLIP [21], domain adaptation methods have taken advantage of its rich vision-language latent space in chiefly two ways: 1. adapting the image embedding space to be more similar to the text embedding space [16] or 2. prompt tuning [1, 2, 15, 29]. Unfreezing CLIP for target adaptation [16] is computationally expensive and risks harming CLIP's strong zero-shot generalizability. Instead of directly fine-tuning CLIP, methods that take the prompt tuning approach capitalize on its generalizability by embedding the text prompts describing the target domain into a space better suited to the target domain. Our method is most inspired by RISE [9] which generates domain-agnostic text embeddings and uses these to search for domain-agnostic image representations. The key difference lies in how we generate text descriptions for our domains: while their approach uses simple hand-crafted domain descriptors, we develop LLM Domain Template, a process to generate descriptions of each domain. In addition, their approach is designed for image classification, allowing them to use CLIP's image encoder directly for downstream classification, whereas we are focusing on segmentation and must grapple with anchoring pixel-specific features to text embeddings from a separately learned space. Jin et al. [12] similarly use domain descriptors for domain adaptive object detection, a task that also requires feature localization, but their method promotes visual context learning instead of using domain-agnostic text embeddings. A comprehensive review of related works can be found in Appendix A.

### 3 Methods

#### 3.1 Preliminaries

In unsupervised domain adaptation (UDA) for semantic segmentation, we are given access to a labeled source dataset  $S = \{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^{N_S}$  of  $N_S$  image and mask

pairs  $(\boldsymbol{x}, \boldsymbol{y})$  that have spatial height and width (H, W). Every pixel in the ground truth mask  $\boldsymbol{y}$  falls into one of C classes:  $\{\boldsymbol{y} \in \mathbb{Z} : 0 \leq \boldsymbol{y} \leq C\}$ . The goal is to learn a model  $\mathcal{F}_{\theta}$  parameterized by  $\theta$  that performs pixel-wise classification well on inputs from an unlabeled target dataset  $T = \{\boldsymbol{x}_i\}_{i=1}^{N_T}$ .

Typically, the segmentation network  $\mathcal{F}_{\theta}$  is trained to minimize the following pixel-wise cross entropy loss  $\mathcal{L}_{ce}$  over each source image sample [17, 26, 31],

$$\mathcal{L}_{ce} = -\sum_{i=1}^{H} \sum_{j=1}^{W} \sum_{c=1}^{C} \mathbb{I}[y_{ijc} = c] \log(\hat{y}_{ijc})$$
(1)

where  $\hat{\boldsymbol{y}} = \mathcal{F}_{\theta}(\boldsymbol{x})$  is the model's predicted discrete distribution over all classes C for the image  $\boldsymbol{x}$  and  $\mathbb{I}[y_{ijc} = c]$  is a binary indicator for whether the pixel at (i, j) belongs to class c. Additionally, many UDA methods optimize auxiliary losses to modify  $\mathcal{F}_{\theta}$ 's latent space [3,8].

In our approach, we introduce a loss that modifies the distance between  $\mathcal{F}_{\theta}$ 's class representations to be the same as in the latent space of a frozen CLIPpretrained text encoder  $\mathcal{G}$ . To the best of our knowledge, this is the first UDA work to apply vision-language representations for semantic segmentation. We first build domain-agnostic text embeddings by averaging embeddings of domainspecific prompts generated by a Large Language Model in our LLM Domain Template process (Section 3.3). Then we extract pixel-wise class features from the image segmentation encoder (Figure 2). Lastly our proposed covariancebased loss CoPT bridges the gap between the image encoder and text encoder latent spaces in order to enforce learning of domain-agnostic features (Section 3.4).

### 3.2 Hand-Crafted Domain-Agnostic Text Embeddings

We have access to a CLIP-pretrained text encoder  $\mathcal{G}$  that takes a string prompt p and outputs a text embedding  $\mathbf{t} = \mathcal{G}(p)$ . In our work we keep  $\mathcal{G}$  frozen and use it purely to extract text embeddings. The prompt template specific to domain d and class c is hand-crafted in [9] as,

$$p_d(c) =$$
"A  of a " (2)

where d = DOMAIN and c = CLASS. For example, when the target domain is photos of objects, one would plug in "photo" to DOMAIN so that "A photo of a car" denotes the target domain description of a car.

A prompt template is crafted for each domain so that, given a domain dand a class c, the text embedding  $\mathbf{t}_d^c = \mathcal{G}(p_d(c))$  contains domain-specific and class-specific information. According to RISE [9], simply averaging the text embeddings of multiple class templates results in a generic text representation of an object. This ensures that the text embedding is representative of a class across multiple domains. To generate a domain-agnostic prompt  $\mathbf{t}^c$  for a class c, the source and target embeddings are averaged together as shown in Figure 3:

$$\boldsymbol{t}^{c} = \frac{1}{2} (\boldsymbol{t}^{c}_{S} + \boldsymbol{t}^{c}_{T}) \tag{3}$$

Such a generalized text embedding of a class is called domain-agnostic text embedding by Bose et al. [1] and we follow this terminology. While we follow RISE's hand-crafted prompt construction method, we modify it to use two domain templates (for source and target) instead of the 80 CLIP templates [21].

So far, prior works have relied on using hand-crafted domain descriptors to embed semantic information specific to a domain. However these descriptors are short and lack details about what contributes to the domain



Fig. 3: The final domain-agnostic text embedding  $\mathbf{t}^{\mathbf{c}}$  for class c is generated by averaging the source and target text embeddings,  $\mathbf{f}_{\mathbf{S}}^{\mathbf{c}}$  and  $\mathbf{f}_{\mathbf{T}}^{\mathbf{c}}$ .

gap, namely the texture and appearance of a domain. We propose a framework called LLM Domain Template to increase the semantic expressiveness of the text embeddings.

#### 3.3 LLM Domain Template

LLM Domain Template uses a large language model, ChatGPT [20], to automatically generate long-form descriptors of the source and target domains. We hypothesize that doing so will lead to better domain-agnostic text embeddings because the average embedding will be closer to the centroid of the domainspecific regions in latent space. Figure 4 shows LLM Domain Template in action. We query ChatGPT with "Can you give detailed descriptions of what makes a <DOMAIN> image look <DOMAIN>?" and get a list of K domain attribute templates. The prompt template specific to attribute k = ATTRIBUTE K generated by ChatGPT is

$$p_k(c) =$$
"A  with " (4)

We fill in the class name and attribute for all templates, then feed them to the text encoder  $\mathcal{G}$  and average them to get the domain-specific text embedding for a class:

$$\boldsymbol{t}_{d}^{c} = \frac{1}{K} \sum_{k=1}^{K} \mathcal{G}(p_{k}(c))$$
(5)

The final domain-agnostic text embedding for a class,  $t^c$ , is generated using Equation 3. The domain-agnostic embeddings are generated once at model initialization and stored in a memory bank.

#### 3.4 CoPT: Covariance-based Pixel-Text Loss

Our goal is to adapt the segmentation encoder's latent space to the target by learning domain-agnostic class representations. We use the domain-agnostic text



**Fig. 4:** LLM Domain Template: An LLM is queried to describe the source and target domain attributes. Here the query is about the synthetic source domain. The resulting templates are formatted with a class name, fed to a frozen CLIP text encoder and averaged to obtain a single text embedding for a class.

embeddings  $t^c$  to structure this representation space using a covariance-based pixel-text loss, or CoPT. CoPT enforces the distance between classes in pixel feature space to be the same as the distance between classes in text embedding space. This requires extracting class pixel features from the encoder given a source domain sample.

Our image segmentation model  $\mathcal{F}_{\theta}$  is composed of an encoder  $\mathcal{E}_{\psi}$  and decoder  $\mathcal{D}_{\rho}$ :  $\mathcal{F}_{\theta} = \mathcal{D}_{\rho}(\mathcal{E}_{\psi})$ .  $\mathcal{E}_{\psi}$  spatially downsamples its input by a factor r, so while input image  $\boldsymbol{x}$  has resolution (H, W), the encoder's features  $\boldsymbol{f} = \mathcal{E}_{\psi}(\boldsymbol{x})$  have resolution  $(\frac{H}{r}, \frac{W}{r})$ . To extract class pixel features from  $\boldsymbol{f}$ , we spatially downsample  $\boldsymbol{x}$ 's corresponding ground truth mask  $\boldsymbol{y}$  to  $(\frac{H}{r}, \frac{W}{r})$  and generate a binary mask  $\boldsymbol{b}^c$  for class c with the same dimensions:

$$\boldsymbol{b}^{c} = \mathbb{I}[\text{downsample}(\boldsymbol{y}) = c]$$
(6)

The mask  $b^c$  is multiplied by f to extract class features and spatial averaging is applied to get the final pixel feature for a class,  $f^c$ :

$$\boldsymbol{f}^{c} = \frac{1}{\left(\frac{H}{r}\right)\left(\frac{W}{r}\right)} \sum_{i=1}^{\frac{H}{r}} \sum_{j=1}^{\frac{W}{r}} b_{i,j}^{c} \boldsymbol{f}_{i,j}$$
(7)

where  $f_{i,j}$  is the pixel feature at location (i, j) and  $b_{i,j}^c$  is a binary value indicating whether the pixel belongs to class c. Figure 2 illustrates the class pixel feature extraction process.

The pixel covariance matrix  $\Sigma_p$  for an input sample  $\boldsymbol{x}$  is calculated by measuring the cosine similarity between all pairs of class pixel features  $\boldsymbol{f}^c$  from  $\boldsymbol{x}$ . A similar process is carried out with the domain-agnostic text embeddings  $\boldsymbol{t}^c$  to get the text covariance matrics  $\Sigma_t$ :

$$\Sigma_{p} = \begin{bmatrix} \operatorname{Var}(\boldsymbol{f}^{1}) & \dots \operatorname{Cov}(\boldsymbol{f}^{1}, \boldsymbol{f}^{C}) \\ \vdots & \ddots & \vdots \\ \operatorname{Cov}(\boldsymbol{f}^{C}, \boldsymbol{f}^{1}) & \dots & \operatorname{Var}(\boldsymbol{f}^{C}) \end{bmatrix}, \Sigma_{t} = \begin{bmatrix} \operatorname{Var}(\boldsymbol{t}^{1}) & \dots \operatorname{Cov}(\boldsymbol{t}^{1}, \boldsymbol{t}^{C}) \\ \vdots & \ddots & \vdots \\ \operatorname{Cov}(\boldsymbol{t}^{C}, \boldsymbol{t}^{1}) & \dots & \operatorname{Var}(\boldsymbol{t}^{C}) \end{bmatrix}$$
(8)

The best segmentation performance is usually achieved by training with large input resolutions, leading to low batch sizes in order to fit images into GPU



Fig. 5: CoPT calculates the difference in pixel and text covariance matrices in order to transform the segmentation encoder's latent space to emulate the domain-agnostic text space.

memory. As a result it is rare that all classes appear in a mini-batch. To address this we construct the text embedding's covariance matrix  $\Sigma_t$  for corresponding classes as they occur in a sample during training. Low batch sizes also lead to high variance in class pixel features that jeopardizes stable training. To reduce variance in gradient updates we store the average pixel feature for each class in a memory bank, and update the bank on each iteration using a linear decay. That is, upon calculation of a pixel feature  $\mathbf{f}^{\mathbf{c}}$  for a class, the class in memory  $\mathbf{f}^{\mathbf{c}}_{\text{current}}$ is updated with decay  $\lambda$  following

$$\mathbf{f}_{new}^{\mathbf{c}} = \lambda * \mathbf{f}_{current}^{\mathbf{c}} + (1 - \lambda) * \mathbf{f}^{\mathbf{c}}$$
(9)

When calculating the pixel covariance matrix, the updated pixel feature in the memory bank  $f_{new}^c$  is used.

CoPT's final objective is to minimize the cosine distance between the covariance matrices  $\Sigma_p$  and  $\Sigma_t$ :

$$\mathcal{L}_{\text{CoPT}} = 1 - \frac{\Sigma_p \cdot \Sigma_t}{\|\Sigma_p\| \cdot \|\Sigma_t\|}$$
(10)

This calculation is illustrated in Figure 5. CoPT can be added to any UDA training scheme that uses source images and their labels.

### 4 Experiments

### 4.1 Datasets

We run experiments on four benchmarks commonly used for UDA for semantic segmentation:  $GTA \rightarrow CS$  [22],  $Synthia \rightarrow CS$  [23],  $CS \rightarrow DZ$  [24] and  $CS \rightarrow ACDC$  [25].  $GTA \rightarrow CS$  and  $Synthia \rightarrow CS$  consider synthetic images as the source domain

9

and real images as the target. This emulates the common industry practice of developing in-house synthetic datasets for the source domain.  $CS \rightarrow DZ$  considers the adaptation of day-time source to night-time target images.  $CS \rightarrow ACDC$ considers adaptation of clear to adverse weather conditions. All benchmarks are of road scenes, which is grounded in the importance of segmentation for autonomous vehicle applications. GTA [22] consists of 24,966 synthetic images of street scenes collected from a video game and comes with pixel-level semantic labels for 19 classes. Cityscapes (CS) [5] is a real street-scene dataset with fine pixel-level annotations across the same 19 classes as GTA for 2,975 training images and 500 validation images. DarkZurich (DZ) [24] consists of 2,416 training images of twilight scenes and 151 testing images, and has labels for the 19 CS classes. ACDC [25] consists of 1600 training, 406 validation and 2000 test images and is labeled with the 19 CS classes. Synthia [23] consists of 9,400 synthetic images collected in a virtual city and contains annotations for 16 of the 19 CS classes.  $GTA \rightarrow CS$  and  $Synthia \rightarrow CS$  use the 500 labeled images from the CS validation set for evaluation. The DZ training set is used as an unlabeled target image source while the 151 images in the testing set are used only for evaluation. Similarly we use the ACDC training set as an unlabeled target image source and evaluate on its 2000 test images. In our experiments we report per-class % Intersection over Union (IOU) as well as the average mIOU across 19 classes in  $GTA \rightarrow CS$ ,  $CS \rightarrow DZ$  and  $CS \rightarrow ACDC$ , and across 16 classes in Synthia $\rightarrow CS$ .

### 4.2 Implementation

For LLM Domain Template, we query ChatGPT [20], a transformer-based LLM trained for next-token-prediction on public internet and private licensed data. We submit one query for each domain following Equation 4. A full list of the query and responses can be found in Appendix B.1. Once templates are formatted with class names, they are passed to a frozen ViT-B/32 CLIP [21] text encoder to generate text embeddings, which takes 5.6 seconds on an NVIDIA RTX A5000. For our memory bank we use a dictionary initialized with empty Tensors of length 512. We perform an ablation experiment, detailed in Appendix D.1 to find the optimal decay  $\lambda$  is 0.5. Our implementation is in PyTorch. The code may be found at https://github.com/cfmata/CoPT.

In our main experiments we implement CoPT on top of MIC [8], the current state of the art for UDA for semantic segmentation. In MIC, the auxiliary losses are pixel-wise cross entropy over source images, masked pseudo-label self-training using an EMA teacher, ImageNet feature distance regularization, and strongly augmented self-training. CoPT is added with weight 1 to the auxiliary losses. We include detailed equations for self-training losses in Appendix B.2 but strongly encourage the reader to refer to [8] for complete details on the auxiliary losses and their hyperparameters. CoPT is optimized by Adam with learning rate 0.00006, betas (0.9, 0.999) and weight decay 0.01. We use batch size 2 per GPU and train for 40,000 iterations, which takes approximately 34 hours on a single NVIDIA RTX A5000. Image crops sized  $1024 \times 1024$  are fed to the model during training.

### 5 Results

#### 5.1 Unsupervised Domain Adaptation for Semantic Segmentation

We compare CoPT to previous UDA methods for semantic segmentation on  $GTA \rightarrow CS$  in Table 1. The methods are trained using the GTA training set as the source domain and evaluated in terms of IOU on the Cityscapes validation set. CoPT gives a 1.6 % boost in mIOU when trained on top of MIC, and shows improvements on most classes, setting a new state of the art on the benchmark. CoPT gives the highest improvement on the train class: the train's large size means it is made up of a larger variety of features than other classes, so it might benefit more from associating the class with the centroid of a wide variety of semantic features. Figure 6 shows qualitative results on Cityscapes from CoPT and two other baselines. It illustrates how CoPT can improve performance on the sidewalk class, which also exhibits high intra-class variation. CoPT struggles most with the traffic sign class: this class has more varied geometric shapes than the other classes and its accurate segmentation depends more on accurate boundary localization rather than on distinguishing domain textures.

To test how CoPT compares to previous methods on a different source dataset, we train it using Synthia synthetic data and test it on the Cityscapes validation set, reporting IOU for all classes in Table 2. Synthia data constitutes a larger domain difference from Cityscapes than GTA, largely due to the camera angles being from above rather than from a driver's perspective. Regardless of this challenge, CoPT is able to surpass MIC's performance as well as previous methods' performance on most classes. CoPT performs particularly well on the motorbike/motorcycle class. However, it struggles most with the sidewalk and road classes. In many Synthia source images, foliage is placed in front of the camera and blocks the sidewalk and roads. Since CoPT relies on downsampled ground truth masks to extract pixel features for a class, when there are thin structures in front of large classes this can lead to mixing class features in the extracted pixel feature, explaining CoPT's struggle on sidewalks and roads.

To show the generalizability of CoPT, we evaluate its performance on CS $\rightarrow$ DZ, a benchmark that considers adaptation from daytime images to nighttime. CoPT outperforms previous methods on overall mIOU, and consistently attains top one to two performance on nearly all categories. It achieves the best performance on the truck class and also performs particularly well on the sidewalk and terrain classes compared to previous methods. This affirms our findings on the GTA $\rightarrow$ CS benchmark in Table 1 that CoPT is improving intra-class understanding, since these classes tend to be large in the image, containing numerous visual features while making up the same class.

To test whether CoPT performs well on another target domain we train it with labeled Cityscapes training images and unlabeled ACDC training images and test it on the adverse conditions in the ACDC test set as the target domain. Results on each condition are presented in Table 4 where CoPT is compared to MIC. CoPT outperforms MIC on average due to its 3.6% improvement on the Fog condition and maintains similar performance on other conditions. Compared

Table 1: Semantic segmentation results reported as % IOU on the GTA $\rightarrow$ CS val set. Bold indicates highest performance and <u>underline</u> indicates second highest. All methods' results are taken from the published paper except for those labeled with  $\dagger$ , which indicates the method was retrained. Our method, built on top of MIC [8], attains state of the art performance.

Method	Road	S.Walk	Build.	Wall	Fence	Pole	T.Light	T.Sign	Veg.	Terr.	$\mathbf{Sky}$	Person	Rider	$\operatorname{Car}$	Truck	Bus	Train	M.Bike	Bike	mIOU
DAFormer [6]	95.7	70.2	89.4	53.5	48.1	49.6	55.8	59.4	89.9	47.9	92.5	72.2	44.7	92.3	74.5	78.2	65.1	55.9	61.8	68.3
SePiCo [28]	95.2	67.8	88.7	41.4	38.4	43.4	55.5	63.2	88.6	46.4	88.3	73.1	49.0	91.4	63.2	60.4	0.0	45.2	60.0	61.0
HRDA [7]	96.4	74.4	91.0	61.6	51.5	57.1	63.9	69.3	91.3	48.4	94.2	79.0	52.9	93.9	84.1	85.7	75.9	63.9	67.5	73.8
PiPa (HRDA) [3]	96.8	76.3	91.6	63.0	57.7	60.0	65.4	<u>72.6</u>	<u>91.7</u>	51.8	94.8	<u>79.7</u>	56.4	94.4	85.9	88.4	78.9	63.5	67.2	<u>75.6</u>
$MIC^{\dagger}$ [8]	97.5	<u>80.3</u>	91.3	60.2	52.5	59.7	64.1	73.1	91.3	49.6	93.9	79.3	54.8	<u>94.6</u>	86.5	86.6	67.2	66.8	<u>68.7</u>	74.6
CoPT (Ours)	97.6	80.9	91.6	<u>62.1</u>	<u>55.9</u>	59.3	66.7	70.5	91.9	53.0	<u>94.4</u>	80.0	55.6	94.7	87.1	88.6	82.1	<u>65.0</u>	68.8	76.1

Table 2: Semantic segmentation results reported as % IOU on the Synthia $\rightarrow$ CS val set. All methods' results are taken from the published paper except for those labeled with  $\dagger$ , which indicates the method was retrained. Our method attains state of the art performance.

Method	Road	S.Walk	Build.	Wall	Fence	Pole	T.Light	T.Sign	Veg.	Sky	Person	Rider	$\operatorname{Car}$	Bus	M.Bike	Bike	mIOU
DAFormer [6]	84.5	40.7	88.4	41.5	6.5	50.0	55.0	54.6	86.0	89.9	73.2	48.2	87.2	53.2	53.9	61.7	60.9
SePiCo [28]	77.0	35.3	85.1	23.9	3.4	38.0	51.0	55.1	85.6	80.5	73.5	46.3	87.6	69.7	50.9	66.5	58.1
HRDA [7]	85.2	47.7	88.8	49.5	4.8	57.2	65.7	60.9	85.3	92.9	79.4	52.8	89.0	64.7	63.9	64.9	65.8
PiPa (HRDA) <sup>†</sup> [3]	86.6	52.7	88.2	46.6	1.4	47.8	60.3	46.2	86.7	93.2	75.5	50.8	86.6	17.0	59.8	63.6	67.3
MIC [8]	86.6	50.5	<u>89.3</u>	47.9	7.8	59.4	66.7	63.4	$\underline{87.1}$	<u>94.6</u>	<u>81.0</u>	58.9	90.1	61.9	67.1	64.3	<u>67.3</u>
CoPT (Ours)	83.4	44.3	90.0	50.4	8.0	60.0	67.0	<u>63.0</u>	87.5	94.8	81.1	58.6	89.7	66.5	68.9	<u>65.0</u>	67.4

**Table 3:** Semantic segmentation results reported as % IOU on the CS $\rightarrow$ DZ test set. Bold indicateshighest performance and <u>underline</u> indicates second highest. All methods' results are taken from thepublication except for those labeled with  $\dagger$ , which indicates the method was retrained.

Method	Road \$	S.Walk	Build.	Wall	Fence	Pole	T.Light	T.Sign	Veg.	Terr.	$\mathbf{Sky}$	Person	Rider	$\operatorname{Car}$	Truck	Bus	Train	M.Bike	Bike	mIOU
DAFormer [6]	93.5	<u>65.5</u>	73.3	39.4	19.2	53.3	44.1	44.0	<u>59.5</u>	34.5	66.6	53.4	52.7	82.1	52.7	<u>9.5</u>	89.3	50.5	38.5	53.8
SePiCo [28]	91.2	61.3	67.0	28.5	15.5	44.7	44.3	41.3	65.4	22.5	80.4	41.3	52.4	71.2	39.3	0.0	39.6	27.5	28.8	45.4
HRDA [7]	90.4	56.3	72.0	39.5	19.5	57.8	52.7	43.1	59.3	29.1	70.5	60.0	58.6	84.0	<u>75.5</u>	11.2	90.5	51.6	40.9	<u>55.9</u>
$MIC^{\dagger}[8]$	89.1	54.0	80.4	53.5	20.3	63.8	54.9	54.2	47.9	<u>37.9</u>	47.6	63.9	58.3	75.2	66.5	7.3	<u>90.7</u>	48.9	45.9	55.8
CoPT (Ours)	92.7	66.1	80.0	$\underline{49.0}$	19.3	63.2	51.7	52.0	50.9	43.1	55.9	61.7	56.5	59.6	79.4	2.96	90.8	<u>50.0</u>	42.5	56.2

Table 4: Semantic segmentation results reported as % IOU on the adverse conditions in CS $\rightarrow$ ACDC test set.  $\dagger$  indicates MIC [8] was retrained. CoPT outperforms MIC on average and especially on fog.

Method	Rain	Snow	Fog	Night	All
$\mathrm{MIC}^{\dagger}$ [8]	74.2	60.7	56.8	57.6	63.3
CoPT (Ours)	74.0	60.6	60.4	56.8	63.7

to the other adverse conditions, fog tends to uniformly change class appearance, whereas a condition like nighttime can drastically change appearance with lights illuminating different object parts and obscuring others, which supports the hypothesis that CoPT improves intra-class understanding.

### 5.2 Ablations

**CoPT Distance Metric** We compare different metrics to calculate the distance between CoPT's pixel and text covariance matrices in Table 5. Using a cosine similarity loss outperforms both L1 and L2 distances. L2 considers the covariance between each pair of classes equally, which may not be ideal since certain class pairs can co-occur more often in target vs. source domains and thus contribute more to the domain gap. L1 is ideal for scenarios involving grid-like space, which is not similar to the covariance matrices since different batches contain different classes. Cosine similarity favors points that share the same direction, which is



Fig. 6: Segmentation predictions on GTA $\rightarrow$ CS validation set. Each column from left to right shows the Cityscapes image, HRDA [7] prediction, MIC [8] prediction, CoPT (our) prediction and the ground truth label. In all examples, CoPT produces fewer spurious predictions in the sidewalk, showing its benefits for classes with large intraclass variation.

conducive to our goal of aligning representations in abstract latent spaces with unknown structure.

LLM Domain Template vs. Hand-Crafted Templates Table 6 shows our evaluation of whether LLM Domain Template, in which an LLM generates domain descriptions, leads to better domain-agnostic features for CoPT than simple hand-crafted domain descriptions. The evaluation is done on the  $\text{GTA}\rightarrow\text{CS}$  benchmark. While it shows that LLM Domain Template outperforms the hand-crafted approach, the slight improvement could indicate that the average hand-crafted template embedding is not very far from the average LLM Domain Template embedding. Averaging templates gives the centroid of the points, and since the domain-specific hand-crafted templates have so far been designed to be quite general already, their centroid will likely also be domain-agnostic. However, this can pose a problem in cases where access to the target domain is not allowed, in which case it is beneficial to have our automated method generate descriptions of the target. This can be accomplished by feeding a target image in addition to the text query in Figure 4 in LLM Domain Template.

**Backbone Method** We test whether CoPT is able to boost performance of two different UDA methods, MIC and HRDA, for semantic segmentation by jointly training these methods with our loss. The results on the GTA $\rightarrow$ CS benchmark in Table 7 show an over 1 % improvement in mIOU across 19 Cityscapes classes. This indicates that CoPT cooperates with multiple methods and is easily incorporated into state of the art training paradigms. CoPT's training scheme is further ablated in Appendix D.2.

**Text Encoder** To test whether CoPT's performance boost is due specifically to CLIP [21] text embeddings or to the use of text embeddings in general, we swap CLIP for Sentence-T5 [19] or Mistral-7B [10] and retrain CoPT on GTA $\rightarrow$ CS in Table 8. CoPT gets state of the art on using both CLIP [21] and Sentence-T5 [19] text embeddings, showing that in general using text embeddings as regularization anchors for pixel feature space benefits the UDA problem. Mistral-

Table 5: Ablation experiment results for CoPT's covariance distance metric are reported as % mIOU on  $GTA \rightarrow CS$  over 19 classes.

Metric	mIOU
L1	75.34
L2	70.18
Cosine Similarity	<b>76.08</b>

 
 Table
 7:
 Ablation experiment results
 for CoPT added to two different UDA methods. Results reported as % mIOU on  $GTA \rightarrow CS$  over 19 classes

Table 6: Ablation experiment results for LLM Domain Templates vs. hand -crafted templates are reported as %mIOU on GTA $\rightarrow$ CS over 19 classes.

Template Type	mIOU
Hand-Crafted	75.95
LLM Generation	76.08

Table 8: CoPT results in % mIOU on  $GTA \rightarrow CS$  val 19 classes using different text encoders.

111 705 0001	10 0100000	•		
Method	mIOU	Δ	Text Encoder	mIO
HRDA	73.8	-	MIC Baseline	74.6
HRDA+CoPT	<b>75.34</b>	+1.54	Mistral	70.8
MIC	74.61	+1.47	Sentence-T5	75.5
MIC+CoPT	<b>76.08</b>		CLIP	<b>76.1</b>

Table 9: CoPT mIOU (%) across 19 classes on  $\text{GTA}\rightarrow\text{CS}$  val set using pixel features extracted from source, pseudo-labeled target, or both domains.

CoPT Pixel Features	mIOU
MIC Baseline	74.6
Source + Target	71.8
Target Only	74.7
Source Only	76.1

7B [10] underperforms because it is trained for next token prediction, leading to worse sentence representations [11] and thus worse representations for the domain prompt templates (see Appendix D.4).

Source vs. Target Pixel Features In our main experiments CoPT is implemented using class pixel features from the source domain. In the ablation shown in Table 9 we test whether applying CoPT using class features from the target domain is beneficial. To extract class pixel features from the target domain, we use the target pseudo label. CoPT achieves state of the art on  $\text{GTA}\rightarrow\text{CS}$ when applied to either source pixel features or target pixel features, showing that adapting the relative distance of class features from each domain to resemble that of text is beneficial. The higher performance on source pixel features is likely due to noise in the target pseudo labels. Due to out-of-memory issues, when applying CoPT to both source and target pixel features we must backpropagate after computing source, which changes the optimization scheme and leads to underperformance.

### 6 Discussion

Based on the results from our benchmark experiments, CoPT gives the biggest performance boosts on semantic classes with large intra-class variations, such as truck, train and bus in GTA $\rightarrow$ CS and wall in Synthia $\rightarrow$ CS. These classes are large in reality, leading them to encompass larger swaths of the image than other classes. They also exhibit varying features, such as the vehicles showing ads on the side, or variation due to windows and paint. Trucks, trains and buses in particular have similar appearance as well, but distinct semantic functions, so it is beneficial for the model to learn more distinct embedding spaces for these classes, which CoPT encourages through the use of text.

CoPT relies on features from the end of the encoder that are spatially downsampled, which loses detailed information. This can explain why CoPT doesn't improve on the traffic sign class in GTA $\rightarrow$ CS, in which different instances of the class have different geometries, requiring the model to use detailed spatial features for accurate boundary segmentation. A naive way to alleviate this problem is to upsample the encoder's features to the original resolution before applying the ground truth mask for each class. However, this leads to out-of-memory errors. Instead, one can upsample crops of the encoder feature to the original resolution and apply the masks to the cropped region, carrying out the rest of the loss as usual. We leave this for future work.

## 7 Conclusion

This work presents the first effort to incorporate vision-language representations into a UDA method for semantic segmentation, a task in which annotations are simultaneously difficult to collect and in dire need. We have introduced a novel covariance-based pixel text loss called CoPT that cooperates with previous UDA methods for semantic segmentation to boost performance, especially on classes with large intra-class variation. CoPT benefits from using domain-agnostic text embeddings based on descriptions of domains generated by LLMs as opposed to hand-crafted ones found in previous methods. CoPT sets a new state of the art based on our experiments on four benchmarks and through our ablations we show that it gives robust improvements to different backbone methods. We hope CoPT inspires future UDA methods to search for novel ways to use domainagnostic properties of alternative modalities.

Acknowledgment This work was supported by Electronics and Telecommunications Research Institute (ETRI) grants funded by the Korean government. [24ZR1100, A Study of Hyper-Connected Thinking Internet Technology by autonomous connecting, controlling and evolving ways] and [24ZB1200, Research of Human-centered autonomous intelligence system original technology]

## References

- Bose, S., Jha, A., Fini, E., Singha, M., Ricci, E., Banerjee, B.: Stylip: Multi-scale style-conditioned prompt learning for clip-based domain generalization. In: WACV (2024)
- Cao, Q., Xu, Z., Chen, Y., Ma, C., Yang, X.: Domain-controlled prompt learning (2023)
- Chen, M., Zheng, Z., Yang, Y., Chua, T.S.: Pipa: Pixel- and patch-wise selfsupervised learning for domain adaptative semantic segmentation. In: ACM MM (2023)
- 4. Chen, R., Rong, Y., Guo, S., Han, J., Sun, F., Xu, T., Huang, W.: Smoothing matters: Momentum transformer for domain adaptive semantic segmentation (2022)
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
- 6. Hoyer, L., Dai, D., Gool, L.V.: Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In: CVPR (2022)
- 7. Hoyer, L., Dai, D., Gool, L.V.: Hrda: Context-aware high-resolution domainadaptive semantic segmentation. In: ECCV (2022)
- Hoyer, L., Dai, D., Wang, H., Gool, L.V.: Mic: Masked image consistency for context-enhanced domain adaptation. In: CVPR (2023)
- Huang, Z., Zhou, A., Lin, Z., Cai, M., Wang, H., Lee, Y.J.: A sentence speaks a thousand images: Domain generalization through distilling clip with language guidance. In: ICCV (2023)
- Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D.S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L.R., Lachaux, M.A., Stock, P., Scao, T.L., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mistral 7b (2023)
- 11. Jiang, T., Huang, S., Luan, Z., Wang, D., Zhuang, F.: Scaling sentence embeddings with large language models (2023)
- 12. Jin, S., Jiang, X., Huang, J., Lu, L., Lu, S.: Llms meet vlms: Boost open vocabulary object detection with fine-grained descriptors. In: ICLR (2024)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: ICCV (2023)
- 14. Kundu, J.N., Kulkarni, A., Singh, A., Jampani, V., Babu, R.V.: Generalize then adapt: Source-free domain adaptive semantic segmentation. In: ICCV (2021)
- Lai, Z., Bai, H., Zhang, H., Du, X., Shan, J., Yang, Y., Chuah, C.N., Cao, M.: Empowering unsupervised domain adaptation with large-scale pre-trained visionlanguage models. In: WACV (2024)
- Lai, Z., Vesdapunt, N., Zhou, N., Wu, J., Huynh, C.P., Li, X., Fu, K.K., Chuah, C.N.: Padclip: Pseudo-labeling with adaptive debiasing in clip for unsupervised domain adaptation. In: ICCV (2023)
- 17. Li, G., Kang, G., Liu, W., Wei, Y., Yang, Y.: Content-consistent matching for domain adaptive semantic segmentation. In: ECCV (2020)
- Lu, Y., Luo, Y., Zhang, L., Li, Z., Yang, Y., Xiao, J.: Bidirectional self-training with multiple anisotropic prototypes for domain adaptive semantic segmentation. In: ACM MM (2022)
- 19. Ni, J., Abrego, G.H., Constant, N., Ma, J., Hall, K., Cer, D., Yang, Y.: Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. In: ACL (2022)

- 16 C. Mata et al.
- 20. OpenAI, :, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., Bello, I., Berdine, J., Bernadett-Shapiro, G., Berner, C., Bogdonoff, L., Boiko, O., Boyd, M., Brakman, A.L., Brockman, G., Brooks, T., Brundage, M., Button, K., Cai, T., Campbell, R., Cann, A., Carey, B., Carlson, C., Carmichael, R., Chan, B., Chang, C., Chantzis, F., Chen, D., Chen, S., Chen, R., Chen, J., Chen, M., Chess, B., Cho, C., Chu, C., Chung, H.W., Cummings, D., Currier, J., Dai, Y., Decareaux, C., Degry, T., Deutsch, N., Deville, D., Dhar, A., Dohan, D., Dowling, S., Dunning, S., Ecoffet, A., Eleti, A., Eloundou, T., Farhi, D., Fedus, L., Felix, N., Fishman, S.P., Forte, J., Fulford, I., Gao, L., Georges, E., Gibson, C., Goel, V., Gogineni, T., Goh, G., Gontijo-Lopes, R., Gordon, J., Grafstein, M., Gray, S., Greene, R., Gross, J., Gu, S.S., Guo, Y., Hallacy, C., Han, J., Harris, J., He, Y., Heaton, M., Heidecke, J., Hesse, C., Hickey, A., Hickey, W., Hoeschele, P., Houghton, B., Hsu, K., Hu, S., Hu, X., Huizinga, J., Jain, S., Jain, S., Jang, J., Jiang, A., Jiang, R., Jin, H., Jin, D., Jomoto, S., Jonn, B., Jun, H., Kaftan, T., Łukasz Kaiser, Kamali, A., Kanitscheider, I., Keskar, N.S., Khan, T., Kilpatrick, L., Kim, J.W., Kim, C., Kim, Y., Kirchner, J.H., Kiros, J., Knight, M., Kokotajlo, D., Łukasz Kondraciuk, Kondrich, A., Konstantinidis, A., Kosic, K., Krueger, G., Kuo, V., Lampe, M., Lan, I., Lee, T., Leike, J., Leung, J., Levy, D., Li, C.M., Lim, R., Lin, M., Lin, S., Litwin, M., Lopez, T., Lowe, R., Lue, P., Makanju, A., Malfacini, K., Manning, S., Markov, T., Markovski, Y., Martin, B., Mayer, K., Mayne, A., McGrew, B., McKinney, S.M., McLeavey, C., McMillan, P., McNeil, J., Medina, D., Mehta, A., Menick, J., Metz, L., Mishchenko, A., Mishkin, P., Monaco, V., Morikawa, E., Mossing, D., Mu, T., Murati, M., Murk, O., Mély, D., Nair, A., Nakano, R., Nayak, R., Neelakantan, A., Ngo, R., Noh, H., Ouyang, L., O'Keefe, C., Pachocki, J., Paino, A., Palermo, J., Pantuliano, A., Parascandolo, G., Parish, J., Parparita, E., Passos, A., Pavlov, M., Peng, A., Perelman, A., de Avila Belbute Peres, F., Petrov, M., de Oliveira Pinto, H.P., Michael, Pokorny, Pokrass, M., Pong, V.H., Powell, T., Power, A., Power, B., Proehl, E., Puri, R., Radford, A., Rae, J., Ramesh, A., Raymond, C., Real, F., Rimbach, K., Ross, C., Rotsted, B., Roussez, H., Ryder, N., Saltarelli, M., Sanders, T., Santurkar, S., Sastry, G., Schmidt, H., Schnurr, D., Schulman, J., Selsam, D., Sheppard, K., Sherbakov, T., Shieh, J., Shoker, S., Shyam, P., Sidor, S., Sigler, E., Simens, M., Sitkin, J., Slama, K., Sohl, I., Sokolowsky, B., Song, Y., Staudacher, N., Such, F.P., Summers, N., Sutskever, I., Tang, J., Tezak, N., Thompson, M.B., Tillet, P., Tootoonchian, A., Tseng, E., Tuggle, P., Turley, N., Tworek, J., Uribe, J.F.C., Vallone, A., Vijayvergiya, A., Voss, C., Wainwright, C., Wang, J.J., Wang, A., Wang, B., Ward, J., Wei, J., Weinmann, C., Welihinda, A., Welinder, P., Weng, J., Weng, L., Wiethoff, M., Willner, D., Winter, C., Wolrich, S., Wong, H., Workman, L., Wu, S., Wu, J., Wu, M., Xiao, K., Xu, T., Yoo, S., Yu, K., Yuan, Q., Zaremba, W., Zellers, R., Zhang, C., Zhang, M., Zhao, S., Zheng, T., Zhuang, J., Zhuk, W., Zoph, B.: Gpt-4 technical report (2024)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML (2021)
- 22. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016)
- Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: CVPR (2016)

- Sakaridis, C., Dai, D., Gool, L.V.: Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In: ICCV (2019)
- 25. Sakaridis, C., Dai, D., Gool, L.V.: Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In: ICCV (2021)
- Tsai, Y.H., Hung, W.C., Schulter, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: CVPR (2018)
- 27. Wang, Z., Liu, X., Suganuma, M., Okatani, T.: Cross-region domain adaptation for class-level alignment (2022)
- 28. Xie, B., Li, S., Li, M., Liu, C.H., Huang, G., Wang, G.: Sepico: Semantic-guided pixel contrast for domain adaptive semantic segmentation. In: IEEE TPAMI (2023)
- 29. Zara, G., Conti, A., Roy, S., Lathuiliere, S., Rota, P., Ricci, E.: The unreasonable effectiveness of large language-vision models for source-free video domain adaptation. In: ICCV (2023)
- Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., Wen, F.: Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In: CVPR (2021)
- Zheng, Z., Yang, Y.: Unsupervised scene adaptation with memory regularization in vivo. In: IJCAI (2020)