

# CONDA: Continual Unsupervised Domain Adaptation Learning in Visual Perception for Self-Driving Cars

Thanh-Dat Truong<sup>1</sup>, Pierce Helton<sup>1</sup>, Ahmed Moustafa<sup>1</sup>, Jackson David Cothren<sup>2</sup>, Khoa Luu<sup>1</sup>

<sup>1</sup>CVIU Lab, University of Arkansas, Fayetteville, AR, USA

<sup>2</sup>Department of Geosciences, University of Arkansas, Fayetteville, AR, USA

{tt032, pchelton, armoustam, jcothre, khoaluu}@uark.edu

## Abstract

Although unsupervised domain adaptation methods have achieved remarkable performance in semantic scene segmentation, these approaches remain impractical in real-world use cases. In practice, the segmentation models may encounter new data that have not been seen yet. Also, the previous data training of segmentation models may be inaccessible due to privacy problems. Therefore, to address these problems, in this work, we propose a **Continual Unsupervised Domain Adaptation (CONDA)** approach that allows the model to continuously learn and adapt with respect to the presence of the new data. Moreover, our proposed approach is designed without the requirement of accessing previous training data. To avoid the catastrophic forgetting problem and maintain the performance of the segmentation models, we present a novel *Bijjective Maximum Likelihood* loss to impose the constraint of predicted segmentation distribution shifts. The experimental results on the benchmark of continual unsupervised domain adaptation have shown the advanced performance of the proposed CONDA method.

## 1. Introduction

Semantic scene segmentation has become one of the most popular research topics in computer vision recently. Its goal is to break down an input image and densely assign each pixel to its corresponding predefined class. There have been new approaches recently based upon deep learning technology that have remarkable results in semantic scene segmentation. [4]. Typically, segmentation is trained on labelled scene data, but annotating the images for semantic segmentation is a time-consuming and expensive process. This is because each and every pixel in the input image must be labelled. A method to reduce the cost of labelling images is to use a simulation to create a large-scale synthetic dataset [47, 49]. Although this strategy saves time, it has a se-

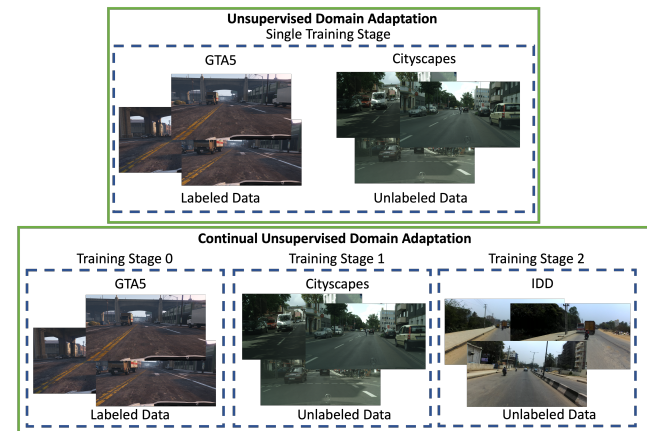


Figure 1. Unsupervised domain Adaptation trains the model on both labeled source data and unlabeled target data simultaneously by a single training stage. Meanwhile, Continual Unsupervised Domain Adaptation first trains on the labeled source data. Then, it continuously adapts the model to the new target domains and at each training stage, the model only accesses data at that stage.

rious trade off when deploying supervised models trained on these synthetic datasets into the real-world data. Particularly, these supervised models trained on the synthetic datasets often perform worse on real images due to a pixel appearance gap between the synthetic and real images and thus are not well-suited for real-image deployment.

To address the problem aforementioned, unsupervised domain adaptation has emerged as a feasible solution. In particular, unsupervised domain adaptation (UDA) aims to learn a model on the large-scale annotated source datasets (referred as the source domain) and adapt to unlabeled target datasets (referred as the target domain) to guarantee its performance on the new domain. Common unsupervised domain adaptation methods try to minimize the distribution discrepancy in the deep representations between the source and target domains. The minimization process is done simultaneously with supervised learning on the

Table 1. **Comparisons in the properties between our proposed approach and other prior methods.** Convolutional Neural Network (CNN), Generative Adversarial Net (GAN), Bijective Network (BiN), Entropy Minimization (EntMin), Segmentation Map (Seg), Depth Map (Depth);  $\ell_{IW}$ : Image-wise Weighting Loss,  $\ell_{CE}$ : Cross-entropy Loss,  $\ell_{Focal}$ : Focal Loss  $\ell_{adv}$ : Adversarial Loss,  $\ell_{Huber}$ : Huber Loss.  $\ell_{square}$ : Maximum Squares Loss,  $\ell_{density}$ : Negative Log-Likelihood Loss,  $\ell_{bmt}$ : Bijective Maximum Likelihood Loss.

Methods	Source-Free Data	Forgetting Aware	Structural Learning	Source Label Require	Target Domain Transfer	Architecture	Designed Loss
AdaptPatch [60]	✗	✗	Weak (Binary label)	Seg	✓	CNN+GAN	$\ell_{adv}$
CBST [73]	✗	✗	—	Seg	✓	CNN	$\ell_{CE}$
ADVENT [64]	✗	✗	Weak (Binary label)	Seg	✓	CNN+GAN	EntMin + $\ell_{adv}$
IntraDA [44]	✗	✗	Weak (Binary label)	Seg	✓	CNN+GAN	EntMin + $\ell_{adv}$
BiMaL [55]	✗	✗	Maximum Likelihood	Seg	✓	CNN + BiN	$\ell_{density}$
SPIGAN [31]	✗	✗	Weak (Binary label)	Seg + Depth	✓	CNN+GAN	$\ell_{adv}$
DADA [65]	✗	✗	Depth-aware Label	Seg + Depth	✓	CNN+GAN	$\ell_{adv} + \ell_{Huber}$
MaxSquare [6]	✗	✗	Weak (Binary label)	Seg	✓	CNN + GAN	$\ell_{square} + \ell_{IW}$
SAC [1]	✗	✗	—	Seg	✓	CNN	$\ell_{CE} + \ell_{Focal}$
<b>Ours</b>	✓	✓	<b>Bijective Maximum Likelihood</b>	—	✓	<b>CNN + BiN</b>	$\ell_{bmt}$

source domain [55, 59, 60, 64, 65]. The discrepancy minimization can be computed at a single level or multiple levels of deep representation using maximum mean discrepancy [19, 35], or adversarial loss within a generative adversarial network framework [8, 22, 23, 59, 61], or contrastive learning [28, 69]. In addition, other domain adaptation approaches utilize the cross entropy loss with pseudo labels with well designed learning strategies [1, 70, 73].

These approaches have shown their potential performance in semantic segmentation with domain adaptation; however, their practical applications are limited. For example, in autonomous driving, vehicles may encounter various urban or highway scenarios with a diversity of environments, e.g., weather, lighting, or geographical conditions, and each of these conditions can be considered as a new target domain. Meanwhile, in the domain adaptation approach, the model is trained on a particular given target domain not able to update itself with the presence of new data. Moreover, the training data, in some cases, is protected during the adaptation phase to preserve privacy. In this work, we address this problem by introducing a continual learning framework for semantic scene segmentation under the domain adaptation setting. The goal of our proposed framework is to sequentially, continuously learn and adapt the model to the new incoming unlabeled target domain while maintaining the performance of the model on previous target domains and alleviating the catastrophic forgetting problem. In contrast to prior domain adaptation approaches, our continual unsupervised domain adaptation does not access the previous training data during the adaptation phase. In other words, only the unlabeled target data is given during the adaptation phase. Figure 1 illustrates the difference between unsupervised domain adaptation and continual unsupervised domain adaptation. Table 1 summarizes the difference between our proposed approach and prior standard domain adaptation approaches.

**Contributions of This Work** In summary, this work presents a novel Continual Unsupervised Domain Adaptation (CONDA) approach to semantic scene segmentation. The contributions in this work can be summarized as follows. Firstly, the problem of continual domain adaptation is formulated in semantic scene segmentation by regularizing the distribution shift of predictions between source and target domains to avoid the catastrophic forgetting problem. As opposed to prior adaptation methods, only data from the target domains are given during the adaptation phase in our continual learning framework. Secondly, given the formulated problem, the distribution shift regularizer is further derived into the Bijective Maximum Likelihood loss that can be used to measure the distribution shift without the demand of the source training data. Then, the bijective maximum likelihood loss is formed using the bijective network so that the loss is able to capture global structural information and represents the distributions of the segmentation. Finally, though our experiments on the benchmark of GTA → Cityscapes → IDD → Mapillary, our proposed approach outperforms the prior methods and achieves state-of-the-art results.

## 2. Related Work

### 2.1. Semantic Segmentation via Deep Learning

Fully Convolutional Networks (FCNs) are the preferred approach for the task of semantic segmentation due to their capacity for learning and high accuracy. Combining FCNs with an encoder-decoder structure produces additional refinements in accuracy. FCNs [4, 34] were first applied to the task of segmentation with multiple convolutional layers followed by spatial pooling. Later approaches [33, 46] combined upsampled, high-level feature maps with low-level feature maps prior to decoding, collecting more information and still producing precise instance borders. Ad-

ditional works [4, 68] improve the performance of models while preserving the field of view by using dilated convolutions. Spatial pyramid pooling has seen success in other recent developments [3, 5]. This method allows the model to acquire contextual information at multiple levels, generating more global information at higher network layers. In a novel FCN architecture, Deeplabv3+ [3] used the encoder-decoder structure in unison with spatial pyramid pooling to produce a faster, stronger network. The latest works use Transformer-based backbones [7, 24, 67] to create more proficient and advanced semantic segmentation networks.

## 2.2. Unsupervised Domain Adaptation via Deep Learning

The common domain adaptation approaches are domain discrepancy minimization [19, 35, 61], adversarial learning [8, 9, 21–23, 59], entropy minimization [39, 44, 64, 71], self-training [25, 57, 73]. Within this work’s scope, UDA is focused on semantic segmentation, and adversarial training is the most commonly employed approach to UDA for semantic segmentation. The adversarial training paradigm is much like generative adversarial networks (GANs) in that they both aim to train a predictive discriminator on the domain of inputs while the segmentation network tries to fool the discriminator. Both the training of the adversarial step and the supervised segmentation task occur simultaneously on the source and target domains. The first instance of a GAN-based UDA approach to segmentation was introduced by Hoffman *et al.* [21]. Later, global and class-wise adaptation learned from the application of adversarial learning on pseudo labels was presented by Chen *et al.* [9]. After considering the spatial distribution difference, [8] proposed a spatial-aware adaptation method to align two domains along with a target-guided distillation loss. A conditional generator that transforms feature maps of the source domain to better match the target domain was learned by Hong *et al.* [23]. Tasi *et al.* [59] learned a consistency of scene layout and local context between target and source domains by using adversarial learning. There exist prior methods that utilize GANs to use source images to synthesize target images [39, 71]. Hoffman *et al.* [22] presented Cycle-Consistent Adversarial Domain Adaptation which aligns at both the pixel and feature level representations. Zhu *et al.* [72] developed the penalization of easy and hard source examples by implementing a conservative loss in the adversarial framework. We *et al.* [66] proposed a DCAN framework that uses channel-wise feature alignment in the segmentation networks. Sakaridis *et al.* [51] proposed a UDA framework on scene understanding that can gradually adapt a segmentation model for increasingly foggy images, i.e. no fog to high fog. Recently, self-supervised approaches [1, 24, 70] have shown their state-of-the-art performance in domain adaptation tasks. Araslanov *et al.* [1] developed

a simple self-supervised framework trained on pseudo labels without the demand of extra training rounds. Zhang *et al.* [70] introduced a self-training approach that is able to denoise pseudo labels and learns structural information by enforcing the consistency between augmentations. Hoyer *et al.* [24] improved the performance of segmentation models by utilizing pseudo labels and introducing a powerful transformer-based backbone. Later, Hoyer *et al.* further improved their approach by using multi-resolution cropped images [25] and masked image consistency learning strategy [26] to enhance contextual learning. Fashes *et al.* [17] presented a new prompt-based approach to zero-shot unsupervised domain adaptation.

## 2.3. Continual Learning

These methods aim to continuously update the model with respect to an incoming streaming data. The simple approach to this task is to re-train or fine-tune the model with the new data. However, under the assumption that the original training data is inaccessible due to the privacy problem, fine-tuning the model on the new updated data could lead to the problem of catastrophic forgetting [13, 14, 18, 30, 36, 38, 42, 43, 48, 54] which is the most challenging problem in continual learning. Catastrophic forgetting refers the existence of performance drops on the original data when the model is updated. This challenge can be alleviated by constraining the updated model to a similar locale of the previous model. This constraint can be done by imposing the regularization on gradients or weights [30, 36], or measuring the probabilities [32], or distilling the intermediary features [13]. Recently, several works [2, 37] developed the continual learning framework for semantic segmentation. However, these approaches are not applicable for the unsupervised domain adaptation setting as these requires the ground truth labels of data during the learning process. [63] presented a continual learning framework where the model is sequentially trained on multiple labeled data domains. [50] proposed a continual learning framework under the unsupervised domain adaptation setting. However, this approach requires storing all data for the purpose of rehearsal. [53] introduced a multi-dead knowledge distillation framework for the continual unsupervised domain but this approach requires access to the source data during the training phase. The recent approaches presented a clustering approach to continual learning [27, 58] In our work, we focus on the continual learning framework under the unsupervised adaptation setting where the source data is not used in the adaptation process.

## 3. The Proposed CONDA Framework

In this section, we present the proposed domain adaptation approach in the context of continual learning. Let  $\mathbf{x}_s \in \mathcal{X}_s$  be the RGB input image of the source domain,  $\hat{\mathbf{y}}_s \in \mathcal{Y}_s$




Figure 2. **The proposed CONDA Framework.** The input images  $\mathbf{x}$  are first forward to the segmentation network  $F$ . Then, the segmentation maps  $\mathbf{y}$  are projected into the deep latent space by the bijective network  $G$  to measure the distribution shift of target predictions compared to the original distribution of source predictions.

be the corresponding segmentation ground truth of  $\mathbf{x}_s$ , and  $F$  be the deep network that maps the input image  $\mathbf{x}$  into the segmentation  $\mathbf{y}$ , i.e  $\mathbf{y} = F(\mathbf{x}, \theta)$ ,  $\theta$  is the deep network parameters. We assume that the network  $F$  has been first well learned on the source domain by supervised learning defined as follows,

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{\mathbf{x}_s, \mathbf{y}_s \in \mathcal{X}_s \times \mathcal{Y}_s} \mathcal{L}_s(F(\mathbf{x}_s, \theta), \hat{\mathbf{y}}_s) \quad (1)$$

where  $\mathcal{L}_s$  is the supervised loss, i.e., the cross-entropy loss, between a predicted segmentation and a ground-truth segmentation. Then, given a well-learned  $F$ , we will continuously learn the model on the incoming target data. **In our setting, we assume that the given source dataset is not used to train during the adaptation phase.** Let  $\mathbf{x} \in \mathcal{X}_t$  be the RGB input image of the incoming target domain  $\mathcal{X}_t$ . We assume there could be multiple incoming target datasets that the model needs to continually adapt to. Let  $q(\mathbf{y})$  represent the distribution of predicted segmentation of the previous model learned on the source domain,  $p(\mathbf{y})$  be the distribution of predicted segmentations of the current updated model.

In general, domain adaptation in the context of continuous learning of the semantic segmentation task can be formed as follows,

$$\theta^* = \arg \min_{\theta} [\mathbb{E}_{\mathbf{x} \sim \mathcal{X}_t} \mathcal{L}_t(\mathbf{y}) + \lambda \mathcal{D}(p(\mathbf{y}), q(\mathbf{y}))] \quad (2)$$

where  $\lambda$  is the hyper-parameter imposing the impact of the distribution shift,  $\mathcal{L}_t$  is the unsupervised loss used in the domain adaptation setting (e.g. entropy loss [64], cross-entropy loss with pseudo labels [1], adversarial loss [65]), and  $\mathcal{D}(p(\mathbf{y}), q(\mathbf{y}))$  is the distance between distributions of source predictions and target predictions, respectively. As shown in Eqn. (2), the constraint  $\mathcal{D}(p(\mathbf{y}), q(\mathbf{y}))$  ensures that the unsupervised adaptation procedure does not lead to catastrophic forgetting. In particular, while the model continuously updates with respect to the incoming target data, it

must not forget the knowledge that has been learned in previous training using the source data. This is ensured by imposing the distribution shift between source prediction and target prediction distributions. Figure 2 illustrates our proposed framework.

### 3.1. Distribution Shift Modeling by Bijective Network

There are several options for  $\mathcal{D}$  to measure the divergence between the two distributions  $p(\mathbf{y})$  and  $q(\mathbf{y})$ . In our paper, we define distance  $\mathcal{D}$  as the Kullback–Leibler (KL) divergence which is a common statistical distance shown in Eqn. (3).

$$\mathcal{D}(p(\mathbf{y}), q(\mathbf{y})) = \int \log \left( \frac{p(\mathbf{y})}{q(\mathbf{y})} \right) p(\mathbf{y}) d\mathbf{y} \quad (3)$$

Eqn. (3) measures the divergence between distributions  $p(\mathbf{y})$  and  $q(\mathbf{y})$ . However, Eqn. (3) requires access to the source dataset during the adaptation phase which conflicts with the setting in our problem. Thus, Eqn. (3) is re-formed as follows,

$$\begin{aligned} & \int \log \left( \frac{p(\mathbf{y}_t)}{q(\mathbf{y}_t)} \right) p(\mathbf{y}) d\mathbf{y} \\ &= \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} \log(p(\mathbf{y})) - \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} \log(q(\mathbf{y})) \\ &\leq -\mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} \log(q(\mathbf{y})) \end{aligned} \quad (4)$$

With any form of the distribution  $p$ , the above inequality always holds as  $p(\mathbf{y}) \in [0, 1]$  and  $\log(p(\mathbf{y})) \leq 0$ . Then, instead of directly computing the KL divergence between  $p(\mathbf{y})$  and  $q(\mathbf{y})$ , we define  $\mathcal{D}_{MaL}$  as the upper bound of Eqn. (3),

$$\mathcal{D}_{MaL}(p(\mathbf{y}), q(\mathbf{y})) = -\mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} \log(q(\mathbf{y})) \quad (5)$$

As Eqn. (5) is an upper bound of Eqn. (3) due to Eqn. (4). Then, minimizing Eqn. (5) also imposes the distance between two distributions  $p(\mathbf{y})$  and  $q(\mathbf{y})$ .



As  $q(\mathbf{y})$  represents the distribution of predicted segmentation of the model learned on the source dataset, we propose to model the distribution  $q(\mathbf{y})$  by the bijective network. Let  $G : \mathbf{Y}_S \rightarrow \mathcal{Z}$  be the bijective network that maps a segmentation to the latent space, i.e.  $\mathbf{z} = G(\mathbf{y}, \theta_G)$ , where  $\mathbf{z} \sim \pi(\mathbf{z})$  is the latent variable,  $\theta_G$  is the set parameters of  $G$ , and  $\pi$  is the prior distribution.

Then, by the change of variable formula, the distribution  $p(\mathbf{y})$  can be formed as follows,

$$\log p(\mathbf{y}) = \log \pi(\mathbf{z}) + \log \left| \frac{\partial G(\mathbf{y}, \theta_G)}{\partial \mathbf{y}} \right|, \quad (6)$$

where  $\left| \frac{\partial G(\mathbf{y}, \theta_G)}{\partial \mathbf{y}} \right|$  is the Jacobian determinant of  $G(\mathbf{y}, \theta_G)$  with respect to  $\mathbf{y}$ . Then, the network  $G$  can be learned by minimizing the negative log-likelihood on the source segmentation maps as follows,

$$\begin{aligned} \theta_G^* &= \arg \min_{\theta_G} \mathbb{E}_{\mathbf{y} \in \mathcal{Y}_s} \left[ -\log(q(\mathbf{y})) \right] \\ &= \arg \min_{\theta_G} \mathbb{E}_{\mathbf{y} \in \mathcal{Y}_s} \left[ \log \pi(\mathbf{z}) + \log \left| \frac{\partial G(\mathbf{y}, \theta_G)}{\partial \mathbf{y}} \right| \right]. \end{aligned} \quad (7)$$

Generally, there are several choices for the prior distribution  $\pi$ . However, the ideal prior distribution should meet two conditions: (i) simplicity in the density estimation, and (ii) easy in sampling. Therefore, taking these conditions into account, we adopt the the Normal distribution as the prior distribution  $\pi$ . It should be noted that the prior distribution is not limited to the Normal distribution; any form of distribution that satisfies the two conditions can also be adopted.

To enhance the modeling capability of the bijective networks  $G$ , we design  $G$  as a multi-scale architecture network in which each scale is designed as an invertible network. Several deep neural structures [11, 12, 15, 16, 29, 41, 56] can be adopted to construct the invertible network at each scale. We assume that the bijective network  $G$  and the segmentation network  $F$  are given during the adaptation phase.

## 4. Experimental Results

In this section, we first review the datasets used in our experiments. Then, we describe our implementation in detail and the evaluation metric used to measure the performance of the segmentation models. Next, we present our experimental results on the standard benchmark of continual unsupervised domain adaptation. Finally, we compare our qualitative results with the baselines.

### 4.1. Datasets and Network Architectures

**GTA5** is a collection of synthetic images and semantic labels at a resolution of  $1914 \times 1052$  pixels. The 24,966 images were collected from the GTA5 game engine using

the communication between the game engine and the graphics hardware. In total, 33 class categories are included in the dataset, but only 7 categories are compatible with the Cityscapes [10], IDD [62], and Mapillary [40] datasets used in our experiments.

**Cityscapes** includes 3,975 images captured from 50 different cities. Each image has a high-quality, semantic, dense pixel annotation of 30 object classes from the urban settings. Cityscapes was created to give researchers more annotated, high quality, and real world data.

**Vistas (Mapillary Vistas Dataset)** [40] is a dataset of various street-level images with pixel-accurate and instance-specific human annotations from areas throughout the world. The data includes 25,000 high-quality images and 124 object categories.

**IDD** [62] contains images from several cities in India. The dataset has 10,000 semantically labeled images. Of the 34 object classes included in the dataset, only 7 classes that are common to SYNTHIA, Cityscapes, and Vistas are used in our experiments.

**Implementation** The DeepLab-V2 [4] architecture with a Resnet-101 [20] backbone is used in all experiments. In our experiments, the image size is set to the resolution of  $640 \times 320$ . We use Exponential Moving Averages (EMA) to update the model. In our work, the unsupervised loss  $\mathcal{L}_t$  is defined as the cross-entropy loss with the pseudo labels computed by the EMA model. The value of  $\lambda$  is set to 0.005.

We adopt the structure of [16, 29, 55] for our generator  $G$ . We used PyTorch [45] to implement the framework. 4 NVIDIA Quadro P800 GPUs with 48GB of VRAM each were utilized during development. The model is optimized by the Stochastic Gradient Descent optimizer where the learning rate is set to  $2.5 \times 10^{-4}$ .

**Evaluation Metric** The performance of semantic segmentation models are often measured by the mean Intersection over Union (mIoU) metric over all classes, expressed in percentage. Follow the evaluation protocol of prior domain adaptation methods [1, 55, 64, 65], we also adopt the mIoU metric in our experiments to measure the performance of different methods.

### 4.2. Quantitative Results

In our experiments, we compare our methods with the baseline (AdvEnt) [64], Multi-Dis. [52], and MTKT [52] on the benchmark of  $\text{GTA5} \rightarrow \text{Cityscapes} \rightarrow \text{IDD} \rightarrow \text{Mapillary}$ . There are two setting in our experiments, i.e., (1) The Multi-target Oracle setting means the model is simultaneous training on all the target domains, and (2) The Continual setting means that the model is trained target domains sequentially. It should be noted that the baseline (AdvEnt) use the source data during the adaptation phase.

**GTA  $\rightarrow$  Cityscapes:** Table 2 displays our method's

	Method	Target	flat	constr.	object	nature	sky	human	vehicle	mIoU	mIoU Avg.
Multi-Target Oracle Setting	Multi-Target Baseline [64]	Cityscapes	93.6	80.6	26.4	78.1	81.5	51.9	76.4	69.8	67.8
		IDD	92.0	54.6	15.7	77.2	90.5	50.8	78.6	65.6	
		Mapillary	89.2	72.4	32.4	73.0	92.7	41.6	74.9	68.0	
	Multi-Dis. [52]	Cityscapes	94.6	80.0	20.6	79.3	84.1	44.6	78.2	68.8	68.2
		IDD	91.6	54.2	13.1	78.4	93.1	49.6	80.3	65.8	
		Mapillary	89.0	72.5	29.3	75.5	94.7	50.3	78.9	70.0	
	MTKT [52]	Cityscapes	94.6	80.7	23.8	79.0	84.5	51.0	79.2	70.4	69.1
		IDD	91.7	55.6	14.5	78.0	92.6	49.8	79.4	65.9	
		Mapillary	90.5	73.7	32.5	75.5	94.3	51.2	80.2	71.1	
Continual Setting	Continual Baseline [64]	Cityscapes	<b>92.9</b>	<b>79.0</b>	18.7	<b>76.9</b>	84.1	47.3	72.9	67.4	67.0
		IDD	<b>91.8</b>	51.1	11.6	<b>79.0</b>	91.6	<b>47.5</b>	<b>72.5</b>	63.6	
		Mapillary	90.3	71.7	30.1	76.1	<b>93.9</b>	50.2	77.3	70.0	
	Ours	Cityscapes	89.4	71.1	<b>32.8</b>	75.8	<b>92.8</b>	<b>48.1</b>	<b>75.6</b>	<b>69.4</b>	<b>68.9</b>
		IDD	88.0	<b>69.4</b>	<b>30.6</b>	73.0	<b>92.8</b>	47.1	62.6	<b>66.2</b>	
		Mapillary	<b>90.4</b>	<b>73.3</b>	<b>33.1</b>	<b>76.4</b>	93.8	<b>51.2</b>	<b>80.1</b>	<b>71.2</b>	

Table 2. Segmentation performance (%) on GTA5 → Cityscapes → IDD → Mapillary.

SOTA performance when compared with prior methods on 7 shared classes from the Cityscapes validation set. In particular, our approach achieved a higher IoU than the previous baseline by +1.9%. Comparing the per class results with the baseline, 3 of the 7 classes saw notable improvements, namely “object” (+14.1%), “sky” (+8.7%), and “vehicle” (+2.9%). The remaining classes did not differ from the baseline by a large degree.

**GTA → Cityscapes → IDD:** In the IDD row of Table 2, the improvements for the 7 shared classes can be seen when adapting from GTA5 → Cityscapes → IDD. Again, our approach improves the baseline IoU results, with a per class improvement in IoU of 18.3%, 19.0%, and 1.2% for the classes of “constr”, “object”, and “sky”, respectively. Furthermore, the overall mIoU accuracy improved by 2.6% from the previous baseline.

**GTA → Cityscapes → IDD → Mapillary:** The final experiment tested the adaptation of the model to the Mapillary dataset. Table 2 displays the results of this experiment. Our approach improved the previous baseline’s mIoU of all classes by an average of 1.2%. The improvement by class was 0.1% for “flat”, 1.6% for “constr”, 3.0% for “object”, 0.3% for “nature”, 1.0% for “sky”, and 2.8% for “vehicle”. Across all experiments, our new approach improved the mIoU average by 1.9%, setting a new SOTA result for the task of continuous learning. It should be noted that without accessing the source training data during the adaptation phase, overall, our approach still gains better results com-




Figure 3. Qualitative Results. We compare our results with the baseline method without baseline method (AdvEnt) [64].

pared to the baseline method.

### 4.3. Qualitative Results

Figure 3 illustrates the qualitative results our experiments. Our new method achieves better results than without adaptation. The continuous learning structure allows our model to continue the learning process, so improvements are made across all classes. The predictions produced by the model more closely matches the ground truth than the baseline without adaptation. The noise is reduced, the edges are clearly defined, and the structure of predictions reflects the objects in the image to a greater degree. Overall, the model maintains the ability to generate accurate predictions and improves these predictions during test time.

## 5. Conclusions

In this paper, we have presented a novel solution to the problem of continual domain adaptation in semantic scene segmentation. By using the proposed Maximum Likelihood Loss to impose the constraint of distribution shift, we have avoided the problem of catastrophic forgetting and allowed the models to continuously update and improve their performance with respect to the new target data. Moreover, the proposed approach does not require the presence of previous training data. Therefore, our approach ensures the privacy of the previous training data. Experiments on the standard benchmark have shown the performance improvements of our method.

**Acknowledgment** This work is supported by NSF Data Science, Data Analytics that are Robust and Trusted (DART) and Google Initiated Research Grant. We also acknowledge the Arkansas High Performance Computing Center for providing GPUs.

## References

- [1] Nikita Araslanov, , and Stefan Roth. Self-supervised augmentation consistency for adapting semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 3, 4, 5
- [2] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulò, Elisa Ricci, and Barbara Caputo. Modeling the background for incremental learning in semantic segmentation. In *Proc. Conf. Comp. Vision Pattern Rec.*, 2020. 3
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 3
- [4] 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. *TPAMI*, 2018. 1, 2, 3, 5
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 3
- [6] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *ICCV*, 2019. 2
- [7] Runfa Chen, Yu Rong, Shangmin Guo, Jiaqi Han, Fuchun Sun, Tingyang Xu, and Wenbing Huang. Smoothing matters: Momentum transformer for domain adaptive semantic segmentation. *CoRR*, 2022. 3
- [8] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *CVPR*, 2018. 2, 3
- [9] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *ICCV*, 2017. 3
- [10] 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*, 2016. 5
- [11] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation, 2015. 5
- [12] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp, 2017. 5
- [13] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Proc. Europ. Conf. Computer Vision*, 2020. 3
- [14] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. Plop: Learning without forgetting for continual semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4040–4050, 2021. 3
- [15] Chi Nhan Duong, Kha Gia Quach, Khoa Luu, T Hoang Ngan Le, Marios Savvides, and Tien D Bui. Learning from longitudinal face demonstration—where tractable deep modeling meets inverse reinforcement learning. *IJCV*, 2019. 5
- [16] Chi Nhan Duong, Thanh-Dat Truong, Khoa Luu, Kha Gia Quach, Hung Bui, and Kaushik Roy. Vec2face: Unveil human faces from their blackbox features in face recognition. In *CVPR*, 2020. 5
- [17] Mohammad Fahes, Tuan-Hung Vu, Andrei Bursuc, Patrick Pérez, and Raoul de Charette. Poda: Prompt-driven zero-shot domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18623–18633, 2023. 3
- [18] Robert French. Catastrophic forgetting in connectionist networks. 1999. 3
- [19] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. 2, 3
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [21] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. FCNs in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv:1612.02649*, 2016. 3
- [22] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 2, 3
- [23] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. Conditional generative adversarial network for structured domain adaptation. In *CVPR*, 2018. 2, 3
- [24] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. DAFormer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *CVPR*, 2022. 3
- [25] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. HRDA: Context-aware high-resolution domain-adaptive semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 3
- [26] Lukas Hoyer, Dengxin Dai, Haoran Wang, and Luc Van Gool. Mic: Masked image consistency for context-enhanced domain adaptation. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11721–11732, 2023. 3
- [27] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5830–5840, 2021. 3
- [28] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*, 2019. 2
- [29] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NIPS*, 2018. 5
- [30] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. 2017. 3
- [31] Kuan-Hui Lee, German Ros, Jie Li, and Adrien Gaidon. SPIGAN: Privileged adversarial learning from simulation. In *ICLR*, 2019. 2
- [32] Z. Li and D. Hoiem. Learning without forgetting. *Proc. Europ. Conf. Computer Vision*, 2016. 3
- [33] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, 2017. 2
- [34] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 2
- [35] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 2, 3
- [36] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. 2017. 3
- [37] Umberto Michieli and Pietro Zanuttigh. Incremental learning techniques for semantic segmentation. 2019. 3
- [38] Umberto Michieli and Pietro Zanuttigh. Incremental learning techniques for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, pages 0–0, 2019. 3
- [39] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *CVPR*, 2018. 3
- [40] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017. 5
- [41] Chi Nhan Duong, Kha Gia Quach, Khoa Luu, Ngan Le, and Marios Savvides. Temporal non-volume preserving approach to facial age-progression and age-invariant face recognition. In *ICCV*, 2017. 5
- [42] Firat Ozdemir and Orcun Goksel. Extending pretrained segmentation networks with additional anatomical structures. *International journal of computer assisted radiology and surgery*, 14:1187–1195, 2019. 3
- [43] Firat Ozdemir, Philipp Fuernstahl, and Orcun Goksel. Learn the new, keep the old: Extending pretrained models with new anatomy and images. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part IV 11*, pages 361–369. Springer, 2018. 3
- [44] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *CVPR*, 2020. 2, 3
- [45] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. 5
- [46] Tobias Pohlen, Alexander Hermans, Markus Mathias, and Bastian Leibe. Full-resolution residual networks for semantic segmentation in street scenes. In *CVPR*, 2017. 2
- [47] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016. 1
- [48] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 1995. 3
- [49] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, 2016. 1
- [50] Mohammad Rostami. Lifelong domain adaptation via consolidated internal distribution. 2021. 3
- [51] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *ECCV*, 2018. 3
- [52] Antoine Saporta, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Multi-target adversarial frameworks for domain adaptation in semantic segmentation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 5, 6
- [53] Antoine Saporta, Arthur Douillard, Tuan-Hung Vu, Patrick Pérez, and Matthieu Cord. Multi-head distillation for continual unsupervised domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*, 2022. 3
- [54] Sebastian Thrun. Lifelong learning algorithms. In *Springer Learning to Learn*, 1998. 3
- [55] Thanh-Dat Truong, Chi Nhan Duong, Ngan Le, Son Lam Phung, Chase Rainwater, and Khoa Luu. Bimal: Bijective maximum likelihood approach to domain adaptation in semantic scene segmentation. In *International Conference on Computer Vision*, 2021. 2, 5
- [56] Thanh-Dat Truong, Chi Nhan Duong, Minh-Triet Tran, Ngan Le, and Khoa Luu. Fast flow reconstruction via robust invertible  $n \times n$  convolution. *Future Internet*, 2021. 5
- [57] Thanh-Dat Truong, Ngan Le, Bhiksha Raj, Jackson Cothren, and Khoa Luu. Freedom: Fairness domain adaptation approach to semantic scene understanding. In *IEEE Confer-*



- ence on *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [58] Thanh-Dat Truong, Hoang-Quan Nguyen, Bhiksha Raj, and Khoa Luu. Fairness continual learning approach to semantic scene understanding in open-world environments. In *NeurIPS*, 2023. 3
- [59] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018. 2, 3
- [60] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative representations. *arXiv:1901.05427*, 2019. 2
- [61] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. 2, 3
- [62] G. Varma, A. Subramanian, Anoop M. Namboodiri, Manmohan Chandraker, and C. V. Jawahar. Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1743–1751, 2019. 5
- [63] Riccardo Volpi, Diane Larlus, and Grégory Rogez. Continual adaptation of visual representations via domain randomization and meta-learning. In *Proc. Conf. Comp. Vision Pattern Rec.*, 2021. 3
- [64] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, 2019. 2, 3, 4, 5, 6
- [65] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *ICCV*, 2019. 2, 4, 5
- [66] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S. Davis. DCAN: Dual channel-wise alignment networks for unsupervised scene adaptation. In *ECCV*, 2018. 3
- [67] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021. 3
- [68] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016. 3
- [69] Xiangyu Yue, Zangwei Zheng, Shanghang Zhang, Yang Gao, Trevor Darrell, Kurt Keutzer, and Alberto Sangiovanni Vincentelli. Prototypical cross-domain self-supervised learning for few-shot unsupervised domain adaptation. In *CVPR*, 2021. 2
- [70] 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. *arXiv preprint arXiv:2101.10979*, 2021. 2, 3
- [71] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 3
- [72] Xinge Zhu, Hui Zhou, Ceyuan Yang, Jianping Shi, and Dahua Lin. Penalizing top performers: Conservative loss for semantic segmentation adaptation. In *ECCV*, 2018. 3
- [73] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018. 2, 3