

Domain Adaptive Video Semantic Segmentation via Cross-Domain Moving Object Mixing

Kyusik Cho Suhyeon Lee Hongje Seong Euntai Kim*

School of Electrical and Electronic Engineering, Yonsei University, Seoul, Korea

{ks.cho, hyeon93, hjseong, etkim}@yonsei.ac.kr

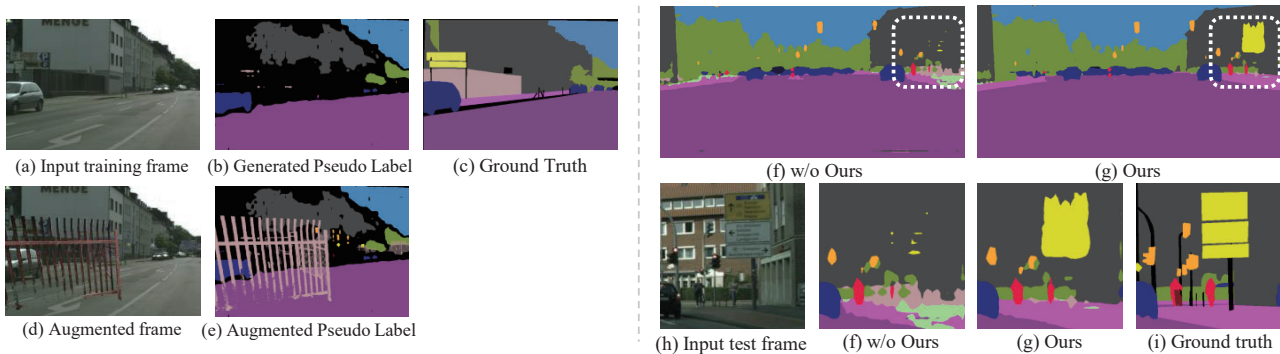


Figure 1. **Motivation: self-training with the biased model.** (a) Given the training sample of the target domain, the model biased toward easy-to-transfer classes generates (b) the pseudo-label without hard-to-transfer classes, such as *sidewalk*, *fence*, and *sign*. (c) The ground truth label is shown for visibility. However, the proposed method pastes several classes from the source domain to (d) the target domain video frames and (e) pseudo-labels. (f) The model trained on the *biased* pseudo-labels cannot detect the *sign* and confused the *sidewalk* with *fence* and *terrain* in the test sample. (g) The model trained with our method detects the classes better and is less noisy in the hard-to-transfer classes.

Abstract

The network trained for domain adaptation is prone to bias toward the easy-to-transfer classes. Since the ground truth label on the target domain is unavailable during training, the bias problem leads to skewed predictions, forgetting to predict hard-to-transfer classes. To address this problem, we propose Cross-domain Moving Object Mixing (CMOM) that cuts several objects, including hard-to-transfer classes, in the source domain video clip and pastes them into the target domain video clip. Unlike image-level domain adaptation, the temporal context should be maintained to mix moving objects in two different videos. Therefore, we design CMOM to mix with consecutive video frames, so that unrealistic movements are not occurring. We additionally propose Feature Alignment with Temporal Context (FATC) to enhance target domain feature discriminability. FATC exploits the robust source domain features, which are trained with ground truth labels, to learn discriminative target domain features in an unsupervised manner by filtering unre-

liable predictions with temporal consensus. We demonstrate the effectiveness of the proposed approaches through extensive experiments. In particular, our model reaches mIoU of 53.81% on VIPER \rightarrow Cityscapes-Seq benchmark and mIoU of 56.31% on SYNTHIA-Seq \rightarrow Cityscapes-Seq benchmark, surpassing the state-of-the-art methods by large margins.

1. Introduction

Video semantic segmentation, a task of classifying every pixel in every video frame [1, 15, 20, 37], is a fundamental problem in machine vision. Training the segmentation network requires extensive pixel-level annotated data, which costs expensive human labor. On the other hand, advances in computer technology have made it possible to create perfectly annotated synthetic datasets [23, 24, 25]. However, segmentation networks trained with synthetic data perform poorly in the real world because of the domain gap.

To handle this issue, many studies of unsupervised domain adaptation (UDA) have been proposed that improve performance in the target domain without any an-

*Corresponding author

notations [10, 22, 35, 36]. These were mainly progressed in image-level, while some advances were recently studied with video data. DA-VSN [8] has taken a successful first step in the domain adaptive video semantic segmentation (DAVSS) through adversarial training and intra-domain temporal consistency regularization. TPS [34] applied self-training for video, and Park *et al.* [27] used both adversarial training and self-training. The mainly used approaches, adversarial training and self-training, are well-known solutions in image-level UDA, and they were effectively exploited by extending the promising solutions in image-level to video data. In this paper, we further exploit the self-training.

In UDA, the self-training on target domain is performed by learning with pseudo-labels that are generated via segmentation networks trained on the source domain. Here, the segmentation networks would generate good quality of the pseudo-labels for easy-to-transfer classes and vice versa. That means, during self-training, the network will be biased toward easy-to-transfer classes, forgetting to predict hard-to-transfer classes. This problem is depicted in Figure 1.

Inspired by recent studies of UDA for image segmentation [29, 33], we design two advanced DAVSS methods to address the bias problem in videos: Cross-domain Moving Object Mixing (CMOM) and Feature Alignment with Temporal Context (FATC).

The first contribution, CMOM, is a data augmentation method that mixes moving objects in the source domain and target domain. To do this, we can simply adopt the image-level mixing approach [29]. However, this will break the temporal context of the video, *e.g.*, an object may suddenly disappear in the next frame or a new object may be presented, making it an unrealistic video. To prevent this problem, CMOM cuts the source domain moving objects from the consecutive video frames and pastes them into the target domain. The second contribution, FATC, is additionally used to mitigate the bias problem during self-training. FATC forces to generate discriminative features from target domain videos by reducing distances between robust source features and semantically weak target domain features. Here, we can also adopt the image-level alignment approach [33], but this cannot exploit a rich cue of the video: we can predict the results of the current frame not only from the current frame but also from the previous frame by warping the results of the previous frame. Our FATC effectively exploits this cue by predicting the target domain segmentation mask of the current frame twice and filtering the noisy results. Then, the remained target domain features are aligned with the source domain features, enabling the network to learn discriminative and robust features in the target domain without accurate ground truth labels.

With two proposed methods, we self-train a semantic segmentation network, and we test on standard DAVSS

benchmarks. Our proposal achieves the state-of-the-arts performances on both VIPER [23] \rightarrow Cityscapes-Seq [4] and SYNTHIA-Seq [25] \rightarrow Cityscapes-Seq benchmarks, showing the effectiveness of our framework. We also conduct extensive ablation studies to demonstrate that our approach facilitates the networks to learn the robust target domain features in videos.

2. Related Work

Video Semantic Segmentation (VSS). VSS is a problem that classifies all pixels in all video frames. The most notable difference between VSS and image semantic segmentation is that VSS can utilize the temporal context of the video. Two primary goals are considered in utilizing the temporal context, one for accurate predictions and the other one for faster predictions [32]. The methods for accurate predictions exploit richer information from multi-frame contexts [5, 20]. On the other hand, the methods for faster prediction save computation costs by reusing the features extracted from the previous frame [26, 37]. Following the previous DAVSS works [8, 34], we adopt ACCEL [15] as a segmentation network. ACCEL modularized the network for segmenting previous and current frames to strike a balance between speed and accuracy, making trade-offs a choice.

Domain Adaptive Semantic Segmentation. The machine learning algorithm optimizes the model on the training data assuming that the training data distribution and the test data distribution are equal. If this assumption is violated, the model performance drops drastically. The goal of domain adaptation is to minimize this performance drop [10]. The most popular methods in unsupervised domain adaptation (UDA) for semantic segmentation are adversarial learning [8, 10, 22, 30, 31] and self-training [13, 35, 36, 39]. Other techniques are also used for domain adaptation: aligning the feature representations of the two domains [13, 28, 33]; augmenting data by changing the source domain image to a target style [3, 18, 35] or mixing two different domain images [17, 29]; and entropy minimization in the target domain [31].

Mixing Image Augmentation. Data augmentation with mixing images are widely used in various image recognition tasks. ClassMix [21] uses the mixing mask generated based on model prediction to create augmented images and pseudo-labels for unlabeled images, and has proven its effectiveness in semi-supervised semantic segmentation task. Ghiasi *et al.* [6] shows that a simple instance-level mixing only with scale jittering is beneficial for instance segmentation. Mixing techniques were also used in UDA [17, 29] to address the bias problem [39] of self-training approach.

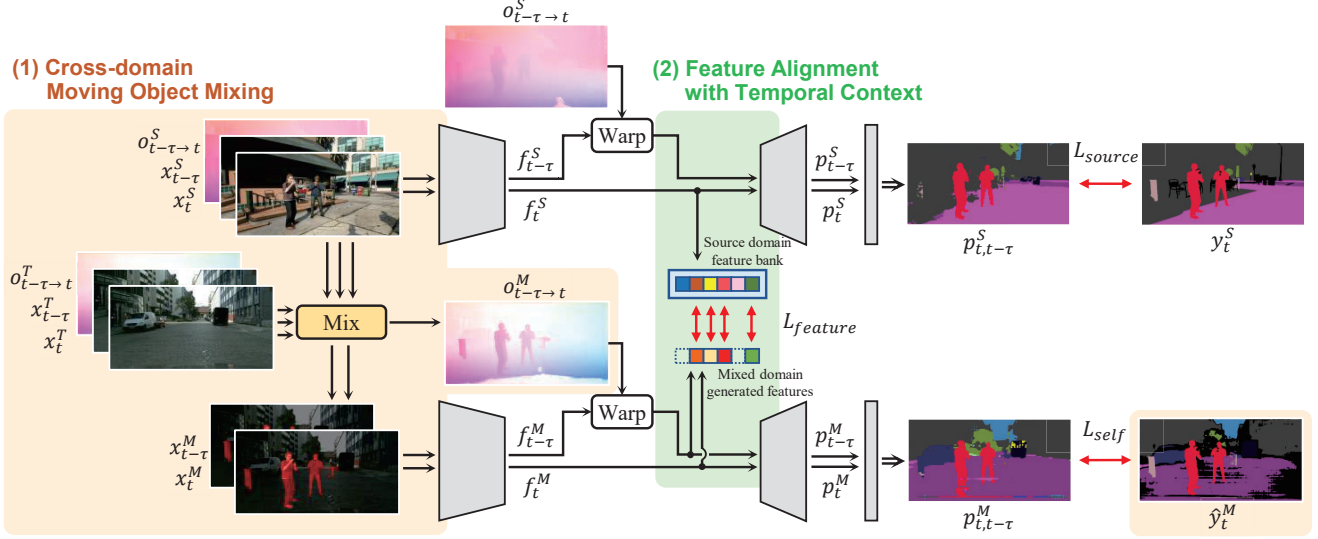


Figure 2. **Illustration of the proposed framework.** Our framework is trained with two proposed methods based on self-training. (1) CMOM generates the mixed domain data sample by mixing data samples from two domains. The mixed domain video frames, optical flow, and pseudo-label replace the target domain data sample during the training process. (2) FATC reduces feature mismatch between the generated mixed domain features and source domain features in the feature bank. The unreliable predictions are excluded by comparing the predictions of the previous frame and the current frame. We train the model with three loss terms: source domain segmentation loss \mathcal{L}_{source} , mixed domain segmentation loss \mathcal{L}_{self} , and feature alignment loss $\mathcal{L}_{feature}$.

Since the pseudo-label is generated from the model’s prediction, the model biased toward easy-to-transfer generates the *biased* pseudo-labels. The mixing approaches address the bias problem by pasting source domain data into the target domain. DACS [29] selects random classes from source images and pastes corresponding pixels to target images. Lee *et al.* [17] mix the tail class objects to further solve the class imbalance problem. All of the image-level mixings, however, have the drawback that they will break the temporal context of the video by ignoring the movement of objects when they are applied to video. Our approach solves the problem by cutting the source domain moving objects from the consecutive video frames and pasting them into the target domain.

Domain Adaptive Video Semantic Segmentation (DAVSS). DAVSS was initially tackled by DA-VSN [8]. DAVSS research mainly focuses on extending image domain adaptation strategies to video using temporal information. For example, DA-VSN [8] uses a sequence of predictions for adversarial learning and enforces consistent prediction in consecutive frames in the target domain. Recently, TPS [34] uses consistency learning between frames in the target domain as the main strategy. Meanwhile, Park *et al.* [27] conduct DAVSS on different architectures and datasets from them, and use both adversarial learning and self-training. Our work also uses self-training, and we further consider the bias problem which has not been explored with video data.

3. Methodology

3.1. Problem Formulation and Overview

Let \mathbb{S} and \mathbb{T} be the source and target domains. In the domain adaptive video semantic segmentation (DAVSS) setting, the source domain video frames $\{x_{1:T}^{\mathbb{S}}\}_{n=1}^{N_{\mathbb{S}}}$ with the corresponding pixel-wise labels $\{y_{1:T}^{\mathbb{S}}\}_{n=1}^{N_{\mathbb{S}}}$ and the unlabeled target domain video frames $\{x_{1:T}^{\mathbb{T}}\}_{n=1}^{N_{\mathbb{T}}}$ are given. T is the length of the video, and $N_{\mathbb{S}}$ and $N_{\mathbb{T}}$ denotes the number of video samples in the source and target datasets, respectively. To exploit temporal information, we adopt FlowNet [14] to produce the source and target domain optical flows $o_{t-\tau \rightarrow t}^{\mathbb{S}}$ and $o_{t-\tau \rightarrow t}^{\mathbb{T}}$ from time $t - \tau$ to t , respectively. In this paper, we further consider the domain \mathbb{M} , a mixture of source and target domains [29].

The overall framework is illustrated in Figure 2. Our framework focuses on the self-training based DAVSS with two proposed methods: Cross-domain Moving Object Mixing (CMOM) and Feature Alignment with Temporal Context (FATC). At each iteration, source domain video frames $\{x_{t-\tau}^{\mathbb{S}}, x_t^{\mathbb{S}}\}$, source domain optical flow $o_{t-\tau \rightarrow t}^{\mathbb{S}}$, mixed domain video frames $\{x_{t-\tau}^{\mathbb{M}}, x_t^{\mathbb{M}}\}$, and mixed domain optical flow $o_{t-\tau \rightarrow t}^{\mathbb{M}}$ enter the our video semantic segmentation (VSS) network. Our network makes source and mixed domain predictions, $p_t^{\mathbb{S}}$ and $p_t^{\mathbb{M}}$ at time t . The frame-level predictions of each domain $p_t^{\mathbb{S}}, p_{t-\tau}^{\mathbb{S}}$ are used to generate a clip-level prediction $p_{t,t-\tau}^{\mathbb{S}}$ with the score fusion module [15]. For training the network, we provide supervision to both source and mixed domains, $y_t^{\mathbb{S}}$ and $\hat{y}_t^{\mathbb{M}}$. All mixed domain

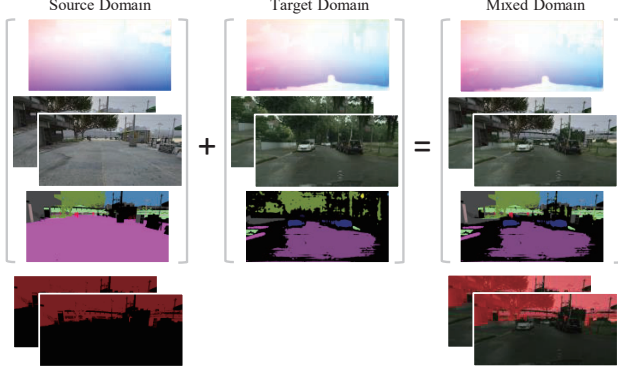


Figure 3. **Illustration of CMOM.** CMOM randomly selects some classes in source domain and cuts them from two consecutive frames. Then, they are pasted into target domain. In order to maintain the temporal context of a video, we create a mask with the same classes in all frames constituting the video and mix not only the RGB frame but also the optical flow. The last row shows the mixing mask in red for visibility.

samples are produced by our proposed CMOM. In the feature space, our FATC excludes noisy features by exploiting the temporal consensus and enhances target domain feature discriminability.

3.2. Cross-Domain Moving Object Mixing

Our method focuses on the self-training. When training the network for DAVSS, due to the absence of the ground truth target labels $\{y_{1:T}^T\}_{n=1}^{N_T}$, the network tends to be biased toward easy-to-transfer classes and forget hard-to-transfer classes. Therefore, the model hard to generate pseudo-labels with hard-to-transfer classes. In the domain adaptive image semantic segmentation, DACS [29] alleviates this problem with a cut-and-paste data augmentation strategy between the source and target domain. DACS data augmentation is formulated as

$$\begin{aligned} x^M &= M \odot x^S + (1 - M) \odot x^T, \\ \hat{y}^M &= M \odot \hat{y}^S + (1 - M) \odot \hat{y}^T, \end{aligned} \quad (1)$$

where M is the pixel selection mask of randomly selected classes, and \odot is the pointwise product. x^S , x^T , and x^M are the source, target, and mixed domain images, respectively. y^S , \hat{y}^T , and \hat{y}^M are the source ground truth label, target pseudo-label, and mixed domain pseudo-label, respectively.

Inspired by [29], we propose a novel data augmentation method CMOM that cuts several objects, including hard-to-transfer classes, in the source domain video clip and pastes them into the target domain video clip. Unlike image-level domain adaptation, the temporal context should be maintained to mix moving objects in two different videos. Therefore, we design CMOM to mix with consecutive video frames, so that unrealistic movements are not occurring.

The process of CMOM is illustrated in Figure 3. We first randomly select three-quarters of the classes that exist in the current source domain label y_t^S . Secondly, we create binary masks $\{M_t, M_{t-\tau}\}$ in which the class pixels selected from the source video labels $\{y_t^S, y_{t-\tau}^S\}$ are 1 and the remainders are 0. Using the masks, the mixed domain video frames $\{x_{t-\tau}^M, x_t^M\}$, pseudo-label \hat{y}_t^M , and optical flow $o_{t-\tau \rightarrow t}^M$ are generated by using the following equations:

$$\begin{aligned} x_{t-\tau}^M &= M_{t-\tau} \odot x_{t-\tau}^S + (1 - M_{t-\tau}) \odot x_{t-\tau}^T, \\ x_t^M &= M_t \odot x_t^S + (1 - M_t) \odot x_t^T, \\ \hat{y}_t^M &= M_t \odot \hat{y}_t^S + (1 - M_t) \odot \hat{y}_t^T, \\ o_{t-\tau \rightarrow t}^M &= M_{t-\tau} \odot o_{t-\tau \rightarrow t}^S + (1 - M_{t-\tau}) \odot o_{t-\tau \rightarrow t}^T. \end{aligned} \quad (2)$$

In this work, we set the τ to 1, and the target pseudo-labels $\{\hat{y}_{1:T}^T\}_{n=1}^{N_T}$ are obtained by using the pre-trained DAVSN model [8] and IAST pseudo-label policy [19] before the self-training phase. This is due to the stability of the network, and they can be replaced by any pre-trained DAVSS model and pseudo-label policy.

This simple data augmentation method gives us several advantages. First, CMOM provides supervision of all classes, addressing the problem of forgetting to predict hard-to-transfer classes due to bias. In addition, the incorrect pseudo-labels generated by erroneous predictions are alleviated by mixing ground truth labels. It helps to overcome the drawbacks of self-training with biased pseudo-labels. Another advantage is that the temporal context of the video is maintained, which can be easily damaged by the naive image-level approach. To maintain the context, we selected the same object in both frames and mixed the optical flow; the movement of each object is accessible with a sharp optical flow. Lastly, CMOM is data augmentation for videos with low computational cost and simplicity. In order to use image-level adaptation strategies in the video, the method must be applied to all frames. Therefore, approaches that require much computation are burdensome to be used in the video. CMOM does not require much computation and can be used easily. We reported the computational cost in ablation experiments.

CMOM is used as target domain data augmentation in our paper. In other words, the mixed domain video frames x^M and pseudo-labels \hat{y}^M are used instead of x^T and \hat{y}^T in the training process.

3.3. Feature Alignment with Temporal Context

Although CMOM helps to enhance the target feature quality by providing supervision of all classes in the mixed domain, the discriminability of the target domain is not as good as that of a fully-supervised trained source domain. In this section, we propose Feature Alignment with Temporal Context (FATC) for enhancing the discriminability

of the target domain. FATC exploits the robust source domain features, which are trained with ground truth labels, to learn discriminative target domain features in an unsupervised manner by filtering unreliable predictions with temporal consensus.

Here, we focus on that there are several recent works that focus on the feature alignment for domain adaptive image semantic segmentation [13, 28, 33], and we can adopt image-level alignment approach to FATC. However, image-level approach cannot exploit a rich cue of the video: we can predict the results of the current frame not only from the current frame but also from the previous frame by warping the results of the previous frame. Focusing on this point, we extend the image-level feature alignment [33] to the DAVSS by performing the feature denoising process.

Our FATC is illustrated in Figure 4. The feature denoising process of FATC excludes noisy features using the temporal context of the video clip. We cannot filter only correct predictions from target domains without ground truth labels, but we can exclude obvious incorrect predictions with video frames. Each video frame is strongly correlated and the corresponding pixel should be predicted to be the same class for the frames. In other words, a consistent prediction between two frames does not guarantee correct answers, but different predictions across frames are always inaccurate.

We check this temporal correspondence to detect and exclude noisy features. Using the optical flow $o_{t-\tau \rightarrow t}^M$, we warp $\{p_{t-\tau}^M, f_{t-\tau}^M\}$ to $\{\tilde{p}_t^M, \tilde{f}_t^M\}$, where $f_t^{M,c}$ indicates the features of each class c , p_t^M is the predictions, and \tilde{f}_t^M and \tilde{p}_t^M are the warped features and predictions from time $t - \tau$ to t , respectively. Then, we generate the class-wise valid region binary mask $v_t^{M,c}$ by comparing p_t^M and $\tilde{p}_{t-\tau}^M$:

$$\begin{aligned} v_t^{M,c} &= \mathbb{1}^C(p_t^M, \tilde{p}_{t-\tau}^M), \\ V_t^{M,c} &= \{v_t^{M,c_1}, v_t^{M,c_2}, \dots, v_t^{M,c_k}\} = \mathcal{T}(v_t^{M,c}). \end{aligned} \quad (3)$$

Here, $\mathbb{1}^C$ is a class-wise mask generating function that returns 1 for pixels where $p_t^M = c$ and $\tilde{p}_{t-\tau}^M = c$; and 0 for otherwise pixels, for $c \in C$. Also, \mathcal{T} is the disconnection detection operator to generate instance-wise features that only works for ‘things’ classes [33], and k is the number of instance-wise masks in the corresponding class. Then, the feature centroid in the mixed domain $\mathcal{F}_t^{M,c}$ is calculated as follows:

$$\mathcal{F}_t^{M,c_k} = \frac{\sum v_t^{M,c_k} f_t^{M,c_k}}{\sum v_t^{M,c_k}}. \quad (4)$$

On the other hand, the ground truth labels for the source domain are available. Therefore we exclude features from misclassified pixels using Y^S , as in [33].

$$\begin{aligned} v_t^{S,c} &= \mathbb{1}^C(p_t^S, y_t^S), \\ V_t^{S,c} &= \{v_t^{S,c_1}, v_t^{S,c_2}, \dots, v_t^{S,c_k}\} = \mathcal{T}(v_t^{S,c}). \end{aligned} \quad (5)$$




Figure 4. **Illustration of FATC.** Feature alignment reduces the mismatch of the features from the two domains. The generated source domain features are stored in the feature bank. The loss function compares each feature in the mixed domain prediction with the closest one in the bank. The noisy features are filtered out by comparing with the ground truth label or warped previous frame prediction in the source or mixed domain, respectively.

Then, the feature centroid in the source domain $\mathcal{F}_t^{S,c}$ is calculated by the same method as the mixed domain.

$$\mathcal{F}_t^{S,c_k} = \frac{\sum v_t^{S,c_k} f_t^{S,c_k}}{\sum v_t^{S,c_k}}. \quad (6)$$

Finally, we can adopt image-level feature alignment method [33] to align the denoised source and mixed domain features in the feature space; create a feature bank that stores \mathcal{F}_t^{S,c_j} , and then align \mathcal{F}_t^{M,c_i} to features in the bank. The feature bank has a fixed size for each class, and the oldest feature is emitted whenever a new $\mathcal{F}_t^{S,c}$ is generated. The mixed domain foreground features generate the loss with the closest \mathcal{F}_t^{S,c_j} in the bank, and i and j are the indexes of each feature sample. The loss for feature alignment is L1 loss formulated as:

$$\mathcal{L}_{feature} = \sum_i \min_j \frac{\|\mathcal{F}_t^{M,c_i} - \mathcal{F}_t^{S,c_j}\|_1}{k}. \quad (7)$$

3.4. Training

The overall training objective is defined as Eq. (8), where λ controls the balance between the loss terms. \mathcal{L}_{source} and \mathcal{L}_{self} are the supervised loss defined through standard cross entropy loss for the source and mixed domain, respectively.

$$\mathcal{L} = \mathcal{L}_{source} + \lambda_M \mathcal{L}_{self} + \lambda_{feature} \mathcal{L}_{feature}. \quad (8)$$

VIPER → Cityscapes-Seq																		
	Methods	Adv. ST	road	side.	buil.	fence	light	sign	vege.	terr.	sky	pers.	car	truck	bus	mot.	bike	mIoU
Image	Source only		56.7	18.7	78.7	6.0	22.0	15.6	81.6	18.3	80.4	59.9	66.3	4.5	16.8	20.4	10.3	37.1
	AdvEnt [31]	✓	78.5	31.0	81.5	22.1	29.2	26.6	81.8	13.7	80.5	58.3	64.0	6.9	38.4	4.6	1.3	41.2
	CBST [39]	✓	48.1	20.2	84.8	12.0	20.6	19.2	83.8	18.4	84.9	59.2	71.5	3.2	38.0	23.8	37.7	41.7
	IDA [22]	✓	78.7	33.9	82.3	22.7	28.5	26.7	82.5	15.6	79.7	58.1	64.2	6.4	41.2	6.2	3.1	42.0
	CRST [38]	✓	56.0	23.1	82.1	11.6	18.7	17.2	85.5	17.5	82.3	60.8	73.6	3.6	38.9	30.5	35.0	42.4
	SVMIn [7]	✓ ✓	51.1	14.3	80.8	11.9	30.9	23.1	83.5	37.7	74.5	59.5	79.7	36.4	53.2	20.0	4.2	44.1
	CrCDA [12]	✓	78.1	33.3	82.2	21.3	29.1	26.8	82.9	28.5	80.7	59.0	73.8	16.5	41.4	7.8	2.5	44.3
	RDA [11]	✓	72.0	25.9	80.8	15.1	27.2	20.3	82.6	31.4	82.2	56.3	75.5	22.8	48.3	19.1	6.7	44.4
	FDA [35]	✓	70.3	27.7	81.3	17.6	25.8	20.0	83.7	31.3	82.9	57.1	72.2	22.4	49.0	17.2	7.5	44.4
Video	DA-VSN [8]	✓	86.8	36.7	83.5	22.9	30.2	27.7	83.6	26.7	80.3	60.0	79.1	20.3	47.2	21.2	11.4	47.8
	TPS [34]	✓	82.4	36.9	79.5	9.0	26.3	29.4	78.5	28.2	81.8	61.2	80.2	39.8	40.3	28.5	31.7	48.9
	Ours	✓	89.0	53.8	86.8	31.0	32.5	47.3	85.6	25.1	80.4	65.1	79.3	21.6	43.4	25.7	40.6	53.8

Table 1. **Quantitative comparisons on VIPER → Cityscapes-Seq.** Results are expressed in per-class IoU and mIoU scores. Our method achieved the highest IoU in most classes, and a significant update to the mIoU compared with the state-of-the-art methods. *Image*: the methods originally proposed for domain adaptive image semantic segmentation. These studies were re-experimented in a structure for video segmentation. *Video*: the methods proposed for domain adaptive video semantic segmentation. *Adv.*: Domain adaptation methods based on adversarial training. *ST*: self-training.

SYNTHIA-Seq → Cityscapes-Seq															
	Methods	Adv.	ST	road	side.	buil.	pole	light	sign	vege.	sky	pers.	rider	car	mIoU
Image	Source only			56.3	26.6	75.6	25.5	5.7	15.6	71.0	58.5	41.7	17.1	27.9	38.3
	AdvEnt [31]	✓		85.7	21.3	70.9	21.8	4.8	15.3	59.5	62.4	46.8	16.3	64.6	42.7
	CBST [39]		✓	64.1	30.5	78.2	28.9	14.3	21.3	75.8	62.6	46.9	20.2	33.9	43.3
	IDA [22]	✓		87.0	23.2	71.3	22.1	4.1	14.9	58.8	67.5	45.2	17.0	73.4	44.0
	CRST [38]		✓	70.4	31.4	79.1	27.6	11.5	20.7	78.0	67.2	49.5	17.1	39.6	44.7
	SVMIn [7]	✓	✓	84.9	0.5	77.9	29.6	7.4	15.0	78.6	73.2	46.9	6.2	73.8	44.9
	CrCDA [12]	✓		86.5	26.3	74.8	24.5	5.0	15.5	63.5	64.4	46.0	15.8	72.8	45.0
	RDA [11]		✓	84.7	26.4	73.9	23.8	7.1	18.6	66.7	68.0	48.6	9.3	68.8	45.1
	FDA [35]		✓	84.1	32.8	67.6	28.1	5.5	20.3	61.1	64.8	43.1	19.0	70.6	45.2
Video	DA-VSN [8]	✓		89.4	31.0	77.4	26.1	9.1	20.4	75.4	74.6	42.9	16.1	82.4	49.5
	TPS [34]		✓	91.2	53.7	74.9	24.6	17.9	39.3	68.1	59.7	57.2	20.3	84.5	53.8
	Ours		✓	90.4	39.2	82.3	30.2	16.3	29.6	83.2	84.9	59.3	19.7	84.3	56.3

Table 2. **Quantitative comparisons on SYNTHIA-Seq → Cityscapes-Seq.**

4. Experiments

4.1. Setup

Architecture. Following the previous works for domain adaptive video segmentation [8, 34], we use ACCEL [15] architecture for video semantic segmentation. The architecture consists of two segmentation branches, an optical flow estimation module, and a score fusion module. Each segmentation branch is designed with ResNet-101 [9] backbone and DeepLab [2] classification head to obtain single-frame prediction. The optical flow estimation module and score fusion module are used to combine each single-frame level prediction into video-level prediction. FlowNet [14] is used as the optical flow estimation module, and the score fusion module is a single 1×1 convolutional layer.

Datasets. Following [8] and [34], we use VIPER [23], SYNTHIA-Seq [25], and Cityscapes-Seq [4] datasets for benchmarking. Cityscapes-Seq is a benchmark set consisting of 5000 video clips captured from real streets. This set is split into 2975, 500, and 1525 clips for training, validation, and testing, respectively. Each video clip consists of 30 frames, and the ground truth label is provided only in a single frame (20th frame). VIPER is a synthetic dataset generated with the game ‘Grand Theft Auto V’. VIPER contains 134K video frames and corresponding segmentation labels. SYNTHIA-Seq is another synthetic dataset containing 8000 frames and corresponding segmentation labels.

Evaluation Protocols. Following the previous DAVSS methods [8, 34], we measure the Intersection over Union (IoU) score on the validation set of Cityscapes-Seq. IoU




Figure 5. **Qualitative result.** We compare our qualitative results with other state-of-the-art methods. The models are trained over VIPER \rightarrow Cityscapes-Seq benchmark. Since the ground truth of Cityscapes-Seq is provided only one frame per 30 frames, we show the same ground truth at time t for all frames.

score is measured for common classes of the source and target domain for evaluation. 15 common classes are selected for VIPER to Cityscapes-Seq task, and 11 common classes are selected for SYNTHIA-Seq to Cityscapes-Seq task.

Training Details. To train our network, We use an SGD optimizer with a momentum of 0.9 and weight decay of 5×10^{-4} . The initial learning rate was set to 5×10^{-4} and decreased according to the polynomial decay with a power of 0.9. We trained our network for 40K iterations. The balancing parameters in Eq. (8) are set to $\lambda_M = 1$, $\lambda_{feature} = 0.01$. The feature bank size is set to 50. Our model is initialized with the pretrained model of DA-VSN [8] and then we start to self-train. For self-training, we generate pseudo-labels offline with our proposed methods and IAST [19] policy that adjusts the threshold of the prediction score for each class and each image. The pseudo-label hyperparameters of IAST [19] are set to $\alpha = 0.2$, $\beta = 0.9$, $\gamma = 8$.

4.2. Comparison with State-of-the-art

In Tables 1 and 2, we compare our framework with the state-of-the-art methods on VIPER \rightarrow Cityscapes-Seq and SYNTHIA-Seq \rightarrow Cityscapes-Seq benchmarks. To our best knowledge, DA-VSN [8] and TPS [34] are the only works, published after peer review, that study domain adaptive video semantic segmentation on the same architectures and datasets as ours. We compare the performance of the proposed method with these baselines and with multiple domain adaptive image segmentation baselines [7, 11, 12, 22, 31, 35, 38, 39]. These baselines are based on various strategies: adversarial training [7, 12, 22, 31], self-training [7, 11, 35, 38, 39], and data augmentation [11, 35]. All methods were tested on the same architecture for video segmentation. As shown in Table 1, our framework, the model trained with CMOM and FATC, achieved mIoU of 53.81% and outperforms all baselines by a large margin. In addition, as shown in Table 2, we achieve the best perfor-

VIPER → Cityscapes-Seq				
Method	ST	CMOM	FATC	mIoU
Baseline				47.85
Only Self-Training	✓			49.89
with CMOM	✓	✓		53.72
with FATC	✓		✓	50.86
Ours	✓	✓	✓	53.81

(a) Ablation study on loss functions. *ST*: self-training

Learning time		
Method	w/o CMOM	w/ CMOM
Time (s/iter)	2.031	2.087

(b) Ablation study on learning time with CMOM.

VIPER → Cityscapes-Seq					
Mix Class	Things	Stuffs	Movable	Stationary	All
mIoU	52.07	52.86	52.54	53.28	53.81

(c) CMOM mixing contents.

VIPER → Cityscapes-Seq					
Mix ratio	0%	25%	50%	75%	100%
mIoU	50.86	53.12	53.59	53.81	40.98

(d) CMOM mixing ratio.

Method	VIPER → Cityscapes-Seq			SYNTHIA-Seq → Cityscapes-Seq		
	TPS	+Ours	Gain	TPS	+Ours	Gain
mIoU	48.9	53.9	+5.0	53.8	58.8	+5.0

(e) Complementary study on TPS [34].

Table 3. Ablation Experiments.

mance with mIoU of 56.31%, surpassing all baselines on SYNTHIA-Seq → Cityscapes-Seq scenario. Our work significantly outperforms state-of-the-art methods in both scenarios, demonstrating the effectiveness of our approach.

In addition, we present qualitative comparisons with the state-of-the-art methods in Figure 5. As shown in the qualitative results, DA-VSN [8] and TPS [34] cannot accurately predict hard-to-transfer classes, *i.e.*, signs (denoted in yellow boxes) and fences (denoted in white boxes), while our approach shows good results. Furthermore, we can observe that our results seem less noisy than the other methods.

4.3. Ablation Experiments

We present several ablation experiments to demonstrate the effectiveness of our method. In this section, we tested on VIPER → Cityscapes-Seq benchmark unless mentioned otherwise.

Computational costs of CMOM. To show the computational cost of CMOM, we measure the training time with and without CMOM. The results are summarized in Table 3b. Compared to “without CMOM”, only 2.8% ($= 2.087/2.031 - 1$) of learning time is additionally required for training with CMOM. Therefore, CMOM does not require much computation but it is effective in training.

Loss Functions. We ablate the proposed loss functions, and the results are given in Table 3a. As shown in the table, naively training with simple pseudo-labels achieves a performance improvement marginally (+2.04%). With either CMOM or FATC, we surpass the simple self-training approach. Furthermore, we achieve significant performance improvement from baseline with both CMOM and FATC by 5.96%.

Mixing Ratio. In Table 3d, we present experimental results with various mixing ratios. We conduct the experiments with mixing ratios of 0%, 25%, 50%, 75%, and 100%. Note that the mixed data at mixing ratios of 0% and 100% are the same as target and source data, respectively.

As shown in Table 3d, we empirically find that the mixing ratio of 75% achieves the best performance. In addition, the experimental results of mixing ratio 25%, 50%, and 75% are significantly improved compared to performance without mixing and achieved similar performance. This means that our approach is valid regardless of the mixing ratio.

Mixing Contents. Things (*e.g.*, person, sign) and stuff (*e.g.*, road, sidewalk) are well-known categories in panoptic segmentation [16]. However, in the context of *moving object mixing*, several classes are movable (*e.g.*, person, car) while others are stationary (*e.g.*, light, sign). In this study, we perform CMOM with things, stuff, movable, or stationary classes. The results are given in Table 3c. Interestingly, we find that mixing with all classes achieves the best performance. We conjecture that every class has movement in the video clip because the data has been collected from moving camera, so every class can be considered movable.

Complementary Study. Table 3e shows the performance of our framework built on TPS [34]. As a result of the experiment, our approach boosted the performance of TPS by +5.0 mIoU in both benchmarks. Experimental results show that our approach is not limited to a specific framework and is well generalized.

5. Conclusion

In this paper, we introduce a novel DAVSS framework. The proposed CMOM and FATC effectively address the bias problem which is caused by self-training in videos. We demonstrate the effectiveness of the proposed methods with comparison experiments and extensive ablation studies on VIPER → Cityscapes-Seq and SYNTHIA-Seq → Cityscapes-Seq benchmarks. We believe that our motivation and insight in this paper can be an essential step to addressing the problem caused by self-training in videos.

Acknowledgements. This work was supported by the KIST Institutional Program (Project No. 2E31571-22-143).

References

- [1] Siddhartha Chandra, Camille Couprie, and Iasonas Kokkinos. Deep spatio-temporal random fields for efficient video segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8915–8924, 2018.
- [2] 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 transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [3] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6830–6840, 2019.
- [4] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [5] Raghudeep Gadde, Varun Jampani, and Peter V Gehler. Semantic video cnns through representation warping. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4453–4462, 2017.
- [6] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2918–2928, 2021.
- [7] Dayan Guan, Jiaxing Huang, Shijian Lu, and Aoran Xiao. Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognition*, 112:107764, 2021.
- [8] Dayan Guan, Jiaxing Huang, Aoran Xiao, and Shijian Lu. Domain adaptive video segmentation via temporal consistency regularization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8053–8064, 2021.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [10] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016.
- [11] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Rda: Robust domain adaptation via fourier adversarial attacking. *arXiv preprint arXiv:2106.02874*, 2021.
- [12] Jiaxing Huang, Shijian Lu, Dayan Guan, and Xiaobing Zhang. Contextual-relation consistent domain adaptation for semantic segmentation. In *The European Conference on Computer Vision*, pages 705–722. Springer, 2020.
- [13] Xinyue Huo, Lingxi Xie, Hengtong Hu, Wengang Zhou, Houqiang Li, and Qi Tian. Domain-agnostic prior for transfer semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7075–7085, 2022.
- [14] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2462–2470, 2017.
- [15] Samvit Jain, Xin Wang, and Joseph E Gonzalez. Accel: A corrective fusion network for efficient semantic segmentation on video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8866–8875, 2019.
- [16] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019.
- [17] Suhyeon Lee, Junhyuk Hyun, Hongje Seong, and Euntai Kim. Unsupervised domain adaptation for semantic segmentation by content transfer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8306–8315, 2021.
- [18] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6936–6945, 2019.
- [19] Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. In *The European Conference on Computer Vision*, pages 415–430. Springer, 2020.
- [20] David Nilsson and Cristian Sminchisescu. Semantic video segmentation by gated recurrent flow propagation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6819–6828, 2018.
- [21] Viktor Olsson, Wilhelm Trane, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1369–1378, 2021.
- [22] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. *arXiv preprint arXiv:2004.07703*, 2020.
- [23] Stephan R Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2213–2222, 2017.
- [24] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016.
- [25] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3234–3243, 2016.

- [26] Evan Shelhamer, Kate Rakelly, Judy Hoffman, and Trevor Darrell. Clockwork convnets for video semantic segmentation. In *European Conference on Computer Vision*, pages 852–868. Springer, 2016.
- [27] Inkyu Shin, Kwanyong Park, Sanghyun Woo, and In So Kweon. Unsupervised domain adaptation for video semantic segmentation. *arXiv preprint arXiv:2107.11052*, 2021.
- [28] Marco Toldo, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation in semantic segmentation via orthogonal and clustered embeddings. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1358–1368, 2021.
- [29] Wilhelm Tranheden, Viktor Olsson, Juliano Pinto, and Lennart Svensson. Dacs: Domain adaptation via cross-domain mixed sampling. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 1379–1389, 2021.
- [30] Y.-H. Tsai, W.-C. Hung, S. Schuster, K. Sohn, M.-H. Yang, and M. Chandraker. Learning to adapt structured output space for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [31] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019.
- [32] Wenguan Wang, Tianfei Zhou, Fatih Porikli, David Crandall, and Luc Van Gool. A survey on deep learning technique for video segmentation. *arXiv preprint arXiv:2107.01153*, 2021.
- [33] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jinjun Xiong, Wen-mei Hwu, Thomas S Huang, and Honghui Shi. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12635–12644, 2020.
- [34] Yun Xing, Dayan Guan, Jiaying Huang, and Shijian Lu. Domain adaptive video segmentation via temporal pseudo supervision. In *European Conference on Computer Vision*, 2022.
- [35] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020.
- [36] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12414–12424, 2021.
- [37] Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2349–2358, 2017.
- [38] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5982–5991, 2019.
- [39] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision*, pages 289–305, 2018.