

Spatio-Temporal Pixel-Level Contrastive Learning-based Source-Free Domain Adaptation for Video Semantic Segmentation

Shao-Yuan Lo^{1*} Poojan Oza² Sumanth Chennupati² Alejandro Galindo² Vishal M. Patel¹

¹Johns Hopkins University ²Amazon

{sylo, vpatel136}@jhu.edu {poojanku, sumchenn, pagh}@amazon.com

Abstract

Unsupervised Domain Adaptation (UDA) of semantic segmentation transfers labeled source knowledge to an unlabeled target domain by relying on accessing both the source and target data. However, the access to source data is often restricted or infeasible in real-world scenarios. Under the source data restrictive circumstances, UDA is less practical. To address this, recent works have explored solutions under the Source-Free Domain Adaptation (SFDA) setup, which aims to adapt a source-trained model to the target domain without accessing source data. Still, existing SFDA approaches use only image-level information for adaptation, making them sub-optimal in video applications. This paper studies SFDA for Video Semantic Segmentation (VSS), where temporal information is leveraged to address video adaptation. Specifically, we propose Spatio-Temporal Pixel-Level (STPL) contrastive learning, a novel method that takes full advantage of spatio-temporal information to tackle the absence of source data better. STPL explicitly learns semantic correlations among pixels in the spatio-temporal space, providing strong self-supervision for adaptation to the unlabeled target domain. Extensive experiments show that STPL achieves state-of-the-art performance on VSS benchmarks compared to current UDA and SFDA approaches. Code is available at: <https://github.com/shaoyuanlo/STPL>

1. Introduction

The availability of large amounts of labeled data has made it possible for various deep networks to achieve remarkable performance on Image Semantic Segmentation (ISS) [2, 4, 30]. However, these deep networks often generalize poorly on target data from a new unlabeled domain that is visually distinct from the source training data. Unsupervised Domain Adaptation (UDA) attempts to mitigate this domain shift problem by using both the labeled source data and un-

*This work was mostly done when S.-Y. Lo was an intern at Amazon.




Figure 1. Comparison of VSS accuracy. Video-based UDA methods [12, 38, 49] outperform image-based UDA methods [33, 51], showing the importance of video-based strategies for the VSS task. Image-based SFDA methods [16, 39] perform lower than the UDA methods, which shows the difficulty of the more restricted SFDA setting. The proposed STPL, even with SFDA, achieves the best accuracy and locates at the top-right corner of the chart (i.e., more restriction, but higher accuracy).

labeled target data to train a model transferring the source knowledge to the target domain [11, 12, 31, 32, 38, 41]. UDA is effective but relies on the assumption that both source and target data are available during adaptation. In real-world scenarios, the access to source data is often restricted (e.g., data privacy, commercial proprietary) or infeasible (e.g., data transmission efficiency, portability). Hence, under these source data restrictive circumstances, UDA approaches are less practical.

To deal with these issues, the Source-Free Domain Adaptation (SFDA) setup, also referred to as Unsupervised Model Adaptation (UMA), has been recently introduced in the literature [6, 26, 27, 52]. SFDA aims to use a source-trained

model (i.e., a model trained on labeled source data) and adapt it to an unlabeled target domain without requiring access to the source data. More precisely, under the SFDA formulation, given a source-trained model and an unlabeled target dataset, the goal is to transfer the learned source knowledge to the target domain. In addition to alleviating data privacy or proprietary concerns, SFDA makes data transmission much more efficient. For example, a source-trained model ($\sim 0.1 - 1.0$ GB) is usually much smaller than a source dataset ($\sim 10 - 100$ GB). If one is adapting a model from a large-scale cloud center to a new edge device that has data with different domains, the source-trained model is far more portable and transmission-efficient than the source dataset.

Under SFDA, label supervision is not available. Most SFDA studies adopt pseudo-supervision or self-supervision techniques to adapt the source-trained model to the target domain [16, 39]. However, they consider only image-level information for model adaptation. In many real-world semantic segmentation applications (autonomous driving, safety surveillance, etc.), we have to deal with temporal data such as streams of images or videos. Supervised approaches that use temporal information have been successful for Video Semantic Segmentation (VSS), which predicts pixel-level semantics for each video frame [19, 22, 28, 46]. Recently, video-based UDA strategies have also been developed and yielded better performance than image-based UDA on VSS [12, 38, 49]. This motivates us to propose a novel SFDA method for VSS, leveraging temporal information to tackle the absence of source data better. In particular, we find that current image-based SFDA approaches suffer from sub-optimal performance when applied to VSS (see Figure 1). To the best of our knowledge, this is the first work to explore video-based SFDA solutions.

In this paper, we propose a novel spatio-temporal SFDA method namely Spatio-Temporal Pixel-Level (STPL) Contrastive Learning (CL), which takes full advantage of both spatial and temporal information for adapting VSS models. STPL consists of two main stages. (1) Spatio-temporal feature extraction: First, given a target video sequence input, STPL fuses the RGB and optical flow modalities to extract spatio-temporal features from the video. Meanwhile, it performs cross-frame augmentation via randomized spatial transformations to generate an augmented video sequence, then extracts augmented spatio-temporal features. (2) Pixel-level contrastive learning: Next, STPL optimizes a pixel-level contrastive loss between the original and augmented spatio-temporal feature representations. This objective enforces representations to be compact for same-class pixels across both the spatial and temporal dimensions.

With these designs, STPL explicitly learns semantic correlations among pixels in the spatio-temporal space, providing strong self-supervision for adaptation to an unlabeled target domain. Furthermore, we demonstrate that STPL is a

non-trivial unified spatio-temporal framework. Specifically, *Spatial-only CL* and *Temporal-only CL* are special cases of STPL, and STPL is better than a naïve combination of them. Extensive experiments demonstrate the superiority of STPL over various baselines, including the image-based SFDA as well as image- and video-based UDA approaches that rely on source data (see Figure 1). The key contributions of this work are summarized as follows:

- We propose a novel SFDA method for VSS. To the best of our knowledge, this is the first work to explore video-based SFDA solutions.
- We propose a novel CL method, namely STPL, which explicitly learns semantic correlations among pixels in the spatio-temporal space, providing strong self-supervision for adaptation to an unlabeled target domain.
- We conduct extensive experiments and show that STPL provides a better solution compared to the existing image-based SFDA methods as well as image- and video-based UDA methods for the given problem formulation.

2. Related work

Video semantic segmentation. VSS predicts pixel-level semantics for each video frame [10, 15, 19, 22, 25, 28], which has been considered a crucial task for video understanding [46]. VSS networks use temporal information, the inherent nature of videos, to pursue more accurate or faster segmentation. For example, FSO [22] employs the dense conditional random field as post-processing to obtain temporally consistent segmentation. NetWarp [10] uses optical flow information to transfer intermediate feature maps of adjacent frames and gains better accuracy. ACCEL [19] integrates predictions of sequential frames via an adaptive fusion mechanism. TDNet [15] extracts feature maps across different frames and merges them by an attention propagation module. ESVS [28] considers the temporal correlation during training and achieves a higher inference speed. These works rely on large densely annotated training data and are sensitive to domain shifts.

Unsupervised domain adaptation. UDA tackles domain shifts by aligning the representations of the two domains [11]. This framework has been widely investigated in ISS. Existing approaches can be categorized into two main streams: adversarial learning-based [3, 9, 41, 43, 45] and self-training-based [17, 33, 51, 53]. Recently, there are several works studying UDA for VSS [12, 38, 49]. DA-VSN [12] presents temporal consistency regularization to minimize temporal discrepancy across different domains and video frames. VAT-VST [38] extends both adversarial learning and self-training techniques to video adaptation. TPS [49] designs temporal pseudo supervision to adapt VSS models from the perspective of consistency training. These UDA approaches rely on labeled source data for adaptation, which is not practical in many real-world scenarios.




Figure 2. Overview of the proposed Spatio-Temporal Pixel-Level (STPL) contrastive learning framework. STPL consists of two main stages. (1) Spatio-temporal feature extraction: First, STPL fuses the RGB and optical flow ($o_{t-1 \rightarrow t}$) modalities to extract spatio-temporal features ($z_{(t-1,t)}, \tilde{z}_{(t-1,t)}$) from both the original and augmented video sequences (X, \tilde{X}). (2) Pixel-level contrastive learning: Next, after passing through a projection head and pseudo pixel-wise feature separation, STPL optimizes the pixel-level contrastive loss between the original and augmented spatio-temporal features (\mathcal{L}^{stpl}). For simplicity, this illustration considers a two-frame video sequence as the input.

Source-free domain adaptation. SFDA, a.k.a. UMA, aims to adapt a source-trained model to an unlabeled target domain without requiring access to the source data [6, 26, 27, 44, 52]. It has been investigated for ISS in recent years [16, 23, 24, 29, 39, 40]. SFDA-SS [29] develops a data-free knowledge distillation strategy for target domain adaptation. UR [39] reduces the uncertainty of target data predictions. HCL [16] presents the historical contrastive learning, which leverages the historical source hypothesis to compensate for the absence of source data. Edge/Feature-Mixup [24] generates mixup domain samples used for both source training and target adaptation. However, the need for modifying source training makes it inflexible, and it is expensive to be scaled to the video level. SFDA for videos is still relatively unexplored.

Contrastive learning. CL has been a successful representation learning technique [5, 13, 20, 21, 34]. The key idea is to create positive and negative sample pairs, then learn discriminative feature representations by maximizing the embedding distance among positive pairs and minimizing that among negative pairs. Recent works [1, 47] further explore pixel-to-pixel contrast for the ISS task, but they need label supervision for training.

3. Proposed method

An overview of the proposed STPL is illustrated in Figure 2. STPL is implemented by two key designs: spatio-temporal feature extraction and pixel-level CL. This section first introduces the detailed designs. Then we demonstrate that STPL is a non-trivial unified spatio-temporal framework.

3.1. Spatio-temporal feature extraction

The input is an unlabeled target video sequence $X = \{x_1, x_2, \dots, x_{t-1}, x_t\}$, where x_t is the current frame. For simplicity, let us consider $X = \{x_{t-1}, x_t\}$, i.e., a video with a current frame and a previous frame. Given X , the VSS network’s encoder E extracts feature representations for each frame: $z_{t-1} = E(x_{t-1})$ and $z_t = E(x_t)$. In addition, we employ FlowNet 2.0 [18] denoted as O , a widely used optical flow estimator, to estimate the optical flow between the previous and the current frames as: $o_{t-1 \rightarrow t} = O(x_{t-1}, x_t)$. **Spatio-temporal fusion block.** Next, we propose a spatio-temporal fusion block F to extract spatio-temporal feature representations from the previous and the current features z_{t-1} and z_t (see Figure 3 (a)). It adopts the estimated optical flow $o_{t-1 \rightarrow t}$ to warp the previous feature z_{t-1} to the propagated feature as: $z'_{t-1} = W(z_{t-1}; o_{t-1 \rightarrow t})$, where W denotes the warping operation. This feature propaga-

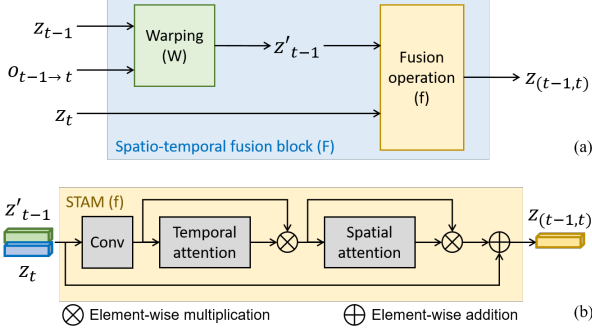


Figure 3. (a) The proposed spatio-temporal fusion block (F). (b) The proposed fusion operation (f): Spatio-Temporal Attention Module (STAM). STAM infers the attention of a spatio-temporal feature along the spatial and temporal dimensions separately, weighting important components in the spatio-temporal space. Details can be found in Appendix A1. Our fusion block is also compatible with various fusion operations.

tion aligns the pixel correspondence between the previous and the current features, which is crucial for the dense prediction task. Then a fusion operation f is used to fuse the cross-frame features into a spatio-temporal feature as: $z_{(t-1,t)} = f(z'_{t-1}, z_t)$.

The fusion operation integrates two input features into one output feature. It can be element-wise addition, concatenation, 1×1 convolution layer, an attention module, or other variants. Inspired by [48], we design a Spatio-Temporal Attention Module (STAM) illustrated in Figure 3 (b). STAM infers the attention of a spatio-temporal feature along the spatial and temporal dimensions separately, weighting important components in the spatio-temporal space. Details can be found in Appendix A1.

Cross-frame augmentation. Meanwhile, we perform cross-frame augmentation [49] that applies randomized spatial transformations T on each input frame to generate an augmented video sequence: $\tilde{X} = T(X) = \{\tilde{x}_{t-1}, \tilde{x}_t\}$. Then we apply the same spatio-temporal feature extraction process on \tilde{X} and extract the augmented spatio-temporal feature $\tilde{z}_{(t-1,t)}$. The augmentation T contains randomized Gaussian blurring and color jittering transformations.

3.2. Pixel-level contrastive learning

With the extracted original and augmented spatio-temporal features $z_{(t-1,t)}$ and $\tilde{z}_{(t-1,t)}$, we propose a new CL method to derive a semantically meaningful self-supervision. Typical CL schemes [5, 20] assume that an input contains only a single semantic category, and need a large batch size to offer sufficient positive/negative pairs for training. Nevertheless, in VSS, the input contains multiple instances, and a large batch size is computationally infeasible. Hence, we propose a method based on a pixel-level CL paradigm that leverages pixel-to-pixel contrast [1, 47], and refer to our

method as Spatio-Temporal Pixel-Level (STPL) CL.

Pseudo pixel-wise feature separation. STPL aims to acquire pixel-level representations that are similar among the same-class pixel samples but distinct among different-class pixel samples. Since we do not have target domain labels, we use our VSS model’s prediction for the input X as pseudo-label \hat{y} . Subsequently, we use \hat{y} to do pixel-wise feature separation. To maintain high-quality pseudo-labels, we set a hyperparameter of confident proportion k to control the proportion of pixels preserved as pseudo-labels. More precisely, the confident pseudo-labels \hat{y}^* are obtained by $\hat{y}^* = \text{top}k(\hat{y}; k) \subset \hat{y}$, where $\text{top}k$ is an operation that returns the k -proportion of the most confident predictions according to their probability scores.

Pixel-to-pixel contrastive loss. To perform CL, we first adopt a projection head H to project our feature representations $z_{(t-1,t)}^h = H(z_{(t-1,t)})$ and $\tilde{z}_{(t-1,t)}^h = H(\tilde{z}_{(t-1,t)})$, similar to SimCLR [5]. According to the generated confident pseudo-labels \hat{y}^* , we denote the confident pixel representation sets in $z_{(t-1,t)}^h$ and $\tilde{z}_{(t-1,t)}^h$ as $z_{(t-1,t)}^{*h} \subset z_{(t-1,t)}^h$ and $\tilde{z}_{(t-1,t)}^{*h} \subset \tilde{z}_{(t-1,t)}^h$, respectively. Next, consider a query confident pixel representation $q \in z_{(t-1,t)}^{*h}$ (i.e., q is a pixel representation in the feature $z_{(t-1,t)}^{*h}$) with a predicted pseudo-label \hat{y}_q^* , we define its positive pair set as:

$$P_q \equiv \{q^+ \in z_{(t-1,t)}^{*h} : \hat{y}_{q^+}^* = \hat{y}_q^*\}, \quad (1)$$

i.e., all the same-class pixels in the augmented feature $\tilde{z}_{(t-1,t)}^{*h}$. Then we define its negative pair set as:

$$N_q \equiv \{q^- \in \tilde{z}_{(t-1,t)}^{*h} : \hat{y}_{q^-}^* \neq \hat{y}_q^*\}, \quad (2)$$

i.e., all the different-class pixels in $\tilde{z}_{(t-1,t)}^{*h}$. We follow SupCon [20] to develop a CL scheme with multiple positive pairs. The complete formulation of the proposed STPL contrastive loss is as follows:

$$\mathcal{L}_q^{stpl} = \frac{-1}{|P_q|} \sum_{q^+ \in P_q} \log \frac{\exp(q \cdot q^+ / \tau)}{\sum_{q^- \in N_q} \exp(q \cdot q^- / \tau)}, \quad (3)$$

where τ is a temperature parameter, and the \cdot symbol denotes the inner product. Finally, the overall objective for the given video sequence input X is defined as:

$$\mathcal{L}^{stpl} = \frac{1}{|z_{(t-1,t)}^{*h}|} \sum_{q \in z_{(t-1,t)}^{*h}} \mathcal{L}_q^{stpl}. \quad (4)$$

This objective enforces the pixel representations in the original spatio-temporal features to be similar to that of the same-class pixels in the augmented features, while being distinct from that of the different-class pixels. This explicitly learns semantic correlations among pixels in the spatio-temporal space and thus can achieve better class discriminability. The proposed STPL provides a strong self-supervision for video adaptation under the SFDA setup.




Figure 4. Illustration of (a) the proposed spatio-temporal contrast \mathcal{L}^{stpl} (Eq. (3), (4)), (b) spatial-only contrast \mathcal{L}^{spa} , and (c) temporal-only contrast \mathcal{L}^{tem} .

3.3. STPL as a unified spatio-temporal framework

We further demonstrate that STPL is a non-trivial unified spatio-temporal framework. Specifically, *Spatial-only CL* and *Temporal-only CL* are special cases of STPL. Moreover, we show that a naïve combination of them is sub-optimal compared to STPL.

Spatial-only contrast. Let us turn off the fusion operation F of the STPL framework with an identity operation. Then, let us allow only the current frame feature z_t , and similarly, only the augmented current frame feature \tilde{z}_t to pass through the fusion block. After the projection head and confident filtering steps, the contrastive loss would be computed between z_t^{*h} and \tilde{z}_t^{*h} instead of the spatio-temporal features $z_{(t-1,t)}^{*h}$ and $\tilde{z}_{(t-1,t)}^{*h}$. That is, in Eq. (3) and Eq. (4), it becomes that $q \in z_t^{*h}$ and $\{q^+, q^-\} \in \tilde{z}_t^{*h}$. This computes contrast between only spatial variations and thus is a spatial-only special case of STPL. We denote this loss as \mathcal{L}^{spa} .

Temporal-only contrast. Let us consider a duplicate copy of the input video as an augmentation (i.e., $\tilde{X} = X$). Next, let us turn off the fusion operation F of STPL, allowing only the current frame feature z_t and the augmented previous frame feature \tilde{z}_{t-1} to pass through the fusion block. Here $\tilde{z}_{t-1} = z_{t-1}$ since $\tilde{X} = X$. Hence, after the projection head and confident filtering steps, the contrastive loss would be computed between z_t^{*h} and z_{t-1}^{*h} . That is, in Eq. (3) and Eq. (4), it becomes that $q \in z_t^{*h}$ and $\{q^+, q^-\} \in z_{t-1}^{*h}$. This computes contrast between only temporal variations and thus is a temporal-only special case of STPL. We denote this loss as \mathcal{L}^{tem} .

Naïve combination. To learn spatio-temporal contrast, a naïve way would be to combine the spatial-only and temporal-only contrastive losses together: $\mathcal{L}^{spa} + \mathcal{L}^{tem}$. Our experiments in Sec. 4.3 show that the naïve combination is sub-optimal compared to STPL. This demonstrates that the proposed STPL is a non-trivial unified spatio-temporal framework. Figure 4 compares the proposed spatio-temporal contrast \mathcal{L}^{stpl} , spatial-only contrast \mathcal{L}^{spa} , and temporal-only contrast \mathcal{L}^{tem} .

4. Experiments

4.1. Experimental setup

Datasets. We evaluate our method on two widely used domain adaptive VSS benchmarks: VIPER [36] \rightarrow Cityscapes-Seq [7] and SYNTHIA-Seq [37] \rightarrow Cityscapes-Seq. VIPER has 133,670 synthetic video frames with a resolution of 1080×1920 . SYNTHIA-Seq consists of 8,000 synthetic video frames with a resolution of 760×1280 . We consider VIPER and Synthia-Seq as source datasets to pre-train source models, respectively. Cityscapes-Seq is a realistic traffic scene dataset. It contains 2,975 training and 500 validation video sequences with a frame resolution of 1024×2048 . We use it as a target dataset. Following [12, 49], we resize the frames of VIPER and Cityscapes-Seq to 760×1280 and 512×1024 , respectively. For evaluations, the output predictions are interpolated to the original size.

Implementation details. Following [12, 49], we employ ACCEL [19] as our VSS network. It includes two segmentation branches, an optical flow estimation branch, and a prediction fusion layer. These branches consist of DeepLabv2 [4] architecture with ResNet-101 [14] backbone, FlowNet [8], and a 1×1 convolution layer, respectively. All the adaptation models are trained by an SGD optimizer with an initial learning rate of $2.5e^{-6}$ and a momentum of 0.9 for 20k iterations. The learning rate decreases along the polynomial decay with a power of 0.9. We set the temperature $\tau = 0.07$ and the confident proportion $k = 0.7$. The mean Intersection-over-Union (mIoU) is used as the evaluation metric. Our experiments are implemented using PyTorch [35].

4.2. Main results

Baselines. Since the proposed STPL is the first SFDA method for VSS, we compare it with multiple related domain adaptation state-of-the-art approaches described as follows. (1) *Image-based UDA*: FDA [51], PixMatch [33] and RDA [17]; (2) *Image-based SFDA*: UR [39] and HCL [16]; and (3) *Video-based UDA*: DA-VSN [12], VAT-VST [38] and TPS [49]. The image-based approaches are applied to videos by using a VSS backbone (ACCEL in our experiments), following the practice of [12, 49]. Furthermore, to fairly assess our STPL, we create the SFDA versions of these video-based UDA approaches as our (4) *Video-based SFDA* baselines. We remove all of their loss terms containing source data while keeping all the loss terms computed from only target data. We use the * symbol to denote these baselines. The results of the source-only and oracle (i.e., trained with target domain labels) models are also reported for reference. For fair comparisons, all four types of baselines use the same VSS backbone and training settings.

VIPER \rightarrow Cityscapes-Seq. Table 1 reports the evaluation results on the VIPER \rightarrow Cityscapes-Seq adaptation benchmark. The proposed STPL outperforms all four types of

Table 1. Quantitative comparisons (%) with multiple types of domain adaptation baselines on VIPER [36] → Cityscapes-Seq [7].

Method	Design	DA	road	side.	buil.	fence	light	sign	vege.	terr.	sky	pers.	car	truck	bus	mot.	bike	mIoU
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
FDA [51] (CVPR'20)	Image	UDA	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
PixMatch [33] (CVPR'21)	Image	UDA	79.4	26.1	84.6	16.6	28.7	23.0	85.0	30.1	83.7	58.6	75.8	34.2	45.7	16.6	12.4	46.7
RDA [17] (ICCV'21)	Image	UDA	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
UR [39] (CVPR'21)	Image	SFDA	84.2	20.1	80.1	11.5	30.7	31.1	82.8	22.1	69.2	59.5	81.0	4.9	52.7	36.6	8.7	45.0
HCL [16] (NeurIPS'21)	Image	SFDA	80.6	34.0	76.8	29.7	20.5	36.3	79.1	19.2	56.3	58.1	73.9	3.4	5.2	20.0	28.9	41.5
DA-VSN [12] (ICCV'21)	Video	UDA	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
VAT-VST [38] (AAAI'22)	Video	UDA	87.1	41.2	82.2	17.1	26.0	33.1	83.2	20.6	70.6	64.3	71.0	11.6	84.1	27.8	11.1	48.7
TPS [49] (ECCV'22)	Video	UDA	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
DA-VSN* [12] (ICCV'21)	Video	SFDA	77.8	32.6	79.6	29.2	37.5	34.7	82.0	22.0	64.1	61.1	76.0	6.6	32.8	32.2	11.4	45.3
VAT-VST* [38] (AAAI'22)	Video	SFDA	48.2	20.4	78.1	28.8	33.1	33.6	81.1	20.0	56.1	58.3	74.7	8.6	73.5	29.7	9.6	43.6
TPS* [49] (ECCV'22)	Video	SFDA	69.9	0.0	77.4	0.0	6.2	14.8	77.5	0.2	47.4	36.9	67.7	0.0	19.3	0.0	0.0	27.8
STPL (Ours)	Video	SFDA	83.1	38.9	81.9	48.7	32.7	37.3	84.4	23.1	64.4	62.0	82.1	20.0	76.4	40.4	12.8	52.5
Oracle	-	-	96.5	76.8	89.2	58.3	49.5	60.0	90.3	37.5	80.5	72.1	92.0	41.6	64.6	63.1	76.2	69.9

Table 2. Quantitative comparisons (%) with multiple types of domain adaptation baselines on SYNTHIA-Seq [37] → Cityscapes-Seq [7].

Method	Design	DA	road	side.	buil.	pole	light	sign	vege.	sky	pers.	rider	car	mIoU
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
FDA [51] (CVPR'20)	Image	UDA	84.1	32.8	67.6	28.1	5.5	20.3	61.1	64.8	43.1	19.0	70.6	45.2
PixMatch [33] (CVPR'21)	Image	UDA	90.2	49.9	75.1	23.1	17.4	34.2	67.1	49.9	55.8	14.0	84.3	51.0
RDA [17] (ICCV'21)	Image	UDA	84.7	26.4	73.9	23.8	7.1	18.6	66.7	68.0	48.6	9.3	68.8	45.1
UR [39] (CVPR'21)	Image	SFDA	83.5	8.0	68.1	16.5	9.9	17.7	62.4	65.1	31.9	15.3	82.3	41.9
HCL [16] (NeurIPS'21)	Image	SFDA	79.0	44.7	78.9	25.4	12.9	36.6	75.2	63.0	49.0	19.5	50.1	48.6
DA-VSN [12] (ICCV'21)	Video	UDA	89.4	31.0	77.4	26.1	9.1	20.4	75.4	74.6	42.9	16.1	82.4	49.5
VAT-VST [38] (AAAI'22)	Video	UDA	82.8	26.5	78.3	23.7	12.8	20.0	78.4	64.5	45.5	16.0	69.6	47.1
TPS [49] (ECCV'22)	Video	UDA	91.2	53.7	74.9	24.6	17.9	39.3	68.1	59.7	57.2	20.3	84.5	53.8
DA-VSN* [12] (ICCV'21)	Video	SFDA	81.0	37.9	68.4	23.7	14.0	27.5	69.8	71.3	46.4	18.7	80.2	49.0
VAT-VST* [38] (AAAI'22)	Video	SFDA	84.8	28.6	72.4	25.6	17.1	32.9	64.5	56.9	50.7	21.9	83.4	49.0
TPS* [49] (ECCV'22)	Video	SFDA	62.6	0.0	69.2	0.2	0.8	14.4	56.6	10.4	4.2	0.2	24.5	22.1
STPL (Ours)	Video	SFDA	87.6	42.5	74.6	27.7	18.5	35.9	69.0	55.5	54.5	17.5	85.9	51.8
Oracle	-	-	96.4	78.1	89.1	43.6	42.3	64.9	90.3	84.4	66.8	50.7	92.7	72.7

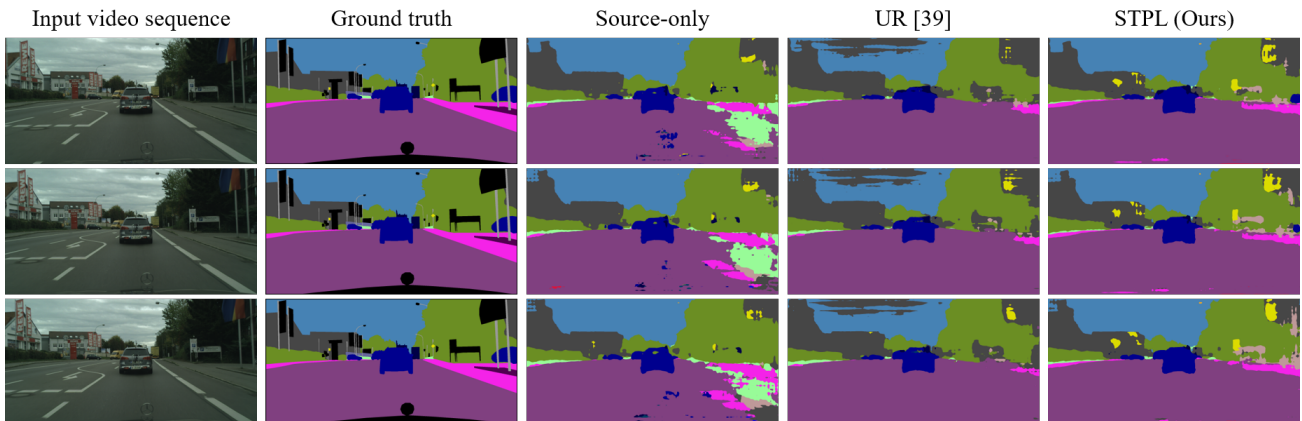


Figure 5. Qualitative results on VIPER [36] → Cityscapes-Seq [7]. The source-only model produces noisy and inconsistent predictions on the road and sidewalk. UR [39], an image-based SFDA method, suffers from inaccurate predictions on the sky and sidewalk. In contrast, the proposed STPL obtains more accurate segmentation results with high temporal consistency across the video sequence.

baselines by decent margins, which is 15.1% higher than the source-only model and 3.6% higher than the best-performing competitor. In particular, its superiority over the image-based SFDA approaches indicates the benefits of a video-based solution and demonstrates the effectiveness of our spatio-temporal strategy for videos. We can also observe that the video-based UDA approaches suffer from performance degradation when applied to SFDA. Whereas, STPL achieves better performance even compared to their UDA results relying on source data.

SYNTHIA-seq \rightarrow Cityscapes-Seq. Table 2 provides the results on the SYNTHIA-Seq \rightarrow Cityscapes-Seq benchmark. Similarly, our STPL is better than most baselines. Although TPS achieves the best accuracy under UDA, this requires accessing source data. Moreover, TPS*'s accuracy dramatically reduces to 22.1% under SFDA, showing that it is not a proper solution when source data are unavailable. Overall, these results clearly demonstrate the superiority of STPL.

Qualitative results. Figure 5 shows examples of qualitative results on VIPER [36] \rightarrow Cityscapes-Seq [7]. The source-only model produces noisy and inconsistent predictions on the road and sidewalk, showing the domain shift effect. UR, an image-based SFDA method, suffers from inaccurate sky predictions and cannot detect the whole sidewalk. In contrast, the proposed STPL obtains more accurate segmentation results with high temporal consistency across the video sequence. This indicates the importance of a video-based strategy for the VSS task and demonstrates our method's effectiveness. The qualitative and quantitative results are consistent.

4.3. Ablation analysis

Objective functions. We conduct an ablation study to validate the effectiveness of our spatio-temporal objective for adaptation. We create several variants for comparison. *Vanilla Self-training* simply computes the cross-entropy loss between predictions and pseudo-labels with a confident threshold. *Duplicate CL* computes the pixel-level contrastive loss between two identical video frames, i.e., the loss described in Sec. 3.2 but uses a duplicate copy as an augmentation and passes through the current frame features only. *Temporal-only CL*, *Spatial-only CL* and *Naïve T+S CL* are described in Sec. 3.3, whose objective functions are \mathcal{L}^{tem} , \mathcal{L}^{spa} and $\mathcal{L}^{tem} + \mathcal{L}^{spa}$, respectively.

As can be seen in Table 3, the simple Duplicate CL achieves higher accuracy than Vanilla Self-training, showing the effectiveness of the pixel-level contrastive loss. Both Temporal-only CL and Spatial-only CL make an improvement over Duplicate CL, which indicates the importance of contrasting with variations. Naïve T+S CL, a naïve combination of the temporal-only and spatial-only contrastive losses, is slightly better than either single loss. The proposed spatio-temporal objective further outperforms Naïve

Table 3. Ablation study of different objective functions on VIPER [36] \rightarrow Cityscapes-Seq [7].

Method / Objective function	mIoU
Source-only	37.1
Vanilla Self-training	45.4 (+8.3)
Duplicate CL	45.7 (+8.6)
Temporal-only CL (\mathcal{L}^{tem})	47.4 (+10.3)
Spatial-only CL (\mathcal{L}^{spa})	51.1 (+14.0)
Naïve T+S CL ($\mathcal{L}^{tem} + \mathcal{L}^{spa}$)	51.4 (+14.3)
STPL (Ours; \mathcal{L}^{stpl})	52.5 (+15.4)

Table 4. Ablation study of different fusion operations f on VIPER [36] \rightarrow Cityscapes-Seq [7].

Fusion operation	mIoU
Element-wise addition	51.4
1×1 convolution layer	51.8
Concatenation	52.3
STAM	52.5

T+S CL, showing that our design can learn more semantically meaningful context from the spatio-temporal space than simply adding the losses of two dimensions together. This demonstrates that our STPL is a non-trivial unified spatio-temporal framework for video adaptation.

Fusion operations. As discussed in Sec. 3.1, our STPL framework is compatible with various fusion operations used to extract spatio-temporal features. Here we consider and compare different fusion operations, such as element-wise addition, 1×1 convolution layer, concatenation, and the proposed STAM module. In Table 4, we can observe that STAM achieves the best performance, showing its effectiveness. On the other hand, adopting any fusion operation can outperform all the baselines in Table 1 and variants in Table 3. This demonstrates that STPL maintains superior performance regardless of the choice of fusion operations.

Feature visualization. Figure 6 provides the t-SNE visualization [42] of the feature space learned for the VIPER \rightarrow Cityscapes-Seq benchmark. For simplicity, we sample four classes (road, traffic light, car, and bicycle) to visualize. Each point in the scatter plots stands for a pixel representation. We compute the intra-class variance σ_{intra} (lower is better) and inter-class variance σ_{inter} (higher is better) of the feature space to provide a quantitative measurement. As can be seen, TPS*, which is originally designed for UDA, has a less discriminative feature space under the SFDA setup. It obtains a higher σ_{intra} and a lower σ_{inter} than the source-trained model. HCL, an image-based SFDA approach, acquires a

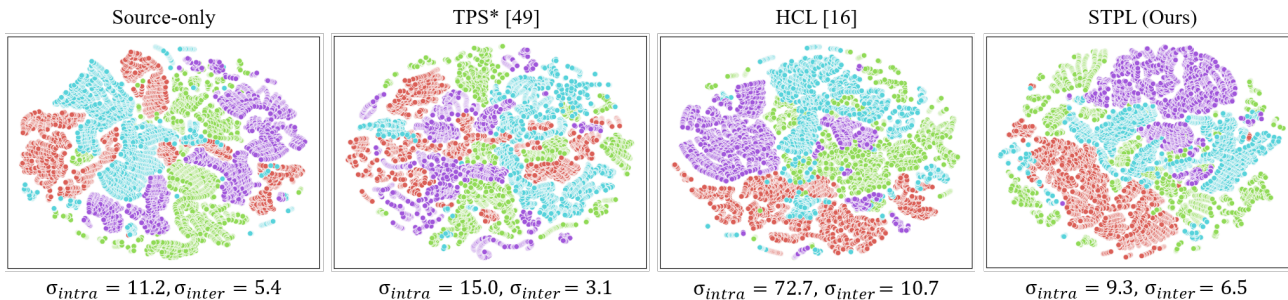


Figure 6. The t-SNE visualization [42] of the feature space learned for VIPER [36] \rightarrow Cityscapes-Seq [7], where each point in the scatter plots stands for a pixel representation. Four classes (road, traffic light, car, and bicycle) are sampled to visualize. σ_{intra} is the intra-class variance (lower is better) and σ_{inter} is the inter-class variance (higher is better) of the feature space. All the methods are evaluated on the same selected video samples. In comparison, the proposed STPL learns the most discriminative feature space, which is reflected by the lowest σ_{intra} and the high σ_{inter} .




Figure 7. The percentage of same-class pixel representations among the k -nearest neighbors in the feature space. STPL achieves higher percentage for every k value, showing that STPL learns a more discriminative and semantically consistent feature space.

higher σ_{inter} , but its σ_{intra} is much higher. In comparison, the proposed STPL learns the most discriminative feature space. Unlike HCL, STPL leverages spatio-temporal information for video adaptation, and the benefit is clearly reflected by the lowest σ_{intra} and the high σ_{inter} . This demonstrates STPL’s ability to learn semantic correlations among pixels in the spatio-temporal space.

Feature space neighborhood. This analysis inspects the neighborhood of the feature space learned by the proposed STPL, which quantitatively measures the discriminability of a feature space [50]. We randomly select several video samples and extract the features at the pixel level. For an unbiased analysis, 500 pixel representations are considered for each semantic class to create a feature analysis set. Next, we query each representation in the set and retrieve the k -nearest neighbors of that representation. Among the retrieved k nearest representations, we inspect the percentage of the same-class representations it contains.

Figure 7 reports the inspection results. For smaller k values, all the methods have similar accuracy, which indicates that their feature spaces have semantically consistent neighbors for query pixel representations. Interestingly, when we increase the k values to retrieve more neighbors, the accuracy differences between the proposed STPL and the other approaches significantly enlarge. In other words, the accuracy of STPL drops much slower than the rest. We can see that for any given k values, STPL has more semantically consistent representations in the neighborhood. This analysis shows that our method effectively learns a discriminative feature space, thereby resulting in better performance.

5. Conclusion

In this paper, we propose STPL, a novel SFDA method for VSS, which takes full advantage of spatio-temporal information to tackle the absence of source data better. STPL explicitly learns semantic correlations among pixels in the spatio-temporal space and provides strong self-supervision for video adaptation. To the best of our knowledge, this is the first work to explore video-based SFDA solutions. Moreover, we demonstrate that STPL is a non-trivial unified spatio-temporal framework. Extensive experiments show the superiority of STPL over various baselines, including the image-based SFDA as well as image- and video-based UDA approaches. Further insights into the proposed method are also provided by our comprehensive ablation analysis.

Limitations. Similar to all the existing SFDA methods, STPL assumes that the source-trained model has learned source knowledge well. A sub-optimal source-trained model would affect adaptation performance. Such limitation of SFDA is an interesting direction for future investigations.

Potential negative social impact. The proposed method may make attackers easier to adapt pre-trained open-source models to malicious uses. To avoid such risk, computer security or defense mechanisms could be incorporated.

References

- [1] Inigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *IEEE/CVF International Conference on Computer Vision*, 2021. 3, 4
- [2] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 1
- [3] Wei-Lun Chang, Hui-Po Wang, Wen-Hsiao Peng, and Wei-Chen Chiu. All about structure: Adapting structural information across domains for boosting semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [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. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 1, 5
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 2020. 3, 4
- [6] Boris Chidlovskii, Stephane Clinchant, and Gabriela Csurka. Domain adaptation in the absence of source domain data. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016. 1, 3
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 5, 6, 7, 8, 11
- [8] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. Flownet: Learning optical flow with convolutional networks. In *IEEE/CVF International Conference on Computer Vision*, 2015. 5
- [9] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. Ssf-dan: Separated semantic feature based domain adaptation network for semantic segmentation. In *IEEE/CVF International Conference on Computer Vision*, 2019. 2
- [10] Raghudeep Gadde, Varun Jampani, and Peter V Gehler. Semantic video cnns through representation warping. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, 2015. 1, 2
- [12] Dayan Guan, Jiaying Huang, Aoran Xiao, and Shijian Lu. Domain adaptive video segmentation via temporal consistency regularization. In *IEEE/CVF International Conference on Computer Vision*, 2021. 1, 2, 5, 6
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 5
- [15] Ping Hu, Fabian Caba, Oliver Wang, Zhe Lin, Stan Sclaroff, and Federico Perazzi. Temporally distributed networks for fast video semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [16] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. In *Conference on Neural Information Processing Systems*, 2021. 1, 2, 3, 5, 6
- [17] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Rda: Robust domain adaptation via fourier adversarial attacking. In *IEEE/CVF International Conference on Computer Vision*, 2021. 2, 5, 6
- [18] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017. 3
- [19] Samvit Jain, Xin Wang, and Joseph E Gonzalez. Accel: A corrective fusion network for efficient semantic segmentation on video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 2, 5
- [20] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *Conference on Neural Information Processing Systems*, 2020. 3, 4
- [21] Donghyun Kim, Yi-Hsuan Tsai, Bingbing Zhuang, Xiang Yu, Stan Sclaroff, Kate Saenko, and Manmohan Chandraker. Learning cross-modal contrastive features for video domain adaptation. In *IEEE/CVF International Conference on Computer Vision*, 2021. 3
- [22] Abhijit Kundu, Vibhav Vineet, and Vladlen Koltun. Feature space optimization for semantic video segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [23] Jogendra Nath Kundu, Akshay Kulkarni, Amit Singh, Varun Jampani, and R Venkatesh Babu. Generalize then adapt: Source-free domain adaptive semantic segmentation. In *IEEE/CVF International Conference on Computer Vision*, 2021. 3
- [24] Jogendra Nath Kundu, Akshay R Kulkarni, Suvaansh Bhambri, Deepesh Mehta, Shreyas Anand Kulkarni, Varun Jampani, and Venkatesh Babu Radhakrishnan. Balancing discriminability and transferability for source-free domain adaptation. In *International Conference on Machine Learning*, 2022. 3
- [25] Jiaying Li, Yikai Zhao, Xingjian He, Xinxin Zhu, and Jing Liu. Dynamic warping network for semantic video segmentation. In *Complexity*, 2021. 2
- [26] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 1, 3

- [27] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, 2020. 1, 3
- [28] Yifan Liu, Chunhua Shen, Changqian Yu, and Jingdong Wang. Efficient semantic video segmentation with per-frame inference. In *European Conference on Computer Vision*, 2020. 2
- [29] Yuang Liu, Wei Zhang, and Jun Wang. Source-free domain adaptation for semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 3
- [30] Shao-Yuan Lo, Hsueh-Ming Hang, Sheng-Wei Chan, and Jing-Jhih Lin. Efficient dense modules of asymmetric convolution for real-time semantic segmentation. In *ACM Multimedia Asia*, 2019. 1
- [31] Shao-Yuan Lo and Vishal M Patel. Exploring adversarially robust training for unsupervised domain adaptation. In *Asian Conference on Computer Vision*, 2022. 1
- [32] Shao-Yuan Lo, Wei Wang, Jim Thomas, Jingjing Zheng, Vishal M Patel, and Cheng-Hao Kuo. Learning feature decomposition for domain adaptive monocular depth estimation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022. 1
- [33] Luke Melas-Kyriazi and Arjun K Manrai. Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 2, 5, 6
- [34] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Conference on Neural Information Processing Systems*, 2019. 5
- [36] Stephan R Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *IEEE/CVF International Conference on Computer Vision*, 2017. 5, 6, 7, 8, 11
- [37] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 5, 6
- [38] Inkyu Shin, Kwanyong Park, Sanghyun Woo, and In So Kweon. Unsupervised domain adaptation for video semantic segmentation. In *AAAI Conference on Artificial Intelligence*, 2022. 1, 2, 5, 6
- [39] Prabhu Teja Sivaprasad and Francois Fleuret. Uncertainty reduction for model adaptation in semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 2, 3, 5, 6
- [40] Serban Stan and Mohammad Rostami. Unsupervised model adaptation for continual semantic segmentation. In *AAAI Conference on Artificial Intelligence*, 2021. 3
- [41] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2
- [42] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. In *Journal of Machine Learning Research*, 2008. 7, 8, 11
- [43] Vibashan VS, Vikram Gupta, Poojan Oza, Vishwanath A Sindagi, and Vishal M Patel. Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2
- [44] Vibashan VS, Poojan Oza, and Vishal M Patel. Instance relation graph guided source-free domain adaptive object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 3
- [45] Vibashan VS, Domenick Poster, Suya You, Shuowen Hu, and Vishal M Patel. Meta-uda: Unsupervised domain adaptive thermal object detection using meta-learning. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022. 2
- [46] 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. 2
- [47] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. In *IEEE/CVF International Conference on Computer Vision*, 2021. 3, 4
- [48] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *European Conference on Computer Vision*, 2018. 4, 11
- [49] Yun Xing, Dayan Guan, Jiaying Huang, and Shijian Lu. Domain adaptive video segmentation via temporal pseudo supervision. In *European Conference on Computer Vision*, 2022. 1, 2, 4, 5, 6
- [50] Shiqi Yang, Yaxing Wang, Joost van de Weijer, Luis Herranz, and Shangling Jui. Generalized source-free domain adaptation. In *IEEE/CVF International Conference on Computer Vision*, 2021. 8
- [51] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 1, 2, 5, 6
- [52] Hao-Wei Yeh, Baoyao Yang, Pong C Yuen, and Tatsuya Harada. Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021. 1, 3
- [53] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *IEEE/CVF International Conference on Computer Vision*, 2019. 2