

Domain Adaptive Video Segmentation via Temporal Pseudo Supervision

Yun Xing^{1*} Dayan Guan^{2*} Jiaxing Huang¹ Shijian Lu^{1†}

¹Nanyang Technological University

²Mohamed bin Zayed University of Artificial Intelligence

Abstract. Video semantic segmentation has achieved great progress under the supervision of large amounts of labelled training data. However, domain adaptive video segmentation, which can mitigate data labelling constraints by adapting from a labelled source domain toward an unlabelled target domain, is largely neglected. We design temporal pseudo supervision (TPS), a simple and effective method that explores the idea of consistency training for learning effective representations from unlabelled target videos. Unlike traditional consistency training that builds consistency in spatial space, we explore consistency training in spatiotemporal space by enforcing model consistency across augmented video frames which helps learn from more diverse target data. Specifically, we design cross-frame pseudo labelling to provide pseudo supervision from previous video frames while learning from the augmented current video frames. The cross-frame pseudo labelling encourages the network to produce high-certainty predictions, which facilitates consistency training with cross-frame augmentation effectively. Extensive experiments over multiple public datasets show that TPS is simpler to implement, much more stable to train, and achieves superior video segmentation accuracy as compared with the state-of-the-art. Code is available at <https://github.com/xing0047/TPS>.

Keywords: Video semantic segmentation, Unsupervised domain adaptation, Consistency training, Pseudo labeling

1 Introduction

Video semantic segmentation [15, 49, 12, 43, 55], which aims to predict a semantic label for each pixel in consecutive video frames, is a challenging task in computer vision research. With the advance of deep neural networks in recent years, video semantic segmentation has achieved great progress [59, 38, 17, 31, 34, 35, 24, 44] by learning from large-scale and annotated video data [4, 11]. However, the

*Equal contribution.

†Corresponding author.




Fig. 1. The proposed temporal pseudo supervision (TPS) handles domain adaptive video segmentation by introducing **Cross-frame Augmentation** and **Cross-frame Pseudo Labelling** for consistency training in target domain. Specifically, the **Cross-frame Pseudo Labelling** obtains one-hot predictions (taken as pseudo labels) for *Previous Frames* and warps the predicted pseudo labels to the current video frames to supervise the learning from the *Augmented Current Frames* that are generated by the **Cross-frame Augmentation**.

annotation in video semantic segmentation involves pixel-level dense labelling which is prohibitively time-consuming and laborious to collect and has become one major constraint in supervised video segmentation. An alternative approach is to resort to synthetic data such as those rendered by game engines where pixel-level annotations are self-generated [56, 22]. On the other hand, video segmentation models trained with such synthetic data often experience clear performance drops [19] while applied to real videos that usually have different distributions as compared with synthetic data.

Domain adaptive video segmentation aims for bridging distribution shifts across different video domains. Though domain adaptive image segmentation has been studied extensively, domain adaptive video segmentation is largely neglected despite its great values in various practical tasks. To the best of our knowledge, DA-VSN [19] is the only work that explores adversarial learning and temporal consistency regularization to minimize the inter-domain temporal discrepancy and inter-frame discrepancy in target domain. However, DA-VSN relies heavily on adversarial learning which cannot guarantee a low empirical error on unlabelled target data [37, 6, 70], leading to negative effects on temporal consistency regularization in target domain. Consistency training is a prevalent

semi-supervised learning technique that can guarantee a low empirical error on unlabelled data by enforcing model outputs to be invariant to data augmentation [68, 60, 53]. It has recently been explored in domain adaptation tasks for guaranteeing a low empirical error on unlabelled target data [1, 62, 48].

Motivated by consistency training in semi-supervised learning, we design a method named temporal pseudo supervision (TPS) that explores consistency training in spatiotemporal space for effective domain adaptive video segmentation. TPS works by enforcing model predictions to be invariant under the presence of cross-frame augmentation that is applied to the unlabelled target-domain video frames as illustrated in Fig. 1. Specifically, TPS introduces cross-frame pseudo labelling that predicts pseudo labels for previous video frames. The predicted pseudo labels are then warped to the current video frames to enforce consistency with the prediction of the augmented current frames. Meanwhile, they also provide pseudo supervision for the domain adaptation model for learning from the augmented current frames. Compared with DA-VSN involving unstable adversarial learning, TPS is simpler to implement, more stable to train and achieve superior video segmentation performance consistently across multiple public datasets.

The major contributions of this work can be summarized in three aspects. First, we introduce a domain adaptive video segmentation framework that addresses the challenge of absent target annotations from a perspective of consistency training. Second, we design an innovative consistency training method that constructs consistency in spatiotemporal space between the prediction of the augmented current video frames and the warped prediction of previous video frames. Third, we demonstrate that the proposed method achieves superior video segmentation performance consistently across multiple public datasets.

2 Related works

2.1 Video Semantic Segmentation

Video Semantic Segmentation is the challenging task of assigning a human-defined category to each pixel in each frame of a given video sequence. To tackle this challenge, the most natural and straightforward solution is to directly apply image segmentation approaches to each frame individually, in which way the model tends to ignore the temporal continuity in the video while training. A great many works explore on leveraging temporal consistency across frames by optical-flow guided feature fusion [72, 17, 34, 44], sequential network based representation aggregation [52] or joint learning of segmentation as well as optical flow estimation [32, 13].

Although video semantic segmentation has achieved great success under a supervised learning paradigm given a large amount of annotated data, pixel-wise video annotations are laborious and usually deficient to train a well-behaved network. Semi-supervised video segmentation aims at exploiting sparsely annotated video frames for segmenting unannotated frames of the same video. To

make better use of unannotated data, a stream of work investigates on learning video segmentation network under annotation efficient settings by exploiting optical flow [51, 72, 52, 13], patch matching [2, 5], motion cues [61, 73], pseudo-labeling [7], or self-supervised learning [66, 39, 33].

To further ease the burden of annotating, a popular line of study explores training segmentation network for real scene with synthetic data that can be automatically annotated, by either adversarial learning [63, 30, 65, 27, 54, 28, 19] or self-training [74, 41, 42, 9, 69, 36, 47, 26, 25, 71, 29, 48], which is known as domain adaptation. For domain adaptive video segmentation, DA-VSN [19] is the only work that addresses the problem by incorporating adversarial learning to bridge the domain gap in temporal consistency. However, DA-VSN is largely constrained by adversarial learning that is unstable during training with high empirical risk. Different from adversarial learning [23, 64, 45, 21, 67, 18], consistency training [68, 60, 48, 1] is widely explored in semi-supervised learning and domain adaptation recently with the benefits of its higher training stability and lower empirical risk. In this work, we propose to address the domain adaptive video segmentation by introducing consistency training across frames.

2.2 Consistency Training

Consistency training is a prevalent semi-supervised learning scheme that regularizes network predictions to be invariant to input perturbations [68, 60, 53, 20, 10]. It intuitively makes sense as the model is supposed to be robust to small changes on inputs. Recent studies that focus on consistency training differ in how and where to set up perturbation. A great many works introduce random perturbation by Gaussian noise [16], stochastic regularization [58, 40] or adversarial noise [50] at input level to enhance consistency training by enlarging sample space. More recently, it has been shown that stronger image augmentation [68, 3, 60] can better improve the consistency training. Conceptually, the strong augmentation on images enriches the sample space of data, which can benefit the semi-supervised learning significantly.

Aside from the effectiveness of consistency training in semi-supervised learning, a line of recent studies explore adapting the strategy in domain adaptation tasks [1, 62, 48]. SAC [1] tackles domain adaptive segmentation by ensuring consistency between predictions from different augmented views. DACS [62] performs augmentation by mixing image patches from the two domains with swapping labels and pseudo labels accordingly. Derived from FixMatch [60] which performs consistency training under the scenario of image classification, PixMatch [48] explores on various image augmentation strategies for domain adaptive image segmentation task. Unlike the aforementioned works involving consistency training in spatial space, we adopt consistency training in spatiotemporal space by enforcing model outputs invariant to cross-frame augmentation at the input level, which is devised to enrich the augmentation set and thus benefit the consistency training on unlabeled target videos.

3 Method

3.1 Background

Consistency training is a prevalent semi-supervised learning technique that enforces consistency between predictions on unlabeled images and the corresponding perturbed ones. Motivated by consistency training in semi-supervised learning, PixMatch [48] presents strong performance on domain adaptive segmentation by exploiting effective data augmentation on unlabeled target images. The idea is based on the assumption that a well-performed model should predict similarly when fed with strongly distorted inputs for unlabeled target data. Specifically, PixMatch performs pseudo labeling to provide pseudo supervision from original images for model training fed with augmented counterparts. As in FixMatch [60], the use of a hard label for consistency training in PixMatch encourage the model to obtain predictions with not only augmentation robustness but also high certainty on unlabeled data. Given a source-domain image $x^{\mathbb{S}}$ and its corresponding ground truth $y^{\mathbb{S}}$, together with an unannotated image $x^{\mathbb{T}}$ from the target domain, the training objective of PixMatch can be formulated as follows:

$$\mathcal{L}_{\text{PixMatch}} = \mathcal{L}_{ce}(\mathcal{F}(x^{\mathbb{S}}), y^{\mathbb{S}}) + \lambda_T \mathcal{L}_{ce}(\mathcal{F}(\mathcal{A}(x^{\mathbb{T}})), \mathcal{P}(\mathcal{F}(x^{\mathbb{T}}), \tau)). \quad (1)$$

where \mathcal{L}_{ce} is the cross-entropy loss, \mathcal{F} and \mathcal{A} denote the segmentation network and the transformation function for image augmentation, respectively. \mathcal{P} represents the operation that selects pseudo labels given a confidence threshold of τ . λ_T is a hyperparameter that controls the trade-off between source and target losses while training.

3.2 Temporal Pseudo Supervision

This work focus on the task of domain adaptive video segmentation. Different from PixMatch [48] that explored consistency training in spatial space for image-level domain adaptation, we propose a Temporal Pseudo Supervision (TPS) method to tackle the video-level domain adaptation by exploring spatio-temporal consistency training. Specifically, TPS introduces cross-frame augmentation for spatio-temporal consistency training to expand the diversity of image augmentation designed for spatial consistency training [48]. For the video-specific domain adaptation problem, we take adjacent frames as a whole in the form of $X_k = \mathcal{S}(x_{k-1}, x_k)$, where \mathcal{S} is a notation for stack operation.

As for cross-frame augmentation in TPS, we apply image augmentation \mathcal{A} defined in Eq. 1 on the current frames $X_k^{\mathbb{T}}$ and such process is treated as performing cross-frame augmentation \mathcal{A}^{cf} on previous frames $X_{k-\eta}^{\mathbb{T}}$, where η is referred to as propagation interval which measures the temporal distance between the previous frames and the current frames. In this way, TPS can construct consistency training in spatiotemporal space by enforcing consistency between predictions on $\mathcal{A}^{cf}(X_{k-\eta}^{\mathbb{T}})$ and $X_k^{\mathbb{T}}$, which is different from PixMatch [48] that enforces

spatial consistency between predictions on $\mathcal{A}(x^\mathbb{T})$ and $x^\mathbb{T}$ (as in Eq. 1). Formally, the cross-frame augmentation \mathcal{A}^{cf} is defined as:

$$\mathcal{A}^{cf}(X_{k-\eta}^\mathbb{T}) = \mathcal{S}(\mathcal{A}(x_{k-1}^\mathbb{T}), \mathcal{A}(x_k^\mathbb{T})). \quad (2)$$

Remark 1 *It is worth highlighting that the image augmentation \mathcal{A} plays a crucial role in consistency training by strongly perturbing inputs to construct unseen views. As for the augmentation set \mathcal{A} , there have been studies [68, 3, 60] presenting that stronger augmentation can benefit the consistency training more. To expand the diversity in image augmentation for the video task, we take the temporal deviation in video as a new kind of data augmentation for the video task and combine it with \mathcal{A} , noted as \mathcal{A}^{cf} . To validate the effectiveness of cross-frame augmentation, we empirically compare TPS (using \mathcal{A}^{cf}) with PixMatch [48] (using \mathcal{A}) in Table 1 and 2.*

With the constructed spatio-temporal space from cross-frame augmentation, TPS performs cross-frame pseudo labelling to provide pseudo supervision from previous video frames for network training fed with augmented current video frames. The cross-frame pseudo labelling has two roles: 1) facilitate the cross-frame consistency training that applies data augmentations across frames; 2) encourage the network to output video predictions with high certainty on unlabeled frames.

Given a video sequence in target domain, we first forward previous video frames $X_{k-\eta}^\mathbb{T}$ through a video segmentation network \mathcal{F} to obtain the previous frame prediction, and use FlowNet [14] to produce the optical flow $o_{k-\eta \rightarrow k}$ estimated from the previous frame $x_{k-\eta}^\mathbb{T}$ and the current frame $x_k^\mathbb{T}$. Subsequently, the obtained previous frame prediction is warped using the estimated optical flow $o_{k-\eta \rightarrow k}$ to ensure the warped prediction is in line with the current frame temporally. We then perform pseudo labeling by utilizing a confidence threshold τ to filter out warped predictions with low confidence. In a nutshell, the process of cross-frame pseudo labelling can be formulated as:

$$\mathcal{P}^{cf}(\mathcal{F}(X_{k-\eta}^\mathbb{T}), o_{k-\eta \rightarrow k}, \tau) = \mathcal{P}(\mathcal{W}(\mathcal{F}(X_{k-\eta}^\mathbb{T}), o_{k-\eta \rightarrow k}), \tau). \quad (3)$$

Remark 2 *we would like to note that the confidence threshold τ is set to pick out high-confident predictions as pseudo labels for consistency training. There exist hard-to-transfer classes in the domain adaptive segmentation task (e.g. light, sign and rider in SYNTHIA-Seq \rightarrow Cityscapes-Seq) that tend to produce low confidence scores as compared to dominant classes, thus more possibly being ignored in pseudo labelling. To retain the pseudo label of hard-to-transfer classes as much as possible, we take 0 as the threshold τ for our experiments and further discussion about the effect of τ in Table 3.*

The training objective of TPS resembles Eq. 1 in both source and target domain except that: 1) instead of feeding single images to the model, TPS takes adjacent video frames as inputs for video segmentation; 2) TPS replaces \mathcal{A} in Eq. 1 with a more diverse version \mathcal{A}^{cf} to enrich the augmentation set by incorporating cross-frame augmentation; 3) in lieu of the straightforward pseudo

labeling in Eq. 1, TPS resorts to cross-frame pseudo labeling that propagates video prediction from previous frames across optical flow $o_{k-\eta \rightarrow k}$ before further step. In a nutshell, given source-domain video frames X^S along with the target-domain video sequence, we formulate our TPS as:

$$\mathcal{L}_{\text{TPS}} = \mathcal{L}_{ce}(\mathcal{F}(X^S), y^S) + \lambda_T \mathcal{L}_{ce}(\mathcal{F}(\mathcal{A}^{cf}(X_{k-\eta}^T)), \mathcal{P}^{cf}(\mathcal{F}(X_{k-\eta}^T), o_{k-\eta \rightarrow k}, \tau)). \quad (4)$$

Remark 3 We should point out that λ_T is set to balance the training between source and target domain as in DA-VSN. In spite of the effectiveness of DA-VSN on domain adaptive video segmentation task, the training process of adversarial learning is inherently unstable with feeding complex or irrelevant cues to the discriminator while training [45]. To alleviate the effect, DA-VSN set λ_T to 0.001 to stabilize the training process whereas compromise the domain adaptation performance. Different from the previous work, we leverage the inherent stability of consistency training and naturally set λ_T to 1.0 for our TPS to treat learning of source and target equally. We further make comparison on the stability of training process between DA-VSN and TPS by visualization in Fig. 3 and explore on the effect of λ_T on the performance in Table 5.

4 Experiments

4.1 Experimental Setting

Datasets. To validate our method, we conduct comprehensive experiments under two challenging synthetic-to-real benchmarks for domain adaptive video segmentation: SYNTHIA-Seq [57] \rightarrow Cityscapes-Seq [11] and VIPER [56] \rightarrow Cityscapes-Seq. As in [19], we treat either SYNTHIA-Seq or VIPER as source-domain data and take Cityscapes-Seq as the target-domain data.

Implementation details. As in [19], we take ACCEL [34] as the video segmentation framework, which is composed of double segmentation branches and an optical flow estimation branch, together with a fusion layer at the output level. Specifically, both branches for segmentation forward a single video frame through Deeplab [8]. Meanwhile, the branch of optical flow estimation [14] produces the corresponding optical flow of the adjacent video frames, which can be further used in a score fusion layer to integrate frame prediction from different views. As regard to the training process, we use SGD as the optimizer with momentum and weight decay set to 0.9 and 5×10^{-4} respectively. The model is trained with a learning rate of 2.5×10^{-4} for 40k iterations. As in [60, 48], we incorporate multiple augmentations in our experiments, including gaussian blur, color jitter and random scaling. The mean intersection-over-union (mIoU) is used to evaluate all methods. For the efficiency of training and inference, we apply bicubic interpolation to resize every video frame in Cityscapes-Seq and VIPER to 512×1024 , 720×1280 , respectively. All the experiments are implemented on a single GPU with 11 GB memory.

Table 1. Quantitative comparisons over the benchmark of SYNTHIA-Seq \rightarrow Cityscapes-Seq: TPS outperforms multiple domain adaptation methods by large margins. These methods include the only domain adaptive video segmentation method [19], the most related domain adaptive segmentation method [48] and other domain adaptive segmentation approaches [65, 75, 54, 74, 30, 27, 69] which serve as baselines. Note that “Source only” denotes the network trained with source-domain data solely

SYNTHIA-Seq \rightarrow Cityscapes-Seq												
Methods	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
AdvEnt [65]	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 [75]	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 [54]	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 [74]	70.4	31.4	79.1	27.6	11.5	20.7	78.0	67.2	49.5	17.1	39.6	44.7
CrCDA [30]	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 [27]	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 [69]	84.1	32.8	67.6	28.1	5.5	20.3	61.1	64.8	43.1	19.0	70.6	45.2
DA-VSN [19]	89.4	31.0	77.4	26.1	9.1	20.4	75.4	74.6	42.9	16.1	82.4	49.5
PixMatch [48]	90.2	49.9	75.1	23.1	17.4	34.2	67.1	49.9	55.8	14.0	84.3	51.0
TPS (Ours)	91.2	53.7	74.9	24.6	17.9	39.3	68.1	59.7	57.2	20.3	84.5	53.8

4.2 Comparison with State-of-the-art

We compare the proposed TPS mainly with the most related methods DA-VSN [19] and PixMatch [48], considering the fact that DA-VSN is current state-of-the-art method on domain adaptive video segmentation (the same task as in this work) and PixMatch is the state-of-the-art method on domain adaptive image segmentation using consistency training (the same learning scheme as in this work). Quantitative comparisons are shown in Table 1 and 2. We note that TPS surpasses DA-VSN by a large margin on the benchmark of both SYNTHIA-Seq \rightarrow Cityscapes-Seq (4.3% in mIoU) and VIPER \rightarrow Cityscapes-Seq (1.1% in mIoU), which presents the superiority of consistency training over adversarial learning for domain adaptive video segmentation. Additionally, we highlight that our method TPS outperforms PixMatch on both benchmarks (a mIoU of 2.8% and 2.2%, respectively) which corroborates the effectiveness of the cross-frame augmentation for consistency training on video-specific task. In addition, we also compare our method with multiple baselines [65, 75, 54, 74, 30, 27, 69] which were originally devised for domain adaptive image segmentation. These baselines are based on adversarial learning [65, 54, 30] and self-training [75, 74, 69, 27]. As in [19], We apply these approaches by simply replacing the image segmentation model with our video segmentation backbone and implement domain adaptation similarly. As presented in Table 4.1 and 4.2, TPS surpasses all baselines by large margins, demonstrating the advantage of our video-specific approach as compared to image-specific ones.

Table 2. Quantitative comparisons over the benchmark of VIPER \rightarrow Cityscapes-Seq: TPS outperforms multiple domain adaptation methods by large margins

VIPER \rightarrow Cityscapes-Seq																
Methods	road	side.	buil.	fencelight	sign	vege.	terr.	sky	pers.	car	truckbus	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
AdvEnt [65]	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 [75]	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 [54]	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 [74]	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
CrCDA [30]	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 [27]	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 [69]	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 [48]	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
DA-VSN [19]	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 (Ours)	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

Furthermore, we present the qualitative result in Fig. 2 to demonstrate the superiority of our method. We point out that despite the impressive adaptation performance of DA-VSN and PixMatch, both approaches are inferior in video segmentation as compared to TPS. As regard to DA-VSN, in spite of its excellence in retaining temporal consistency, the learnt network using DA-VSN produces less accurate segmentation (e.g. sidewalk in Fig. 2). Such outcome demonstrates the superiority of consistency training over adversarial learning in minimizing empirical error. As for PixMatch, we notice that the performance of learnt network with PixMatch is unsatisfying on retaining temporal consistency, which corroborates the necessity of introducing cross-frame augmentation in consistency training. Based on the observation of qualitative results, we conclude that TPS performs better in either keeping temporal consistency or producing accurate segmentation, which is in accordance with the quantitative result in Table 1.

4.3 Ablation Studies

We perform extensive ablation studies to better understand why TPS can achieve superior performance on video adaptive semantic segmentation. All the ablation studies are performed on the benchmark of SynthiaSeq \rightarrow Cityscapes, where TPS achieves a mIoU of 53.8% under the default setting. We present complete ablation results and concrete analysis, including the propagation interval η in Eq. 2 the confidence threshold τ in Eq. 3, and the balancing parameter λ_T in Eq. 4.

Propagation Interval. The propagation interval η in Eq. 2 represents temporal variance between previous and current frames in cross-frame augmentation. We note that increasing propagation interval η will expand temporal variance




Fig. 2. Qualitative comparison of TPS with the state-of-the-art over domain adaptive video segmentation benchmark “SYNTHIA-Seq \rightarrow Cityscapes-Seq”: TPS produces much more accurate segmentation as compared to “source only”, indicating the effectiveness of our approach on addressing domain adaptation issue. Moreover, TPS generates better segmentation than PixMatch and DA-VSN as shown in rows 4-5, which is consistent with our quantitative result. Best viewed in color.

and thus enrich cross-frame augmentation. We present our result of the ablation study on propagation interval in Table 3. Despite all results surpassing current methods in Table 1, we note that the network suffers from a performance drop while increasing propagation interval, especially on the segmentation of small objects, which can be ascribed to the increased warping error caused by propagating video prediction with optical flow.

Confidence Threshold. The confidence threshold τ in Eq. 3 is closely related to the quality of the produced pseudo labels. A common solution is to set a confidence threshold $\tau \in (0, 1)$ to filter out the low-confident predictions while

Table 3. Results of TPS with different propagation interval η : TPS achieves the best performance when $\eta = 1$. For the classes of small objects (*e.g.*, pole, light, sign, person and rider), the performance may suffer from warping error while increasing η

SYNTHIA-Seq \rightarrow Cityscapes-Seq												
η	road	side.	buil.	pole	light	sign	vege.	sky	pers.	rider	car	mIoU
3	88.9	49.5	75.4	23.4	14.1	31.6	73.5	61.0	54.3	15.2	82.2	51.7
2	91.2	52.1	74.9	19.2	14.2	31.7	71.1	61.6	55.9	19.0	84.5	52.3
1	91.2	53.7	74.9	24.6	17.9	39.3	68.1	59.7	57.2	20.3	84.5	53.8

pseudo labelling whereas retains high-confident ones. Despite its potential effectiveness in retaining the quality of pseudo labels, the consistency training in TPS tends to suffer from the inherent class-imbalance distribution in a real-world dataset (target domain), which prevents the network to produce high confidence scores for some hard-to-transfer classes. To explore the effect of the threshold τ on the performance of TPS, we perform relevant experiments and present our results in Table 4. We note that the best result is obtained when τ is set to 0. We highlight that the segmentation on hard-to-transfer classes in our task (*e.g.* pole, light, sign and rider) suffers from performance drops as expected while confidence threshold τ is adopted when pseudo labeling.

Table 4. Results of TPS with different confidence threshold τ : The best result is obtained when $\tau = 0$. It can be noticed that the hard-to-transfer classes (*e.g.*, pole, light, sign, rider) experience performance drop while setting $\tau > 0$ to filter out low-confident predictions when pseudo labeling

SYNTHIA-Seq \rightarrow Cityscapes-Seq												
τ	road	side.	buil.	pole	light	sign	vege.	sky	pers.	rider	car	mIoU
0.50	91.1	54.0	76.5	23.7	14.1	34.5	71.7	59.7	56.4	18.5	84.3	53.1
0.25	88.1	48.1	77.2	21.2	16.2	38.5	74.1	64.1	57.6	17.4	86.0	53.5
0.00	91.2	53.7	74.9	24.6	17.9	39.3	68.1	59.7	57.2	20.3	84.5	53.8

Table 5. Parameter analysis on the balancing weight λ_T . We observe that either prioritizing training process on source or target domain degrades the segmentation performance

SYNTHIA-Seq \rightarrow Cityscapes-Seq						
λ_T	0.1	0.2	0.5	1.0	1.5	2.0
TPS (Ours)	50.0	51.2	52.6	53.8	53.4	53.3

Balancing Weight. The balancing weight λ_T in Eq. 4 contributes to our solution by balancing training process between source and target domain nicely. Both supervised learning in source domain with dense annotations and consistency training in target domain should be taken good care of. We present our result of ablation study on λ_T in Table 5. As presented in Table 5, the best result is retrieved while λ_T is set to 1.0. We can observe that all results of various λ_T surpass the result of previous work DA-VSN (achieved a mIoU of 49.5 in Table 1) on the benchmark of SYNTHIA-Seq \rightarrow Cityscapes-Seq, which demonstrates the superiority of consistency training in TPS.




Fig. 3. Target losses from TPS and DA-VSN for two domain adaptation benchmarks: (a) SYNTHIA-Seq \rightarrow Cityscapes-Seq and (b) VIPER \rightarrow Cityscapes-Seq. We point out that the degradation of target loss in TPS is more stable than that in DA-VSN for both two benchmarks. Best viewed in color.

4.4 Discussion

Training stability. To compare the training stability of DA-VSN with TPS on two benchmarks, we visualize the target-domain training processes of both DA-VSN and TPS by calculating the target losses for every 20 iterations. As illustrated in Fig. 3, the decay of target loss with TPS is much less noisy than in DA-VSN, along with lower empirical error on average in target domain on both benchmarks, indicating the effectiveness of consistency training on the domain adaptive video segmentation task. In contrast, the target loss in DA-VSN degrades more unsteadily and harder to converge due to the adversarial learning module in DA-VSN, and such negative effect is stronger under the scenario of SYNTHIA-Seq \rightarrow Cityscapes-Seq. The performance differences between benchmarks can be explained by the fact that SYNTHIA-Seq has larger domain gap with Cityscapes-Seq than VIPER, and we also point out that the notable advance on the benchmark of SYNTHIA-Seq \rightarrow Cityscapes-Seq brought by TPS further demonstrates the superiority of consistency training over adversarial learning




Fig. 4. Visualization of temporal feature representations in the target domain via t-SNE [46] (different colors represent different categories): the proposed TPS surpasses Source Only, PixMatch [48] and DA-VSN [19] clearly with higher inter-class variance and lower intra-class variance. Note that we obtain the temporal features by stacking features extracted from two consecutive frames as in [19], and perform PCA with whitening on the obtained temporal features to retrieve principal components with unit component-wise variances. The visualization is based on the domain adaptive video segmentation benchmark SYNTHIA-Seq \rightarrow Cityscapes-Seq. Best viewed in color.

approach on bridging larger domain gap between different video distribution. This merit is important for real-world applications, since real scenarios could be very different from pre-built synthetic environment.

Feature Visualization. To delve deeper and investigate on the effectiveness of TPS, we visualize the target-domain video representation with t-SNE [46] presented in Fig. 4, together with visualization for source only, PixMatch and DA-VSN for comparison. We observe that TPS outperforms source-only training by a large margin, which reveals the outstanding adaptation performance of our consistency-training-based approach. Furthermore, we also spot that TPS surpasses the previous works on domain adaptive video segmentation task by achieving largest inter-class variance while keeping smallest intra-class variance, which is a proper indicator that the upstream class-wise representation from TPS are more distinguishable.

Table 6. Complementary Study on TPS: the proposed TPS can be easily integrated with the state-of-the-art work DA-VSN [19] with a clear performance gain over two challenging domain adaptation benchmarks for video segmentation

	SYNTHIA-Seq \rightarrow Cityscapes-Seq			VIPER \rightarrow Cityscapes-Seq		
Method	Base	+TPS	Gain	Base	+TPS	Gain
DA-VSN	49.5	55.1	+5.6	47.8	50.2	+2.4

Complementary Study. We further conduct experiments to explore if TPS complements the domain adaptive video segmentation network DA-VSN [19] by performing additional cross-frame consistency training on target-domain data. The results of our complementary study are summarized in Table 6. It can be observed that the integration of TPS improves the performance of DA-VSN by a large margin over two benchmarks, indicating that consistency training in TPS complements the adversarial learning in DA-VSN productively. Moreover, TPS complements with DA-VSN [19] by surpassing “TPS only” (achieved a mIoU of 53.8 and 48.9 in Table 1 and 2 respectively), which proves that the effects of adversarial learning and consistency training on the domain adaptive video segmentation task are orthogonal.

5 Conclusion

This paper proposes a temporal pseudo supervision method that introduces cross-frame augmentation and cross-frame pseudo labeling to address domain adaptive video segmentation from the perspective of consistency training. Specifically, cross-frame augmentation is designed to expand the diversity of image augmentation in traditional consistency training and thus effectively exploit unlabeled target videos. To facilitate consistency training with cross-frame augmentation, cross-frame pseudo labelling provides pseudo supervision from previous video frames for network training fed with augmented current video frames, where the introduction of pseudo labeling encourages the network to output video predictions with high certainty. Comprehensive experiments demonstrate the effectiveness of our method in domain adaption for video segmentation. In the future, we will investigate how the idea of temporal pseudo supervision perform in other video-specific tasks with unlabeled data, such as semi-supervised video segmentation and domain adaptive action recognition.

Acknowledgement

This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from Singapore Telecommunications Limited (Singtel), through Singtel Cognitive and Artificial Intelligence Lab for Enterprises.

References

1. Araslanov, N., Roth, S.: Self-supervised augmentation consistency for adapting semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15384–15394 (2021)
2. Badrinarayanan, V., Galasso, F., Cipolla, R.: Label propagation in video sequences. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 3265–3272. IEEE (2010)
3. Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. arXiv preprint arXiv:1911.09785 (2019)
4. Brostow, G.J., Shotton, J., Fauqueur, J., Cipolla, R.: Segmentation and recognition using structure from motion point clouds. In: European conference on computer vision. pp. 44–57. Springer (2008)
5. Budvytis, I., Sauer, P., Roddick, T., Breen, K., Cipolla, R.: Large scale labelled video data augmentation for semantic segmentation in driving scenarios. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 230–237 (2017)
6. Chen, C., Xie, W., Huang, W., Rong, Y., Ding, X., Huang, Y., Xu, T., Huang, J.: Progressive feature alignment for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 627–636 (2019)
7. Chen, L.C., Lopes, R.G., Cheng, B., Collins, M.D., Cubuk, E.D., Zoph, B., Adam, H., Shlens, J.: Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In: European Conference on Computer Vision. pp. 695–714. Springer (2020)
8. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: 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)
9. Chen, M., Xue, H., Cai, D.: Domain adaptation for semantic segmentation with maximum squares loss. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2090–2099 (2019)
10. Chen, X., Yuan, Y., Zeng, G., Wang, J.: Semi-supervised semantic segmentation with cross pseudo supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2613–2622 (2021)
11. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
12. Couprie, C., Farabet, C., LeCun, Y., Najman, L.: Causal graph-based video segmentation. In: 2013 IEEE International Conference on Image Processing. pp. 4249–4253. IEEE (2013)
13. Ding, M., Wang, Z., Zhou, B., Shi, J., Lu, Z., Luo, P.: Every frame counts: joint learning of video segmentation and optical flow. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 10713–10720 (2020)
14. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2758–2766 (2015)

15. Floros, G., Leibe, B.: Joint 2d-3d temporally consistent semantic segmentation of street scenes. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 2823–2830. IEEE (2012)
16. French, G., Mackiewicz, M., Fisher, M.: Self-ensembling for visual domain adaptation. arXiv preprint arXiv:1706.05208 (2017)
17. Gadde, R., Jampani, V., Gehler, P.V.: Semantic video cnns through representation warping. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (Oct 2017)
18. Guan, D., Huang, J., Lu, S., Xiao, A.: Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognition* **112**, 107764 (2021)
19. Guan, D., Huang, J., Xiao, A., Lu, S.: Domain adaptive video segmentation via temporal consistency regularization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8053–8064 (2021)
20. Guan, D., Huang, J., Xiao, A., Lu, S.: Unbiased subclass regularization for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9968–9978 (2022)
21. Guan, D., Huang, J., Xiao, A., Lu, S., Cao, Y.: Uncertainty-aware unsupervised domain adaptation in object detection. *IEEE Transactions on Multimedia* (2021)
22. Hernandez-Juarez, D., Schneider, L., Espinosa, A., Vázquez, D., López, A.M., Franke, U., Pollefeys, M., Moure, J.C.: Slanted stixels: Representing san francisco’s steepest streets. arXiv preprint arXiv:1707.05397 (2017)
23. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. arXiv preprint arXiv:1711.03213 (2017)
24. Hu, P., Caba, F., Wang, O., Lin, Z., Sclaroff, S., Perazzi, F.: Temporally distributed networks for fast video semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8818–8827 (2020)
25. Huang, J., Guan, D., Xiao, A., Lu, S.: Cross-view regularization for domain adaptive panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10133–10144 (2021)
26. Huang, J., Guan, D., Xiao, A., Lu, S.: Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. *Advances in Neural Information Processing Systems* **34**, 3635–3649 (2021)
27. Huang, J., Guan, D., Xiao, A., Lu, S.: Rda: Robust domain adaptation via fourier adversarial attacking. arXiv preprint arXiv:2106.02874 (2021)
28. Huang, J., Guan, D., Xiao, A., Lu, S.: Multi-level adversarial network for domain adaptive semantic segmentation. *Pattern Recognition* **123**, 108384 (2022)
29. Huang, J., Guan, D., Xiao, A., Lu, S., Shao, L.: Category contrast for unsupervised domain adaptation in visual tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1203–1214 (2022)
30. Huang, J., Lu, S., Guan, D., Zhang, X.: Contextual-relation consistent domain adaptation for semantic segmentation. In: European conference on computer vision. pp. 705–722. Springer (2020)
31. Huang, P.Y., Hsu, W.T., Chiu, C.Y., Wu, T.F., Sun, M.: Efficient uncertainty estimation for semantic segmentation in videos. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 520–535 (2018)
32. Hur, J., Roth, S.: Joint optical flow and temporally consistent semantic segmentation. In: European Conference on Computer Vision. pp. 163–177. Springer (2016)
33. Jabri, A., Owens, A., Efros, A.A.: Space-time correspondence as a contrastive random walk. *Advances in Neural Information Processing Systems* (2020)

34. Jain, S., Wang, X., Gonzalez, J.E.: Accel: A corrective fusion network for efficient semantic segmentation on video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8866–8875 (2019)
35. Kim, D., Woo, S., Lee, J.Y., Kweon, I.S.: Video panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9859–9868 (2020)
36. Kim, M., Byun, H.: Learning texture invariant representation for domain adaptation of semantic segmentation. arXiv preprint arXiv:2003.00867 (2020)
37. Kumar, A., Sattigeri, P., Wadhawan, K., Karlinsky, L., Feris, R., Freeman, B., Wornell, G.: Co-regularized alignment for unsupervised domain adaptation. *Advances in Neural Information Processing Systems* **31** (2018)
38. Kundu, A., Vineet, V., Koltun, V.: Feature space optimization for semantic video segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3168–3175 (2016)
39. Lai, Z., Lu, E., Xie, W.: Mast: A memory-augmented self-supervised tracker. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6479–6488 (2020)
40. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. arXiv preprint arXiv:1610.02242 (2016)
41. Li, Y., Yuan, L., Vasconcelos, N.: Bidirectional learning for domain adaptation of semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6936–6945 (2019)
42. Lian, Q., Lv, F., Duan, L., Gong, B.: Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
43. Liu, B., He, X.: Multiclass semantic video segmentation with object-level active inference. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4286–4294 (2015)
44. Liu, Y., Shen, C., Yu, C., Wang, J.: Efficient semantic video segmentation with per-frame inference. In: European Conference on Computer Vision. pp. 352–368. Springer (2020)
45. Luo, Y., Liu, P., Guan, T., Yu, J., Yang, Y.: Significance-aware information bottleneck for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6778–6787 (2019)
46. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
47. Mei, K., Zhu, C., Zou, J., Zhang, S.: Instance adaptive self-training for unsupervised domain adaptation. In: Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16. pp. 415–430. Springer (2020)
48. Melas-Kyriazi, L., Manrai, A.K.: Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12435–12445 (2021)
49. Miksik, O., Munoz, D., Bagnell, J.A., Hebert, M.: Efficient temporal consistency for streaming video scene analysis. In: ICRA. pp. 133–139. IEEE (2013)
50. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **41**(8), 1979–1993 (2018)
51. Mustikovela, S.K., Yang, M.Y., Rother, C.: Can ground truth label propagation from video help semantic segmentation? In: European Conference on Computer Vision. pp. 804–820. Springer (2016)

52. Nilsson, D., Sminchisescu, C.: Semantic video segmentation by gated recurrent flow propagation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6819–6828 (2018)
53. Ouali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12674–12684 (2020)
54. Pan, F., Shin, I., Rameau, F., Lee, S., Kweon, I.S.: Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. arXiv preprint arXiv:2004.07703 (2020)
55. Patraucean, V., Handa, A., Cipolla, R.: Spatio-temporal video autoencoder with differentiable memory. arXiv preprint arXiv:1511.06309 (2015)
56. Richter, S.R., Hayder, Z., Koltun, V.: Playing for benchmarks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2213–2222 (2017)
57. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3234–3243 (2016)
58. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems* **29**, 1163–1171 (2016)
59. Shelhamer, E., Rakelly, K., Hoffman, J., Darrell, T.: Clockwork convnets for video semantic segmentation. In: European Conference on Computer Vision. pp. 852–868. Springer (2016)
60. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. arXiv preprint arXiv:2001.07685 (2020)
61. Tokmakov, P., Alahari, K., Schmid, C.: Weakly-supervised semantic segmentation using motion cues. In: European Conference on Computer Vision. pp. 388–404. Springer (2016)
62. Tranheden, W., Olsson, V., Pinto, J., Svensson, L.: Dacs: Domain adaptation via cross-domain mixed sampling. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1379–1389 (2021)
63. Tsai, Y.H., Hung, W.C., Schuster, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7472–7481 (2018)
64. Tsai, Y.H., Sohn, K., Schuster, S., Chandraker, M.: Domain adaptation for structured output via discriminative patch representations. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1456–1465 (2019)
65. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2517–2526 (2019)
66. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2566–2576 (2019)
67. Xiao, A., Huang, J., Guan, D., Zhan, F., Lu, S.: Transfer learning from synthetic to real lidar point cloud for semantic segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 2795–2803 (2022)

68. Xie, Q., Dai, Z., Hovy, E., Luong, T., Le, Q.: Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems* **33**, 6256–6268 (2020)
69. Yang, Y., Soatto, S.: Fda: Fourier domain adaptation for semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4085–4095 (2020)
70. Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., Wen, F.: Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12414–12424 (2021)
71. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision* pp. 1–15 (2021)
72. Zhu, X., Xiong, Y., Dai, J., Yuan, L., Wei, Y.: Deep feature flow for video recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2349–2358 (2017)
73. Zhu, Y., Sapra, K., Reda, F.A., Shih, K.J., Newsam, S., Tao, A., Catanzaro, B.: Improving semantic segmentation via video propagation and label relaxation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8856–8865 (2019)
74. Zou, Y., Yu, Z., Liu, X., Kumar, B., Wang, J.: Confidence regularized self-training. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5982–5991 (2019)
75. Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 289–305 (2018)