

PseudoProp: Robust Pseudo-Label Generation for Semi-Supervised Object Detection in Autonomous Driving Systems

Shu Hu^{1*}, Chun-Hao Liu^{2†}, Jayanta Dutta², Ming-Ching Chang³, Siwei Lyu¹, Naveen Ramakrishnan⁴

¹University at Buffalo, SUNY, USA ²Bosch Center for Artificial Intelligence, USA

³University at Albany, SUNY, USA ⁴Amazon, USA

{shuhu, siweilyu}@buffalo.edu, {Chun-Hao.Liu, Jayanta.Dutta}@us.bosch.com

mchang2@albany.edu, rnaveen83@gmail.com

Abstract

Semi-supervised object detection methods are widely used in autonomous driving systems, where only a fraction of objects are labeled. To propagate information from the labeled objects to the unlabeled ones, pseudo-labels for unlabeled objects must be generated. Although pseudo-labels have proven to improve the performance of semi-supervised object detection significantly, the applications of image-based methods to video frames result in numerous miss or false detections using such generated pseudo-labels. In this paper, we propose a new approach, PseudoProp, to generate robust pseudo-labels by leveraging motion continuity in video frames. Specifically, PseudoProp uses a novel bidirectional pseudo-label propagation approach to compensate for misdetection. A feature-based fusion technique is also used to suppress inference noise. Extensive experiments on the large-scale Cityscapes dataset demonstrate that our method outperforms the state-of-the-art semi-supervised object detection methods by 7.4% on mAP⁷⁵.

1. Introduction

In autonomous driving system design and development, it is common to collect multiple video sequences and only label key frames to train a deep neural network (DNN) based object detector. However, the efficacy of the detector may be limited by the size of the human-annotated dataset. Therefore, we have witnessed the *tour de force* of modern DNNs with semi-supervised learning (SSL) in the past that have been applied to autonomous driving [8, 20]. SSL uses available human-annotated data to guide the model training with unlabeled data. One dominant idea in SSL is pseudo-labeling, where pseudo-labels of unlabeled data are repeatedly generated by a pre-trained model. The model is then updated by training on a mixture of pseudo-labels and human-annotated data. Since the pre-trained model can generate highly confident pseudo-labels, SSL-based models can improve the performance of both image [27, 32] and video [8] object detection. Notwithstanding this tremendous success, pseudo-labels generated by conventional SSL-based object detection models from unlabeled data are not all reliable. Hence they cannot be directly applied to the training procedure of the detector network to improve performance [1, 3]. In particular, misdetections and false detections can easily appear in the pseudo-labels, due to the performance bottleneck of the selected pre-trained object detector. In this paper, we propose to leverage motion cues to gather useful information among sequential frames for robust pseudo-label generation. Our hypothesis is that motion continuity can effectively improve the quality of pseudo-labels for the critical task of object detection in autonomous driving systems.

*Work done while interning at Bosch Center for Artificial Intelligence

†Corresponding author

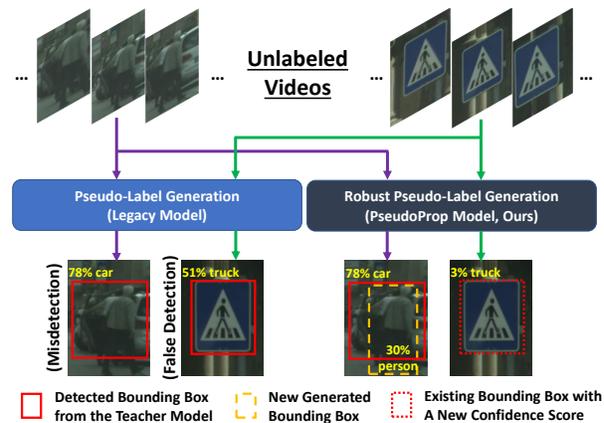


Figure 1. An illustrative example comparing the proposed PseudoProp model for robust pseudo-label generation and the legacy semi-supervised object detection, where both models leverage the teacher-student training framework. Images are from the Cityscapes dataset [9].

beled data are repeatedly generated by a pre-trained model. The model is then updated by training on a mixture of pseudo-labels and human-annotated data. Since the pre-trained model can generate highly confident pseudo-labels, SSL-based models can improve the performance of both image [27, 32] and video [8] object detection. Notwithstanding this tremendous success, pseudo-labels generated by conventional SSL-based object detection models from unlabeled data are not all reliable. Hence they cannot be directly applied to the training procedure of the detector network to improve performance [1, 3]. In particular, misdetections and false detections can easily appear in the pseudo-labels, due to the performance bottleneck of the selected pre-trained object detector. In this paper, we propose to leverage motion cues to gather useful information among sequential frames for robust pseudo-label generation. Our hypothesis is that motion continuity can effectively improve the quality of pseudo-labels for the critical task of object detection in autonomous driving systems.

Despite the idea being intuitive, motion cue is often overlooked in the design of SSL-based object detectors for au-

onomous driving. Most existing SSL-based object detection methods worked on single images individually [25, 27, 31, 32], thus the relationship among images is not considered thoroughly. A few object detection works [8, 20] leverage SSL-based methods on videos to generate pseudo-labels during the training, where the original labeled data are mostly composed of sparse video frames [9]. Each frame from the video can be viewed as an image, and then image-based SSL models can be applied for object detection. However, such methods suffer from unwanted misdetections and false detections.

In this paper, we propose a novel effective algorithmic model by leveraging motion information for robust pseudo-label generation; our model can effectively improve SSL-based object detection for autonomous driving. Our model is named **PseudoProp**, as it exploits motion as a unique property from the autonomous driving data to robustly propagate pseudo labels. Fig. 1 overviews PseudoProp and compares it against a legacy baseline. We adopt a teacher-student framework in PseudoProp, where a teacher model annotates pseudo-labels and a student model learns and benefits from the pseudo-labels.

We developed a similarity-aware weighted boxes fusion (SWBF) in PseudoProp based on a novel bidirectional pseudo-label propagation (BPLP) to make pseudo-labels more robust. BPLP can generate better pseudo-labels to ease the misdetection problems. On the other hand, BPLP might also generate too many redundant bounding boxes and inevitably introduce false positives due to the exhaustive forward and backward motion prediction. To this end, we propose an approach to reduce confidence scores of falsely transferred bounding boxes, based on the similarity between their extracted features. With this similarity check, we adapt the weighted boxes fusion (WBF) [28] originally designed for bounding boxes reduction. Fig. 2 explains how PseudoProp works in details. PseudoProp can alleviate the misdetection problem and significantly reduce the confidence scores for falsely detected objects. Experiments are performed on the large-scale Cityscapes dataset [9], which demonstrates the effectiveness of PseudoProp on generating robust pseudo-labels for image-based SSL object detection.

Note that we do not aim to develop a full-scale object tracking algorithm, where the goal is to find the motion trajectories for objects in the video. Instead, we focus on how best to associate motion objects across frames to generate pseudo-labels to improve SSL object detection. Our work is distinct from the conventional video object detection methods, which execute only one round of detection on all video frames.

The main contributions of our work can be summarized as follows:

1. We present a novel framework – PseudoProp for robust pseudo-label generation for per-image object detection

based on motion propagation and SSL.

2. The proposed SWBF method based on the BPLP approach can solve the misdetection problem and significantly reduce the confidence scores of the false positives in the generated pseudo-labels.
3. Experiments on the Cityscapes dataset demonstrate the effectiveness of our model in generating robust pseudo-labels and boosting SSL object detection performance.

2. Related Work

This section surveys relevant works of semi-supervised learning, pseudo-label generation, and video motion prediction.

2.1. SSL with Pseudo-Label Generation

Object detection is widely used in autonomous driving and video surveillance systems. Deep learning methods [10–13, 15, 24, 34, 35] have become *de facto* for object detection because of their dominant performance and scalability. However, training deep neural networks requires a large amount of annotated data, and to this end *Semi-Supervised Learning (SSL)* has growing popularity in generating or augmenting annotations for training powerful networks.

SSL has been widely applied to computer vision tasks such as object detection [25, 32] and semantic segmentation [8, 22, 33]. One important idea in this domain is pseudo-labeling [1, 3]. For object detection, the pseudo-labels are the bounding boxes of objects in unlabeled data repeatedly generated by a pre-trained model. Most SSL-based object detection methods focus on images. For example, STAC [27] and Humble teacher [32]. There are few existing works [8, 20] on object detection by leveraging SSL-based methods on videos to generate pseudo-labels on unlabeled data. In [20], the authors assume that the training videos contain only sparsely labeled bounding boxes and apply a traditional detector (Exemplar-SVM) instead of deep learning-based models for object detection. This results in lower performance for the final model. In addition, the teacher-student framework in Naive-Student [8], can be applied to object detection and semantic segmentation on videos. However, this work does not consider the relationship among frames in the same video. Therefore, the generated pseudo-labels may include many misdetections and false detections.

2.2. Video Motion

The spatially-displaced convolution network (SDC-Net) in [23] can predict future video frames based on a two-stage process of first estimating motion then predicting frames. A

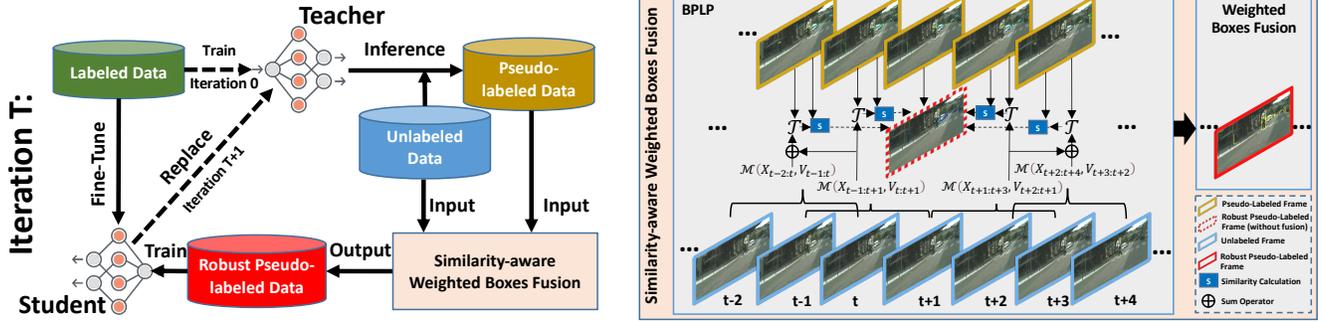


Figure 2. Overview of our proposed PseudoProp model for robust pseudo-label generation in semi-supervised object detection. (Left) Teacher-student framework for semi-supervised learning. The dash lines represent operations (‘Train’ and ‘Replace’) only working on specific iterations (Iteration 0 and Iteration T+1). (Right) Details of the proposed similarity-aware weighted boxes fusion (SWBF). The figure is better viewed in color.

follow-up work [38] is proposed to enhance semantic segmentation via ground truth label propagation, where the semantic labels are propagated through video motion built on top of the SDC-Net. The spatial-temporal algorithm of [6] achieves fast object detection on videos through motion-assisted supervised learning. In [18], the authors propose to generate pseudo-labels and adopt those to train a multi-object tracking model with unlabeled videos. Their pseudo bounding boxes are derived from hallucination of videos, where the videos are generated by motion transformations to simulate various effects. A deep neural tracker is trained with hard example mining.

There is a large literature on video object tracking [2, 4, 18]. However, a full-blown tracking algorithm is not required in this study, since generating pseudo-labels does not require creating accurate trajectories of objects in every video frame. Our method is more efficient and pertinent to the requirement of the SSL tasks.

3. Methodology

The PseudoProp model contains two parts: (1) a teacher-student framework for training semi-supervised object detector (§ 3.1), (2) Motion Prediction (§ 3.2) and (3) the *similarity-aware weighted boxes fusion* (§ 3.3). Fig. 2 illustrates detailed components of the parts.

3.1. Teacher-Student based SSL

The teacher-student framework [8, 36], which starts with the idea of knowledge distillation [14] has been widely applied in SSL. In this paper, we use a state-of-the-art teacher-student architecture from [8] for video object detection, but our method introduces motion propagation to generate robust pseudo-labels.

Given the labeled training data $\mathcal{D}_L = \{(\tilde{X}_i, \tilde{Y}_i)\}_{i=1}^n$, where n is the size of labeled data. \tilde{X}_i denotes a video frame, and \tilde{Y}_i is the corresponding human annotations (a set of bounding boxes) of \tilde{X}_i . Let $\mathcal{D}_U = \{X_i\}_{i=1}^m$ be an un-

labeled dataset, where m is the size of unlabeled data. \mathcal{D}_U is extracted from multiple video sequences with no human annotation. The human-annotated dataset \mathcal{D}_L is exploited to train a teacher network θ_1 by using loss \mathcal{L} for object detection, where \mathcal{L} consists of conventional classification and regression losses for bounding box prediction. Therefore, we have

$$\theta_1^* = \arg \min_{\theta_1} \frac{1}{n} \sum_{(\tilde{X}_i, \tilde{Y}_i) \in \mathcal{D}_L} \mathcal{L}(\tilde{Y}_i, f_{\theta_1}(\tilde{X}_i)),$$

where θ_1^* is the trained teacher network with a prediction function f .

We apply θ_1^* to generate (or update) the pseudo-labels for all unlabeled data in \mathcal{D}_U . Therefore, $Y_i = f_{\theta_1^*}(X_i)$, where Y_i is a set of pseudo-labels (bounding boxes) of the unlabeled data X_i . Next, we propose the *similarity-aware weighted boxes fusion* (SWBF) based on a motion prediction model and a noise-resistant pseudo-label fusion model, to enhance the quality of the generated pseudo-labels. This can be represented as $\bar{Y}_i = \text{SWBF}(Y_i), \forall X_i \in \mathcal{D}_U$, where \bar{Y}_i is a set of high-quality pseudo-labels after performing SWBF on Y_i .

A student network is subsequently trained with the pseudo-labeled frames using the same loss function \mathcal{L} . Thus, we have

$$\theta_2^* = \arg \min_{\theta_2} \frac{1}{m} \sum_{X_i \in \mathcal{D}_U} \mathcal{L}(\bar{Y}_i, f_{\theta_2}(X_i)).$$

Since the pseudo-labeled data are noisy, the trained student network cannot achieve high performance yet. The student network is next fine-tuned on \mathcal{D}_L before evaluated on the validation or test dataset. This way,

$$\theta_2^{**} = \arg \min_{\theta_2^*} \frac{1}{n} \sum_{(\tilde{X}_i, \tilde{Y}_i) \in \mathcal{D}_L} \mathcal{L}(\tilde{Y}_i, f_{\theta_2^*}(\tilde{X}_i)).$$

Finally, we replace the teacher $f_{\theta_1^*}$ with the student $f_{\theta_2^{**}}$ and iterate the procedure again until termination.

3.2. Motion Prediction

To estimate motion from unlabeled video frames, we adopt the SDC-Net [23] to predict the motion vector (du, dv) on each pixel (u, v) per frame X_t at time t . SDC-Net is proposed to predict a video frame X_{t+1} based on past frame observations as well as their estimated optical flow. It can be trained easily using consecutive frames without providing any manual labels. A later work [38] was proposed to improve the SDC-Net by using video frame reconstruction instead of frame prediction, *i.e.*, to apply bidirectional frames to reconstruct the current frame. The predicted frame \hat{X}_{t+1} and its corresponding predicted pseudo-labels \hat{Y}_{t+1} can be formulated as:

$$\begin{aligned}\hat{X}_{t+1} &= \mathcal{B}(\mathcal{M}(X_{t-\tau:t+1}, V_{t-\tau+1:t+1}), X_t), \\ \hat{Y}_{t+1} &= \mathcal{T}(\mathcal{M}(X_{t-\tau:t+1}, V_{t-\tau+1:t+1}), Y_t),\end{aligned}\quad (1)$$

where $X_{t-\tau:t+1}$ are frames from time $t - \tau$ to $t + 1$, $V_{t-\tau+1:t+1}$ are the corresponding optical flows from time $t - \tau + 1$ to $t + 1$, \mathcal{M} is a convolutional neural network (CNN) to predict per-pixel motion vector (du, dv) on X_t , \mathcal{B} is a bilinear sampling operation to interpolate the motion-translated frame into the predicted frame, and \mathcal{T} is a floor operation for deriving pseudo-labels from motion prediction. We adopt the pre-trained optical flow estimation model FlowNet2 [16] to generate V , and this video frame reconstruction approach is used for \mathcal{M} . We select $\tau = 1$ throughout all experiments unless specified otherwise. Once the motion vectors on all pixels are available, we use \mathcal{T} to predict (u, v) in Y_t as $(\lfloor u + du \rfloor, \lfloor v + dv \rfloor)$ in \hat{Y}_{t+1} , where $\lfloor \cdot \rfloor$ is the floor operation.

3.3. Similarity-aware Weighted Boxes Fusion (SWBF)

In this section, we first propose a bidirectional pseudo-label propagation method to generate candidate pseudo-labels according to the motion predictions. Then we propose a robust fusion method to generate final pseudo-labels.

Bidirectional Pseudo-Label Propagation (BPLP). Since the predicted pseudo-labels from the teacher model may contain many false negatives, *e.g.*, humans are misdetected as in Fig. 1, we apply the motion prediction in Eq. (1) to propagate pseudo-label prediction. However, such motion prediction can only predict frames and labels in one time step. To make the predicted pseudo-labels more robust at time $t + 1$, we propose the *bidirectional pseudo-label propagation* to generate pseudo-label proposals across frames, via interpolations from both forward propagation from existing labels and time-reversed backward propagation. We also apply different propagation lengths $k \in \mathbb{Z}^+$. Specifically,

$$\bar{Y}_{t+1} = Y_{t+1} \cup \hat{Y}_{t+1}, \quad \hat{Y}_{t+1} = \bigcup_{i \in K} \hat{Y}_{t+1}^i, \quad (2)$$

$$\hat{Y}_{t+1}^i = \mathcal{T}\left(\sum_{j \in J} \mathcal{M}(X_{t-j:t-j+2}, V_{t+1-j:o}), Y_{t+1-i}\right), \quad (3)$$

where $K = \{\pm 1, \dots, \pm(k-1), \pm k\}$, $i \in K$, as well as $J = \begin{cases} \{1, \dots, i\}, & \text{if } i > 0 \\ \{i, \dots, -1\}, & \text{if } i < 0 \end{cases}$, $o = \begin{cases} t+2-j, & \text{if } i > 0 \\ t-j, & \text{if } i < 0 \end{cases}$, where \pm indicates forward and backward propagation. In Eq. (2), Y_{t+1} is the pseudo-label set of the unlabeled frame X_{t+1} from the teacher model prediction. \hat{Y}_{t+1} is a set containing pseudo-labels from the past and future frames after using motion propagation from Eq. (3). \hat{Y}_{t+1}^i is the pseudo-label set from Y_{t+1-i} . We also show details on calculating \hat{Y}_{t+1} in the Appendix. Next, we calculate \bar{Y}_{t+1} for frame X_{t+1} by taking the union of Y_{t+1} and \hat{Y}_{t+1} .

Robust Fusion. BPLP with different settings of k can generate many candidates of pseudo-labels, which may include additional false positives (FP), which we categorize into two types. For *Type-A FP*, refer to an example in Fig. 3(a), where a person is detected at time t and $t + 2$, but not detected at $t + 1$ due to occlusion by a tree. In this case, BPLP can generate two bounding boxes at $t + 1$ however with low confidence scores caused by the occlusion. For *Type-B FP*, refer to Fig. 3(b), where a billboard is mistakenly detected as a truck at time $t + 1$ with a high confidence score. Even worse, the number of candidate pseudo-labels (bounding boxes) increases as the value of k increases. Thus, many redundant bounding boxes \bar{Y}_{t+1} can be predicted in frame X_{t+1} .

We propose a similarity-based approach to reduce the confidence scores of these false pseudo-label predictions. We define $Y_{t+1-i} := \{(L_{t+1-i}^z, P_{t+1-i}^z, S_{t+1-i}^z)\}_{z=1}^{|\hat{Y}_{t+1-i}|}$, where L_{t+1-i}^z , P_{t+1-i}^z , S_{t+1-i}^z are the class, positions, confidence score of the z -th bounding box in Y_{t+1-i} , respectively, and $|\cdot|$ denotes the number of bounding boxes in the set. Similarly, we define $\hat{Y}_{t+1}^i := \{(\hat{L}_{t+1}^{i,z}, \hat{P}_{t+1}^{i,z}, \hat{S}_{t+1}^{i,z})\}_{z=1}^{|\hat{Y}_{t+1}^i|}$. Note that $L_{t+1-i}^z = \hat{L}_{t+1}^{i,z}$, $\forall z$, because (1) we do not change the bounding box class during the propagation and (2) $\hat{P}_{t+1}^{i,z}$ (assume it is inside the frame) can be obtained from P_{t+1-i}^z by applying \mathcal{T} from Eq. (3). Recall that previously $S_{t+1-i}^z = \hat{S}_{t+1}^{i,z}$, $\forall z$, which results in Type-A FPs. Therefore, we introduce a similarity function $\text{sim}(\cdot)$ based on $\hat{P}_{t+1}^{i,z}$ and P_{t+1-i}^z to estimate the bounding box confidence score when transitioned from S_{t+1-i}^z to $\hat{S}_{t+1}^{i,z}$. To calculate this similarity, we first crop images at frame X_{t+1-i} and X_{t+1} according to the positions P_{t+1-i}^z and $\hat{P}_{t+1}^{i,z}$, respectively. Then we use a pre-trained neural network to extract the high-level feature representatives from the cropped images. The similarity is obtained by comparing these two high level feature representatives:

$$\hat{S}_{t+1}^{i,z} = S_{t+1-i}^z \cdot \text{sim}\left(\mathcal{C}(P_{t+1}^{i,z}), \mathcal{C}(P_{t+1-i}^z)\right), \quad (4)$$

where $\mathcal{C}(\cdot)$ is a function extracting the high-level feature

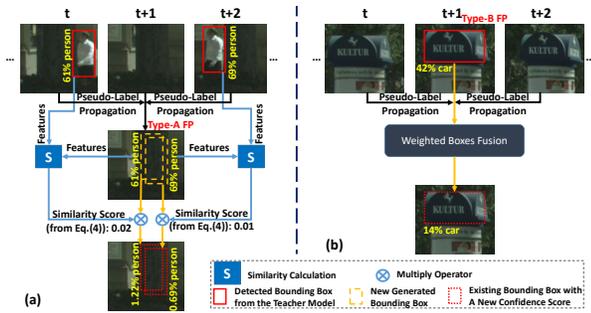


Figure 3. Examples of (a) Type-A FP and (b) Type-B FPs produced from BPLP that can be solved by the proposed SWBF method.

representatives from the cropped images based on the box positions, and $\text{sim}(\cdot)$ is a similarity function. The reason we adopt a feature-based approach for similarity calculation is that we prefer assigning similar scores to objects within the same class before and after pseudo-label propagation. The use of such similarity scores can effectively reduce Type-A FPs. Fig. 3(a) shows an illustrative example.

Although this similarity screening can reduce the confidence score for some Type-A FPs, it can not address Type-B FPs. To further filter out redundant prediction boxes, we adapt weighted boxes fusion (WBF) [28]. WBF also reduces the confidence scores of the Type-B FP boxes by averaging the localization and confidence scores of predictions from all sources (previous, current, and future frames) of the same object. Fig. 3(b) shows an illustrative example.

Before using WBF, we split \bar{Y}_{t+1} into d parts according to the bounding boxes classes, where d is the total number of classes in \bar{Y}_{t+1} . We define $\bar{Y}_{t+1,c} \subseteq \bar{Y}_{t+1}$ as a subset for the c -th class. For each subset, i.e. $\bar{Y}_{t+1,c}$, we briefly introduce the fusion procedures as follows:

1. We first divide bounding boxes from $\bar{Y}_{t+1,c}$ into different clusters. For each cluster, the intersection over union (IoU) of each two bounding boxes should be greater than a user-defined threshold Thr (in our experiments, $\text{Thr} = 0.5$ is close to an optimal threshold).
2. For boxes in each cluster r , we calculate their average confidence score C_r , and the weighted average for the positions using $C_r = \frac{1}{B} \sum_{l=1}^B C_r^l$ and $P_r = \frac{\sum_{l=1}^B C_r^l \cdot P_r^l}{\sum_{l=1}^B C_r^l}$, where B is the total number of boxes in the cluster r . C_r^l and P_r^l are the confidence score and the position of the l -th box in the cluster r , respectively.
3. The above two steps can reduce the redundant bounding boxes. However, it cannot solve the Type-B FP problems. To reduce the confidence score of false de-

tected boxes, we re-scale C_r by

$$C_r = C_r \cdot \frac{\min(B, |K| + 1)}{|K| + 1}, \quad (5)$$

where $|K|$ is the size of the set K . The interpretation is that, if a small number of sources can provide pseudo-labels on an object, this detection is most likely a false detection. An illustrative example is shown in Figure 3(b).

4. Finally, $\bar{Y}_{t+1,c}$ only contains the averaged bounding box information (c, P_r, C_r) from each cluster.

Therefore, the final \bar{Y}_{t+1} only contains the updated $\bar{Y}_{t+1,c}$ from each class. The details of this fusion method can be found in Appendix. The pseudo-code of our proposed SWBF method for \bar{Y}_{t+1} is described in Algorithm 1.

Algorithm 1: Similarity-aware Weighted Boxes Fusion (SWBF) for \bar{Y}_{t+1} .

Input: k, d, Thr, X, Y (from the teacher model prediction).
Output: The new pseudo-labels \bar{Y}_{t+1} on X_{t+1}

- 1 **for** $i \in \{\pm 1, \dots, \pm(k-1), \pm k\}$ **do**
- 2 Create or update \hat{Y}_{t+1} based on Eq.(2), Eq.(3), and Eq.(4).
- 3 **end**
- 4 Create \bar{Y}_{t+1} based on Eq.(2).
- 5 **for** $c = 1, \dots, d$ **do**
- 6 $\bar{Y}_{t+1,c} \leftarrow \text{WBF}(\bar{Y}_{t+1,c}, \text{Thr})$.
- 7 **end**
- 8 Update \bar{Y}_{t+1} based on $\bar{Y}_{t+1,c}, \forall c$.
- 9 **return** \bar{Y}_{t+1}

4. Experiments

We evaluate PseudoProp for SSL-based object detection in autonomous driving applications. PseudoProp can be applied in any video dataset of autonomous driving with sequential frames. Our experiments are performed on the popular large-scale Cityscapes dataset [9], as it fits our scenario well. Due to space limitations, we only present significant results and leave additional results in the supplementary material.

4.1. Experimental Settings

Datasets and Evaluation Metrics. The Cityscapes dataset [9] contains diverse street-views recorded from 50 cities in Germany. We use the annotated 2,975 training images as our training set and the annotated validation 500 images as our test set. Each image is selected from the 20-th

frame of a 30-frame video snippet. Therefore, the training video set contains 2,975 videos. For each training video, we estimate motion vectors and randomly select 3 frames (excluding the frames already in the training set) without replacement as the pseudo-labeled frames. To generate robust pseudo-labels for the student model, we first use the pseudo-labels of the selected frames with confidence scores > 0.4 produced from the teacher model. Note that the threshold of 0.4 yields nearly the best performance in our experiments, and is also widely used in the literature [30]. This procedure is important to suppress noisy labels as in [39]. We next apply SWBF on these noisy labels and obtain the robust pseudo-labels. We then use these frames with robust pseudo-labels to construct the pseudo-labeled sets with $1\times$, $2\times$, $3\times$ sizes of the original training set. We report the mean average precision (mAP), mAP with IoU 0.5 (mAP⁵⁰), and mAP with IoU 0.75 (mAP⁷⁵) [17] as the object detection evaluation results.

Computing Infrastructure. The algorithm in this paper is implemented with Python 3.6, and it is trained and tested on an Intel(R) Xeon(R) Gold 6150 CPU @2.70GHz with 128GB RAM, and one NVIDIA Tesla V100 GPU with 32GB VRAM. The TensorFlow version is 2.5.0 for EfficientDet-D1 (object detection). The PyTorch version is 1.7.0 for SDC-Net (motion prediction).

Teacher and Student Models. We train a deep neural network object detector as our initial teacher model. We adopt the EfficientDet-D1 [30] as the teacher network, with backbone pre-trained on ImageNet [26] and the whole network fine-tuned on Cityscapes with batch size 8. The maximum number of epochs is 180. Random image horizontal flip and scaling are applied as our data augmentation strategy. We also adopt the stochastic gradient descent optimizer and a cosine decay learning rate scheduler in the training loop. The learning rate is set to 0.08, after 1 epoch warmup with an initial learning rate 0.008. After fine-tuning, we obtained 0.355 mAP⁵⁰ performance for Cityscapes on the test set, which is close to the state-of-the-art performance [7, 19].

EfficientDet-D1 is also adopted as our student network. We train it with the pseudo-labeled data with a maximal 180 number of epochs. After that, we fine-tune the student model on the training dataset. The maximum number of epochs is also 180. Note that SWBF is a *post-processing* method to generate pseudo-labels. Any object detector producing detection boxes can be integrated with our method to take advantage of the improved pseudo-labels. Since the pseudo-labels generation is only used in the training procedure, the inference speed of PseudoProp is the same as EfficientDet-D1.

Motion Prediction and Feature Extraction. A pre-trained SDC-Net [23] is used to predict pseudo-labels according to the motion vectors presented in Eq. (1). For cal-

culating the similarity in Eq. (4) between two cropped images, we use EfficientNet-B1 [29], which is also the backbone of EfficientDet-D1 for feature extraction. Cosine similarity is used to calculate the value of $\text{sim}(\cdot)$. Feature values are normalized into $[0, 1]$ to ensure $\text{sim}(\cdot) \in [0, 1]$. Other similarity functions can also be applied.

Comparisons. We compare PseudoProp with two existing SSL object detectors based on pseudo-label generation.

- **Naive-Student** [8]: The original Naive-Student model uses a teacher-student model with a test-time augmentation. However, their test-time augmentation does not fit directly to object detection. We also use the same selected frames and pseudo-labels with confidence scores higher than 0.4 to construct the pseudo-labeled datasets. Naive-Student can be viewed as a variant of PseudoProp without SWBF.
- **VideoProp** [38]: It was originally proposed to improve the semantic segmentation. Here we use it to increase the pseudo-labeled data size via their label propagation. Since this method can only generate pseudo-labels according to the ground truth (GT), we follow the approach in their original paper by only considering the 19-th, 20-th, 21-th, and 22-th frames in each video. Specifically, for each training video, we use the 20-th, 21-th frames, and GT labels (from 20-th frame) to predict the labels in the 19-th frame. Similarly, we use the 19-th, 20-th frames, and GT labels to predict the labels in the 21-th frames. We also reconstruct the 21-th frame and combined it with the 19-th and 20-th frames to predict pseudo-labels for the 22-th frame. The evaluation pseudo-labeled set is thus constructed for experimental comparison.

Since PseudoProp is a pure image-based object detector, we also compare the inference performance with state-of-the-art supervised object detection models SSD and DSP-Net [7] as baselines.

4.2. Results

General Performance. We set 2 iterations for the Naive-Student and PseudoProp for the setting of $1\times$ pseudo-label size. We also test model performance on the size of pseudo-labeled data with $2\times$ and $3\times$ settings. Since VideoProp can also generate pseudo-labeled data according to the GT labeled data, we proposed to combine $1\times$ pseudo-labeled data from PseudoProp and $1\times$ ($2\times$) pseudo-labeled data from VideoProp as a new $2\times$ ($3\times$) pseudo-labeled data. We then train the student models on these new datasets and name them PseudoProp*. A similar approach is applied to the Naive-Student model (Naive-Student*). The performance of all models is shown in Table 1. Observe that PseudoProp and PseudoProp* achieve superior

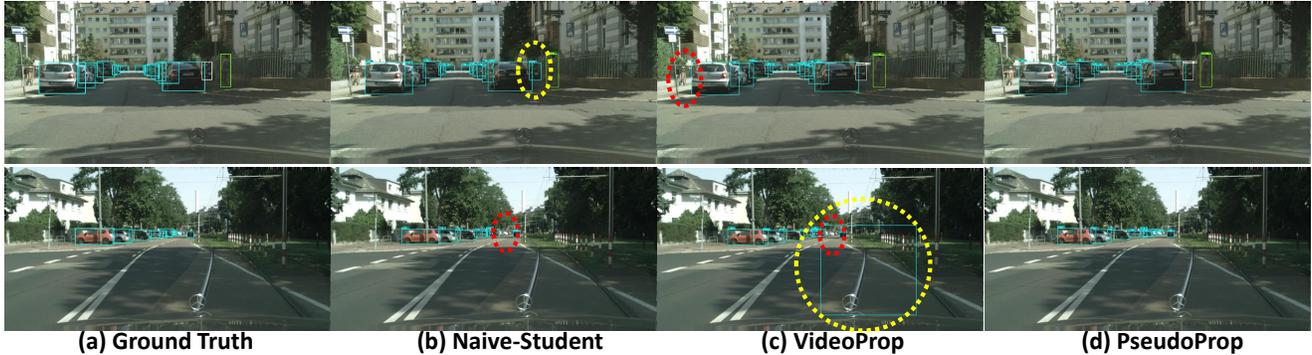


Figure 4. Visual comparison among the (a) ground truth, (b) Naive-Student, (c) VideoProp, and (d) the proposed PseudoProp in the Cityscapes evaluation. Red and yellow dotted ovals indicate the objects misdetections and falsely detected, respectively.

Ratio	Models	Training Methods	Test Set		
			mAP (%)	mAP ⁵⁰ (%)	mAP ⁷⁵ (%)
1×	EfficientDet-D1	s	19.0	35.5	17.2
	SSD	s	-	36.7	-
	DSPNet	s	-	36.9	-
	VideoProp	ss	21.7	40.3	19.9
	Naive-Student (iteration 1)	ss	20.8	39.0	18.8
	Naive-Student (iteration 2)	ss	22.2	40.8	20.3
	PseudoProp (iteration 1, ours)	ss	21.6	40.4	19.9
PseudoProp (iteration 2, ours)	ss	22.6 (+1.8%)	41.4 (+1.5%)	20.9 (+3.0%)	
2×	VideoProp	ss	21.9	43.0	19.6
	Naive-Student (iteration 1)	ss	21.2	38.9	19.6
	Naive-Student* (iteration 1)	ss	22.8	43.3	19.8
	PseudoProp (iteration 1, ours)	ss	21.7	41.0	20.2
	PseudoProp* (iteration 1, ours)	ss	23.2 (+1.8%)	44.4 (+2.5%)	20.9 (+5.6%)
3×	VideoProp	ss	22.3	42.0	19.8
	Naive-Student (iteration 1)	ss	21.0	39.7	18.7
	Naive-Student* (iteration 1)	ss	23.1	43.2	21.5
	PseudoProp (iteration 1, ours)	ss	21.7	40.0	19.8
	PseudoProp* (iteration 1, ours)	ss	23.1 (+0%)	43.9 (+1.6%)	23.1 (+7.4%)

Table 1. Comparison of mAP, mAP⁵⁰, and mAP⁷⁵ of different object detection models in the Cityscapes evaluation when using 1×, 2×, and 3× pseudo-labeled data and k=1. The improved percentage is also reported when compared with the best baseline result under the same conditions. “-” represents that no performance is provided in the existing works. “s” stands for supervised and “ss” for semi-supervised training. All “ss”-based models are implemented based on the EfficientDet-D1. Best results are shown in bold.

performance in all settings. They achieve larger performance gains in mAP⁷⁵, which implies our generated bounding boxes are more accurate than the others. Fig. 4 shows qualitative comparison results.

Discussions. We next discuss observations and analysis of our experiments. First, SSL models outperform supervised learning-based models. This is because SSL models use not only the original labeling but also high-quality pseudo-labeled data for training. Second, PseudoProp improves the pseudo-label quality of the ordinary Naive-Student thanks to the SWBF. Comparing VideoProp and PseudoProp (iteration 2) in the 1× setting, we find the improved performance of the motion-based model in the teacher-student architecture. PseudoProp is more general and flexible than VideoProp, as VideoProp only generates

pseudo-labels near the GT. Third, performance for all SSL-based models can be improved by increasing the pseudo-labeled data size. Fourth, PseudoProp* achieves the best performance, as the most high-quality pseudo-labeled data propagated from the GT are used. SWBF generates more random pseudo-labels, which increases the diversity of the data. When the generated data ratio increases 2× to 3×, PseudoProp performance decreases slightly, and a reason is that more noisy data were used for training. Finally, the inference time of PseudoProp is no different than any teacher model.

4.3. Ablation Study

Performance on Different Score Thresholds. After using SWBF to generate the robust pseudo-labels, we set a threshold to remove noisy pseudo-labels based on their confidence scores before passing them to the student model. We compare Naive-Student and PseudoProp using 4 different thresholds {0, 0.1, 0.2, 0.3} and report the results in Fig. 5(a). Observe that even without setting the threshold, PseudoProp outperforms Naive-Student on all evaluation metrics. This means SWBF generates more reliable pseudo-labels. Furthermore, PseudoProp can be more robust by tuning the threshold *w.r.t.* the observed mAP⁵⁰ performance.

Performance on Low-Data Regime (training with only a small amount of labeled data). We explore the model performance with different amounts of training data together with a *fixed* amount of pseudo-labeled data. This way, we can understand how many labeled data points are needed to fine-tune our model that can effectively speed up the whole training process. Therefore, we randomly extract small sets of labeled data with different sizes such as 500, 1000, 2000 from the original training set. Then we compare Naive-Student and PseudoProp models on each of these sets when performing the fine-tuning of the student model. We show the comparison results in Fig. 5(b). Note that PseudoProp outperforms Naive-Student in all evaluation metrics, even in the low-data regime.

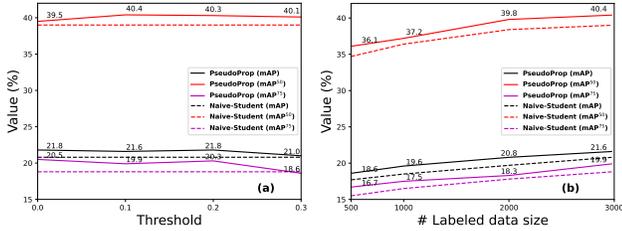


Figure 5. Comparison of mAP, mAP⁵⁰, and mAP⁷⁵ of Naive-Student and PseudoProp models on the Cityscapes dataset with 1 iteration when using 1× pseudo-labeled data, $k = 1$, and (a) different thresholds and (b) different size of labeled data.

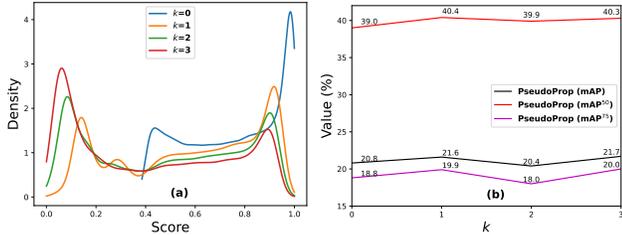


Figure 6. (a) Confidence score distributions with different k . (b) Comparison of mAP, mAP⁵⁰, and mAP⁷⁵ of Naive-Student ($k=0$) and PseudoProp models ($k=1,2,3$) on the Cityscapes dataset with 1 iteration, 1× pseudo-labeled data, and threshold 0.1.

Performance on Different Motion Propagation Length k . We compare the model performance with different k values, which control the BPLP time steps. First, in Figure 6(a), we explore the pseudo-label confidence score distributions with different k . Note that $k=0$ represents the Naive-Student. As discussed before, we select pseudo-labels with confidence scores larger than 0.4 to remove the noisy labels. Therefore, the density curve starts from 0.4 for $k=0$. In addition, since we use a similarity function in Eq.(4) and different k for the fusion method to reduce the false positives, we should expect that our confidence scores for pseudo-labels are smaller than the original ones. Fig. 6(a) verifies this phenomenon in that the curves shift to the left as k increases. The more neighboring frames we use to propagate pseudo-labels, the more likely the fused final confidence score would decrease. In other words, the confidence score for non-robust pseudo-labels would be suppressed further.

Observe in Fig. 6(b) that PseudoProp outperforms Naive-Student model ($k=0$) in mAP⁵⁰ by setting $k > 0$. We also observe that increasing k does not always increase performance. This can be explained by the fact that larger k also induces more noise (false positives) from neighboring frames. Nevertheless, our method achieves the best performance in mAP and mAP⁷⁵ when setting $k=3$. Since both metrics consider higher requirements for predicting accurate bounding box positions, the proposed SWBF yields better estimation for bounding box coordinates when fusing with a larger k . The advantage of accurate bounding box coordinates outweighs the introduced noise.

Fusion Methods	Test Set		
	mAP (%)	mAP ⁵⁰ (%)	mAP ⁷⁵ (%)
NMS	21.0	39.7	19.1
NMW	21.0	39.8	19.1
SNMS	21.2	39.8	19.3
WBF	21.0	39.6	19.1
SWBF (ours)	21.6 (+1.9%)	40.4 (+1.5%)	19.9 (+3.1%)

Table 2. Comparison of mAP, mAP⁵⁰, and mAP⁷⁵ for PseudoProp model with different fusion methods in the Cityscapes evaluation. The hyperparameters are 1 iteration, 1× pseudo-labeled data, $k = 1$, and threshold 0.1. Best results are shown in bold.

Performance on Different Fusion Methods. A similarity-based WBF approach is used in PseudoProp for bounding boxes fusion. The following ablation study verifies that our fusion module is indeed better than the state-of-the-art fusion methods [28], including non-maximum suppression (NMS) [21], soft-NMS (SNMS) [5], non-maximum weighed (NMW) [37], and the original WBF [28]. We replace our fusion module with each of the comparison methods and then evaluate performance on the test dataset. Table 2 shows results of this experiment. It is clear that our proposed SWBF method achieves the best performance. Note the performance of the original WBF is very close to other state-of-the-art methods. By adding the proposed similarity modification on top of WBF, SWBF outperforms all comparison methods.

5. Conclusion

In this paper, we develop the PseudoProp model to generate robust pseudo-labels that can effectively improve semi-supervised, per-image based object detection. We propose a BPLP method to resolve the misdetection problem in the pseudo-labels. In addition, the proposed similarity approach combined with the WBF method can effectively suppress the confidence scores of the falsely detected bounding boxes. Experimental evaluations on the Cityscapes dataset demonstrate that PseudoProp can improve not only traditional teacher-student based but also motion-based semi-supervised models. Our generated pseudo-labels are reliable for model training, which is validated qualitatively and quantitatively.

Future Works. There are several important future directions to improve PseudoProp. First, we would consider jointly training an end-to-end model for object detection and motion prediction in our semi-supervised framework. Second, motion prediction error can be taken into account, in which we assume perfect prediction in this work. Third, since the fusion mechanism naturally benefits from soft labels, we believe that incorporating soft labels into PseudoProp will further improve its performance. Fourth, further evaluation can be performed on quantifying how the feature extraction module affects PseudoProp performance.

References

- [1] Eric Arazo and et al. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *IJCNN*, pages 1–8. IEEE, 2020. 1, 2
- [2] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 941–951, 2019. 3
- [3] David Berthelot and et al. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. 1, 2
- [4] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Know your surroundings: Exploiting scene information for object tracking. In *European Conference on Computer Vision*, pages 205–221. Springer, 2020. 3
- [5] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. Soft-nms—improving object detection with one line of code. In *ICCV*, 2017. 8
- [6] Kai Chen, Jiaqi Wang, Shuo Yang, Xingcheng Zhang, Yuanjun Xiong, Chen Change Loy, and Dahua Lin. Optimizing video object detection via a scale-time lattice. In *CVPR*, June 2018. 3
- [7] Liangfu Chen and et al. Driving scene perception network: Real-time joint detection, depth estimation and semantic segmentation. In *WACV*, 2018. 6
- [8] Liang-Chieh Chen and et al. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *ECCV*, 2020. 1, 2, 3, 6
- [9] Marius Cordts, Mohamed Omran, and et al. The Cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 1, 2, 5
- [10] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NeurIPS*, pages 379–387, 2016. 2
- [11] Sorin Grigorescu et al. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3):362–386, 2020. 2
- [12] Hui Guo, Shu Hu, Xin Wang, Ming-Ching Chang, and Siwei Lyu. Eyes tell all: Irregular pupil shapes reveal gan-generated faces. *ICASSP*, 2022. 2
- [13] Hui Guo, Shu Hu, Xin Wang, Ming-Ching Chang, and Siwei Lyu. Robust attentive deep neural network for exposing gan-generated faces. *IEEE ACCESS*, 2022. 2
- [14] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [15] Shu Hu, Yuezun Li, and Siwei Lyu. Exposing gan-generated faces using inconsistent corneal specular highlights. In *ICASSP*, pages 2500–2504. IEEE, 2021. 2
- [16] E. Ilg and et al. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 4
- [17] Tsung-Yi Lin and et al. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 6
- [18] Daniel McKee, Bing Shuai, Andrew Berneshawi, Manchen Wang, Davide Modolo, Svetlana Lazebnik, and Joseph Tighe. multi-object tracking with hallucinated and unlabeled videos. In *CVPR workshop*, 2021. 3
- [19] Claudio Michaelis and et al. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *Machine Learning for Autonomous Driving Workshop at NeurIPS*, 2019. 6
- [20] Ishan Misra and et al. Watch and learn: Semi-supervised learning for object detectors from video. In *CVPR*, 2015. 1, 2
- [21] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *ICPR*, 2006. 8
- [22] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015. 2
- [23] Fitsum A. Reda and et al. SDC-Net: Video prediction using spatially-displaced convolution. In *ECCV*, 2018. 2, 4, 6
- [24] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [25] Chuck Rosenberg and et al. Semi-supervised self-training of object detection models. In *WACV*, volume 1, pages 29–36, 2005. 2
- [26] Olga Russakovsky and et al. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 6
- [27] Kihyuk Sohn and et al. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020. 1, 2
- [28] Roman Solovyev and et al. Weighted boxes fusion: Ensembling boxes from different object detection models. *IVC*, 2021. 2, 5, 8
- [29] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 6
- [30] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *CVPR*, 2020. 6
- [31] Yuxing Tang and et al. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *CVPR*, 2016. 2
- [32] Yihe Tang and et al. Humble teachers teach better students for semi-supervised object detection. In *CVPR*, 2021. 1, 2
- [33] Caiyong Wang et al. Nir iris challenge evaluation in non-cooperative environments: Segmentation and localization. In *IJCB*, pages 1–10. IEEE, 2021. 2
- [34] Xin Wang, Hui Guo, Shu Hu, Ming-Ching Chang, and Siwei Lyu. Gan-generated faces detection: A survey and new perspectives. *arXiv preprint arXiv:2202.07145*, 2022. 2
- [35] Shifeng Zhang et al. Single-shot refinement neural network for object detection. In *CVPR*, pages 4203–4212, 2018. 2
- [36] Xujiang Zhao, Feng Chen, Shu Hu, and Jin-Hee Cho. Uncertainty aware semi-supervised learning on graph data. *NeurIPS*, 2020. 3
- [37] Huajun Zhou et al. Cad: Scale invariant framework for real-time object detection. In *ICCVW*, 2017. 8
- [38] Yi Zhu and et al. Improving semantic segmentation via video propagation and label relaxation. In *CVPR*, 2019. 3, 4, 6
- [39] Yuliang Zou and et al. PseudoSeg: Designing pseudo labels for semantic segmentation. In *ICLR*, 2021. 6