

Naive-Student: Leveraging Semi-Supervised Learning in Video Sequences for Urban Scene Segmentation

Liang-Chieh Chen¹, Raphael Gontijo Lopes¹, Bowen Cheng²,
Maxwell D. Collins¹, Ekin D. Cubuk¹, Barret Zoph¹,
Hartwig Adam¹, Jonathon Shlens¹

¹ Google Research, ² UIUC

Abstract. Supervised learning in large discriminative models is a mainstay for modern computer vision. Such an approach necessitates investing in large-scale human-annotated datasets for achieving state-of-the-art results. In turn, the efficacy of supervised learning may be limited by the size of the human annotated dataset. This limitation is particularly notable for image segmentation tasks, where the expense of human annotation is especially large, yet large amounts of unlabeled data may exist. In this work, we ask if we may leverage semi-supervised learning in unlabeled video sequences and extra images to improve the performance on urban scene segmentation, simultaneously tackling semantic, instance, and panoptic segmentation. The goal of this work is to avoid the construction of sophisticated, learned architectures specific to label propagation (*e.g.*, patch matching and optical flow). Instead, we simply predict pseudo-labels for the unlabeled data and train subsequent models with both human-annotated and pseudo-labeled data. The procedure is iterated for several times. As a result, our Naive-Student model, trained with such simple yet effective iterative semi-supervised learning, attains state-of-the-art results at all three Cityscapes benchmarks, reaching the performance of 67.8% PQ, 42.6% AP, and 85.2% mIOU on the test set. We view this work as a notable step towards building a simple procedure to harness unlabeled video sequences and extra images to surpass state-of-the-art performance on core computer vision tasks.

Keywords: semi-supervised learning, pseudo label, semantic segmentation, instance segmentation, panoptic segmentation




Fig. 1. Naive-Student: semi-supervised learning in video sequences for scene segmentation. We iteratively train on human-annotated frames (1 out of 30 in each Cityscapes video sequence), and generate pseudo-labels for the other unlabeled video frames (left). Segmentation performances (val set) improve at each iteration (right).

1 Introduction

Significant advances in computer vision due to deep learning [40,66,26,10] have been tempered by the fact that these advances have been accrued through supervised learning on large-scale, human-annotated datasets [49,69]. The paradigm of supervised learning requires the expenditure of a large amount of resources to manually label static images – whether through the development of specialized annotation tools [5,70,8], or the amount of human hours for the annotation itself [49,69]. Such an approach does not scale effectively to comprehensively label real-time video frames (but see [65,2,7]). More importantly, supervised training is rather sample-inefficient as many examples are required for good generalization [29,39]. Ideally, one would expect and hope that a training method may be able to learn in a more self-supervised manner particularly on video – much as presumed to occur in human visual learning [41,82].

The limitations of supervised learning is most pronounced in the task of image segmentation [22]. Human annotation of static images for segmentation is particularly expensive, requiring, for instance, 90 minutes per image [17] or 22 worker hours per 1,000 mask segmentations [49]. In the case of self-driving cars, the annotation of video is a critical supervised learning problem [24,75], and in turn has fostered an industry of specialized companies for data annotation.

In contrast, recent findings on the benefits of pre-training on ImageNet for segmentation [47] indicate that current segmentation approaches may benefit from large-scale image classification datasets. This direction has been further pursued by [74,12] on an extremely large image classification dataset [30]. Additionally, many segmentation methods [92,9] apply transfer learning by pre-training on augmented segmentation datasets [27,49,55] and then fine-tuning on the target datasets [21,17]. Likewise, other works attempt to exploit label propagation in video to improve segmentation. However, these methods require building specialized modules to propagate labels across video frames [54,23,56,95].

In this work, we leverage both unlabeled video frames and extra unlabeled images to improve the urban scene segmentation evaluated in terms of semantic segmentation, instance segmentation, and panoptic segmentation. Importantly, we do not require any specialized methods for propagating label information across video frames, such as optical flow [54,23,56], patch matching [4,6], or learned motion vector [95]. Instead, we propose to employ a simple iterative semi-supervised learning procedure. At each iteration, the model from the previous iteration generates pseudo-labels for unlabeled video frames (Figure 1). Specifically, a pseudo-label is generated through a distillation across multiple augmentations applied to each unlabeled video frame. Subsequent iterations of the training procedure train on the original labeled data as well as the newly pseudo-labeled data. Our model, trained with such a simple yet effective method, simultaneously sets new state-of-the-art results on the Cityscapes urban scene segmentation [17], achieving 67.8% PQ, 42.6% AP, and 85.2% mIOU on test set. We hope that such an iterative semi-supervised learning may provide more label-efficient methods for developing a machine learning solution to segmentation.

2 Related Works

Our method is related to both *self-training* [71,88,67,20,25,91], where the predictions of a model on unlabeled data is used to train the model, and *semi-supervised learning* [68,44,57,64,86], where additionally extra human-annotated data is available to guide the training with unlabeled data. In particular, our model is trained with some human-annotated images and abundant pseudo-labeled [42,72,34,3] video sequences.

Semi-supervised learning has been widely applied to several computer vision tasks, including semantic segmentation [57,19,59,31,80,35,81,73,96], object detection [68,76,64], instance segmentation [35,60], panoptic segmentation [45], human pose estimation [58], person re-identification [93], multi-object tracking and segmentation [77,62], and so on. A comprehensive literature survey is beyond the scope of this work, and thus we focus on comparing our proposed method with the most related ones.

Our proposed iterative semi-supervised learning is similar to the work by Papandreou *et al.* [57], STC [80], Simple-Does-It [35], the work by Li *et al.* [45], and Noisy-Student [84]. In particular, our iterative semi-supervised learning is similar to the Expectation-Maximization method by Papandreou *et al.* [57] which alternates between estimating the latent pixel labels (*i.e.*, pseudo labels) and optimizing the network parameters with bounding box or image-level annotations. Similarly, Li *et al.* [45] generate pseudo labels for panoptic segmentation by exploiting both fully-annotated and weakly-annotated images, where bounding boxes for ‘thing’ classes and image-level tags for ‘stuff’ classes are provided. However, unlike those two works, we do not exploit any weakly-annotated data. Additionally, we do not sort the images by the annotation difficulty and do not exploit any other assistance, such as saliency maps, as in STC [80]. Simple-Does-It [35] adopts a complicated de-noising procedure to clean the pseudo labels, while we simply use the outputs from a neural network. Finally, following Noisy-Student [84], we employ a stronger Student network in the subsequent iterations, but we do not employ any noisy data augmentation (*i.e.*, RandAugment [18]).

When generating pseudo labels, we employ a simple test-time augmentation, *i.e.*, multi-scale inputs and left-right flips, a common strategy used by segmentation models [12,92], which bears a similarity to Data-Distillation [64]. However, our framework is deployed in an iterative manner, and we exploit unlabeled video sequences for scene segmentation, simultaneously tackling semantic, instance, and panoptic segmentation. Additionally, we do not set a threshold as [64] to remove false positives, avoiding tuning of another hyper-parameter.

Video sequences have also been exploited in semi-supervised learning for semantic segmentation. Human-annotated ground-truth labels of certain frames in a video sequence could be propagated to other unlabeled frames via patch matching [4,6] or optical flow [54,23,94,53,56]. Recently, Zhu *et al.* [95] generate pseudo-labeled video sequences by jointly propagating the image-label pair with learned motion vectors, and demonstrate promising results. Similarly, our method also exploits unlabeled video sequences. However, our method is much simpler since we do not employ any label-propagation modules (*e.g.*, patch matching [4,6],

Algorithm 1 Iterative semi-supervised learning for urban scene segmentation.

Labeled data: n pairs of image x_i and corresponding human annotation y_i

Unlabeled data: m images collected from multiple video sequences or extra images with no human annotations $\{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_m\}$.

Step 1: Train a Teacher network θ_t (with prediction function f) on the manually labeled images by minimizing the total loss \mathcal{L} for scene segmentation.

$$\theta_t^* = \arg \min_{\theta_t} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(y_i, f(x_i, \theta_t))$$

where $\mathcal{L} = \lambda_{\text{sem}} \mathcal{L}_{\text{sem}} + \lambda_{\text{heatmap}} \mathcal{L}_{\text{heatmap}} + \lambda_{\text{offset}} \mathcal{L}_{\text{offset}}$ in our framework.

Step 2: Generate pseudo-labels \tilde{y}_i for unlabeled images with test-time augmentations (*i.e.*, multi-scale inputs and left-right flips).

$$\tilde{y}_i = f(\text{Aug}(\tilde{x}_i), \theta_t^*), \forall i = 1, \dots, m$$

where $\text{Aug}(\cdot)$ is test-time augmentation.

Step 3: Train an *equal or larger* Student network θ_s on pseudo-labeled images $(\tilde{x}_i, \tilde{y}_i)$ with the same objective.

$$\theta_s^* = \arg \min_{\theta_s} \frac{1}{m} \sum_{i=1}^m \mathcal{L}(\tilde{y}_i, f(\tilde{x}_i, \theta_s))$$

Step 4: Fine-tune the Student network θ_s^* from step 3 on the manually labeled image annotations (x_i, y_i) using the same objective.

$$\theta_s^{**} = \arg \min_{\theta_s^*} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(y_i, f(x_i, \theta_s^*))$$

Step 5: Return to step 2 but employ the Student network θ_s^{**} as a Teacher until reaching desired number of iterations.

optical flow [54,23,56], or motion vectors [95]) but instead directly generate the pseudo labels for each video frame.

3 Methods

Alg. 1 gives an overview of our proposed iterative semi-supervised learning for scene segmentation. Suppose two sets of images are given, where one contains human annotations and the other does not. The human-annotated images are exploited to train a Teacher network using the loss function for scene segmentation. Pseudo-labels for those un-annotated images are then generated by the Teacher network with a test-time augmentation function. A Student network is subsequently trained with the pseudo-labeled images using the same loss function for scene segmentation. The Student network is then fine-tuned on human-labeled images before evaluating on the validation set or test set. Finally, one could optionally replace the Teacher network with the Student network and iterate the procedure again. Our method, dubbed *Naive-Student*, is motivated by Noisy-Student [84] where we adopt a stronger Student network in the following iterations, but we do not inject noise (*i.e.*, RandAugment [18]) to the Student. Our algorithm is illustrated in Fig. 2. We elaborate on the details below.

The Loss for Scene Segmentation: Our core building block is the state-of-the-art bottom-up panoptic segmentation model, Panoptic-DeepLab [14], which improves the semantic segmentation model DeepLabv3+ [13] by incorporating another class-agnostic instance segmentation prediction. Its instance segmentation prediction involves a simple instance center prediction as well as the offset regression from each pixel to its corresponding center. As a result, the total

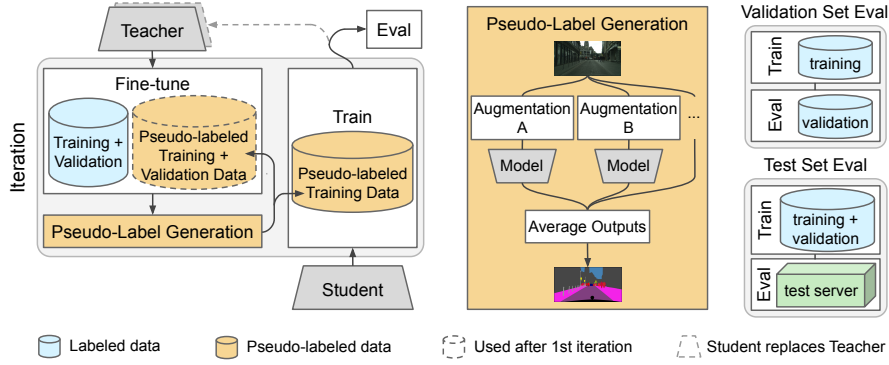


Fig. 2. Overview of our proposed iterative semi-supervised learning for scene segmentation. The Teacher network is trained with all available human-annotated images (and extra pseudo-labeled images after 1st iteration), and then generates pseudo-labels for all the unlabeled images with a simple test-time augmentation (*i.e.*, multi-scale inputs and left-right flips). The Student network is subsequently trained with the pseudo-labeled data, and optionally replaces the Teacher network in following iterations. Before evaluating the validation or test set performance, the Student network is fine-tuned on the human-annotated images. Note that the validation set is only exploited by the Teacher in order to generate high-quality pseudo-labels, and the Student has no access to it. Additionally, the final test set results are evaluated on a fair test server where the annotations are held-out.

loss function \mathcal{L} for scene segmentation boils down to three loss functions: softmax cross entropy loss \mathcal{L}_{sem} for semantic segmentation, mean squared error loss $\mathcal{L}_{heatmap}$ for instance center prediction, and L_1 loss \mathcal{L}_{offset} for offset regression. In our algorithm, the Teacher and the Student networks are trained with the same total loss function \mathcal{L} .

Pseudo-Label Generation: After training the Teacher network on all human-annotated images (and all pseudo-labeled images after iteration 1), we generate (or update) the pseudo labels for all un-annotated images with a test-time augmentation function $Aug(\cdot)$. We simply use the common test-time augmentations, *i.e.*, multi-scale inputs and left-right flips. We only generate hard pseudo labels (*i.e.*, a one-hot distribution) in order to save disk space when processing large resolution images (*e.g.*, Cityscapes image size is 1024×2048).

Ego-Car Region in Pseudo Labels: Cityscapes images are collected (or recorded) with a driving vehicle. A part of the vehicle, called "ego-car" region, is thus visible in all frames of a video sequence. This region is ignored during evaluating the model performance. However, we find that assigning a random pseudo label value to those regions will confuse models during training. To handle this problem, we adopt a simple solution by exploiting the prior that Cityscapes images are all well-calibrated and the ego-car regions are in the same locations for images collected from the same sequence. Since we have access to the only one human-annotated image from a 30-frame sequence, we propagate this ego-

car region information to the other 29 frames in the same sequence and assign them with void label (*i.e.*, no loss back-propagation for those regions).

A Better Network Backbone for Scene Segmentation: The efficient backbone Xception-71 (X-71) [16,63,13] is adopted in the Teacher network at the first iteration in our iterative semi-supervised learning algorithm. In the next iteration, a stronger backbone should be used to generate pseudo labels with a better quality. In this work, we modify the powerful Wide ResNet-38 (WR-38) [83,90] for scene segmentation. In particular, we remove the last residual block B7 in WR-38 [83] and repeat the residual block B6 two more times, resulting in our proposed WR-41. Additionally, we adopt drop path [32,50] (with a constant survival probability 0.8) and multi-grid scheme [12,79] in the last three residual blocks (with unit rate $\{1, 2, 4\}$, same as [12]). As a result, the proposed WR-41 attains better performance than X-71 in the fully supervised setting.

4 Experiments

We conduct experiments on the popular Cityscapes dataset [17], which consists of a large and diverse set of street-view video sequences recorded from 50 cities primarily in Germany. From the video sequences, 5000 images are provided with high-quality pixel-wise annotations in which 2975, 500, and 1525 images are used for training, validation, and test, respectively. Each image is selected from the 20th frame of a 30-frame video snippet. Additionally, another 20000 images are accompanied with coarse annotations. We define each dataset split below.

train-fine: Training set (2,975 images) with fine pixel-wise annotations.

val-fine: Validation set (500 images) with fine pixel-wise annotations.

test-fine: Test set (1,525 images) where the fine pixel-wise annotations are held-out, and the evaluation is performed on a fair test server.

train-extra: Extra 20,000 images with coarse annotations. Our proposed method is not limited to video sequences, and thus we also generate pseudo-labels for this set, instead of using the provided coarse annotations.

train-sequence: The video sequences where the train-fine set is selected from. This set contains $2975 \times 30 = 89,250$ frames.

val-sequence: The video sequences where the val-fine set is selected from. This set contains $500 \times 30 = 15,000$ frames.

Furthermore, one could merge training and validation splits (*e.g.*, trainval-fine is merged from train-fine and val-fine, and similarly for trainval-sequence).

Experimental Setup: We report mean intersection-over-union (mIOU), average precision (AP), and panoptic quality (PQ) to evaluate the semantic, instance, and panoptic segmentation results, respectively.

The state-of-art bottom-up panoptic segmentation model, Panoptic-DeepLab [15], is included in our proposed iterative semi-supervised learning pipeline. Panoptic-DeepLab is a simple framework and simultaneously produces semantic, instance, and panoptic segmentation results without the need to fine-tune on each task. We adopt the same training protocol as [15] when using Panoptic-DeepLab. For example, our models are trained using TensorFlow [1] on 32 TPUs.

We use the ‘poly’ learning rate policy [52] with an initial learning rate of 0.001 for Xception-71 (X-71) backbone [16,63] and 0.0001 for our proposed Wide ResNet-41 (WR-41) [28,90,83], respectively. During training, the batch normalization [33] is fine-tuned, random scale data augmentation and Adam [36] optimizer without weight decay are adopted. On Cityscapes, we employ training crop size equal to 1025×2049 with batch size 32, and 180K training iterations. Similar to other works on panoptic segmentation [38,43,37,85,61,87,78], we re-assign to void label all ‘stuff’ segments whose areas are smaller than a threshold of 4096. Additionally, we employ multi-scale inference (scales equal to $\{0.5, 0.75, 1, 1.25, 1.5, 1.75, 2\}$ for Cityscapes) and left-right flipped inputs, to further improve the performance for test server evaluation.

4.1 Urban Scene Segmentation Results

In this subsection, we summarize our main results on the Cityscapes dataset.

Cityscapes val-fine set: In our iterative semi-supervised learning framework, at each iteration, all data splits, including Mapillary Vistas [55] and Cityscapes trainval-fine, (also trainval-sequence and train-extra after 1st iteration), are exploited for the Teacher networks in order to generate better pseudo-labels, while the Student networks are always initialized from the Mapillary Vistas pretrained checkpoint (unless it is specified that it is initialized from previous iterations). In Tab. 1, we report the validation set results. At iteration 0, we employ the state-of-art Panoptic-DeepLab with Xception-71 (validation set results from [15] are shown in the table for comparison) as the Teacher network to generate pseudo-labels for train-sequence and train-extra splits which are subsequently used to train our Student network using the proposed Wide ResNet-41 as backbone. As a result, at iteration 1, we improve over the Panoptic-DeepLab (X-71) baseline by a margin of 3.8% PQ, 3.1% AP, and 3.2% mIOU. The Student network is then selected as the new Teacher network after fine-tuning on all the available data splits (*i.e.*, trainval-sequence, train-extra, and trainval-fine). At iteration 2, by training with the better quality pseudo-labels, we observe an additional improvement of 1.3% PQ, 1.5% AP, and 0.8% mIOU for the new Student network. Additionally, one could further slightly improve the performance by initializing the Student network from iteration 1, as shown in the last row.

Cityscapes test-fine set: In Tab. 2, we report our Cityscapes test set results. As shown in the table, our single model simultaneously ranks 1st at all three Cityscapes benchmarks. In particular, for the panoptic segmentation benchmark, our model outperforms Panoptic-DeepLab (X-71) [15] by 2.3% PQ, Li *et al.* [46] by 4.5% PQ, and Seamless [61] by 5.2% PQ. For the instance segmentation benchmark, our model outperforms PolyTransform [48] by 2.5% AP, Panoptic-DeepLab (X-71) [15] by 3.6% AP, and PANet [51] by 6.2% AP. Finally, for the competitive semantic segmentation benchmark, our model outperforms Panoptic-DeepLab (X-71) [15] by 1.0% mIOU, OCR [89] by 1.5% mIOU, and Zhu *et al.* [95] by 1.7% mIOU.

Visualization of generated pseudo labels: We observe visually subtle differences between iteration 1 and iteration 2, since both Teachers yield high-

Table 1. Iterative semi-supervised training systematically improves urban scene segmentation results. Results presented on Cityscapes validation data. The Students are pretrained on ImageNet [69] and Mapillary Vistas [55]. Baseline Xception-71 (X-71) at iteration 0 is obtained from [15], while Wide-ResNet-41 (WR-41) is *modified* from [83]. All Teacher networks have been trained on ImageNet, Mapillary Vistas, and Cityscapes. Labeled data is from Cityscapes train-fine set. Pseudo-Labeled data is from Cityscapes train-sequence and train-extra sets. [†] indicates that model was initialized from the checkpoint in the previous iteration. The text color indicates how the Student is selected as the new Teacher (*e.g.*, the Student X-71 at iteration 0 becomes the Teacher at iteration 1 after fine-tuning).

Itr	Architecture		Training Set		Validation Set		
	Student	Teacher	Labeled	Pseudo-Labeled	PQ (%)	AP (%)	mIOU (%)
0	X-71	-	✓		65.3	38.8	82.5
1	WR-41	X-71		✓	69.1	41.9	85.7
2	WR-41	WR-41		✓	70.4	43.4	86.5
2	WR-41 [†]	WR-41		✓	70.8	44.3	86.7

Table 2. Iterative semi-supervised learning achieves state-of-the-art urban scene segmentation results. Results presented on Cityscapes *test-fine* set. **C**: Cityscapes coarse annotation. **V**: Cityscapes video. **MV**: Mapillary Vistas. Note that we do not exploit the train-extra coarse annotations, but instead we generate pseudo-labels for them.

Model	Extra Data	Training Method	PQ (%)	AP (%)	mIOU (%)
Naive-Student (ours)	C, V, MV	iterative semi-supervised	67.8	42.6	85.2
Seamless [61]	MV	supervised	62.6	-	-
Li <i>et al.</i> [46]	COCO	supervised	63.3	-	-
Panoptic-DeepLab (X-71) [15]	MV	supervised	65.5	39.0	84.2
PANet [51]	COCO	supervised	-	36.4	-
PolyTransform [48]	COCO	supervised	-	40.1	-
Zhu <i>et al.</i> [95]	C, V, MV	semi-supervised	-	-	83.5
OCR [89]	C, MV	supervised	-	-	83.7

quality results. To further look into the minor differences, we zoom-in some generated pseudo-labels in Fig. 3. As shown in the figure, the Teacher at iteration 2 generates slightly better pseudo-labels along the thin and small objects.

Visualization of segmentation results: In Fig. 4, we visualize some segmentation results obtained by the Student network on val-fine set.

4.2 Ablation Studies

In this subsection, we provide ablation studies on several design choices. Xception-71 is used as the backbone if not specified.

Training iterations: First, we verify that the performance improvement does not solely result from longer training iterations, but from the extra large pseudo-labeled images. We train Panoptic-DeepLab [15] with 60K iterations on Cityscapes train-fine set and obtain a PQ of 62.9%. We increase the training iterations to 120K iterations, but do not observe any improvement (62.7% PQ)

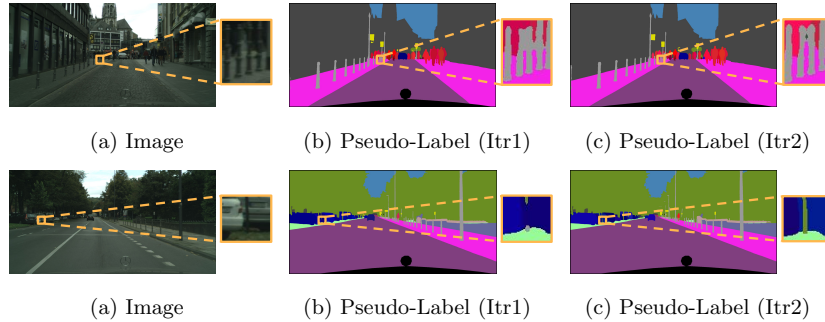


Fig. 3. Generated pseudo-labels by Teacher improves qualitatively with more iterations. We only observe subtle difference between pseudo-labeled video frames at iteration 1 and iteration 2. The results from iteration 2 capture better thin objects, as zoomed-in in the yellow regions.



Fig. 4. Segmentation results by Student on Cityscapes val set.

(*i.e.*, performance saturates after 60K iterations). On the other hand, our proposed Naive-Student attains a better performance with 180K iterations (65.3% PQ) when trained with the larger train-sequence set.

Design choices for the Teacher to generate pseudo-labels: When generating the pseudo-labels, there are four factors involved in our design, namely (1) assignment of void label to the ego-car region, (2) employment of test-time augmentation, (3) more Cityscapes human-labeled images, and (4) Mapillary Vistas pretraining. This “Void-Ego-Car” design, or factor (1), improves 0.7% PQ, 0.5% AP, and 0.4% mIOU. We think the wrongly generated labels in the ego car region slightly affect the model training. Without employing the test-time augmentation, (*i.e.*, multi-scale inference and left-right flipping), when generating the pseudo-labels, the performance drops by 0.8% PQ, 1.1% AP, and 0.9% mIOU. Excluding more Cityscapes human-labeled images for fine-tuning the Teacher network degrades the performance by 1.4% PQ, 1.2% AP, and 1.3% mIOU. Finally, if the Teacher network is not pretrained on the Mapillary Vistas dataset, the performance decreases by 2.2% PQ, 2.2% AP, and 1.5% mIOU.

Table 3. Design choices for the Teacher to generate pseudo-labels. Void-Ego-Car: Assignment of void label to ego-car regions. **Test-Aug:** Test-time augmentation (*i.e.*, multi-scale inputs and left-right flips). **Val-Fine:** Inclusion of Cityscapes val-fine set for training the Teacher network. **MV-Pretrained:** Employment of a pretrained checkpoint on Mapillary Vista for the Teacher network. Results presented on Cityscapes validation set. Note the Student network has no access to the validation set.

Pseudo-Label Generation Scheme				Validation Set Results		
Void-Ego-Car	Test-Aug	Val-Fine	MV-Pretrained	PQ (%)	AP (%)	mIOU (%)
✓	✓	✓	✓	67.5	39.8	83.7
	✓	✓	✓	66.8	39.3	83.3
✓		✓	✓	66.7	38.7	82.8
✓	✓		✓	66.1	38.6	82.4
✓	✓	✓		65.3	37.6	82.2

Table 4. Design choices for training the Student. We experiment with different training splits for training the Student. Results presented on Cityscapes validation set.

Training Set for Student			Validation Set Results		
Train-Fine	Train-Sequence	Train-Extra	PQ (%)	AP (%)	mIOU (%)
✓			63.1	35.2	80.1
	✓		65.3	37.6	82.2
✓	✓		65.2	37.3	82.0
	✓	✓	66.9	40.2	84.2

Design choices for training the Student: In Tab. 4, we report the results when training the Student network with different training set splits. The baseline Student network, trained with Cityscapes train-fine, attains the performance of 63.1% PQ, 35.2% AP, and 80.1% mIOU. Using the pseudo-labeled train-sequence, the performance is improved by 2.2% PQ, 2.4% AP, and 2.1% mIOU. Mixing human-labeled train-fine and pseudo-labeled train-sequence slightly degrades the performance. We think it is because of the inconsistent annotations between human-labeled and pseudo-labeled images, since train-fine is a subset of train-sequence. Finally, adding more pseudo-labeled images (train-sequence and train-extra) improves the result to 66.9% PQ, 40.2% AP, and 84.2% mIOU.

Vary human-labeled images and fix pseudo-labeled images: In Fig. 5, we explore the semi-supervised setting with different amounts of human-labeled images but fixed amount of pseudo-labeled images. In particular, the Teacher network has *only* been trained with different numbers of Cityscapes train-fine images (*i.e.*, no other human-labeled images, such as Mapillary Vistas). The generated pseudo-labels (on Cityscapes train-sequence and train-extra) are used to train another Student network. Both Teacher and Student networks employ the Xception-71 as backbone. For comparison, we also show the performance of the supervised setting with the same amount of human-labeled images. As shown in the figure, we observe (1) the semi-supervised learning setting consistently improves over the fully supervised setting in all three metrics (PQ, AP, and mIOU) as more human-labeled images are exploited, (2) when using only 40%

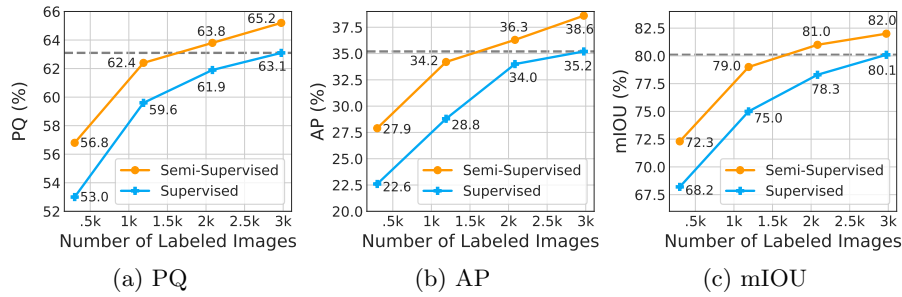


Fig. 5. Semi-supervised learning with a fraction of the original labels may match supervised segmentation performance. Results presented on Cityscapes validation data. We vary the numbers of human-annotated images from train-fine set. With only 40% of labeled data, our semi-supervised learning method attains 98.9% PQ, 97.2% AP, and 98.6% mIOU of the performance of their fully supervised counterparts.

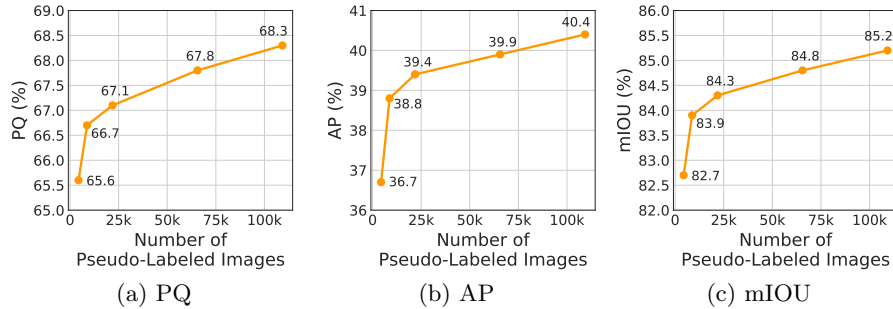


Fig. 6. Increasing the amount of unlabeled data improves segmentation performance for PQ, AP, and mIOU. Results presented on Cityscapes validation data.

of the human-labeled images, our semi-supervised learning method could reap 98.9%, 97.2%, and 98.6% performance from its fully supervised counterparts in PQ, AP, and mIOU, respectively, and (3) when using 100% of the human-labeled images, our semi-supervised learning method attains 65.2% PQ, 38.6% AP, and 82% mIOU, comparable to the fully supervised counterpart with a Mapillary Vistas pretrained checkpoint (65.3% PQ, 38.8% AP, and 82.5% mIOU in [15]).

Fix human-labeled images and vary pseudo-labeled images: In Fig. 6, we explore the semi-supervised setting with different amounts of pseudo-labeled images. In particular, the Teacher network will generate different numbers of pseudo-labeled images for training the Student network. As shown in the figure, we observe consistent improvement in all three metrics when more and more pseudo-labeled images are included in the training.

Training method: In Tab. 5, we experiment with the effect of different training methods: supervised, semi-supervised, and iterative semi-supervised learning. We employ our most powerful backbone, WR-41, attempting to push the envelope of performance. We observe a significant improvement of semi-

Table 5. Comparisons with different training methods. Results presented on Cityscapes validation set. The proposed WR-41 is used as the network backbone. Semi-supervised learning (*i.e.*, only iterate once in our framework) significantly improves the performance, while iterative semi-supervised learning further improves the results.

Training Method	PQ (%)	AP (%)	mIOU (%)
supervised	64.0	38.0	80.7
semi-supervised	69.1	42.4	85.9
<i>iterative</i> semi-supervised	70.4	43.4	86.1

Table 6. Transfer learning from Cityscapes to Mapillary Vistas. Results presented on Mapillary Vistas validation data. We experiment with the network pretrained on ImageNet [69], Cityscapes [17] labeled data (trainval-fine set), and Cityscapes pseudo-labeled data (trainval-sequence, and train-extra).

Training Set		Val Set		
Labeled	Pseudo-Labeled	PQ (%)	AP (%)	mIOU (%)
		37.1	16.7	56.2
✓		36.9	16.5	55.3
✓	✓	37.8	17.0	56.2

supervised learning over supervised learning by 5.1% PQ, 4.4% AP, and 5.2% mIOU, mostly because of the small Cityscapes dataset. Adopting the iterative semi-supervised learning further improves the performance by 1.3% PQ, 1% AP, and 0.2% mIOU. We think there is more room for improving PQ and AP, since the mIOU result is starting to be saturated, as demonstrated in the public leader-board (*i.e.*, differences between top-performing models are about 0.1%).

Transfer learning from Cityscapes to Mapillary Vistas: The transfer learning from the large-scale Mapillary Vistas to Cityscapes has been shown to be effective in the literature [95, 61, 15] and in our work as well, since both datasets contain street-view images. In Tab. 6, we experiment with the other direction of transfer learning from Cityscapes to Mapillary Vistas. The baseline model with WR-41 backbone, pretrained only on ImageNet [69], attains the performance of 37.1% PQ, 16.7% AP, and 56.2% mIOU. If we pretrain the model on the original Cityscapes trainval-fine set, we observe a slight degradation. Interestingly, when we further pretrain the model on the generated pseudo-labels, we observe a small amount of 0.7% improvement in PQ. We think the improvement gained from Cityscapes pretraining is marginal, mainly because the Cityscapes images are mostly taken in Germany, while Mapillary Vistas contains more diverse images.

4.3 Modified Wide ResNet-38: WR-41

In this subsection, we report the experimental results with our modified Wide ResNet-38 [90, 83], called WR-41, on both ImageNet [69] and Cityscapes [17].

ImageNet-1K val set: In Tab. 7, we report the results on the ImageNet-1K validation set. As shown in the table, our TensorFlow re-implementation of wide ResNet-38 (WR-38), proposed in [83], attains 20.36% Top-1 error, which

Table 7. *Single-model* error rates on ImageNet-1K validation set.

Backbone	Drop-Path	Top-1 Error	Top-5 Error
WR-38 [83] (our TF imp.)		20.36%	5.11%
WR-38 [83] (our TF imp.)	✓	19.89%	4.98%
WR-41		20.08%	4.93%
WR-41	✓	19.41%	4.68%

Table 8. Adopting Panoptic-DeepLab with our WR-41 achieves better segmentation accuracy on Cityscapes val set with fewer model parameters and fewer M-Adds than with our TensorFlow re-implemented WR-38.

Backbone	Drop-Path	Multi-Grid	Params (M)	M-Adds (B)	PQ (%)	AP (%)	mIOU (%)
WR-38 [83] (our TF imp.)			173.75	3486.32	62.4	35.3	79.1
WR-38 [83] (our TF imp.)	✓		173.75	3486.32	62.6	36.7	79.5
WR-38 [83] (our TF imp.)	✓	✓	173.75	3493.80	63.1	37.4	80.1
WR-41			147.27	3238.81	63.0	35.5	80.0
WR-41	✓		147.27	3238.81	63.6	36.8	80.4
WR-41	✓	✓	147.27	3246.40	64.0	38.0	80.7

Table 9. Effect of test-time augmentation on Cityscapes val set. MV: Mapillary Vistas pretrained. Flip: Left-right flips. MS: Multi-scale inputs.

Method	MV	Flip	MS	PQ (%)	AP (%)	mIOU (%)
Panoptic-DeepLab (WR-41)				64.0	38.0	80.7
Panoptic-DeepLab (WR-41)		✓		64.5	39.3	81.2
Panoptic-DeepLab (WR-41)		✓	✓	65.0	40.7	81.4
Panoptic-DeepLab (X-71) [15]	✓	✓	✓	64.1	38.5	81.5
Panoptic-DeepLab (WR-41)	✓			66.5	41.5	83.4
Panoptic-DeepLab (WR-41)	✓	✓		66.7	41.8	83.7
Panoptic-DeepLab (WR-41)	✓	✓	✓	67.3	43.4	83.8
Panoptic-DeepLab (X-71) [15]	✓	✓	✓	67.0	42.5	83.1

Table 10. Panoptic-DeepLab with proposed WR-41 backbone on Cityscapes test set. MV: Mapillary Vistas pretrained.

Method	MV	PQ (%)	AP (%)	mIOU (%)
Panoptic-DeepLab (X-71) [15]		62.3	34.6	79.4
Panoptic-DeepLab (WR-41) (ours)		63.7	36.5	81.5
Panoptic-DeepLab (X-71) [15]	✓	65.5	39.0	84.2
Panoptic-DeepLab (WR-41) (ours)	✓	66.5	40.6	84.5

is slightly worse than the one reported in the original paper. We think there are some differences between the deep learning libraries. Note however that our main focus is on the segmentation results, while ImageNet is only used for pretraining. Our proposed WR-41 achieves a slightly better performance. Employing the drop path [32] with a constant survival probability 0.8 improves the performance.

Cityscapes val set: In Tab. 8, we report the Cityscapes validation set results when using Panoptic-DeepLab [15] with WR-38 (our TensorFlow re-

implementation) and WR-41 as backbones. As shown in the table, we observe (1) using drop path [32] (constant survival probability 0.8) consistently improves the performance in both backbones, (2) the performance could be further improved by adopting the multi-grid scheme proposed in [12] (where the unit rates in the last two or three residual blocks are set to (1, 2) or (1, 2, 4) for WR-38 and WR-41, respectively), (3) using WR-41 as backbone slightly improves over WR-38, and (4) Panoptic-DeepLab with WR-41 as backbone is slightly faster (and with slightly fewer parameters) than with WR-38 because the ASPP module [11] is added on the last feature map with 2048 channels (instead of 4096 channels). Additionally, the GPU inference times (Tesla V100-SXM2) on a 1025×2049 input for WR-38 and WR-41 are 437.9 ms and 396.5 ms, respectively.

In Tab. 9, we report the effect of using test-time augmentation (*i.e.*, multi-scale inputs and left-right flips) and pretraining on Mapillary Vistas, when using Panoptic-DeepLab with WR-41 as network backbone. The performance consistently improved with test-time augmentation and pretraining on Mapillary Vistas. Additionally, adopting Panoptic-DeepLab with WR-41 slightly outperforms Panoptic-DeepLab with X-71 as reported in [15].

Cityscapes test set: In Tab. 10, we report the Cityscapes test set results when using Panoptic-DeepLab [15] with our modified WR-41. Without extra data, our Panoptic-DeepLab (WR-41) outperforms Panoptic-DeepLab (X-71) by 1.4% PQ, 1.9% AP, and 2.1% mIOU. With Mapillary Vistas pretraining, our Panoptic-DeepLab (WR-41) outperforms Panoptic-DeepLab (X-71) by 1.0% PQ and 1.6% AP, and 0.3% mIOU.

5 Conclusion

In this work, we have described an iterative semi-supervised learning method that significantly improves the performance of urban scene segmentation on Cityscapes, simultaneously tackling semantic, instance, and panoptic segmentation. This semi-supervised learning procedure effectively harnesses both unlabeled video frames and extra unlabeled images to improve the predictive performance of the model without the creation of additional architectures and learned modules. Namely, pseudo-labeled data garnered through a simple data augmentation (*i.e.*, multi-scale inputs and left-right flips) suffices to boost performance on supervised learning tasks. As a result, our model sets the new state-of-art performance at all three Cityscapes benchmarks without the need to fine-tune or any special design on each task. We hope our simple yet effective learning scheme could establish a baseline procedure to harness the abundant unlabeled video sequences and extra images for computer vision tasks.

Acknowledgments We would like to thank the support from Google Mobile Vision and Brain.

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation (2016) [6](#)
2. Abu-El-Haija, S., Kothari, N., Lee, J., Natsev, P., Toderici, G., Varadarajan, B., Vijayanarasimhan, S.: YouTube-8M: A large-scale video classification benchmark. arXiv:1609.08675 (2016) [2](#)
3. Arazo, E., Ortego, D., Albert, P., O'Connor, N.E., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. arXiv:1908.02983 (2019) [3](#)
4. Badrinarayanan, V., Galasso, F., Cipolla, R.: Label propagation in video sequences. In: CVPR (2010) [2](#), [3](#)
5. Bell, S., Upchurch, P., Snavely, N., Bala, K.: OpenSurfaces: A richly annotated catalog of surface appearance. ACM Transactions on Graphics (2013) [2](#)
6. Budvytis, I., Sauer, P., Roddick, T., Breen, K., Cipolla, R.: Large scale labelled video data augmentation for semantic segmentation in driving scenarios. In: ICCV Workshop (2017) [2](#), [3](#)
7. Caba Heilbron, F., Escorcia, V., Ghanem, B., Carlos Niebles, J.: ActivityNet: A large-scale video benchmark for human activity understanding. In: CVPR (2015) [2](#)
8. Castrejon, L., Kundu, K., Urtasun, R., Fidler, S.: Annotating object instances with a polygon-RNN. In: CVPR (2017) [2](#)
9. Chen, L.C., Collins, M., Zhu, Y., Papandreou, G., Zoph, B., Schroff, F., Adam, H., Shlens, J.: Searching for efficient multi-scale architectures for dense image prediction. In: NeurIPS (2018) [2](#)
10. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected CRFs. In: ICLR (2015) [2](#)
11. 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 TPAMI (2017) [14](#)
12. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv:1706.05587 (2017) [2](#), [3](#), [6](#), [14](#)
13. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (2018) [4](#), [6](#)
14. Cheng, B., Collins, M.D., Zhu, Y., Liu, T., Huang, T.S., Adam, H., Chen, L.C.: Panoptic-DeepLab. In: ICCV COCO + Mapillary Joint Recognition Challenge Workshop (2019) [4](#)
15. Cheng, B., Collins, M.D., Zhu, Y., Liu, T., Huang, T.S., Adam, H., Chen, L.C.: Panoptic-DeepLab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In: CVPR (2020) [6](#), [7](#), [8](#), [11](#), [12](#), [13](#), [14](#)
16. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: CVPR (2017) [6](#), [7](#)
17. 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: CVPR (2016) [2](#), [6](#), [12](#)

18. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical data augmentation with no separate search. arXiv:1909.13719 (2019) 3, 4
19. Dai, J., He, K., Sun, J.: Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: ICCV (2015) 3
20. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV (2015) 3
21. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (VOC) challenge. IJCV **88**(2), 303–338 (2010) 2
22. Forsyth, D.A., Ponce, J.: Computer vision: a modern approach. Prentice Hall Professional Technical Reference (2002) 2
23. Gadde, R., Jampani, V., Gehler, P.V.: Semantic video CNNs through representation warping. In: ICCV (2017) 2, 3, 4
24. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. The International Journal of Robotics Research (2013) 2
25. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: CVPR (2018) 3
26. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR (2014) 2
27. Hariharan, B., Arbelaez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: ICCV (2011) 2
28. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) 7
29. Hénaff, O.J., Razavi, A., Doersch, C., Eslami, S., Oord, A.v.d.: Data-efficient image recognition with contrastive predictive coding. arXiv:1905.09272 (2019) 2
30. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) 2
31. Hong, S., Noh, H., Han, B.: Decoupled deep neural network for semi-supervised semantic segmentation. In: NeurIPS (2015) 3
32. Huang, G., Sun, Y., Liu, Z., Sedra, D., Weinberger, K.Q.: Deep networks with stochastic depth. In: ECCV (2016) 6, 13, 14
33. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: ICML (2015) 7
34. Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Label propagation for deep semi-supervised learning. In: CVPR (2019) 3
35. Khoreva, A., Benenson, R., Hosang, J., Hein, M., Schiele, B.: Simple does it: Weakly supervised instance and semantic segmentation. In: CVPR (2017) 3
36. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: ICLR (2015) 7
37. Kirillov, A., Girshick, R., He, K., Dollár, P.: Panoptic feature pyramid networks. In: CVPR (2019) 7
38. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: CVPR (2019) 7
39. Kornblith, S., Shlens, J., Le, Q.V.: Do better imagenet models transfer better? In: CVPR (2019) 2
40. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NeurIPS (2012) 2
41. Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J.: Building machines that learn and think like people. Behavioral and brain sciences (2017) 2
42. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: ICML Workshop (2013) 3

43. Li, J., Raventos, A., Bhargava, A., Tagawa, T., Gaidon, A.: Learning to fuse things and stuff. arXiv:1812.01192 (2018) 7
44. Li, L.J., Fei-Fei, L.: Optimol: automatic online picture collection via incremental model learning. IJCV (2010) 3
45. Li, Q., Arnab, A., Torr, P.H.: Weakly-and semi-supervised panoptic segmentation. In: ECCV (2018) 3
46. Li, Q., Qi, X., Torr, P.H.: Unifying training and inference for panoptic segmentation. arXiv:2001.04982 (2020) 7, 8
47. Li, Y., Qi, H., Dai, J., Ji, X., Wei, Y.: Fully convolutional instance-aware semantic segmentation. In: CVPR (2017) 2
48. Liang, J., Homayounfar, N., Ma, W.C., Xiong, Y., Hu, R., Urtasun, R.: Poly-transform: Deep polygon transformer for instance segmentation. arXiv:1912.02801 (2019) 7, 8
49. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014) 2
50. Liu, C., Chen, L.C., Schroff, F., Adam, H., Hua, W., Yuille, A., Fei-Fei, L.: Auto-DeepLab: Hierarchical neural architecture search for semantic image segmentation. In: CVPR (2019) 6
51. Liu, S., Qi, L., Qin, H., Shi, J., Jia, J.: Path aggregation network for instance segmentation. In: CVPR (2018) 7, 8
52. Liu, W., Rabinovich, A., Berg, A.C.: Parsenet: Looking wider to see better. arXiv:1506.04579 (2015) 7
53. Luc, P., Neverova, N., Couprie, C., Verbeek, J., LeCun, Y.: Predicting deeper into the future of semantic segmentation. In: ICCV (2017) 3
54. Mustikovela, S.K., Yang, M.Y., Rother, C.: Can ground truth label propagation from video help semantic segmentation? In: ECCV (2016) 2, 3, 4
55. Neuhold, G., Ollmann, T., Bulò, S.R., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV (2017) 2, 7, 8
56. Nilsson, D., Sminchisescu, C.: Semantic video segmentation by gated recurrent flow propagation. In: CVPR (2018) 2, 3, 4
57. Papandreou, G., Chen, L.C., Murphy, K.P., Yuille, A.L.: Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In: ICCV (2015) 3
58. Papandreou, G., Zhu, T., Chen, L.C., Gidaris, S., Tompson, J., Murphy, K.: PersonLab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In: ECCV (2018) 3
59. Pathak, D., Krahenbuhl, P., Darrell, T.: Constrained convolutional neural networks for weakly supervised segmentation. In: ICCV (2015) 3
60. Pinheiro, P.O., Collobert, R., Dollár, P.: Learning to segment object candidates. In: NeurIPS (2015) 3
61. Porzi, L., Bulò, S.R., Colovic, A., Kotschieder, P.: Seamless scene segmentation. In: CVPR (2019) 7, 8, 12
62. Porzi, L., Hofinger, M., Ruiz, I., Serrat, J., Bulò, S.R., Kotschieder, P.: Learning multi-object tracking and segmentation from automatic annotations. In: CVPR (2020) 3
63. Qi, H., Zhang, Z., Xiao, B., Hu, H., Cheng, B., Wei, Y., Dai, J.: Deformable convolutional networks – COCO detection and segmentation challenge 2017 entry. ICCV COCO Challenge Workshop (2017) 6, 7
64. Radosavovic, I., Dollár, P., Girshick, R., Gkioxari, G., He, K.: Data distillation: Towards omni-supervised learning. In: CVPR (2018) 3

65. Real, E., Shlens, J., Mazzocchi, S., Pan, X., Vanhoucke, V.: YouTube-BoundingBoxes: A large high-precision human-annotated data set for object detection in video. In: CVPR (2017) [2](#)
66. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: NeurIPS (2015) [2](#)
67. Riloff, E., Wiebe, J.: Learning extraction patterns for subjective expressions. In: EMNLP (2003) [3](#)
68. Rosenberg, C., Hebert, M., Schneiderman, H.: Semi-supervised self-training of object detection models. WACV/MOTION (2005) [3](#)
69. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. IJCV (2015) [2](#), [8](#), [12](#)
70. Russell, B.C., Torralba, A., Murphy, K.P., Freeman, W.T.: LabelMe: a database and web-based tool for image annotation. IJCV (2008) [2](#)
71. Scudder, H.: Probability of error of some adaptive pattern-recognition machines. IEEE Transactions on Information Theory (1965) [3](#)
72. Shi, W., Gong, Y., Ding, C., MaXiaoyu Tao, Z., Zheng, N.: Transductive semi-supervised deep learning using min-max features. In: ECCV (2018) [3](#)
73. Souly, N., Spampinato, C., Shah, M.: Semi supervised semantic segmentation using generative adversarial network. In: ICCV (2017) [3](#)
74. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: ICCV (2017) [2](#)
75. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. arXiv:1912.04838 (2019) [2](#)
76. Tang, Y., Wang, J., Gao, B., Dellandréa, E., Gaizauskas, R., Chen, L.: Large scale semi-supervised object detection using visual and semantic knowledge transfer. In: CVPR (2016) [3](#)
77. Voigtlaender, P., Krause, M., Osep, A., Luiten, J., Sekar, B.B.G., Geiger, A., Leibe, B.: Mots: Multi-object tracking and segmentation. In: CVPR (2019) [3](#)
78. Wang, H., Zhu, Y., Green, B., Adam, H., Yuille, A., Chen, L.C.: Axial-DeepLab: Stand-alone axial-attention for panoptic segmentation. arXiv:2003.07853 (2020) [7](#)
79. Wang, P., Chen, P., Yuan, Y., Liu, D., Huang, Z., Hou, X., Cottrell, G.: Understanding convolution for semantic segmentation. arXiv:1702.08502 (2017) [6](#)
80. Wei, Y., Liang, X., Chen, Y., Shen, X., Cheng, M.M., Feng, J., Zhao, Y., Yan, S.: Stc: A simple to complex framework for weakly-supervised semantic segmentation. IEEE TPAMI (2016) [3](#)
81. Wei, Y., Xiao, H., Shi, H., Jie, Z., Feng, J., Huang, T.S.: Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In: CVPR (2018) [3](#)
82. Wu, J., Yildirim, I., Lim, J.J., Freeman, B., Tenenbaum, J.: Galileo: Perceiving physical object properties by integrating a physics engine with deep learning. In: NeurIPS (2015) [2](#)
83. Wu, Z., Shen, C., Van Den Hengel, A.: Wider or deeper: Revisiting the ResNet model for visual recognition. Pattern Recognition (2019) [6](#), [7](#), [8](#), [12](#), [13](#)
84. Xie, Q., Hovy, E., Luong, M.T., Le, Q.V.: Self-training with noisy student improves imagenet classification. arXiv:1911.04252 (2019) [3](#), [4](#)
85. Xiong, Y., Liao, R., Zhao, H., Hu, R., Bai, M., Yumer, E., Urtasun, R.: UPSNet: A unified panoptic segmentation network. In: CVPR (2019) [7](#)
86. Yalniz, I.Z., J'egou, H., Chen, K., Paluri, M., Mahajan, D.: Billion-scale semi-supervised learning for image classification. arXiv 1905.00546 (2019) [3](#)

87. Yang, T.J., Collins, M.D., Zhu, Y., Hwang, J.J., Liu, T., Zhang, X., Sze, V., Papandreou, G., Chen, L.C.: DeeperLab: Single-shot image parser. arXiv:1902.05093 (2019) 7
88. Yarowsky, D.: Unsupervised word sense disambiguation rivaling supervised methods. In: ACL (1995) 3
89. Yuan, Y., Chen, X., Wang, J.: Object-contextual representations for semantic segmentation. arXiv:1909.11065 (2019) 7, 8
90. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: BMVC (2016) 6, 7, 12
91. Zhai, X., Oliver, A., Kolesnikov, A., Beyer, L.: S4l: Self-supervised semi-supervised learning. In: ICCV (2019) 3
92. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR (2017) 2, 3
93. Zheng, Z., Zheng, L., Yang, Y.: Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In: ICCV (2017) 3
94. Zhu, X., Xiong, Y., Dai, J., Yuan, L., Wei, Y.: Deep feature flow for video recognition. In: CVPR (2017) 3
95. 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: CVPR (2019) 2, 3, 4, 7, 8, 12
96. Zhu, Y., Zhang, Z., Wu, C., Zhang, Z., He, T., Zhang, H., Manmatha, R., Li, M., Smola, A.: Improving semantic segmentation via self-training. arXiv:2004.14960 (2020) 3