

Self-supervised Augmentation Consistency for Adapting Semantic Segmentation

Nikita Araslanov¹Stefan Roth^{1,2}¹Department of Computer Science, TU Darmstadt² hessian.AI

Abstract

We propose an approach to domain adaptation for semantic segmentation that is both practical and highly accurate. In contrast to previous work, we abandon the use of computationally involved adversarial objectives, network ensembles and style transfer. Instead, we employ standard data augmentation techniques – photometric noise, flipping and scaling – and ensure consistency of the semantic predictions across these image transformations. We develop this principle in a lightweight self-supervised framework trained on co-evolving pseudo labels without the need for cumbersome extra training rounds. Simple in training from a practitioner’s standpoint, our approach is remarkably effective. We achieve significant improvements of the state-of-the-art segmentation accuracy after adaptation, consistent both across different choices of the backbone architecture and adaptation scenarios.

1. Introduction

Unsupervised domain adaptation (UDA) is a variant of semi-supervised learning [6], where the available unlabelled data comes from a different distribution than the annotated dataset [4]. A case in point is to exploit synthetic data, where annotation is more accessible compared to the costly labelling of real-world images [59, 60]. Along with some success in addressing UDA for semantic segmentation [67, 69, 80, 91], the developed methods are growing increasingly sophisticated and often combine style transfer networks, adversarial training or network ensembles [39, 46, 68, 77]. This increase in model complexity impedes reproducibility, potentially slowing further progress.

In this work, we propose a UDA framework reaching state-of-the-art segmentation accuracy (measured by the Intersection-over-Union, IoU) without incurring substantial training efforts. Toward this goal, we adopt a simple semi-supervised approach, *self-training* [12, 42, 91], used in recent works only in conjunction with adversarial training or

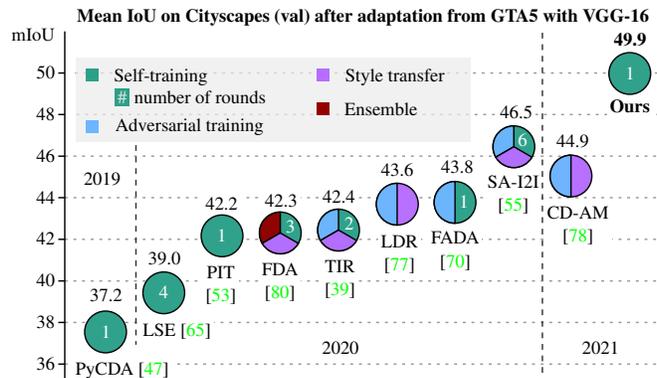


Figure 1. **Results preview.** Unlike much recent work that combines multiple training paradigms, such as adversarial training and style transfer, our approach retains the modest single-round training complexity of self-training, yet improves the state of the art for adapting semantic segmentation by a significant margin.

network ensembles [17, 39, 54, 70, 80, 86, 87]. By contrast, we use self-training *standalone*. Compared to previous self-training methods [9, 43, 65, 91, 92], our approach also sidesteps the inconvenience of multiple training rounds, as they often require expert intervention between consecutive rounds. We train our model using co-evolving pseudo labels end-to-end without such need.

Our method leverages the ubiquitous *data augmentation* techniques from fully supervised learning [11, 85]: photometric jitter, flipping and multi-scale cropping. We enforce *consistency* of the semantic maps produced by the model across these image perturbations. The following assumption formalises the key premise:

Assumption 1. Let $f : \mathcal{I} \rightarrow \mathcal{M}$ represent a pixelwise mapping from images \mathcal{I} to semantic output \mathcal{M} . Denote $\rho_\epsilon : \mathcal{I} \rightarrow \mathcal{I}$ a photometric image transform and, similarly, $\tau_{\epsilon'} : \mathcal{I} \rightarrow \mathcal{I}$ a spatial similarity transformation, where $\epsilon, \epsilon' \sim p(\cdot)$ are control variables following some predefined density (e.g., $p \equiv \mathcal{N}(0, 1)$). Then, for any image $I \in \mathcal{I}$, f is *invariant* under ρ_ϵ and *equivariant* under $\tau_{\epsilon'}$, i.e. $f(\rho_\epsilon(I)) = f(I)$ and $f(\tau_{\epsilon'}(I)) = \tau_{\epsilon'}(f(I))$.

Next, we introduce a training framework using a *momentum network* – a slowly advancing copy of the original model.

Code is available at <https://github.com/visinf/da-sac>.

The momentum network provides stable, yet recent targets for model updates, as opposed to the fixed supervision in model distillation [15, 86, 87]. We also re-visit the problem of long-tail recognition in the context of generating pseudo labels for self-supervision. In particular, we maintain an *exponentially moving class prior* used to discount the confidence thresholds for those classes with few samples and increase their relative contribution to the training loss. Our framework is simple to train, adds moderate computational overhead compared to a fully supervised setup, yet sets a new state of the art on established benchmarks (*cf.* Fig. 1).

2. Related Work

Most of the work on scene adaptation for semantic segmentation has been influenced by a parallel stream of work on domain adaptation (DA) and semi-supervised learning for image classification [23, 24, 27, 45, 50]. The main idea behind these methods is to formulate an upper bound on the target risk using the so-called $\mathcal{H}\Delta\mathcal{H}$ -divergence [3]. In a nutshell, it defines the discrepancy between the marginals of the source and target data by means of a binary classifier. In the following, we briefly review implementation variants of this idea in the context of semantic segmentation.

Learning domain-invariant representations. Adversarial feature alignment follows the GAN framework [24, 26] and minimises the gap between the source and target feature representations in terms of some distance (*e.g.*, Wasserstein in [41]). The discriminator can be employed at multiple scales [15, 67, 77] and use local spatial priors [83]; it can be conditional [33] and class-specific [22, 52], or align the features of ‘hard’ and ‘easy’ target samples [56]. Often, self-supervised losses, such as entropy minimisation [69], or a ‘conservative loss’ [90] assist in this alignment.

The alternative to adversarial feature alignment are more interpretable constraints, such as feature priors [51], bijective source-target association [37] or aligning the domains directly in the image space with style transfer [89] used either alone [74] or, most commonly, jointly with adversarial feature alignment [8, 16, 25, 55, 78, 79, 82]. One issue with style translation is to ensure semantic consistency despite the changes in appearance. To address this, Hoffman *et al.* [32] use semantic and cycle-consistency losses, while Yang *et al.* [77] reconstruct the original image from its label-space representation.

These methods tend to be computationally costly and challenging to train, since they require concurrent training of one or more independent networks, *e.g.* discriminators or style transfer networks. Although Yang and Soatto [80] obviate the need for style networks by incorporating the phase of a Fourier-transformed target image into a source sample, multiple networks have to be trained, each with its own pre-defined phase band.

Features	PIT [53]	LDR [77]	SA-I2I [55]	IAST [54]	RPT [83]	Ours
Adversarial training		✓	✓	✓	✓	
1-round training	✓	✓	(6)	(3)	(3)	✓
SOTA-VGG		✓	✓			✓
SOTA-ResNet				✓	✓	✓

Table 1. **Relation to state of the art.** Previous work reaches the state of the art in terms of IoU either with VGG-16 (SOTA-VGG) or ResNet-101 (SOTA-ResNet). Our framework uses neither adversarial training nor multiple training rounds (given in parentheses), yet outperforms the state of the art consistently in both cases.

Self-training on pseudo labels. As a more computationally lightweight approach, self-training seeks high-quality pseudo supervision coming in the form of class predictions with high confidence. Our work belongs to this category. Most of such previous methods pre-compute the labels ‘offline’, used subsequently to update the model, and repeat this process for several rounds [43, 65, 91, 92]. More recent frameworks following this strategy have a composite nature: they rely on adversarial (pre-)training [14, 20, 87], style translation [17, 80] or both [54, 46, 39, 70, 73].

Training on co-evolving pseudo labels can be computationally unstable, hence requires additional regularisation. Chen *et al.* [13] minimise the entropy with improved behaviour of the gradient near the saturation points. Using fixed representations, be it from a ‘frozen’ network [15, 87], a fixed set of global [53] or self-generated local labels [47, 68, 81], further improves training robustness.

Overconfident predictions [28] have direct consequences for the quality of pseudo labels. Zou *et al.* [92] attain some degree of confidence calibration via regularising the loss with prediction smoothing akin to temperature scaling [28]. Averaging the predictions of two classifiers [86], or using Dropout-based sampling [7, 88], achieves the same goal.

Spatial priors. Different from DA for classification, the characteristic feature of adaptation methods for segmentation is the use of spatial priors. Local priors have been enforced patch-wise [15, 47, 68] and in the form of pre-computed super-pixels [81, 83]. Although global spatial priors have also been used [91], their success hinges on the similarity of the semantic layout in the current benchmarks.

Relation to our approach. As shown in Table 1, our work streamlines the training process. First, we do not use adversarial training, as feature invariance alone does not guarantee *label* invariance [36, 84]. Second, we train our model with co-evolving pseudo labels in one round. Our framework bears resemblance to the noisy mean teacher [76] and combines consistency regularisation [2, 61, 64, 75] with self-ensembling [40, 66]. Similar approaches have been explored in medical imaging [44, 58] and concurrent UDA work [71], albeit limited in the scope of admissible aug-

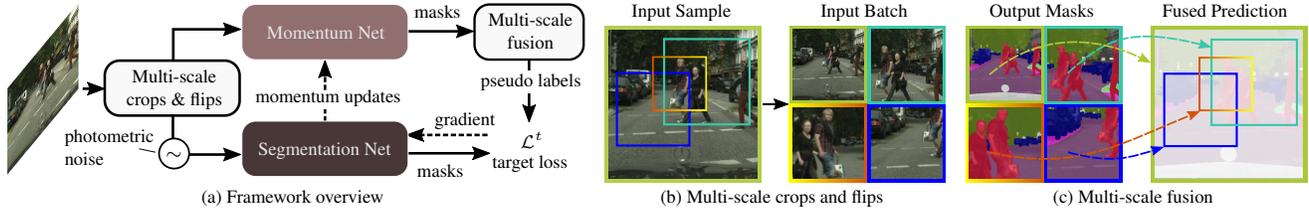


Figure 2. **Overview.** The segmentation network in our framework (a) maintains a slow copy of itself, the momentum network, which provides stable targets for self-supervision. In addition to encouraging semantic invariance w.r.t. the photometric noise, we facilitate consistent predictions across multiple scales and flips by first (b) feeding random multi-scale crops and flips to the momentum network and then (c) fusing the predictions by simple averaging to produce the pseudo-supervision targets.

mentations. We leverage photometric invariance, scale and flip equivariance [72] to extract high-fidelity pseudo supervision instead of more computationally expensive sampling techniques [38]. Contrary to [65], we find that scale alone is not predictive of the label quality, hence we average the predictions produced at multiple scales and flips. This parallels uncertainty estimation using test-time augmentation [1], but at training time [5].

3. Self-Supervised Augmentation Consistency

3.1. Framework overview

Shown in Fig. 2a, our framework comprises a segmentation network, which we intend to adapt to a target domain, and its slowly changing copy updated with a momentum, a *momentum network*. To perform self-supervised scene adaptation, we first supply a batch of random crops and horizontal flips from a sample image of the target domain to both networks. For each pixel we average the predictions (*i.e.* semantic masks) from the momentum network after the appropriate inverse spatial transformation. We then create a pseudo ground truth by selecting confident pixels from the averaged map using thresholds based on running statistics, which are capable of adapting to individual samples. Finally, the segmentation network uses stochastic gradient descent to update its parameters w.r.t. these pseudo labels.

Our approach closely resembles the mean teacher framework [23, 66] and temporal ensembling [35, 40]. However, as we will show empirically, the ensembling property itself plays only an auxiliary role. More importantly, akin to the critic network in reinforcement learning [48] and the momentum encoder in unsupervised learning [30], our momentum network provides stable targets for self-supervised training of the segmentation network. This view allows us to focus on the target-generating process, detailed next.

3.2. Batch construction

For each sampled target image, we generate N crops with random scales, flips and locations, but preserving the aspect ratio. We re-scale the crops as well as the original image to a fixed input resolution $h \times w$ and pass them as the in-

put to the networks. Fig. 2b demonstrates this process. Following the noisy student model in image classification [76], the input to the segmentation network additionally undergoes a photometric augmentation: we add random colour jitter and smooth the images with a Gaussian filter at random. The momentum network, on the other hand, receives a ‘clean’ input, *i.e.* without such augmentations. This is to encourage model invariance to photometric perturbations.

3.3. Self-supervision

Multi-scale fusion. We re-project the output masks from the momentum network back to the original image canvas of size $h \times w$, as illustrated in Fig. 2c. For each pixel, the overlapping areas average their predictions. Note that some pixels may lie outside the crops, hence contain the result of a single forward pass with the original image. We keep these predictions intact. The merged maps are then used to extract the pseudo masks for self-supervision.

A short long-tail interlude. Handling rare classes (*i.e.* classes with only a few training samples) is notoriously difficult in recognition [29]. For semantic segmentation, we here distinguish between the classes with low image-level (*e.g.*, “truck”, “bus”) and pixel-level (*e.g.*, “traffic light”, “pole”) frequency. While generating self-supervision, we take special care of these cases and encourage (i) lower thresholds for selecting their pseudo labels, (ii) increased contributions to the gradient with a focal loss, and (iii) employ importance sampling. We describe these in detail next.

Sample-based moving threshold. Most previous work with self-training employs multi-round training that requires interrupting the training process and re-generating the pseudo labels [43, 46, 54, 65, 91]. One of the reasons is the need to re-compute the thresholds for filtering the pseudo labels for supervision, which requires traversing the predictions for the complete target dataset with the model parameters fixed. In pursuit of our goal of enabling end-to-end training without expert intervention, we take a different approach and compute the thresholds on-the-go. As the main ingredient, we maintain an *exponentially moving class prior*. In detail, for each softmax prediction of the momentum network, we first compute a prior estimate of

the probability that a pixel in sample n belongs to class c as

$$\chi_{c,n} = \frac{1}{hw} \sum_{i,j} m_{c,n,i,j}, \quad (1)$$

where $m_{c,n,i,j}$ is the mask prediction for class c (with resolution $h \times w$). We keep an exponentially moving average after each training iteration t with a momentum $\gamma_\chi \in [0, 1]$:

$$\chi_c^{t+1} = \gamma_\chi \chi_c^t + (1 - \gamma_\chi) \chi_{c,n}. \quad (2)$$

Our sample-based moving threshold $\theta_{c,n}$ takes lower values when the moving prior $\chi_c \approx 0$ (*i.e.* for long-tail classes), but is bounded from above as $\chi_c \rightarrow 1$. We define it as

$$\theta_{c,n} = \zeta (1 - e^{-\chi_c/\beta}) m_{c,n}^*, \quad (3)$$

where β and ζ are hyperparameters and $m_{c,n}^*$ is the predicted peak confidence score for class c , *i.e.*

$$m_{c,n}^* = \max_{i,j} m_{c,n,i,j}. \quad (4)$$

Fig. 3 plots Eq. (3) as a function of the moving class prior χ_c for a selection of β . For predominant classes (*e.g.*, “road”), the exponential term has nearly no effect; the threshold is static w.r.t. the peak class confidence, *i.e.* $\theta_{c,n} \approx \zeta m_{c,n}^*$. However, for long-tail classes such that $\chi_c \approx \beta$, the threshold is lower than this upper bound, hence more pixels for these classes are selected for supervision. To obtain the pseudo labels, we apply the threshold $\theta_{c,n}$ to the peak predictions of the merged output from the momentum network:

$$\hat{m}_{n,i,j} = \begin{cases} c^* & m_{c^*,n,i,j} > \theta_{c,n} \\ \text{ignore} & \text{otherwise,} \end{cases} \quad (5)$$

where $c^* = \arg \max_c m_{c,n,i,j}$ is the dominant class for that pixel. Note that the pixels with confidence values lower than the threshold, as well as non-dominant predictions, will be ignored in the self-supervised loss.

Focal loss with confidence regularisation. Our loss function incorporates a focal multiplier [49] to further increase the contribution of the long-tail classes in the gradient signal. Unlike previous work [49, 65], however, our moving class prior χ_c regulates the focal term:

$$\mathcal{L}_n^t(\bar{m}, m \mid \phi) = -m_{c^*,n} (1 - \chi_{c^*})^\lambda \log(\bar{m}_{c^*,n}), \quad (6)$$

where \bar{m} is the prediction of the segmentation network with parameters ϕ , the pseudo label c^* derives from \hat{m} in Eq. (5) and λ is a hyperparameter of the focal term. Recall that low values of χ_c signify a long-tail category, hence should have a higher weight. High values of λ (*i.e.* > 1) increase the relative weighting on the long-tail classes, while setting $\lambda = 0$ disables the focal term. Note that we also regularise our loss with the confidence value of the momentum network, $m_{c^*,n}$ (Eq. 4). In case of an incorrect pseudo label, we expect this confidence to be low and to regularise the training owing to its calibration with the multi-scale fusion. We minimise the loss in Eq. (6), applied for each pixel, w.r.t. ϕ .

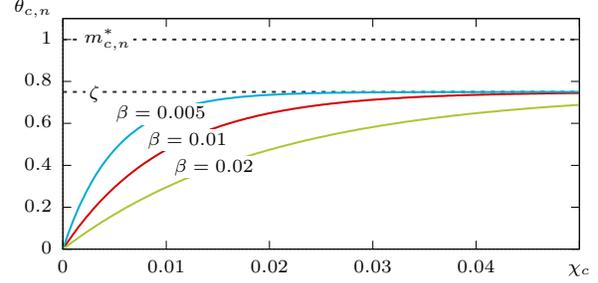


Figure 3. **Sample-based moving threshold.** Our thresholding scheme has two hyperparameters, ζ and β . In this example, $m_{c,n}^* = 1$ and $\zeta = 0.75$. Predominant classes (*e.g.*, “road”) have $\chi_c \gg 0$, hence their threshold approximates $\zeta m_{c,n}^*$. Long-tail classes (*e.g.*, “traffic light”) have $\chi_c \approx 0$ and their thresholds are further reduced with a steepness controlled by β (see Eq. 3).

3.4. Training

Pre-training with source-only loss. Following [47, 83], we use Adaptive Batch Normalisation (ABN) [45] to jump-start our model on the segmentation task by minimising the cross-entropy loss on the source data only. In our experiments, we found it unnecessary to re-compute the mean and the standard deviation only at the end of the training. Instead, in pre-training we alternate batches of source and target images, but ignore the loss for the latter. For a target batch, this implies updating the running mean and the standard deviation in the Batch Normalisation (BN) [34] layers and leaving the remaining model parameters untouched.

Importance sampling. Our loss function in Eq. (6) accounts for long-tail classes with a high image frequency (*e.g.*, “traffic light”, “pole”), and may not be effective for the classes appearing in only few samples (*e.g.*, “bus”, “train”). To alleviate this imbalance, we use importance sampling [21] and increase the sample frequency of these long-tail classes. We minimise the *expected* target loss by re-sampling the target images using the density p_t :

$$\min_{\phi} \mathbb{E}_{n \sim p_t} [\mathcal{L}_n^t(\phi)]. \quad (7)$$

To obtain p_t , we use our pre-trained segmentation network and pre-compute $\chi_{c,n}$, the class prior estimate, for each image n using Eq. (1). At training time, we (*i*) sample a semantic class c uniformly, and then (*ii*) obtain a target sample l with probability

$$\hat{\chi}_{c,l} = \frac{\chi_{c,l}}{\sum_n \chi_{c,n}}. \quad (8)$$

This two-step sampling process ensures that all images have non-zero sample probability owing to the prevalent classes for which $\hat{\chi}_{c,l} > 0$ for all l (*e.g.*, “road” in urban scenes).

Joint target-source training. We train the segmentation network with stochastic gradient descent using the cross-entropy loss for the source and our focal loss for the target

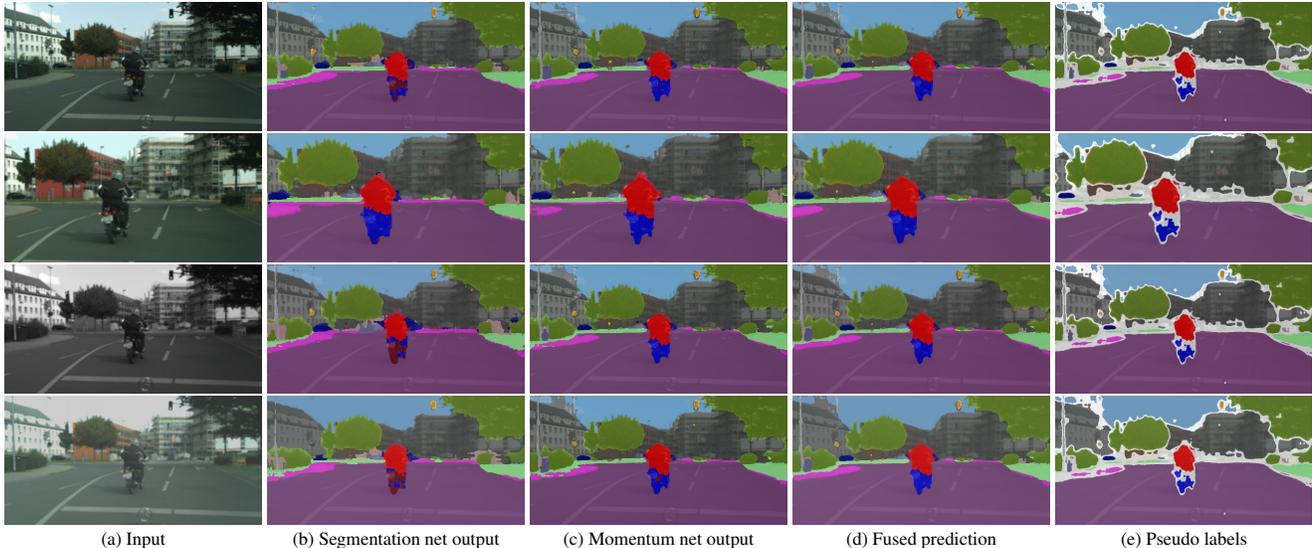


Figure 4. **Self-supervision example.** In this image sample (a) and its crops, the segmentation network (b) tends to mistake the “motorcycle” for a “bicycle”. The momentum network (c) improves on this prediction, but may still produce an inconsistent labelling. Averaging the predictions over multiple scales (d) corrects this inconsistency, allowing to produce high-precision pseudo labels (e) for self-supervision.

data sampled from p_t , as defined by Eqs. (6) and (7). Fig. 4 illustrates the synthesis of pseudo labels. We periodically update the parameters ψ of the momentum network as

$$\psi_{t+1} = \gamma_\psi \psi_t + (1 - \gamma_\psi) \phi, \quad (9)$$

where ϕ are the parameters of the segmentation network. γ_ψ regulates the pace of the updates: low values result in faster, but unstable training, while high γ_ψ leads to a premature and suboptimal convergence. We keep γ_ψ moderate, but update the momentum network only every T iterations.

4. Experiments

Datasets. In our experiments we use three datasets. The Cityscapes dataset [18] contains 2048×1024 images from real-world traffic scenes, split into 2975 images for training and 500 for validation. The GTA5 dataset [59] contains 24 966 synthetic scenes with resolution 1914×1052 and pixelwise annotation aided by the GTA5 game engine. We also use the SYNTHIA-RAND-CITYSCAPES subset of the SYNTHIA dataset [60], which contains 9400 synthetic images with resolution 1280×760 and a semantic annotation compatible with Cityscapes.

Setup. We adopt the established evaluation protocol from previous work [47, 67, 69]. The synthetic traffic scenes from GTA5 [59] and SYNTHIA [60] serve as the source data, and the real images from the Cityscapes dataset as the target (obviously ignoring the available semantic labels). This results in two domain adaptation scenarios depending on the choice of the source data: GTA5 \rightarrow Cityscapes and SYNTHIA \rightarrow Cityscapes. As in previous work, at training time we only use the training split of the Cityscapes dataset

and report the results on the validation split. We measure the segmentation accuracy with per-class Intersection-over-Union (IoU) and its average, the mean IoU (mIoU).

4.1. Implementation details

We implement our framework in PyTorch [57]. We adopt DeepLabv2 [10] as the segmentation architecture, and evaluate our method with two backbones, ResNet-101 [31] and VGG16 [63], following recent work [39, 67, 68, 69, 73]. Both backbones initialise from the models pre-trained on ImageNet [19]. We first train the models with ABN [45] (cf. Sec. 3.4), implemented via SyncBN [57], on multi-scale crops resized to 640×640 and a batch size of 16. Next, training proceeds with the self-supervised target loss (cf. Sec. 3.3) and the BatchNorm layers [34] frozen. The batch size of 16 comprises 8 source images and 8 target images at resolution 1024×512 , which is a common practice [70, 80]. The target batch contains only two image samples along with 3 random crops each (*i.e.* $N = 3$ in Sec. 3.2), downsampled up to a factor of 0.5. As the photometric noise, we use colour jitter, random blur and greyscaling (see Appendix B for details). The optimisation uses SGD with a constant learning rate of 2.5×10^{-4} , momentum 0.9 and weight decay of 5×10^{-4} . We accumulate the gradient in alternating source-target forward passes to keep the memory footprint in check. Since the focal term in Eq. (6) reduces the target loss magnitude w.r.t. the source loss, we scale it up by a factor of 5 (2 for VGG-16). We train our VGG-based framework on two TITAN X GPUs (12GB), while the ResNet-based variant requires four. This is a substantially reduced requirement compared to recent work (*e.g.*,

Method	road	sidew	build	wall	fence	pole	light	sign	veg	terr	sky	pers	ride	car	truck	bus	train	moto	bicy	mIoU
<i>Backbone: VGG-16</i>																				
CyCADA [32]	85.2	37.2	76.5	21.8	15.0	23.8	22.9	21.5	80.5	31.3	60.7	50.5	9.0	76.9	17.1	28.2	4.5	9.8	0.0	35.4
ADVENT [69]	86.9	28.7	78.7	28.5	25.2	17.1	20.3	10.9	80.0	26.4	70.2	47.1	8.4	81.5	26.0	17.2	18.9	11.7	1.6	36.1
CBST [91]	90.4	50.8	72.0	18.3	9.5	27.2	8.6	14.1	82.4	25.1	70.8	42.6	14.5	76.9	5.9	12.5	1.2	14.0	28.6	36.1
PyCDA [47]	86.7	24.8	80.9	21.4	27.3	30.2	26.6	21.1	86.6	28.9	58.8	53.2	17.9	80.4	18.8	22.4	4.1	9.7	6.2	37.2
PIT [53]	86.2	35.0	82.1	31.1	22.1	23.2	29.4	28.5	79.3	31.8	81.9	52.1	23.2	80.4	29.5	26.9	30.7	20.5	1.2	41.8
FDA [80]	86.1	35.1	80.6	30.8	20.4	27.5	30.0	26.0	82.1	30.3	73.6	52.5	21.7	81.7	24.0	30.5	29.9	14.6	24.0	42.2
LDR [77]	90.1	41.2	82.2	30.3	21.3	18.3	33.5	23.0	84.1	37.5	81.4	54.2	24.3	83.0	27.6	32.0	8.1	29.7	26.9	43.6
FADA [70]	92.3	51.1	83.7	33.1	29.1	28.5	28.0	21.0	82.6	32.6	85.3	55.2	28.8	83.5	24.4	37.4	0.0	21.1	15.2	43.8
CD-AM [78]	90.1	46.7	82.7	34.2	25.3	21.3	33.0	22.0	84.4	41.4	78.9	55.5	25.8	83.1	24.9	31.4	20.6	25.2	27.8	44.9
SA-I2I [55]	91.1	46.4	82.9	33.2	27.9	20.6	29.0	28.2	84.5	40.9	82.3	52.4	24.4	81.2	21.8	44.8	31.5	26.5	33.7	46.5
Baseline (ours)	81.5	28.6	79.5	23.2	21.1	31.3	28.2	18.5	75.6	14.9	72.2	58.0	17.1	81.1	19.7	26.3	13.7	12.9	2.1	37.1
SAC (ours)	90.0	53.1	86.2	33.8	32.7	38.2	46.0	40.3	84.2	26.4	88.4	65.8	28.0	85.6	40.6	52.9	17.3	13.7	23.8	49.9
<i>Backbone: ResNet-101</i>																				
PyCDA [†] [47]	90.5	36.3	84.4	32.4	28.7	34.6	36.4	31.5	86.8	37.9	78.5	62.3	21.5	85.6	27.9	34.8	18.0	22.9	49.3	47.4
CD-AM [78]	91.3	46.0	84.5	34.4	29.7	32.6	35.8	36.4	84.5	43.2	83.0	60.0	32.2	83.2	35.0	46.7	0.0	33.7	42.2	49.2
FADA [70]	92.5	47.5	85.1	37.6	32.8	33.4	33.8	18.4	85.3	37.7	83.5	63.2	39.7	87.5	32.9	47.8	1.6	34.9	39.5	49.2
LDR [77]	90.8	41.4	84.7	35.1	27.5	31.2	38.0	32.8	85.6	42.1	84.9	59.6	34.4	85.0	42.8	52.7	3.4	30.9	38.1	49.5
FDA [80]	92.5	53.3	82.4	26.5	27.6	36.4	40.6	38.9	82.3	39.8	78.0	62.6	34.4	84.9	34.1	53.1	16.9	27.7	46.4	50.5
SA-I2I [55]	91.2	43.3	85.2	38.6	25.9	34.7	41.3	41.0	85.5	46.0	86.5	61.7	33.8	85.5	34.4	48.7	0.0	36.1	37.8	50.4
PIT [53]	87.5	43.4	78.8	31.2	30.2	36.3	39.9	42.0	79.2	37.1	79.3	65.4	37.5	83.2	46.0	45.6	25.7	23.5	49.9	50.6
IAST [54]	93.8	57.8	85.1	39.5	26.7	26.2	43.1	34.7	84.9	32.9	88.0	62.6	29.0	87.3	39.2	49.6	23.2	34.7	39.6	51.5
RPT [†] [83]	89.2	43.3	86.1	39.5	29.9	40.2	49.6	33.1	87.4	38.5	86.0	64.4	25.1	88.5	36.6	45.8	23.9	36.5	56.8	52.6
Baseline (ours)	80.2	29.3	76.8	23.8	21.9	37.7	35.4	21.1	79.8	21.3	75.0	59.5	17.5	83.5	22.4	33.4	13.0	30.7	12.3	40.8
SAC (ours)	90.4	53.9	86.6	42.4	45.1	48.5	42.7	87.4	40.1	86.1	67.5	29.7	88.5	49.1	54.6	9.8	26.6	45.3	53.8	

([†]) denotes the use of PSPNet [85] instead of DeepLabv2 [10].

Table 2. **Per-class IoU (%) comparison** on GTA5 \rightarrow Cityscapes adaptation, evaluated on the Cityscapes validation set.

FADA [70] requires 4 Tesla P40 GPUs with 24GB memory). Note that the momentum network is always in evaluation mode, has gradient tracking disabled, hence adds only around 35% memory overhead. For the momentum network, we fix $\gamma_\psi = 0.99$ and $T = 100$ in all our experiments. For the other hyperparameters, we use $\gamma_\chi = 0.99$, $\zeta = 0.75$, $\beta = 10^{-3}$ and $\lambda = 3$. Appendix C.2 provides further detail on hyperparameter selection, as well as a sensitivity analysis of our framework w.r.t. ζ and β . The inference follows the usual procedure of a single forward pass through the segmentation network at the original image resolution without any post-processing.

4.2. Comparison to state of the art

We compare our approach to the current state of the art on the two domain adaptation scenarios: GTA5 \rightarrow Cityscapes in Table 2 and SYNTHIA \rightarrow Cityscapes in Table 3. For a fair comparison, all numbers originate from single-scale inference. In both cases, our approach, denoted as SAC (“Self-supervised Augmentation Consistency”), substantially outperforms our baseline (*i.e.* the source-only loss model with ABN, see Sec. 3.4), and, in fact, sets a new state of the art in terms of mIoU. Importantly, while the ranking of previous works depends on the backbone choice and the source data, we reach the top rank consistently in all settings.

GTA5 \rightarrow Cityscapes (Table 2). Our method achieves a clear improvement over the best published results [55, 83] of +3.4% and +1.2% using the VGG-16 and ResNet-101 backbones, respectively. Note that RPT [83] and SA-I2I [55] have a substantially higher model complexity. RPT [83] uses PSPNet [85], which has a higher upper bound than DeepLabv2 in a fully supervised setup (*e.g.*, +5.7% IoU on PASCAL VOC [85]); it requires extracting superpixels and training an encoder-decoder LSTM, thus increasing the model capacity and the computational overhead. SA-I2I [55] initialises from a stronger baseline, BDL [46], and relies on a style transfer network and adversarial training. While both RPT [83] and SA-I2I [55] require multiple rounds of training, 3 and 6 (from BDL [46]), respectively, we train with the target loss in a single pass. Notably, compared to the previous best approach for VGG with a ResNet evaluation, SA-I2I [55], our improvement with ResNet-101 is substantial, +3.4%, and is comparable to the respective margin on VGG-16.

SYNTHIA \rightarrow Cityscapes (Table 3). Here, the result is consistent with the previous scenario. Our approach attains state-of-the-art accuracy for both backbones, improving by 7.6% and 1.4% with VGG-16 and ResNet-101 backbones over the best results previously published [55, 83]. Again, our method with ResNet-101 outperforms the previous best method with full evaluation, PyCDA [47], by 5.9% IoU.

Method	road	sidew	build	wall	fence	pole	light	sign	veg	sky	pers	ride	car	bus	moto	bicy	mIoU ¹³	mIoU
<i>Backbone: VGG-16</i>																		
PyCDA [47]	80.6	26.6	74.5	2.0	0.1	18.1	13.7	14.2	80.8	71.0	48.0	19.0	72.3	22.5	12.1	18.1	—	35.9
PIT [53]	81.7	26.9	78.4	6.3	0.2	19.8	13.4	17.4	76.7	74.1	47.5	22.4	76.0	21.7	19.6	27.7	—	38.1
FADA [70]	80.4	35.9	80.9	2.5	0.3	30.4	7.9	22.3	81.8	83.6	48.9	16.8	77.7	31.1	13.5	17.9	—	39.5
FDA [80]	84.2	35.1	78.0	6.1	0.4	27.0	8.5	22.1	77.2	79.6	55.5	19.9	74.8	24.9	14.3	40.7	—	40.5
CD-AM [78]	73.0	31.1	77.1	0.2	0.5	27.0	11.3	27.4	81.2	81.0	59.0	25.6	75.0	26.3	10.1	47.4	—	40.8
LDR [77]	73.7	29.6	77.6	1.0	0.4	26.0	14.7	26.6	80.6	81.8	57.2	24.5	76.1	27.6	13.6	46.6	—	41.1
SA-I2I [55]	79.1	34.0	78.3	0.3	0.6	26.7	15.9	29.5	81.0	81.1	55.5	21.9	77.2	23.5	11.8	47.5	—	41.5
Baseline (ours)	60.7	26.9	67.1	8.3	0.0	33.5	11.9	18.3	66.4	70.4	52.1	16.1	64.6	15.5	11.5	26.4	39.1	34.4
SAC (ours)	77.9	38.6	83.5	15.8	1.5	38.2	41.3	27.9	80.8	83.0	64.3	21.2	78.3	38.5	32.6	62.1	56.2	49.1
<i>Backbone: ResNet-101</i>																		
ADVENT [69]	85.6	42.2	79.7	8.7	0.4	25.9	5.4	8.1	80.4	84.1	57.9	23.8	73.3	36.4	14.2	33.0	—	41.2
PIT [53]	83.1	27.6	81.5	8.9	0.3	21.8	26.4	33.8	76.4	78.8	64.2	27.6	79.6	31.2	31.0	31.3	—	44.0
PyCDA [†] [47]	75.5	30.9	83.3	20.8	0.7	32.7	27.3	33.5	84.7	85.0	64.1	25.4	85.0	45.2	21.2	32.0	—	46.7
CD-AM [78]	82.5	42.2	81.3	—	—	—	18.3	15.9	80.6	83.5	61.4	33.2	72.9	39.3	26.6	43.9	52.4	—
FDA [80]	79.3	35.0	73.2	—	—	—	19.9	24.0	61.7	82.6	61.4	31.1	83.9	40.8	38.4	51.1	52.5	—
LDR [77]	85.1	44.5	81.0	—	—	—	16.4	15.2	80.1	84.8	59.4	31.9	73.2	41.0	32.6	44.7	53.1	—
SA-I2I [55]	87.7	49.7	81.6	—	—	—	19.3	18.5	81.1	83.7	58.7	31.8	73.3	47.9	37.1	45.7	55.1	—
FADA [70]	84.5	40.1	83.1	4.8	0.0	34.3	20.1	27.2	84.8	84.0	53.5	22.6	85.4	43.7	26.8	27.8	—	45.2
IAST [54]	81.9	41.5	83.3	17.7	4.6	32.3	30.9	28.8	83.4	85.0	65.5	30.8	86.5	38.2	33.1	52.7	—	49.8
RPT [†] [83]	88.9	46.5	84.5	15.1	0.5	38.5	39.5	30.1	85.9	85.8	59.8	26.1	88.1	46.8	27.7	56.1	—	51.2
Baseline (ours)	63.9	25.9	71.0	11.0	0.2	36.9	7.6	20.0	72.9	75.5	46.7	16.7	74.5	15.8	20.8	21.7	41.0	36.3
SAC (ours)	89.3	47.2	85.5	26.5	1.3	43.0	45.5	32.0	87.1	89.3	63.6	25.4	86.9	35.6	30.4	53.0	59.3	52.6

([†]) denotes the use of PSPNet [85] instead of DeepLabv2 [10]. mIoU¹³ is the average IoU over 13 classes (*i.e.* excluding “wall”, “fence” and “pole”).

Table 3. **Per-class IoU (%) comparison** on SYNTHIA → Cityscapes adaptation, evaluated on the Cityscapes validation set.

Δ	mIoU	Configuration
-8.0	41.9	No augmentation consistency
-6.4	43.5	No momentum net ($\gamma_\psi = 0, T = 1$)
-3.9	46.0	No photometric noise
-2.6	47.3	No multi-scale fusion
-2.4	47.5	No focal loss ($\lambda = 0$)
-1.9	48.0	Min. entropy fusion (<i>vs.</i> averaging)
-1.7	48.2	No class-based thresholding ($\beta \rightarrow 0$)
-1.6	48.3	No confidence regularisation
-1.5	48.4	No importance sampling
-0.6	49.3	No horizontal flipping
0.0	49.9	Full framework (VGG-16)

Table 4. **Ablation study.** We use the GTA5 → Cityscapes setting with the VGG-based model to study the effect of the components of our framework by individually removing each. We report the mean IoU for the Cityscapes validation split.

Remarkably, in both settings our approach is more accurate or competitive with many recent works [65, 70, 77] even when using a weaker backbone, *i.e.* VGG-16 instead of ResNet-101. This is significant, as these improvements are not due to increased training complexity or model capacity, in contrast to these previous works. Additional results, including the evaluation on Cityscapes *test*, are shown in Appendices C and D.

4.3. Ablation study

To understand what makes our framework effective, we conduct an ablation study using the GTA5 → Cityscapes

setting with the VGG-16 backbone. We independently switch off each component and report the results in Table 4. We find that two components, augmentation consistency and the momentum network, play a crucial role. Disabling the momentum network leads to a 6.4% IoU decrease, while abolishing augmentation consistency leads to a drop of 8.0% IoU. Recall that augmentation consistency comprises three augmentation techniques: photometric noise, multi-scale fusion and random flipping. We further assess their individual contributions. Training without the photometric jitter deteriorates the IoU more severely, by 3.9%, compared to disabling the multi-scale fusion (−2.6%) or flipping (−0.6%). We hypothesise that encouraging model robustness to photometric noise additionally alleviates the inductive bias inherited from the source domain to rely on strong appearance cues (*e.g.*, colour and texture), which can be substantially different from the target domain.

Following the intuition that high-confidence predictions should be preferred [65], we study an alternative implementation of the multi-scale fusion. For overlapping pixels, instead of averaging the predictions, we pool the prediction with the minimum entropy. The accuracy drop by 1.9% is somewhat expected. Averaging predictions via data augmentation has previously been shown to produce well-calibrated uncertainty estimates [1]. This is important for our method, since it relies on the confidence values to select the predictions for use in self-supervision. Importance sampling contributes 1.5% IoU to the total accuracy. This

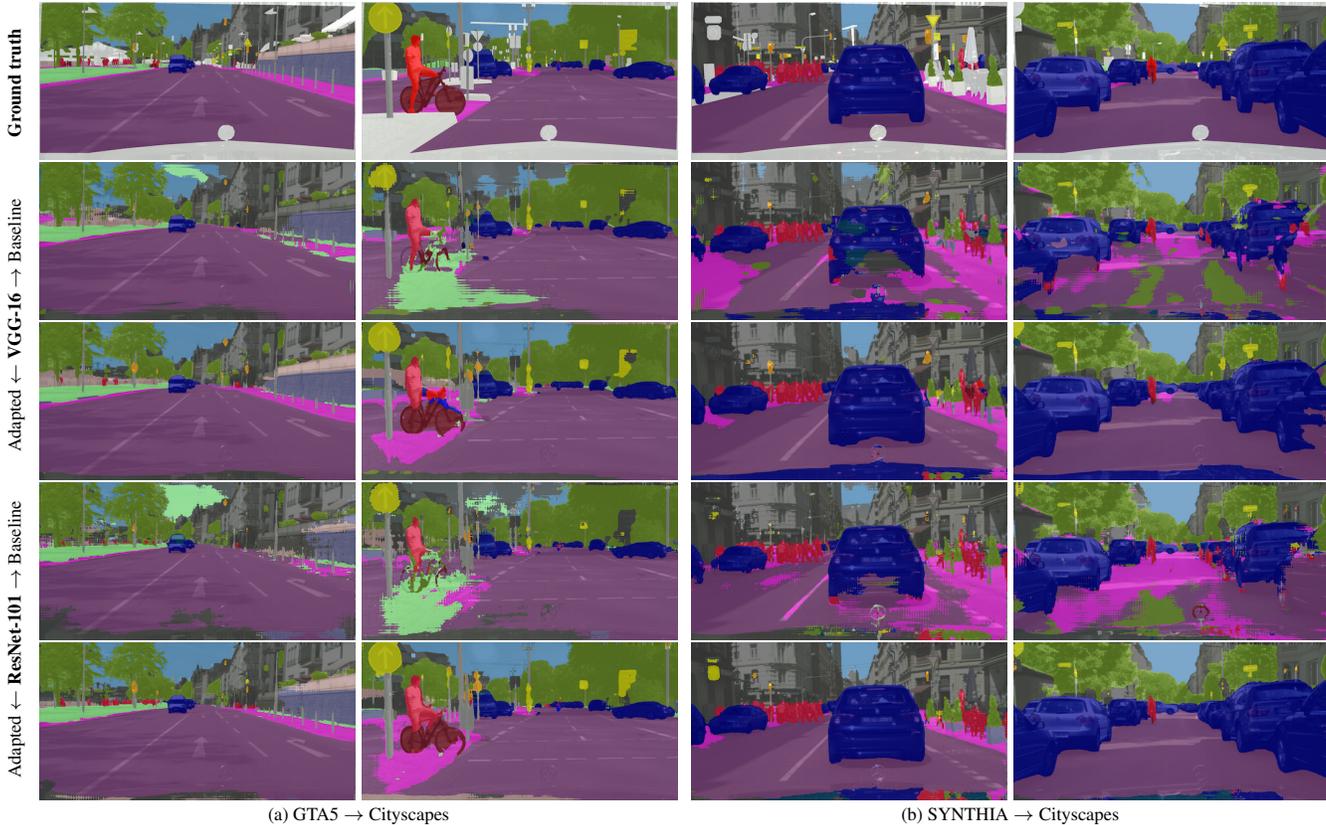


Figure 5. **Qualitative examples.** Our approach rectifies an appreciable amount of erroneous predictions from the baseline.

is surprisingly significant despite that our estimates $\chi_{c,t}$ are only approximate (*cf.* Sec. 3.4), but the overall benefit is in line with previous work [29]. Recall from Eq. (3) that our confidence thresholds are computed per class to encourage lower values for long-tail classes. Disabling this scheme is equivalent to setting $\beta \rightarrow 0$ in Eq. (3), which reduces the mean IoU by 1.7%. This confirms our observation that the model tends to predict lower confidences for the classes occupying only few pixels. Similarly, the loss in Eq. (6) without the focal term ($\lambda = 0$) and confidence regularisation ($m_{c^*,n} = 1$) are 2.4% and 1.6% IoU inferior. This is a surprisingly significant contribution at a negligible computational cost.

4.4. Qualitative assessment

Fig. 5 presents a few qualitative examples, comparing our approach to the naive baseline (*i.e.* source-only loss with ABN). Particularly prominent are the refinements of the classes “road”, “sidewalk” and “sky”, but even small-scale elements improve substantially (*e.g.*, “person”, “fence” in the leftmost column). This is perhaps not surprising, owing to our multi-scale training and the thresholding technique, which initially ignores incorrectly predicted pixels in self-supervision (as they initially tend to have low confidence). Remarkably, the segment boundaries tend to

align well with the object boundaries in the image, although our framework *has no explicit encoding of spatial priors*, which was previously deemed necessary [15, 68, 81, 83]. We believe that enforcing semantic consistency with data augmentation makes our method less prone to the contextual bias [62], often blamed for coarse boundaries.

5. Conclusion

We presented a simple and accurate approach for domain adaptation of semantic segmentation. With ordinary augmentation techniques and momentum updates, we achieve state-of-the-art accuracy, yet make no sacrifice of the modest training or model complexity. No components of our framework are strictly specialised; they build on a relatively weak and broadly applicable assumption (*cf.* Sec. 1). Although this work focuses on semantic segmentation, we are keen to explore the potential of the proposed techniques for adaptation of other dense prediction tasks, such as optical flow, monocular depth, panoptic and instance segmentation, or even compositions of these multiple tasks.

Acknowledgements. This work has been co-funded by the LOEWE initiative (Hesse, Germany) within the emergenCITY center. Calculations for this research were partly conducted on the Lichtenberg high performance computer of TU Darmstadt.

References

- [1] Murat Seckin Ayhan and Philipp Berens. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. In *MIDL*, 2018. 3, 7
- [2] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *NIPS*, pp. 3365–3373, 2014. 2
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Mach. Learn.*, 79(1-2):151–175, 2010. 2
- [4] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *NIPS*, pp. 137–144, 2006. 1
- [5] David Berthelot, Nicholas Carlini, Ian J. Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. MixMatch: A holistic approach to semi-supervised learning. In *NeurIPS*, pp. 5050–5060, 2019. 3
- [6] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, pp. 92–100, 1998. 1
- [7] Minjie Cai, Feng Lu, and Yoichi Sato. Generalizing hand segmentation in egocentric videos with uncertainty-guided model adaptation. In *CVPR*, pp. 14380–14389, 2020. 2
- [8] Wei-Lun Chang, Hui-Po Wang, Wen-Hsiao Peng, and Wei-Chen Chiu. All about structure: Adapting structural information across domains for boosting semantic segmentation. In *CVPR*, pp. 1900–1909, 2019. 2
- [9] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D. Collins, Ekin D. Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *ECCV*, vol. IX, pp. 695–714, 2020. 1
- [10] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2018. 5, 6, 7
- [11] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, vol. VII, pp. 833–851, 2018. 1
- [12] Minmin Chen, Kilian Q. Weinberger, and John Blitzer. Co-training for domain adaptation. In *NIPS*, pp. 2456–2464, 2011. 1
- [13] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *ICCV*, pp. 2090–2099, 2019. 2
- [14] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *ICCV*, pp. 2011–2020, 2017. 2
- [15] Yuhua Chen, Wen Li, and Luc Van Gool. ROAD: Reality oriented adaptation for semantic segmentation of urban scenes. In *CVPR*, pp. 7892–7901, 2018. 2, 8
- [16] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. CrDoCo: Pixel-level domain transfer with cross-domain consistency. In *CVPR*, pp. 1791–1800, 2019. 2
- [17] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with GAN-based data augmentation for domain adaptation in semantic segmentation. In *ICCV*, pp. 6829–6839, 2019. 1, 2
- [18] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *CVPR*, pp. 3213–3223, 2016. 5
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A large-scale hierarchical image database. In *CVPR*, pp. 248–255, 2009. 5
- [20] Jiahua Dong, Yang Cong, Gan Sun, Yuyang Liu, and Xiaowei Xu. CSCL: Critical semantic-consistent learning for unsupervised domain adaptation. In *ECCV*, vol. VIII, pp. 745–762, 2020. 2
- [21] Arnaud Doucet, Nando de Freitas, and Neil J. Gordon. An introduction to sequential Monte Carlo methods. In *Sequential Monte Carlo Methods in Practice*, pp. 3–14. Springer, 2001. 4
- [22] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. SSF-DAN: Separated semantic feature based domain adaptation network for semantic segmentation. In *ICCV*, pp. 982–991, 2019. 2
- [23] Geoffrey French, Michal Mackiewicz, and Mark H. Fisher. Self-ensembling for visual domain adaptation. In *ICLR*, 2018. 2, 3
- [24] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.*, 17(1):2096–2030, 2016. 2
- [25] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. DLOW: Domain flow for adaptation and generalization. In *CVPR*, pp. 2477–2486, 2019. 2
- [26] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, pp. 2672–2680, 2014. 2
- [27] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *NIPS*, pp. 529–536, 2004. 2
- [28] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *ICML*, vol. 70, pp. 1321–1330, 2017. 2
- [29] Agrim Gupta, Piotr Dollár, and Ross B. Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, pp. 5356–5364, 2019. 3, 8
- [30] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pp. 9726–9735, 2020. 3

- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016. 5
- [32] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *ICML*, vol. 80, pp. 1994–2003, 2018. 2, 6
- [33] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. Conditional generative adversarial network for structured domain adaptation. In *CVPR*, pp. 1335–1344, 2018. 2
- [34] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, vol. 37, pp. 448–456, 2015. 4, 5
- [35] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *UAI*, pp. 876–885, 2018. 3
- [36] Fredrik D. Johansson, David A. Sontag, and Rajesh Ranganath. Support and invertibility in domain-invariant representations. In *AISTATS*, vol. 89, pp. 527–536, 2019. 2
- [37] Guoliang Kang, Yunchao Wei, Yi Yang, Yueting Zhuang, and Alexander G. Hauptmann. Pixel-level cycle association: A new perspective for domain adaptive semantic segmentation. In *NeurIPS*, pp. 3569–3580, 2020. 2
- [38] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *NIPS*, pp. 5574–5584, 2017. 3
- [39] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. In *CVPR*, pp. 12972–12981, 2020. 1, 2, 5
- [40] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2, 3
- [41] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. Sliced Wasserstein discrepancy for unsupervised domain adaptation. In *CVPR*, pp. 10285–10295, 2019. 2
- [42] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshops*, vol. 3, 2013. 1
- [43] Guangrui Li, Guoliang Kang, Wu Liu, Yunchao Wei, and Yi Yang. Content-consistent matching for domain adaptive semantic segmentation. In *ECCV*, vol. XIV, pp. 440–456, 2020. 1, 2, 3
- [44] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, and Pheng-Ann Heng. Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model. In *BMVC*, 2018. 2
- [45] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109–117, 2018. 2, 4, 5
- [46] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *CVPR*, pp. 6936–6945, 2019. 1, 2, 3, 6
- [47] Qing Lian, Lixin Duan, Fengmao Lv, and Boqing Gong. Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In *ICCV*, pp. 6757–6766, 2019. 1, 2, 4, 5, 6, 7
- [48] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *ICLR*, 2016. 3
- [49] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(2):318–327, 2020. 4
- [50] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, pp. 1647–1657, 2018. 2
- [51] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *CVPR*, pp. 6778–6787, 2019. 2
- [52] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, pp. 2507–2516, 2019. 2
- [53] Fengmao Lv, Tao Liang, Xiang Chen, and Guosheng Lin. Cross-domain semantic segmentation via domain-invariant interactive relation transfer. In *CVPR*, pp. 4333–4342, 2020. 1, 2, 6, 7
- [54] Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. In *ECCV*, vol. XXVI, pp. 415–430, 2020. 1, 2, 3, 6, 7
- [55] Luigi Musto and Andrea Zinelli. Semantically adaptive image-to-image translation for domain adaptation of semantic segmentation. In *BMVC*, 2020. 1, 2, 6, 7
- [56] Fei Pan, Inkyu Shin, François Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *CVPR*, pp. 3763–3772, 2020. 2
- [57] Adam Paszke, Sam Gross, and Francisco Massa et al. PyTorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pp. 8024–8035, 2019. 5
- [58] Christian S. Perone, Pedro L. Ballester, Rodrigo C. Barros, and Julien Cohen-Adad. Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. *NeuroImage*, 194:1–11, 2019. 2
- [59] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, vol. II, pp. 102–118, 2016. 1, 5
- [60] Germán Ros, Laura Sellart, Joanna Materzynska, David Vázquez, and Antonio M. López. The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pp. 3234–3243, 2016. 1, 5
- [61] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NIPS*, pp. 1163–1171, 2016. 2
- [62] Rakshith Shetty, Bernt Schiele, and Mario Fritz. Not using the car to see the sidewalk – quantifying and controlling the effects of context in classification and segmentation. In *CVPR*, pp. 8218–8226, 2019. 8

- [63] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5
- [64] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. FixMatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, pp. 596–608, 2020. 2
- [65] Muhammad Subhani and Mohsen Ali. Learning from scale-invariant examples for domain adaptation in semantic segmentation. In *ECCV*, vol. XXII, pp. 290–306, 2020. 1, 2, 3, 4, 7
- [66] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, pp. 1195–1204, 2017. 2, 3
- [67] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, pp. 7472–7481, 2018. 1, 2, 5
- [68] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *ICCV*, pp. 1456–1465, 2019. 1, 2, 5, 8
- [69] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. ADVENT: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, pp. 2517–2526, 2019. 1, 2, 5, 6, 7
- [70] Haoran Wang, Tong Shen, Wei Zhang, Lingyu Duan, and Tao Mei. Classes matter: A fine-grained adversarial approach to cross-domain semantic segmentation. In *ECCV*, vol. XIV, pp. 642–659, 2020. 1, 2, 5, 6, 7
- [71] Kaihong Wang, Chenhongyi Yang, and Margrit Betke. Consistency regularization with high-dimensional non-adversarial source-guided perturbation for unsupervised domain adaptation in segmentation. *AAAI*, 2021. 2
- [72] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, pp. 12272–12281, 2020. 3
- [73] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogério Feris, Jinjun Xiong, Wen-Mei Hwu, Thomas S. Huang, and Honghui Shi. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *CVPR*, pp. 12632–12641, 2020. 2, 5
- [74] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gökhan Uzunbas, Tom Goldstein, Ser-Nam Lim, and Larry S. Davis. DCAN: Dual channel-wise alignment networks for unsupervised scene adaptation. In *ECCV*, vol. V, pp. 535–552, 2018. 2
- [75] Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, pp. 6256–6268, 2020. 2
- [76] Qizhe Xie, Minh-Thang Luong, Eduard H. Hovy, and Quoc V. Le. Self-training with noisy student improves imagenet classification. In *CVPR*, pp. 10684–10695, 2020. 2, 3
- [77] Jinyu Yang, Weizhi An, Sheng Wang, Xinliang Zhu, Chaochao Yan, and Junzhou Huang. Label-driven reconstruction for domain adaptation in semantic segmentation. In *ECCV*, vol. XXVII, pp. 480–498, 2020. 1, 2, 6, 7
- [78] Jinyu Yang, Weizhi An, Chaochao Yan, Peilin Zhao, and Junzhou Huang. Context-aware domain adaptation in semantic segmentation. In *WACV*, pp. 514–524, 2021. 1, 2, 6, 7
- [79] Yanchao Yang, Dong Lao, Ganesh Sundaramoorthi, and Stefano Soatto. Phase consistent ecological domain adaptation. In *CVPR*, pp. 9008–9017, 2020. 2
- [80] Yanchao Yang and Stefano Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *CVPR*, pp. 4084–4094, 2020. 1, 2, 5, 6, 7
- [81] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *ICCV*, pp. 2039–2049, 2017. 2, 8
- [82] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, and Tao Mei. Fully convolutional adaptation networks for semantic segmentation. In *CVPR*, pp. 6810–6818, 2018. 2
- [83] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Dong Liu, and Tao Mei. Transferring and regularizing prediction for semantic segmentation. In *CVPR*, pp. 9618–9627, 2020. 2, 4, 6, 7, 8
- [84] Han Zhao, Remi Tachet des Combes, Kun Zhang, and Geoffrey J. Gordon. On learning invariant representations for domain adaptation. In *ICML*, vol. 97, pp. 7523–7532, 2019. 2
- [85] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, pp. 6230–6239, 2017. 1, 6, 7
- [86] Zhedong Zheng and Yi Yang. Unsupervised scene adaptation with memory regularization in vivo. In *IJCAI*, pp. 1076–1082, 2020. 1, 2
- [87] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *Int. J. Comput. Vis.*, 129(4):1106–1120, 2021. 1, 2
- [88] Qianyu Zhou, Zhengyang Feng, Guangliang Cheng, Xin Tan, Jianping Shi, and Lizhuang Ma. Uncertainty-aware consistency regularization for cross-domain semantic segmentation. *arXiv:2004.08878 [cs.CV]*, 2020. 2
- [89] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pp. 2242–2251, 2017. 2
- [90] Xinge Zhu, Hui Zhou, Ceyuan Yang, Jianping Shi, and Dahua Lin. Penalizing top performers: Conservative loss for semantic segmentation adaptation. In *ECCV*, vol. VII, pp. 568–583, 2018. 2
- [91] Yang Zou, Zhiding Yu, B.V.K. Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, vol. III, pp. 297–313, 2018. 1, 2, 3, 6
- [92] Yang Zou, Zhiding Yu, Xiaofeng Liu, B.V.K. Vijaya Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, pp. 5981–5990, 2019. 1, 2