Learning Instance-Specific Adaptation for Cross-Domain Segmentation

Yuliang Zou¹ Zizhao Zhang² Chun-Liang Li² Han Zhang³ Tomas Pfister² Jia-Bin Huang⁴

¹Virginia Tech ²Google Cloud AI ³Google Brain ⁴University of Maryland, College Park



Fig. 1: Cross-domain segmentation. Models trained with the standard recipe on source domain data perform poorly on unseen target domains. On the contrary, the proposed method significantly improves upon off-the-shelf pre-trained models, without accessing the target domain at training time or parameter optimization at test-time.

Abstract. We propose a test-time adaptation method for cross-domain image segmentation. Our method is simple: Given a new unseen instance at test time, we adapt a pre-trained model by conducting instance-specific BatchNorm (statistics) calibration. Our approach has two core components. First, we replace the manually designed BatchNorm calibration rule with a learnable module. Second, we leverage strong data augmentation to simulate random domain shifts for learning the calibration rule. In contrast to existing domain adaptation methods, our method does not require accessing the target domain data at training time or conducting computationally expensive test-time model training/optimization. Equipping

our method with models trained by standard recipes achieves significant improvement, comparing favorably with several state-of-the-art domain generalization and one-shot unsupervised domain adaptation approaches. Combining our method with the domain generalization methods further improves performance, reaching a new state of the art. Our project page is https://yuliang.vision/InstCal/.

1 Introduction

Deep neural networks have shown impressive results in many computer vision applications. However, these models suffer from inevitable performance drops when deployed in out-of-distribution environments due to domain shift [3]. For example, segmentation models trained on sunny images may perform poorly on foggy or rainy scenes [9]. Improving the cross-domain performance of deep vision models has thus received considerable attention in recent years.

Domain adaptation. One straightforward approach for reducing the domain shift is to collect diverse labeled data in the target domain of interest for supervised fine-tuning. However, collecting sufficient annotated data in the target domain could be expensive or infeasible (e.g., in continuously changing environments). This is particularly challenging for many dense prediction tasks such as image segmentation as it requires *dense* (pixel-wise) labels. Unsupervised domain adaptation (UDA) [17,32,50,53] is an alternative route for reducing the domain gap by using unlabeled target data. However, UDA methods require accessing target domain data for model training before deployment. Such assumptions may not hold as we are not able to anticipate what scenarios the model would encounter (e.g., different weather conditions) and therefore cannot collect the unlabeled data accordingly. One-shot UDA [4,31] relaxes the constraint by requiring only one target example for model training. The model can thus use the first example encountered in the unseen target domain as the training example. However, the adaptation procedure often requires thousands of training steps [31], hindering its applicability as a plug-and-play module when deployed on new target domains. These UDA methods also require access to source data *during* the adaptation process, which may be unrealistic at test time.

Domain generalization (DG) [1,27,34,60] overcomes the above limitations by learning invariant representations using multiple source domains to improve model robustness on unseen or continuously changing environments. Recent approaches [40,59] relax the constraint by requiring one single source domain only. However, as Dubey et al. [14] points out, the optimal model learned from training domains may be far from being optimal for an unseen target domain.

Test-time adaptation approaches have been proposed to tackle exactly the same problem. These methods can be roughly categorized into two groups: 1) optimizing model parameters at test time with a proxy task [2,49], prediction pseudo-label [29], or entropy regularizations [54], and 2) BatchNorm calibration [22,36,46]. These approaches can be applied to update models along with observing each target test data, thus observing the entire test distribution. Alternatively, they can be used to create an *instance-specific* model for each test

2

example individually. Despite the flexibility, the optimization-based methods all require time-consuming backprop computation to update model parameters. **Our work.** In this paper, we present a simple test-time adaptation method for cross-domain segmentation (Figure 1). Building upon BatchNorm calibration methods [36,46], we propose to learn *instance-specific* calibration rules using strong data augmentations to simulate various pseudo source domains. Our approach offers several advantages. First, compared with existing work [36,46] with *manually determined* calibration rules that require time-consuming grid searches and may not transfer to different models, our approach is *data-driven* and instance-specific. Second, unlike other test-time adaptation methods [2,29,49,54], our work does not involve expensive gradient-based optimization for updating model parameters at test time. Third, our method learns to calibrate BatchNorm statistics with *one* single instance (i.e., without accessing to a batch of samples). We validate our proposed methods on cross-domain semantic and panoptic segmentation tasks on several benchmarks. Our experiments show a sizable boost over existing adaptation methods.

Contributions. In summary, we make the following contributions:

- We propose a simple instance-specific test-time adaptation method and show its applicability to off-the-shelf segmentation models containing BatchNorm.
- We conduct a detailed ablation study and analysis to validate our design choices in semantic segmentation. Applying the optimal configuration to the more complex panoptic segmentation task leads to promising performance.
- When combined with the models pre-trained by standard recipes, our method compares favorably with state-of-the-art one-shot UDA methods and domain generalizing semantic segmentation methods. Our approach can also be combined with existing DG methods to improve the performance further.

2 Related Work

Domain adaptation. Models trained on one (source) domain often suffers from a severe performance drop when processing samples from unseen (target) domains. Domain adaptation methods aim to mitigate this issue by adapting a pre-trained model using samples from target domains. Unsupervised domain adaptation (UDA) methods show promising results by leveraging unlabeled target data. These UDA techniques include 1) domain invariant learning, 2) generative models, and 3) self-training. Domain invariant learning methods learn invariant features for the source and target domains by imposing an adversarial loss [17,32,50,53]. minimizing the domain distribution distance (e.g., MMD) [30,51] or correlation distance [48]. Applying data augmentation with generative models can also reduce domain gap using image-to-image translation [5,47,61], style transfer [15], or hybrid methods that integrate with domain invariance learning methods [7,21]. Self-training methods [62,63] select confident/reliable target data predictions and convert them into pseudo labels. These methods then iterate the fine-tuning and pseudo-labeling procedures until convergence. While we have observed remarkable progress in UDA, pre-collected target domain data requirement makes it less

practical. Recently, one-shot UDA methods [4,31] have been proposed to tackle this problem. Instead of training on *many* unlabeled target data, these approaches require only *one* unlabeled target data. However, these methods require timeconsuming offline training before deploying on the target domain. In contrast, our proposed method efficiently adapts the model by calibrating BatchNorm on each target example *on the fly*, without offline training on each target domain separately. Models trained with our method can be easily applied to many different unseen domains.

Domain generalization. Instead of adapting models to using target domain data, domain generalization [1,27,34,60] aims to train a model on source domains that are *generalizable* to unseen target domains by encouraging the networks to learn domain-invariant representations. However, these approaches require multiple source domains for training, which poses additional challenges in (labeled) data collection and restricts their feasibility in practical usage. To mitigate the data collection issues, single domain generalization [40] trains models on one single source domain only by either exploiting strong data augmentation strategies [40,52,59] to diversify the source domain training data, or performing feature whitening operations or normalization [9,16,23,39] to remove domain-specific information during training. Similar to these methods, our method also exploits data augmentation to diversify one single source domain data. However, instead of enforcing models to learn domain-invariant features, we encourage models to calibrate BatchNorm on each unseen target data at test time, by training them on diverse pseudo domains generated with strong data augmentation strategies in training time. Our proposed method can also complement (single) domain generalization approaches to improve the performance further.

Test-time adaptation. Depending on the use of online training/optimization, test-time adaptation methods can be divided into two groups. First, *optimization-free* test-time adaptation methods mostly focus on calibrating the running statistics inside BatchNorm layers [22,25,35,36,46] because these feature statistics carry domain-specific information [28]. However, these methods either directly replace the running statistics with current input batch statistics or mix the running statistics and current input batch statistics with a *pre-defined* calibration rule. In contrast, we propose to learn the instance-specific BatchNorm calibration rule from source domain data. Second, *test-time optimization* methods adapt the model parameters using a training objective such as entropy minimization [54], pseudo-labeling [29], or self-supervised proxy tasks [10,49]. Our experiments show that integrating our method with test-time optimization boosts performance.

3 Learning Instance-Specific BatchNorm Calibration

Our method applies to off-the-shelf pre-trained segmentation models containing BatchNorm layers [24], a reasonable assumption in most modern CNN models. In section 3.1, we first review BatchNorm and recent test-time calibration techniques. We then introduce our method (unconditional and conditional BatchNorm calibration) in section 3.2.



Fig. 2: Overview. (a) At training time, we learn the BatchNorm calibration rule (equation 8) by training only the newly initialized parameters on the strongly-augmented source domain data; (b) At test time, we conduct instance-specific BatchNorm calibration using the learned calibration rule. Note that our method does not perform test-time training or optimization, and thus the model parameters are fixed after training.

3.1 Background

A brief review of BatchNorm. BatchNorm has been empirically shown to stabilize model training and improve model convergence speed [55], making it an essential component in most modern CNN models. The inputs to each BatchNorm layer are CNN features $x \in \mathbb{R}^{B \times C \times H \times W}$, where *B* denotes the batch size, *C* denotes the number of feature channels, $H \times W$ denotes the spatial size. BatchNorm conducts normalization followed by an affine transformation on the inputs *x* to get outputs $y \in \mathbb{R}^{B \times C \times H \times W}$

$$y = \frac{x - \mu}{\sqrt{\sigma^2 + \epsilon}} \times \gamma + \beta, \tag{1}$$

where ϵ is a small constant for numerical stability, $\gamma \in \mathbb{R}^C$ and $\beta \in \mathbb{R}^C$ are learnable parameters. We generalize the tensor operations by assuming broadcasting when the dimensions are not exactly the same. Note that the definition of μ and σ^2 differs in training and test time. In training time, μ and σ^2 are set as the input batch statistics.¹

$$\mu_B = \operatorname{mean}(x, \operatorname{axis} = (B, H, W)) \tag{2}$$

$$\sigma_B^2 = \operatorname{var}(x, \operatorname{axis} = (B, H, W)) \tag{3}$$

At test time, μ and σ^2 are set to *population statistics* $\mu_{pop}, \sigma_{pop}^2$, accumulated in training time using exponential moving averaging

$$\mu_{pop,t} = (1 - \alpha) \times \mu_{pop,t-1} + \alpha \times \mu_B \tag{4}$$

$$\sigma_{pop,t}^2 = (1 - \alpha) \times \sigma_{pop,t-1}^2 + \alpha \times \sigma_B^2 \tag{5}$$

¹ Following the "NamedTensor" practice [43], this computes the statistics over the B, H, W dimensions and return vectors with dimension C.

where α is a scalar called momentum, and the default value is 0.1 (in PyTorch convention). Note that the population statistics update happens during training in every feed-forward step.

Manual BatchNorm calibration. Despite the empirical success in in-domain testing, models with BatchNorm layers suffer from a significant performance drop when testing on out-of-distribution data. One potential reason is that the population statistics within the BatchNorm layers carry *domain-specific* information [28], and thus these statistics are not suitable for normalizing inputs from a different domain. Recent studies [28,46,54] show that, by calibrating the population statistics with input statistics, the cross-domain performance can be significantly improved:

$$y = \frac{x - ((1 - m) \times \mu_{pop} + m \times \mu_{ins})}{\sqrt{\left((1 - m) \times \sigma_{pop}^2 + m \times \sigma_{ins}^2\right) + \epsilon}} \times \gamma + \beta,$$
(6)

where *m* indicates calibration strength, each method has a different empirically specified value, $\mu_{ins} \in \mathbb{R}^{B \times C}$ and $\sigma_{ins}^2 \in \mathbb{R}^{B \times C}$ indicate instance mean and variance. Note that the above calibration step happens in every BatchNorm layer, and thus the input features in later BatchNorm layers will be increasingly more calibrated.

In our study (Table 1(a)), we show that simply setting calibration strength m as the default momentum value 0.1 can improve overall cross-domain performance. However, we find several potential issues. First, the calibration strength m is specified *empirically*, requiring a grid search to obtain the optimal value. Nevertheless, the optimal value for one setting might not well transfer to other settings with different pre-trained models or target domains of interest. Second, the calibration strength m is a scalar. However, different feature channels may encode different semantic information [56]. Therefore, we may use different calibration strengths for different feature channels. Third, calibrating the mean and variance with the same strength leads to sub-optimal results.

3.2 The proposed method

Learning to calibrate BatchNorm (InstCal-U). To address the aforementioned issues, we propose to learn the calibration strengths *during training* instead of manually specifying them at test time. For simplicity, we define the following function

$$f_c(a, b, \mathbf{m}) = (\vec{1} - \mathbf{m}) \times a + \mathbf{m} \times b \tag{7}$$

The proposed calibration and normalization process can thus be written as

$$y = \frac{x - f_c \left(\mu_{pop}, \mu_{ins}, \mathbf{m}_{\mu}\right)}{\sqrt{f_c \left(\sigma_{pop}^2, \sigma_{ins}^2, \mathbf{m}_{\sigma}\right) + \epsilon}} \times \gamma + \beta, \tag{8}$$

where $\mathbf{m}_{\mu} \in \mathbb{R}^{C}$ and $\mathbf{m}_{\sigma} \in \mathbb{R}^{C}$ are two learnable parameters. More specifically, we initialize two learnable parameters \mathbf{m}_{μ} and \mathbf{m}_{σ} with the default momentum value 0.1.

Given an off-the-shelf model, we convert all BatchNorm layers into the instancespecific calibrated format in equation 8. Note that we only train the newly initialized calibration parameters \mathbf{m}_{μ} and \mathbf{m}_{σ} , and we keep the other learnable parameters (including γ and β) fixed.

Using training data in the source domain, we train parameters \mathbf{m}_{μ} and \mathbf{m}_{σ} on a diverse set of domains. Our intuition is that, by exposing the model to diverse (simulated) domains, we implicitly constrain the learnable calibration parameters \mathbf{m}_{μ} and \mathbf{m}_{σ} to be robust and invariant to unseen target domains. However, since we only have one single source domain, we need to generate multiple pseudo domains based on the source domain. Instead of adopting complex generative models to generate pseudo domains, we find that applying appropriate strong data augmentation during training leads to promising results. We explore three different augmentation strategies: RandAugment [12], AugMix [20], and DeepAugment [19], and empirically find that DeepAugment performs the best (Table 1(b)). The details of augmentations are in supplementary materials.

Learning to conditionally calibrate BatchNorm (InstCal-C). While the goal is to learn the BatchNorm calibration parameters so that the models can adapt to unseen domains at test time, the learnable parameters \mathbf{m}_{μ} and \mathbf{m}_{σ} are *fixed* after training. We propose an optional module to enable *conditional* calibration to increase the flexibility.

Instead of directly learning the parameters \mathbf{m}_{μ} and \mathbf{m}_{σ} , we propose to learn a set of parameters $\mathbf{m}_{\mu,i}$ and $\mathbf{m}_{\sigma,i}$ for mean and variance, respectively. These parameters can be viewed as the *basis* of calibration rules. We will use two lightweight MLPs (one for each statistic) to predict the coefficients to combine the basis to get the actual calibration strength for each test example, given the concatenation of instance and population statistics. Take \mathbf{m}_{μ} as an example, the computation step can be written as follows

$$\{c_{\mu,i}\}_{1}^{K} = \mathbf{Softmax}\left(g_{\mu}\left(\mathbf{Concat}(\mu_{pop}, \mu_{ins})\right)\right)$$
(9)

$$\mathbf{m}_{\mu} = \sum_{i}^{K} c_{\mu,i} \mathbf{m}_{\mu,i} \tag{10}$$

where $c_{\mu,i}$ is a scalar, **Concat**(·) is the channel-wise concatenation operation, and $g_{\mu}(\cdot)$ is a small 2-layer MLP. The computation of \mathbf{m}_{σ} is similar.

Thanks to the redesign, we further increase the learnable instance-specific BatchNorm calibration flexibility by setting the calibration rule to be conditional on the input features. As a result, the calibration rule is now *dynamically changing* according to different test target samples, while the inference process is still done within one forward pass. As shown in section 4, in general, the performance of instance-specific calibration (InstCal-C) improves upon the unconditional calibration (InstCal-U) on synthetic-to-real settings where significant domain shifts exist.

4 Experimental Results

We mainly validate and analyze our method using the semantic segmentation tasks. In section 4.8, we also apply the proposed method to panoptic segmentation and observe promising results.

4.1 Experimental setup

We conduct experiments on the public semantic segmentation benchmarks: GTA5 [41], SYNTHIA [42], Cityscapes [11], BDD100k [57], Mapillary [37], and WildDash2 [58] datasets. The GTA5 and SYNTHIA datasets are synthetic, while the others are real-world datasets. For both synthetic datasets, we split the data following Chen et al. [6]. For the WildDash2 dataset, we only evaluate the 19 classes overlapping with Cityscapes and ignore the remaining classes. We evaluate model performance using the standard mean intersection-over-union (mIoU) metric. We provide the implementation details in the supplementary material.

4.2 Ablation study

We use GTA5 as the source domain for ablation experiments and Cityscapes as the unseen target domain. We use the DeepLabv2 model with a ResNet-101 backbone.

Table 1: Ablation study. We show results from a DeepLabv2 model with a ResNet-101 backbone. We train models on the GTA5 dataset and treat the Cityscapes dataset as the unseen target domain for evaluation.

(a) Calib	oration param	eters to	learn	(b) Di	fferent	augmen	tations
Stra	tegy 1	nIoU (%)		Augmentat	ion	InstCal-U	InstCal-C
Pre- m =	trained 0.1, fixed	$35.7 \\ 40.1$		Default	N1	39.7	40.9
$\mathbf{m}\in\mathbf{m}\in\mathbf{m}$	$\mathbb{R} \text{ (scalar)} \\ \mathbb{R}^C \text{ (vector)} \\ \in \mathbb{R}^C, \mathbf{m}_{\sigma} \in \mathbb{R}^C $	39.8 41.1 41.5		RandAugm DeepAugm	ont [12] ent [19]	40.6 41.1 41.5	41.3 40.0 42.2
(c) Not en	ough to pre-tr	ain with	strong aug	. (d) Nu	umber c	of basis f	or InstCal-C
A	ugmentation m	IoU (%)			#basis n	nIoU (%)	
De	efault	35.7			2	40.9	
A	ugMix [20]	35.9			4	41.6	
Ra	andAugment [12]	37.9			8	42.2	
De	eepAugment [19]	31.7			16	40.7	

What calibration parameters should we learn? We first conduct experiments to study what calibration parameters should be learned. As shown in Table 1(a), suppose we directly learn a scalar parameter shared by the mean and

variance. The performance is worse than using a default value of calibration strength (0.1) to calibrate BatchNorm. Learning a *vector* parameter works much better than a single *scalar* and outperforms the baseline calibration. Separating the learned vector for mean and variance leads to further improved performance. **Which data augmentation strategy should we use?** As we mentioned in section 3.2, since we only require one source domain for model training, we need to use strong data augmentation to simulate a diverse set of training domains. In this experiment, we study the impact of data augmentation methods.

Table 1(b) shows that using the default weak augmentation, e.g., random scaling, cropping, the performance is even worse than the default baseline. While RandAugment [12] and AugMix [20] work well for InstCal-U or InstCal-C separately, these two augmentation strategies do not work well in both variants. Our results show that DeepAugment [19] achieves the best overall performance. We thus adopt DeepAugment as our default strong data augmentation strategy.

Pre-training with strong data augmentation is not sufficient. In the previous study, we show that the selection of strong data augmentation is critical. One may wonder if pre-training with strong data augmentation without the proposed adaptation (InstCal-U and InstCal-C) is sufficient for performance improvement. Table 1(c) shows that pre-training models using strong data augmentations do not achieve the models trained with our proposed adaptation methods. AugMix [20] and RandAugment [12] can improve the performance over the baseline with standard weak augmentation, but not as significant as using them in InstCal-U or InstCal-C. If we directly use DeepAugment [19] for model pre-training, the performance even drops significantly. The results suggest that it is necessary to apply strong data augmentation, but we need to use them in the InstCal-U/InstCal-C training stage instead of simply using them during pre-training.

Number of basis for conditional calibration. In Table 1(d), we study the impact of number of basis (K in equation 9) for InstCal-C. Using eight basis leads to the best results among several options.

4.3 Comparison with other test-time adaptation methods

We conduct experiments using the DeepLabv2 model with a ResNet-101 backbone. We first construct a baseline using a default value (m = 0.1) to calibrate BatchNorm statistics. We compare with one optimization-based approach (TENT [54]) and two BatchNorm calibration based methods (AdaptiveBN [46] and PT-BN [36]). We use the same protocols to separately conduct test-time adaptation on multiple unseen target domains. (i.e., setting test batch size to 1 and adapting to each test example individually).

Note that AdaptiveBN [46], PT-BN [36], and our baseline share the same formulation (equation 8) but using different m values. As shown in Table 2, while both the simple baseline and AdaptiveBN [46] show improved results, PT-BN [36] even hurts the pre-trained performance in many cases. TENT [54] also shows strong results in some of the test settings, but with the price of significantly increased computation time. In contrast, the proposed InstCal-U



Fig. 3: Manually set calibration strength *m*. We show results from a DeepLabv2 model with a ResNet-101 backbone. The source domain is the GTA5 dataset.

and InstCal-C outperform these test-time adaptation methods in most settings. We also note that InstCal-U performs better in real-world cross-domain settings, while InstCal-C achieves more promising results in synthetic-to-real settings.

In addition to setting the calibration strength to the default value (0.1), we also experiment with different values. We try from 0.0 to 1.0 with a step size of 0.1, and visualize the results in Figure 3. As we can see, using scalar as strength to calibrate BatchNorm is highly sensitive to selecting the values. On the contrary, the proposed InstCal-U and InstCal-C consistently perform well across unseen target domains.

Table 2: Generalizing across multiple domains. We show results from DeepLabv2 models with a ResNet-101 backbone. The baseline uses calibration strength m = 0.1. "C" indicates Cityscapes, "B" indicates BDD100k, "M" indicates Mapillary, and "W" indicates WildDash2. The best performance is in **bold** and the second best is <u>underlined</u>.

	Source: GTA5					Source: Cityscapes			
Method	С	В	М	W	Avg.	В	Μ	W	Avg.
Pre-trained	35.7	32.9	41.1	27.4	34.3	41.2	49.5	33.9	41.5
Baseline $(m = 0.1)$	40.1	37.4	45.3	32.0	38.7	42.8	51.9	37.6	44.1
AdaptiveBN [46]	39.0	36.3	44.3	30.8	37.6	43.0	52.4	37.2	44.2
PT-BN [36]	33.9	34.3	40.5	27.8	34.1	34.4	39.1	28.8	34.1
TENT [54]	38.1	37.8	44.7	32.5	38.3	44.2	<u>52.3</u>	36.8	44.4
InstCal-U (Ours)	41.5	<u>39.4</u>	<u>46.0</u>	34.4	<u>40.3</u>	45.1	52.2	40.3	45.9
InstCal-C (Ours)	42.2	40.2	46.8	35.3	41.1	<u>44.3</u>	51.5	<u>39.3</u>	45.0

4.4 Analysis

For the following studies, we use DeepLabv2 models with a ResNet-101 backbone. **Improvement on in-domain performance.** We test if our models can improve the performance for *source domain*. We do so by evaluating the trained model on the *test split* of the source data. We report in Table **3**(a) that our learned BatchNorm calibration (both unconditional and conditional) still achieve sizable performance gain.



Table 3: Analysis. Results are from DeepLabv2 models with a ResNet-101 backbone.(a) In-domain performance(b) Input batch statistics

Input batch statistics v.s. instance statistics As mentioned in section 3.2, we compute input instance statistics for *each* test example, instead of computing the batch statistics for mixing statistics across different examples within a minibatch. We validate this design choice by replacing instance statistics with batch statics using different batch sizes during test time. Table 3(b) shows that the performance drops as we increase the batch size, even though the test examples come from the same target distribution. We conjecture the performance will worsen if the mini-batch contains test examples from multiple target domains. Thus, we stick to using the input instance statistics.

Improvement on model calibration. We compute the expected calibration error (ECE) [18] for the pre-trained model, baseline update (m = 0.1), PT-BN [36], and the proposed InstCal-U/InstCal-C. As shown in Table 3(c), calibrating BatchNorm statistics indeed reduces model calibration error.

Compatible with test-time optimization. We incorporate the optimizationbased method, TENT [54], into our methods, by optimizing the prediction entropy at test-time. Following TENT [54], we only optimize the weight γ and bias β in BatchNorm layers. Moreover, we conduct instance-specific adaptation. Table 3(d) shows the complementary nature of these two strategies.

Running time. We test the inference speed on Cityscapes on a single V100 GPU. The pre-trained model takes 39 ms to process each testing sample $(1024 \times 512$ resolution). The BatchNorm calibration method [36] induce a 60 ms overhead. Our method increases the inference time by 58 ms (for InstCal-U) and 149 ms (InstCal-C).

4.5 Comparison with one-shot unsupervised domain adaptation

This section compares the proposed method with recent state-of-the-art oneshot UDA methods. One-shot UDA methods adapt source domain pre-trained

models on one single unlabeled target example offline. In contrast, InstCal-U and InstCal-C adapt pre-trained models on the fly on each test example individually. Conceptually, one-shot UDA methods and InstCal-U/InstCal-C use the same amount of data for adaptation. However, one-shot UDA methods usually require time-consuming offline training, and thus it is impossible to adapt models on each target example separately. So these methods only adapt the models using one single unlabeled (training) example and then deploy the adapted model at test-time without adaptation. As shown in Table 4, simply augmenting pre-trained models with InstCal-U/InstCal-C, compares favorably with recent one-shot UDA methods and even outperforms the state of the arts by a large margin in Synthia \rightarrow Cityscapes setting.

Table 4: Comparison with one-shot unsupervised domain adaptation. All results are from modified DeepLabv2 models (specific for domain adaptation). The best performance is in **bold** and the second best is underlined.

	${\rm GTA5}{\rightarrow}{\rm Cityscapes}$	$Synthia {\rightarrow} City scapes$				
Method	mIoU	mIoU (13-class)	mIoU (16-class)			
CLAN [32]	37.7	40.4	-			
AdvEnt [53]	36.1	39.9	-			
CBST [62]	37.1	38.5	-			
OST [4]	42.3	42.8	-			
ASM [31]	43.2	40.7	34.6			
Source-only pre-trained	36.2	36.2	31.6			
+ InstCal-U (Ours)	42.4	43.5	37.7			
+ InstCal-C (Ours)	42.2	44.1	38.1			

4.6 Comparison with domain generalizing segmentation

This section compares our InstCal-U/InstCal-C with recent domain generalizing (DG) semantic segmentation approaches. We use the DeepLabv3+ model with a ResNet-50 backbone. As shown in Table 5, upgrading non-DG pre-trained *weak* models with InstCal-U/InstCal-C compares favorably with these *strong* domain generalizing segmentation methods across different testing settings. Our method even outperforms all the methods except ISW [9] by a large margin.

Note that our method and these domain generalizing methods complement each other. Thus, we can also incorporate our methods on top of these domain generalizing segmentation methods. As shown in Figure 4, the proposed method consistently improves the performance of these methods. Our method can even improve the strong ISW [9] approach and achieve a new state of the art.

4.7 Backbone network agnostic

In previous sections, we have shown the proposed method can improve pre-trained model performance on multiple unseen target domains. However, we only conduct

Table 5: Comparison with state-of-the-art domain generalizing semantic segmentation methods. We show results from a DeepLabv3+ model with a ResNet-50 backbone. "C" indicates Cityscapes, "B" indicates BDD100k, and "M" indicates Mapillary. The best performance is in **bold** and the second best is <u>underlined</u>.



Fig. 4: Combining domain generalization with our method. In addition to the model pre-trained with the standard recipe (non-DG, labeled as "Default"), we choose three DG methods: SW [39], IBN-Net [38], and ISW [9]. All methods use DeepLabv3+ models with a ResNet-50 backbone, trained on the GTA5 dataset.

experiments using the ResNet backbones. In this section, we use the DeepLabv3+ model with ShuffleNetV2 [33] and MobileNetV2 [45] as backbones, to demonstrate our methods is network-agnostic. As shown in Table 6, the proposed methods can also improve pre-trained models with these backbones by a large margin, outperforming the recent state of the arts.

4.8 Panoptic segmentation results

In this section, we directly apply InstCal-U/InstCal-C to an even more challenging task, panoptic segmentation [26]. We start with the off-the-shelf models from Panoptic-DeepLab [8] and train these models on the Cityscapes dataset. We test the models on Foggy Cityscapes [44], which inserts synthetic fog into the original Cityscapes clear images with three strength levels (0.005, 0.01, and 0.02). We adopt panoptic quality (PQ), mean intersection-over-union (mIoU), and mean average precision (mAP) as the evaluation metrics. As shown in Table 7, InstCal-U/InstCal-C greatly improves off-the-shelf Panoptic-DeepLab performance on out-of-distribution foggy scenes by a large margin, validating the proposed method

Table 6: The proposed module is backbone network agnostic. We show results from a DeepLabv3+ model with ShuffleNetV2 and MobileNetV2 as backbones. These models are trained on the GTA5 dataset. "C" indicates Cityscapes, "B" indicates BDD100k, and "M" indicates Mapillary. The best performance is in **bold** and the second best is <u>underlined</u>.

	S	huffle	NetV	2	Ν	lobile	NetV	NetV2	
Method	С	В	Μ	Avg.	С	В	Μ	Avg.	
IBN-Net [38] ISW [9]	$\begin{array}{c} 27.1\\ 31.0 \end{array}$	31.8 32.1	$34.9 \\ 35.3$	$31.3 \\ 32.8$	$30.1 \\ 30.9$	27.7 <u>30.1</u>	$27.1 \\ 30.7$	$28.3 \\ 30.6$	
Non-DG pre-trained + InstCal-U (Ours) + InstCal-C (Ours)	25.7 <u>35.8</u> 35.9	22.1 31.1 30.8	28.3 36.4 <u>35.4</u>	25.4 34.4 <u>34.0</u>	27.1 <u>37.2</u> 37.8	27.5 31.2 30.0	27.3 34.5 <u>33.9</u>	27.3 34.3 <u>33.9</u>	

is universally applicable to different image segmentation tasks without further tuning. We also provide visual results in Figure 1 and supplementary material.

Table 7: Panoptic segmentation results. We show results on two Panoptic-DeepLab model variants (w/ and w/o depthwise separable convolution). The best performance is in **bold** and the second best is <u>underlined</u>.

		Synthetic fog strength								
		0.005			0.01			0.02		
Method	w/ DSConv	\mathbf{PQ}	mIoU	mAP	\mathbf{PQ}	mIoU	mAP	\mathbf{PQ}	mIoU	mAP
Pre-trained	×	53.3	72.2	25.3	45.0	64.9	18.8	32.6	52.8	11.6
+ InstCal-U	×	56.6	75.7	28.9	51.1	71.9	<u>24.3</u>	42.8	64.5	18.5
+ InstCal-C	×	56.6	75.7	28.5	51.2	71.9	24.4	<u>42.4</u>	64.8	<u>18.0</u>
Pre-trained	\checkmark	53.0	73.2	24.5	45.3	66.5	18.3	33.1	54.8	11.5
+ InstCal-U	\checkmark	55.5	76.0	27.2	49.1	71.3	21.8	40.4	63.2	16.2
+ InstCal-C	\checkmark	55.6	76.3	27.6	<u>48.9</u>	71.6	22.4	40.9	64.1	16.8

5 Discussions

This paper proposes a simple learning-based test-time adaptation method for cross-domain segmentation. The proposed method is learned to perform *instance-specific* BatchNorm calibration during training, without time-consuming test-time parameter optimization. As a result, our method is efficient and effective, demonstrating competitive performance across multiple cross-domain image segmentation settings.

Limitations. Currently, we conduct calibration for every BatchNorm layers. It will be interesting to study which layer is more important and thus only calibrate specific layers to increase inference speed. And it will be interesting to extend our method to other normalization layers (e.g., LayerNorm for Vision Transformers [13]) and other challenging tasks. We leave these as future work.

References

- Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: NeurIPS (2018) 2, 4
- Bartler, A., Bühler, A., Wiewel, F., Döbler, M., Yang, B.: Mt3: Meta test-time training for self-supervised test-time adaption. arXiv preprint arXiv:2103.16201 (2021) 2, 3
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. Machine learning **79**(1), 151–175 (2010)
 2
- Benaim, S., Wolf, L.: One-shot unsupervised cross domain translation. In: NeurIPS (2018) 2, 4, 12
- Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: CVPR (2017) 3
- Chen, M., Xue, H., Cai, D.: Domain adaptation for semantic segmentation with maximum squares loss. In: ICCV (2019) 8
- Chen, Y.C., Lin, Y.Y., Yang, M.H., Huang, J.B.: Crdoco: Pixel-level domain transfer with cross-domain consistency. In: CVPR (2019) 3
- 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) 13
- Choi, S., Jung, S., Yun, H., Kim, J.T., Kim, S., Choo, J.: Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In: CVPR (2021) 2, 4, 12, 13, 14
- Cohen, T., Shulman, N., Morgenstern, H., Mechrez, R., Farhan, E.: Self-supervised dynamic networks for covariate shift robustness. arXiv preprint arXiv:2006.03952 (2020) 4
- 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) 8
- 12. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPR Workshop (2020) 7, 8, 9
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021) 14
- Dubey, A., Ramanathan, V., Pentland, A., Mahajan, D.: Adaptive methods for real-world domain generalization. In: CVPR (2021) 2
- Dundar, A., Liu, M.Y., Wang, T.C., Zedlewski, J., Kautz, J.: Domain stylization: A strong, simple baseline for synthetic to real image domain adaptation. arXiv preprint arXiv:1807.09384 (2018) 3
- Fan, X., Wang, Q., Ke, J., Yang, F., Gong, B., Zhou, M.: Adversarially adaptive normalization for single domain generalization. In: CVPR (2021) 4
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. JMLR 17(1), 2096–2030 (2016) 2, 3
- Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: ICML (2017) 11
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: ICCV (2021) 7, 8, 9

- 16 Y. Zou et al.
- Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple data processing method to improve robustness and uncertainty. In: ICLR (2020) 7, 8, 9
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: ICML (2018) 3
- Hu, X., Uzunbas, G., Chen, S., Wang, R., Shah, A., Nevatia, R., Lim, S.N.: Mixnorm: Test-time adaptation through online normalization estimation. arXiv preprint arXiv:2110.11478 (2021) 2, 4
- Huang, L., Zhou, Y., Zhu, F., Liu, L., Shao, L.: Iterative normalization: Beyond standardization towards efficient whitening. In: CVPR (2019) 4, 13
- Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML (2015) 4
- Khurana, A., Paul, S., Rai, P., Biswas, S., Aggarwal, G.: Sita: Single image test-time adaptation. arXiv preprint arXiv:2112.02355 (2021)
- Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: CVPR (2019) 13
- Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: ICCV (2019) 2, 4
- Li, Y., Wang, N., Shi, J., Liu, J., Hou, X.: Revisiting batch normalization for practical domain adaptation. In: ICLR Workshop (2017) 4, 6
- 29. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: ICML (2020) 2, 3, 4
- Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: ICML (2015) 3
- Luo, Y., Liu, P., Guan, T., Yu, J., Yang, Y.: Adversarial style mining for one-shot unsupervised domain adaptation. In: NeurIPS (2020) 2, 4, 12
- Luo, Y., Zheng, L., Guan, T., Yu, J., Yang, Y.: Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In: CVPR (2019) 2, 3, 12
- Ma, N., Zhang, X., Zheng, H.T., Sun, J.: Shufflenet v2: Practical guidelines for efficient cnn architecture design. In: ECCV (2018) 13
- Matsuura, T., Harada, T.: Domain generalization using a mixture of multiple latent domains. In: AAAI (2020) 2, 4
- Mirza, M.J., Micorek, J., Possegger, H., Bischof, H.: The norm must go on: Dynamic unsupervised domain adaptation by normalization. arXiv preprint arXiv:2112.00463 (2021) 4
- Nado, Z., Padhy, S., Sculley, D., D'Amour, A., Lakshminarayanan, B., Snoek, J.: Evaluating prediction-time batch normalization for robustness under covariate shift. arXiv preprint arXiv:2006.10963 (2020) 2, 3, 4, 9, 10, 11
- Neuhold, G., Ollmann, T., Rota Bulo, S., Kontschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV (2017) 8
- Pan, X., Luo, P., Shi, J., Tang, X.: Two at once: Enhancing learning and generalization capacities via ibn-net. In: ECCV (2018) 13, 14
- 39. Pan, X., Zhan, X., Shi, J., Tang, X., Luo, P.: Switchable whitening for deep representation learning. In: ICCV (2019) 4, 13
- 40. Qiao, F., Zhao, L., Peng, X.: Learning to learn single domain generalization. In: CVPR (2020) 2, 4
- 41. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016) 8

Learning Instance-Specific Adaptation for Cross-Domain Segmentation

- Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: CVPR (2016) 8
- 43. Rush, A.: Tensor considered harmful, http://nlp.seas.harvard.edu/ NamedTensor 5
- Sakaridis, C., Dai, D., Van Gool, L.: Semantic foggy scene understanding with synthetic data. IJCV 126(9), 973–992 (Sep 2018) 13
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: CVPR (2018) 13
- 46. Schneider, S., Rusak, E., Eck, L., Bringmann, O., Brendel, W., Bethge, M.: Improving robustness against common corruptions by covariate shift adaptation. In: NeurIPS (2020) 2, 3, 4, 6, 9, 10
- Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR (2017) 3
- Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: ECCV (2016) 3
- 49. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., Hardt, M.: Test-time training with self-supervision for generalization under distribution shifts. In: ICML (2020) 2, 3, 4
- 50. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR (2017) 2, 3
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv preprint arXiv:1412.3474 (2014) 3
- Volpi, R., Namkoong, H., Sener, O., Duchi, J.C., Murino, V., Savarese, S.: Generalizing to unseen domains via adversarial data augmentation. In: NeurIPS (2018) 4
- Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: CVPR (2019) 2, 3, 12
- 54. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. In: ICLR (2021) 2, 3, 4, 6, 9, 10, 11
- Wu, Y., Johnson, J.: Rethinking" batch" in batchnorm. arXiv preprint arXiv:2105.07576 (2021) 5
- 56. Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., Lipson, H.: Understanding neural networks through deep visualization. In: ICML Workshop (2014) 6
- 57. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: CVPR (2020) 8
- Zendel, O., Honauer, K., Murschitz, M., Steininger, D., Dominguez, G.F.: Wilddash - creating hazard-aware benchmarks. In: ECCV (2018) 8
- 59. Zhao, L., Liu, T., Peng, X., Metaxas, D.: Maximum-entropy adversarial data augmentation for improved generalization and robustness. In: NeurIPS (2020) 2, 4
- Zhao, S., Gong, M., Liu, T., Fu, H., Tao, D.: Domain generalization via entropy regularization. In: NeurIPS (2020) 2, 4
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017) 3
- Zou, Y., Yu, Z., Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV (2018) 3, 12
- Zou, Y., Yu, Z., Liu, X., Kumar, B., Wang, J.: Confidence regularized self-training. In: ICCV (2019) 3