

# Domain Bridge for Unpaired Image-to-Image Translation and Unsupervised Domain Adaptation

Fabio Pizzati<sup>1,2</sup>, Raoul de Charette<sup>2</sup>, Michela Zaccaria<sup>3</sup>, Pietro Cerri<sup>1</sup>

<sup>1</sup>VisLab, <sup>2</sup>Inria, <sup>3</sup>University of Parma

{fabio.pizzati, raoul.de-charette}@inria.fr

michela.zaccaria@studenti.unipr.it

pcerri@ambarella.com

## Abstract

Image-to-image translation architectures may have limited effectiveness in some circumstances. For example, while generating rainy scenarios, they may fail to model typical traits of rain as water drops, and this ultimately impacts the synthetic images realism. With our method, called domain bridge, web-crawled data are exploited to reduce the domain gap, leading to the inclusion of previously ignored elements in the generated images. We make use of a network for clear to rain translation trained with the domain bridge to extend our work to Unsupervised Domain Adaptation (UDA). In that context, we introduce an online multimodal style-sampling strategy, where image translation multimodality is exploited at training time to improve performances. Finally, a novel approach for self-supervised learning is presented, and used to further align the domains. With our contributions, we simultaneously increase the realism of the generated images, while reaching on par performances with respect to the UDA state-of-the-art, with a simpler approach.

## 1. Introduction

GANs have demonstrated great ability to learn image-to-image (i2i) translation from paired [15] or unpaired images [41, 13] in different domains (e.g. summer/winter, clear/rainy, etc.). The latter relies on cycle-consistency or style/content disentanglement to learn complex mapping in an unsupervised manner, producing respectively a single translation of the source image or a multimodal translation [13]. This unsupervised i2i translation opened a wide range of applications especially for autonomous driving, for which it may be *virtually* impossible to acquire the same scene in different domains (e.g. say clear/rainy) due to the dynamic nature of the scenes.

Instead, image-to-image translation can be used to gen-

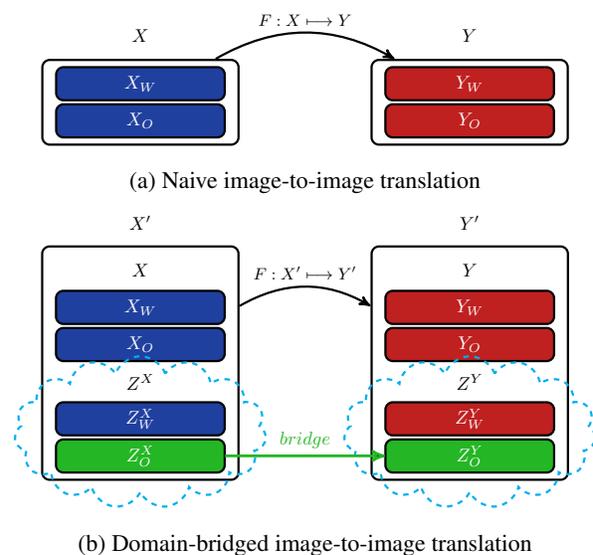


Figure 1: Naive image-to-image translation (Fig. 1a) learns the  $X \rightarrow Y$  domain mapping. Conversely, our *domain bridge* (Fig. 1b) completes source and target domains with automatically retrieved web-crawled data ( $Z^X, Z^Y$ ) which share common characteristics, thus easing the image-to-image translation task.

erate realistic image synthesis exploitable for both domain adaptation and performance evaluation [10, 17], without additional human-labeling effort. However, state-of-the-art i2i translation fails in some situations. For example, while translating clear to rain the networks tend to change only the global appearance of the scene (wetness, puddle, etc.) ignoring essential traits as drops on the windshield or reflections. Ultimately, this greatly impacts the realism of the generated images.

In this paper, we present a simple domain-bridging technique (Fig. 1b) which, opposite to the standard i2i transla-

tion (Fig. 1a), benefits of additional sub-domains retrieved automatically from web-crawled data. Our method produces qualitatively significantly better results, especially when the source and target domains are known to be far since the bridge ease the learning of the mapping. We apply our i2i methodology to the case of clear  $\rightarrow$  rainy images showing that domain bridging leads the translation to preserve drops on the synthetic images, and extend our work to Unsupervised Domain Adaptation (UDA) for which we make novel contributions too (Fig. 2) and demonstrate that all together we perform on par with the most recent UDA methods while being much simpler.

We make three main contributions in our paper:

- **i2i:** we propose a novel *domain-bridge* (Sec. 3.1) to augment automatically the source and target domains and ease i2i mapping,
- **i2i with UDA:** *online multimodal style-sampling* (Sec. 3.2) is applied for UDA, thus increasing the translation diversity,
- **UDA:** we propose novel *Weighted Pseudo Label* (Sec. 3.2) to benefit from self-supervision without the need of offline processing as for the original Pseudo Label [16].

## 2. Related work

**Image-to-image translation.** Early work for image-to-image translation has been done in [15], where an adversarial-based method has been proposed. The training process required paired samples of the same scene in two different domains. In [41], instead, cycle consistency is exploited to perform image-to-image translation on unpaired images. [18] supposes the existence of a shared latent space between images in two domains, and exploits it to perform translations across both domains using a single GAN. Recently, a lot of efforts have been dedicated to achieve multimodal translation [42, 13, 19]. Some others, instead, make use of additional information, such as bounding boxes [33], instance segmentation maps [22], or semantic segmentation and depth maps [3], to increase the translation quality and diversity. In [8], intermediate domains are used to achieve better performances in domain adaptation.

**Synthetic rain modeling.** To synthesize rain on images, Garg and Nayar [7] first proposed to rely on the accurate drop oscillation and photometry modeling which has been extended in [5, 34] for rain streaks and [29] for stationary drops. All assume the impractical complete 3D knowledge of the scene geometry and light sources. Circumventing such limitations, the recent physics-based rendering [9, 12] rely on the rain/fog layers decomposition to simulate the complete rain visual effect. In [9], the rain dynamics and

photometry is approximated from physics models to realistically augment image or sequences with controllable amount of rain, which further allow benchmarking vision algorithms in rain. Alternately, [25] proposed to synthesize photorealistic raindrops on a screen using a simplified refraction model and approximating their dynamics with metaballs. While priors works do produce high quality visual results, none of them face visual characteristics of rainy images, as wetness and reflections.

**Domain adaptation for semantic segmentation.** Most methods for domain adaptation are based on adversarial training, as it regularizes the feature extraction process, making it robust to the domain shift [11, 2, 21, 32, 38, 39, 35, 20, 27]. Complementary approaches, instead, connect the two domains with pixel-level regularization, making use of image-to-image translation GANs [23, 26]. Some recent works combine the two to obtain better results [10, 17]. Others do not use adversarial training at all: for example, Zou et al. [43] exploit self-supervised learning and pseudo-labels only, while [40] make use of mutual learning. Finally, some approaches have been recently introduced specialized on night time adaptation [28] or transfer learning for adverse weather [6, 30].

## 3. Method

Our methodology aims to translate clear images to rainy images reaching both high-qualitative images for both qualitative evaluation and usability to train semantic segmentation networks in rainy weather. Thus, our innovations are spread between image-to-image translation (Sec. 3.1) and Unsupervised Domain Adaptation for semantic segmentation (Sec. 3.2).

### 3.1. Image-to-image translation (i2i)

Image-to-image translation GAN networks learn to approximate the mapping  $F : X \mapsto Y$  using adversarial training, from two sets of representative images in each domain, denoted here  $A$  and  $B$ . Each image in both sets can be interpreted as a sampling from a probability distribution  $P$  associated with the image domain [18]. Formally,  $\forall a \in A, a \sim P_X(x); \forall b \in B, b \sim P_Y(y)$ . GANs are well-known for their instability at training. For the latter to succeed, the network needs to be fed with representative sets of images, so that it can extract the common domain characteristics. Even though, some domain gap may be difficult to model for the network, resulting in a loss of characteristic features of the target domain. This may be caused by a significant domain shift or by the lack of data.

Some minor image details still have a significant perceptual impact. This is the case for rain images, where even a few drops on the lens play an important role in *sensing* the

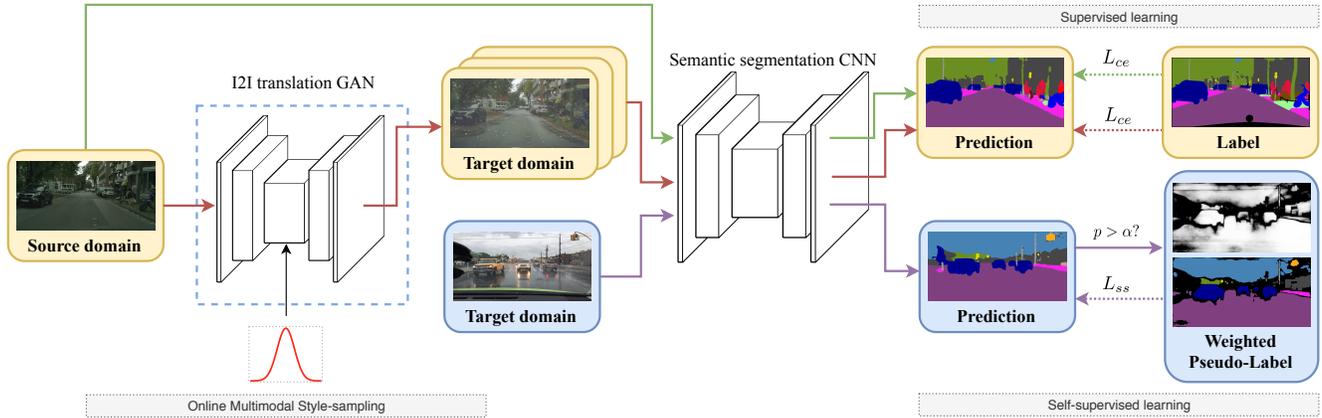


Figure 2: Overview of our pipeline for unsupervised domain adaptation. The blue dashed square means that the GAN parameters are frozen. The Image-to-image translation network is trained offline with our domain bridge strategy. Different line colors refer to different probabilities for one path to be executed. Loss functions are denoted with dotted lines.



Figure 3: Rainy images from the Berkeley Deepdrive dataset [36]. Drops on the windshield or reflections help us perceive that it is raining.

rain, as it can be seen in Fig. 3. State-of-the-art networks may ignore some fundamental elements as drops, lens artifacts or reflections, and this ultimately impacts the realism of generated images. We argue that some characteristics changes (drops, artifacts, etc.) are ignored because they are relatively minor compared to other characteristics changes (e.g. wetness, puddles, etc.), and demonstrate the training may benefit from bridging to ease domain mapping.

**Domain bridging.** Studying the specific case of adverse weather conditions, it is possible to formalize a generic domain  $K$  as the union of finer-grained domains, such as  $K = \{K_W, K_O\}$ . In it,  $K_W$  represents the sub-domains typical of weather, e.g. the presence of precipitations, road wetness, and many more.  $K_O$ , instead, is composed of sub-domains unrelated to the weather. Examples are the scenario, the city, and the illumination. Thus, it is possible to represent  $X$  and  $Y$  as

$$\begin{aligned} X &= \{X_W, X_O\}, \\ Y &= \{Y_W, Y_O\}. \end{aligned} \quad (1)$$

Generally, only the joint probability distribution  $P_K(k) = P_{K_W, K_O}(k)$ ,  $k \in K$  is estimable, as we have no knowl-

edge about the marginal probability distributions  $P_{K_W}(k)$ ,  $P_{K_O}(k)$ .

On one hand, we hypothesize that it is possible to obtain a more stable image-to-image translation if the differences between the two datasets are minimized. On the other, we have to obtain a GAN that produces an effective transformation, so it is necessary to model correctly all relevant features of adverse weather. To simultaneously reach both objectives, images collected from web-crawled videos are added to the  $A$  and  $B$  datasets, obtaining two new training sets  $A'$  and  $B'$ , with respective domains  $X'$  and  $Y'$ , which aims as bridging the gap between the initial domain  $X$  and  $Y$ . This is illustrated in Fig. 1.

Our intuition is that adding samples with reasonable criteria will reduce the Kullback-Leibler divergence between the probability distributions  $P_{X'_O}(x)$  and  $P_{Y'_O}(y)$ , with respect to  $KL(P_{X'_O}(x), P_{Y'_O}(y))$ . As a consequence, the translation model will be more focused on weather-related characteristics and more stable during training.

Once the main hypotheses are formalized, the approach on how to select new images is needed. Let  $Z^x$  and  $Z^y$  be two images sets. As before, we have

$$\begin{aligned} Z^x &= \{Z^x_W, Z^x_O\}, \\ Z^y &= \{Z^y_W, Z^y_O\}. \end{aligned} \quad (2)$$

We choose  $Z^x$  and  $Z^y$  in order to have

$$\begin{aligned} \max |Z^x_W \cap X_W|, \\ \max |Z^y_W \cap Y_W|, \\ \max |Z^x_O \cap Z^y_O|, \end{aligned} \quad (3)$$

where  $|\cdot|$  is the set cardinality. Hence, it is possible to identify two image sets  $C$  and  $D$  such as  $\forall c \in C, c \sim P_{Z^x}(z); \forall d \in D, d \sim P_{Z^y}(z)$ .



Figure 4: Examples of multimodal style-sampling from our Domain-bridged i2i (Sec. 3.1). Note the consistency of style regardless of the source image.

It is now possible to train the image-to-image translation network on  $A'$  and  $B'$  defined as:

$$\begin{aligned} A' &= A \cup C, \\ B' &= B \cup D. \end{aligned} \quad (4)$$

Adding new images, the differences in the global appearance of the two domains is minimized, while the weather-related domain shift remains constant.

In other words, our approach consists in selecting new image samples, with weather conditions corresponding to those in the original dataset, and join them to the existing data. The newly-added images are required to share some domains unrelated to the weather. Retrieving images from the same location and with the same setup ensures that.

In practice, we retrieved these additional samples from web-crawled videos using domain-related keywords search (details in Sec. 4.1.1). The same bridging can be applied automatically to other domain shifts, though as the domain differences become less semantically evident, human expertise may be required to properly select the  $C$  and  $D$  datasets.

We use MUNIT [13] as backbone for our image-to-image translation network, as it allows disentanglement of style and content, which will be of high interest for us.

### 3.2. Unsupervised Domain Adaptation (UDA)

Similar to previous works [10, 17], we use our i2i network with domain-bridge to translate images from pre-

labeled *clear weather* datasets and learn semantic segmentation in rain in an unsupervised fashion. We follow the standard UDA practice which is to alternately train in a supervised manner from source (*clear*) images with labels and train in a self-supervised manner from target (*rain*) images without labels. Our entire UDA methodology (depicted in Fig. 2) brings two novel contributions: a) we use multi-modal clear→rain translations as additional supervised learning, b) we introduce Weighted Pseudo Label - a differentiable extension of Pseudo Label [16] - to align source and target without any offline process.

**Online Multimodal Style-sampling (OMS).** The standard strategy for UDA with i2i networks is to learn from the offline translation of the whole dataset [10, 17].

We instead propose to use the multimodal capacity of MUNIT i2i to generate multiple target styles (i.e. *rain appearances*) for each input image. Styles are sampled during training time. In this way, even if the source image content remains unaltered, it will be possible for the segmentation network to learn different representations of the same scene in the target domain, ultimately leading towards wider diversity and thus more robust detection. Different styles for the same image modify, among other factors, the position and size of drops on the windshield, and the intensity of reflections. This is visible in Fig. 4 showing three arbitrarily sampled styles, where Style 3 consistently produces images that resemble heavy rain.

**Weighted Pseudo-Labels (WPL).** Pseudo Label [16] was proposed as a self-supervised loss to further align source and target distributions. The principle is to self-train a network on target (here, *rain*) whenever the prediction confidence is above some threshold, thus reinforcing the network beliefs.

Most often for UDA, thresholds are calculated offline as the median per-class confidence dataset-wise [17, 43]. This requires storing all predictions for the whole dataset, which is cumbersome. To overcome this, thresholds may be estimated online image-wise or batch-wise [14]. It has to be noted that thresholds are critical since pseudo-labeling can harm global performances if thresholds were underestimated<sup>1</sup> or have limited impact if overestimated [16].

We instead propose Weighted Pseudo-Labels (WPL) which estimates a global threshold  $\alpha$  within the network optimization process. The general principle of WPL is to weight the self-supervised cross-entropy using learned threshold  $\alpha$ , thus acting as continuous pseudo-labeling. Not only WPL does not require offline processing, but it is aware of the global network confidence thus leading to better results. In detail, let  $x$  be an input image and  $\hat{x}_u$  the pseudo

<sup>1</sup>In such case, the ratio of wrong pixels over pseudo-labeled pixels will be too high and lead to incorrect self-supervision.

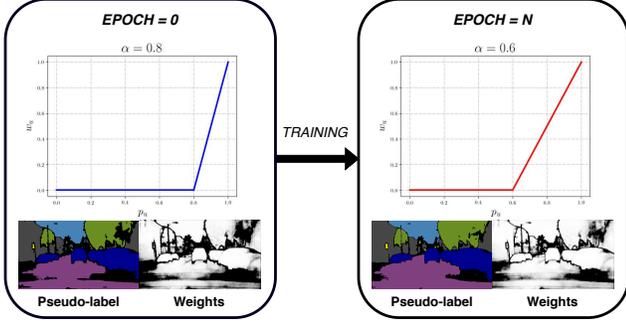


Figure 5: Analysis of the effect of  $\alpha$  optimization. During training the Weighted Pseudo Label expands from high confidence pixels only (left) to lower ones (right).

label of pixel  $u$ , such that

$$\hat{x}_u = \begin{cases} \operatorname{argmax}_q f_u(x) & \text{if } \max(f_u(x)) \geq \alpha \\ \text{None} & \text{otherwise} \end{cases}, \quad (5)$$

where  $f_u(x)$  refers to the class probabilities of  $u$  predicted by the network  $f$ , and  $\operatorname{argmax}_q$  is the best class prediction. In its original implementation [16],  $\hat{x}_u$  is directly used to weight the cross-entropy self-supervision. Instead, we weight this with a weight matrix  $W$  of the same size than  $x$ :

$$w_u = \begin{cases} \frac{\max(f_u(x)) - \alpha}{1 - \alpha} & \text{if } \max(f_u(x)) \geq \alpha \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

The complete loss for WPL is thus defined as the weighted sum of cross-entropy loss  $L_w$  and a balancing loss  $L_b$ :

$$L_{ss} = \sigma L_w + \gamma L_b, \quad (7)$$

where  $\sigma$  and  $\gamma$  are loss weights and cross-entropy loss is:

$$L_w = - \sum_{q \in Q} w_u \hat{x}_{u,q} \log(p_u). \quad (8)$$

$\hat{x}_u$  is the one-hot encoding of pseudo-label as in Eq. 5 and  $Q$  is the set of classes. In this way, predictions where the network is uncertain are weighted less in the network pseudo-label based training. To avoid that the self-supervised contribution remains set to zero by the optimizer,  $L_b$  is required as a balancing loss:

$$L_b = \log^2(1 - \alpha). \quad (9)$$

The optimization of  $\alpha$  leads to a pseudo-label expansion within the training process. Fig. 5 is an illustration of the growing process during training. For the first iteration (Fig. 5, left), the  $L_w$  term prevails over  $L_b$ , pushing  $\alpha$  towards 1 thus including in the pseudo-label only pixels with high confidence. With the minimization of  $L_w$  (Fig. 5,

right),  $L_b$  becomes gradually more important, leading the network to simultaneously include lower confidence pixels inside the pseudo-label, and increasing the informative potential of higher-confidence labels. Note that for numerical stability, we assume  $\alpha = \operatorname{sigmoid}(\beta)$  and estimate  $\beta$ .

**Losses.** To balance the self-supervised WPL contribution with the supervised learning in segmentation, we employ a probability-based approach where pseudo-label is applied only if a uniformly sampled variable  $p_{pl} \in U(0, 1)$  is above a predefined threshold  $p_{tp}$ . Hence, the complete UDA loss function is:

$$L(x_a, \hat{x}_a, x_b) = L_{ce}(f(x_a), \hat{x}_a) + L_{ss}(f(x_b))[p_{pl} > p_{tp}] \quad (10)$$

if we train on source data + target, and

$$L(x_a, \hat{x}_a, x_b) = L_{ce}(f(g(x_a)), \hat{x}_a) + L_{ss}(f(x_b))[p_{pl} > p_{tp}] \quad (11)$$

if we train on translated images + target. In Eq. 10 and 11,  $x_a \in A$  is source image with label  $\hat{x}_a$ ,  $x_b \in B$  target image,  $L_{ce}$  the cross-entropy loss,  $f$  our segmentation network,  $g$  our bridged-GAN, and  $[\cdot]$  are the Iverson brackets.

## 4. Experiments

We now evaluate the performance of both our i2i proposal (Sec. 4.2) and our UDA proposal (Sec. 4.3) on the clear→rain problem using clear/rain datasets recorded with different setups.

### 4.1. Experimental settings

#### 4.1.1 Datasets

For i2i and UDA, we use the german dataset Cityscapes [4] as source (*clear*), and a subset of the American Berkeley DeepDrive [36] (BDD) as target (*rain*). The bridge dataset, only used for the i2i, is a collection of Youtube videos. We now detail each dataset.

**Cityscapes.** We train on Cityscapes training set with 2975 pixel-wise annotated images, and evaluate on their validation set with 500 images. While we train on crops, we evaluate on full-size images, i.e.  $2048 \times 1024$ . The *trainExtra* set, with 19997 images, is also included in the domain bridge to further reduce the domain shift.

**BDD-rainy.** We use the coarse weather annotations of BDD together with daylight annotation to obtain a subset we call BDD-rainy (i.e. rainy+daylight), i.e.  $1280 \times 720$ . For training the rainy+daylight is extracted from the 100k split, while for validation only the 10k split is used. Obviously, duplicates present in both splits are removed. It has to



Figure 6: Samples from the bridge dataset in different weather conditions. Note that the acquisition setup (camera position, optics, etc.) remains unaltered.

Network	LPIPS	IS	Network	mIoU
Real images	0.7137	-	Baseline	31.67
CycleGAN [41]	0.1146	1.15	CycleGAN [41]	35.09
MUNIT [13]	<b>0.3534</b>	<b>1.92</b>	MUNIT [13]	20.78
MUNIT-Bridged	0.2055	1.69	MUNIT-Bridged	<b>35.18</b>

(a) GAN metrics                      (b) Semantic Segmentation

Table 1: Quantitative evaluation of translated image realism, diversity, and semantic segmentation effectiveness.

be noted that, while daylight annotation is accurate, weather annotation is approximate and "rainy" images may either be taken during or *after* a rain event, thus with or without drops on the windshield. This further increases complexity.

**Bridge dataset.** 5 sequences ( $1280 \times 720$ ) were extracted from a single Youtube channel with keywords "driving" (2 videos) for clear weather and "driving rain"/"driving heavy rain" (2/1 video) for rainy scenarios. The choice of using videos from a unique channel further reduces domain gaps, ensuring the same acquisition setup. Some samples are shown in Fig. 6. Also to maximize image diversity, videos are uniformly sub-sampled into  $2 \times 6026$  clear weather images and  $3 \times 9294$  rainy images, leading to a total of 39934 frames.

#### 4.1.2 Networks details

**Image-to-image translation** During training, the images are downsampled to be 720 pixels in height, and cropped to  $480 \times 480$  resolution. The network is trained for 200k iterations, with batch size 1. Adam is used as optimizer, with learning rate  $1e-4$ ,  $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$ .

**Segmentation.** We use Light-weight Refinenet [24] with Resnet-101 as backbone, pretrained on the full-size Cityscapes dataset. The refining is achieved by training for 100 epochs on  $512 \times 512$  crops, after downscaling images to  $1024 \times 512$  for GPU memory constraints. We employ data augmentation for the training process, with random rescaling between a factor 0.5 and 2, and random horizontal

flipping. The batch size used is 6. We use the SGD optimizer with different learning rates for the encoder ( $1e-4$ ) and the decoder ( $1e-3$ ). The momentum is set to 0.9, and the learning rate is divided by 2 every 33 epochs. When pseudo-labels are added to the training, we further refine the network for 70 additional epochs, with constant learning rate divided by 10 with respect to the initial values. The  $\alpha$  parameter is initialized to 0.8 and estimated by SGD as well, with learning rate 0.01 and momentum 0.9.

#### 4.2. Bridged image-to-image translation

We evaluate our bridged i2i (Sec. 3.1) on the Cityscapes to BDD-rainy translation task, and compare results against the recent CycleGAN [41] and MUNIT [42]. As stated, our i2i uses a MUNIT based and is referred to as MUNIT-bridged. It is trained on the bridged versions of the two datasets. Training follow details from Sec. 4.1.2, except for CycleGAN that follows the original implementation<sup>2</sup>. We argue - like others - that GAN metrics aren't appropriate for such evaluation. Thus, we report qualitative evaluation and segmentation task evaluation, together with usual GAN metrics.

**Qualitative evaluation.** Fig. 7 shows randomly selected samples from the Cityscapes validation set<sup>3</sup>. It is visible that both CycleGAN and original MUNIT method fails at modeling the rain appearance, probably due to the large domain gap. In detail, CycleGAN brings no realistic changes to the scene appearance, only adjusting color-levels in the image. Original MUNIT, instead, seems to have collapsed and fails to produce significant outputs, probably due to instability related to the domain gap. Conversely, our MUNIT-bridged model is the only one able to add realistic traits of rain in the synthetic images, thanks to the domain bridge.

**Quantitative evaluation.** We compute GAN metrics following usual practices from [13] and report results in Tab. 1a. The LPIPS distance [37, 33] measures the image diversity [13], while the Inception Score evaluates both quality and diversity [31]. In detail, LPIPS is the average on 19 paired translation of 100 images, and we report the diversity of real data in the target dataset as *upper bound*. Inception Score uses the InceptionV3 network previously trained to classify source and target images.

Overall, we successfully improve performances over CycleGAN in both metrics, but original MUNIT has significantly higher performance. However, the images generated by MUNIT are evidently unrealistic (cf. Fig. 7) and thus

<sup>2</sup>[41] claims that best performances are obtained keeping constant the learning rate for half the training process (100k iterations in this case) to  $2e-4$  and then linearly decreasing to 0.

<sup>3</sup>For MUNIT and our method, MUNIT-bridged, we also randomize the style.

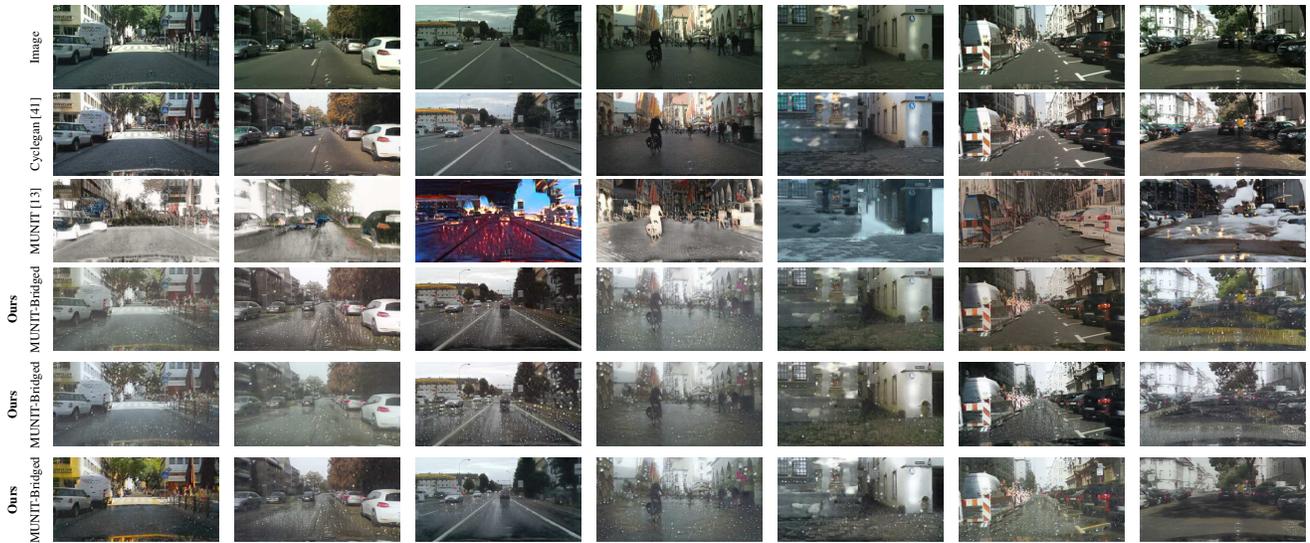


Figure 7: Qualitative comparison between state-of-the-art architectures for i2i in the clear  $\rightarrow$  rain transformation.

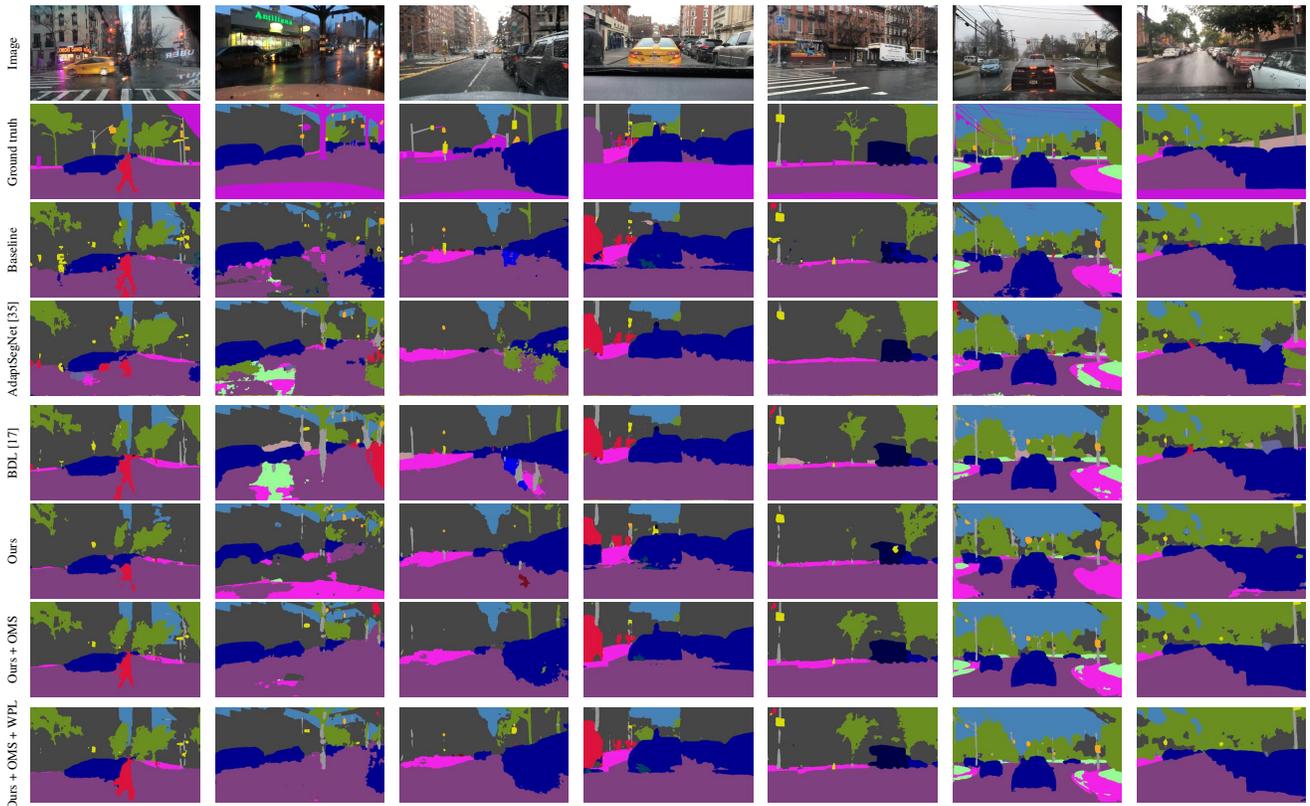


Figure 8: Comparison of our method with the state-of-the-art for semantic segmentation UDA.

we argue that GAN metrics are unreliable which is in fair alignment with [1] advocating that Inception Score is uncorrelated with image quality.

For a more comprehensive evaluation, we train a seg-

mentation task on GAN translated clear  $\rightarrow$  rain images, and evaluate the standard mIoU metric on real rain images, reporting results in Tab. 1b. Note that for a fair comparison, we only sample a single style for MUNIT-based mod-

Method	mIoU	road	sidewalk	building	wall	fence	pole	t light	t sign	veg	terrain	sky	person	rider	car	truck	bus	train	m. bike	bike
Baseline	31.67	77.40	39.95	61.20	12.01	24.76	23.68	13.21	24.11	58.33	27.18	78.86	24.73	12.78	63.34	24.01	28.43	0.00	4.76	2.90
AdaptSegNet [35]	33.44	82.23	39.85	62.06	9.84	17.73	20.39	10.91	22.47	66.30	22.81	76.54	32.24	38.49	68.95	13.08	30.31	0.00	17.97	3.26
BDL [17]	39.60	83.18	<b>48.78</b>	<b>73.93</b>	<b>30.87</b>	<b>27.33</b>	26.03	<b>15.10</b>	26.05	72.63	26.08	<b>88.01</b>	28.59	26.59	<b>76.37</b>	43.31	<b>50.11</b>	0.00	7.38	2.10
Ours	35.18	79.00	37.23	62.36	8.60	14.78	20.98	11.94	22.92	68.02	13.11	82.55	38.96	<b>44.61</b>	72.34	29.10	39.40	0.00	19.32	3.16
Ours + OMS	39.72	82.53	44.51	69.97	20.29	22.91	<b>28.93</b>	14.02	29.17	74.32	<b>28.98</b>	83.53	36.75	32.80	71.29	43.03	46.34	0.00	<b>21.80</b>	3.54
Ours + OMS + WPL	<b>40.04</b>	<b>84.03</b>	44.09	70.51	24.10	23.02	28.31	14.08	<b>30.07</b>	<b>75.31</b>	27.89	83.49	<b>39.10</b>	33.63	74.70	<b>48.60</b>	49.34	0.00	6.77	<b>3.80</b>

Table 2: State-of-the-art comparison. OMS refers to Online Multimodal Style-sampling. WPL is the Weighted Pseudo Labels strategy.

Pseudo-labels	mIoU target
None	39.77
Batch-wise	38.23
WPL (Ours)	<b>40.04</b>

Table 3: Comparison of various Pseudo Labeling strategies: Batch-wise, with our WPL, or with None.

els, and report results when only trained on clear images as *baseline*. If the domain gap were reduced by the GAN translations, an improvement should be visible.

Instead, from the table, training on the original MUNIT-translated dataset leads to a significant performance decrease disproving the high GAN metrics. Finally, our method outperforms CycleGAN by a little margin although CycleGAN fails to produce good quality images. Conversely, our method simultaneously reduces the domain shift and increases realism, thus it also eases performances evaluation on synthetic data.

### 4.3. Unsupervised Domain Adaptation

We now evaluate our UDA contributions encompassing our i2i translation methodology and compare with AdaptSegNet [35] and BDL [17], the best found recent works. We do not compare with less recent methods as CyCADA [10] since the approaches we are evaluating have already demonstrated to guarantee superior performances in UDA [35, 17]. For fair comparison and given architectural similarities, BDL was adapted to work with the same segmentation network, data augmentation policy and hyperparameters detailed in Sec. 4.1.2.

Quantitative results are shown in Tab. 2 where *Ours* refers to UDA with only our domain-bridge i2i translation, *Ours+OMS* using also our Online Multimodal Style-sampling, and *Ours+OMS+WPL* using also our Weighted Pseudo Label. *baseline* refers to the training without any UDA. Overall, our methodology performs on par (+0.44) with BDL the best state-of-the-art, using a much simpler domain adaptation method, and significantly better (+6.6) than AdaptSegNet. Studying the contributions of our OMS and WPL contributions, all components are necessary to reach the best performances. Qualitative evaluation on the target dataset is shown in Fig. 8, and in fair alignment with quantitative metrics.

**Weighted Pseudo Labels.** We evaluate the effectiveness of our WPL proposal and report results in Tab. 3, comparing similar training with either *WPL* (Ours), *batch-wise* Pseudo-Label<sup>4</sup>, or *None*. For all, the training is performed using as target the whole BDD100k train set (removing duplicates from 10k split) together with the rainy sequences from Domain-bridge dataset, resulting in over 90k images. Performance is reported on target BDD-rainy. We do not compare against offline Pseudo Label, as this would be impractical with such a big dataset, and this evaluation is partly encompassed in BDL comparison (cf. Tab. 2). For WPL, we empirically set  $\gamma = 1$ ,  $\sigma = 0.005$  (Eq. 7) to balance contributions and  $p_{tp} = 0.75$  (Eq. 10 & 11).

From results in Tab. 3, *WPL* performs the best and *batch-wise* Pseudo Label third. In fact, the performance decrease for *batch-wise* (compared to no self-supervision) may be explained since best batch pixels are used as pseudo-labels, thus possibly implying some incorrect self-supervision in case of low batch accuracy. Instead, our *WPL* boosts the mIoU on the target dataset which is expected due to its expansion behavior.

## 5. Conclusions

In this work, we introduced a novel approach to generate realistic rainy images with an i2i network, while preserving traits of adverse weather that are typically ignored by state-of-the-art architectures. Then, we extended our system and demonstrated its performances in UDA for semantic segmentation, and with a simple pipeline we obtained on par performances with respect to the state-of-the-art. Finally, we introduced a novel pseudo labeling strategy that works with an unlimited number of images, and has an optimizable weight parameter used to guide region growing. For future work, we plan to extend our pseudo labeling approach with class-wise thresholds, and to extend the method to different adverse weather conditions.

## References

- [1] S. T. Barratt and R. Sharma. A note on the inception score. In *ICML Workshops*, 2018.

<sup>4</sup>For batch-wise Pseudo-Label implementation, we compute optimal threshold per class and per batch.

- [2] M. Biassetton, U. Michieli, G. Agresti, and P. Zanuttigh. Un-supervised domain adaptation for semantic segmentation of urban scenes. In *CVPR Workshops*, 2019.
- [3] Y. Chen, W. Li, X. Chen, and L. V. Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *CVPR*, 2019.
- [4] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [5] C. Creus and G. A. Patow. R4: Realistic rain rendering in realtime. *Computers & Graphics*, 2013.
- [6] D. Dai, C. Sakaridis, S. Hecker, and L. Van Gool. Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *IJCV*, 2019.
- [7] K. Garg and S. K. Nayar. Photorealistic rendering of rain streaks. In *SIGGRAPH*, 2006.
- [8] R. Gong, W. Li, Y. Chen, and L. V. Gool. Dlow: Domain flow for adaptation and generalization. In *CVPR*, 2019.
- [9] S. S. Halder, J.-F. Lalonde, and R. de Charette. Physics-based rendering for improving robustness to rain. In *ICCV*, 2019.
- [10] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2017.
- [11] J. Hoffman, D. Wang, F. Yu, and T. Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016.
- [12] X. Hu, C.-W. Fu, L. Zhu, and P.-A. Heng. Depth-attentional features for single-image rain removal. In *CVPR*, 2019.
- [13] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018.
- [14] A. Iscen, G. Toliyas, Y. Avrithis, and O. Chum. Label propagation for deep semi-supervised learning. In *CVPR*, 2019.
- [15] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [16] D.-H. Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. 2013.
- [17] Y. Li, L. Yuan, and N. Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *CVPR*, 2019.
- [18] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017.
- [19] M.-Y. Liu, X. Huang, A. Mallya, T. Karras, T. Aila, J. Lehtinen, and J. Kautz. Few-shot unsupervised image-to-image translation. In *ICCV*, 2019.
- [20] Y. Luo, L. Zheng, T. Guan, J. Yu, and Y. Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, 2019.
- [21] U. Michieli, M. Biassetton, G. Agresti, and P. Zanuttigh. Adversarial learning and self-teaching techniques for domain adaptation in semantic segmentation. *arXiv preprint arXiv:1909.00781*, 2019.
- [22] S. Mo, M. Cho, and J. Shin. Instagan: Instance-aware image-to-image translation. In *ICLR*, 2019.
- [23] Z. Murez, S. Kolouri, D. Kriegman, R. Ramamoorthi, and K. Kim. Image to image translation for domain adaptation. In *CVPR*, 2018.
- [24] V. Nekrasov, C. Shen, and I. D. Reid. Light-weight refinenet for real-time semantic segmentation. In *BMVC*, 2018.
- [25] H. Porav, T. Bruls, and P. Newman. I can see clearly now: Image restoration via de-raining. *arXiv preprint arXiv:1901.00893*, 2019.
- [26] P. Z. Ramirez, A. Tonioni, and L. Di Stefano. Exploiting semantics in adversarial training for image-level domain adaptation. In *IPAS*, 2018.
- [27] P. Z. Ramirez, A. Tonioni, S. Salti, and L. Di Stefano. Learning across tasks and domains. In *ICCV*, 2019.
- [28] E. Romera, L. M. Bergasa, K. Yang, J. M. Alvarez, and R. Barea. Bridging the day and night domain gap for semantic segmentation. In *IV*, 2019.
- [29] P. Rousseau, V. Jolivet, and D. Ghazanfarpour. Realistic real-time rain rendering. *Computers & Graphics*, 2006.
- [30] C. Sakaridis, D. Dai, and L. Van Gool. Semantic foggy scene understanding with synthetic data. *IJCV*, 2018.
- [31] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *NeurIPS*, 2016.
- [32] S. Sankaranarayanan, Y. Balaji, A. Jain, S. Nam Lim, and R. Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *CVPR*, 2018.
- [33] Z. Shen, M. Huang, J. Shi, X. Xue, and T. Huang. Towards instance-level image-to-image translation. In *CVPR*, 2019.
- [34] N. Tatarchuk. Artist-directable real-time rain rendering in city environments. In *EUROGRAPHICS Workshops*, 2006.
- [35] Y.-H. Tsai, W.-C. Hung, S. Schuler, K. Sohn, M.-H. Yang, and M. Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018.
- [36] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. In *CVPR*, 2017.
- [37] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [38] Y. Zhang, P. David, and B. Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *ICCV*, 2017.
- [39] Y. Zhang, Z. Qiu, T. Yao, D. Liu, and T. Mei. Fully convolutional adaptation networks for semantic segmentation. In *CVPR*, 2018.
- [40] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu. Deep mutual learning. In *CVPR*, 2018.
- [41] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.
- [42] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. In *NeurIPS*, 2017.
- [43] Y. Zou, Z. Yu, B. Vijaya Kumar, and J. Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018.