Guided Curriculum Model Adaptation and Uncertainty-Aware Evaluation for Semantic Nighttime Image Segmentation

Christos Sakaridis[1], Dengxin Dai[1], and Luc Van Gool[1,2]

[1]ETH Zürich, [2]KU Leuven

Abstract

Most progress in semantic segmentation reports on daytime images taken under favorable illumination conditions. We instead address the problem of semantic segmentation of nighttime images and improve the state-of-the-art, by adapting daytime models to nighttime without using nighttime annotations. Moreover, we design a new evaluation framework to address the substantial uncertainty *of semantics in nighttime images. Our central contributions are: 1) a curriculum framework to gradually adapt semantic segmentation models from day to night via labeled synthetic images and unlabeled real images, both for progressively darker times of day, which exploits cross-time-of-day correspondences for the real images to guide the inference of their labels; 2) a novel uncertainty-aware annotation and evaluation framework and metric for semantic segmentation, designed for adverse conditions and including image regions* beyond human recognition capability *in the evaluation in a principled fashion; 3) the* Dark Zurich *dataset, which comprises 2416 unlabeled nighttime and 2920 unlabeled twilight images with correspondences to their daytime counterparts plus a set of 151 nighttime images with fine pixel-level annotations created with our protocol, which serves as a first benchmark to perform our novel evaluation. Experiments show that our guided curriculum adaptation significantly outperforms state-of-the-art methods on real nighttime sets both for standard metrics and our uncertainty-aware metric. Furthermore, our uncertainty-aware evaluation reveals that selective invalidation of predictions can lead to better results on data with ambiguous content such as our nighttime benchmark and profit safety-oriented applications which involve invalid inputs.*

1. Introduction

The state of the art in semantic segmentation is rapidly improving in recent years. Despite the advance, most methods are designed to operate at daytime, under favorable illumination conditions. However, many outdoor applications require robust vision systems that perform well at all times of day, under challenging lighting conditions, and in bad weather [22]. Currently, the popular approach to solving perceptual tasks such as semantic segmentation is to train deep neural networks [20, 38, 44] using large-scale human annotations [6, 9, 23]. This supervised scheme has achieved great success for daytime images, but it scales badly to adverse conditions. In this work, we focus on semantic segmentation at nighttime, both at the method level and the evaluation level.

At the method level, this work adapts semantic segmentation models from daytime to nighttime, without annotations in the latter domain. To this aim, we propose a new method called Guided Curriculum Model Adaptation (GCMA). The underpinnings of GCMA are threefold: power of time, power of place, and power of data. **Time**: environmental illumination changes continuously from daytime to nighttime. This enables adding intermediate domains between the two to smoothly transfer semantic knowledge. This idea is found to be effective in [8, 26]; we extend it by adding two more modules. **Place**: images taken over different time but with the same 6D camera pose share a large portion of content. The shared content can be used to guide the knowledge transfer process from a favorable condition (daytime) to an adverse condition (nighttime). We formalize this observation and propose a solution for large-scale application. **Data**: GCMA takes advantage of the powerful image translation techniques to stylize large-scale real annotated daytime datasets to darker target domains in order to perform standard supervised learning.

The adversity of nighttime poses further challenges for perceptual tasks compared to daytime. The extracted features become corrupted due to visual hazards [41] such as underexposure, noise, and motion blur. The degradation of affected input regions is often so intense that they are rendered *indiscernible*, i.e. determining their semantic content is impossible even for humans. We term such regions as

invalid for the task of semantic segmentation. A robust model should predict with high *uncertainty* on invalid regions while still being confident on valid (discernible) regions, and a sound evaluation framework should reward such behavior. The above requirement is particularly significant for safety-oriented applications such as autonomous cars, since having the vision system declare a prediction as invalid can help the downstream driving system avoid the fatal consequences of this prediction being false, e.g. when a pedestrian is missed.

To this end, we design a generic uncertainty-aware annotation and evaluation framework for semantic segmentation in adverse conditions which explicitly distinguishes invalid from valid regions of input images, and apply it to nighttime. On the annotation side, our novel protocol leverages privileged information in the form of daytime counterparts of the annotated nighttime scenes, which reveal a large portion of the content of invalid regions. This allows to reliably label invalid regions and to indeed *include* invalid regions in the evaluation, contrary to existing semantic segmentation benchmarks [6] which completely exclude them from evaluation. Moreover, apart from the standard class-level semantic annotation, each image is annotated with a mask which designates its invalid regions. On the evaluation side, we allow the *invalid* label in predictions and adopt from [40] the principle that for invalid pixels with legitimate semantic labels, both these labels and the *invalid* label are considered correct predictions. However, this principle does not cover the case of valid regions. We address this by introducing the concept of false invalid predictions. This enables calculation of *uncertainty-aware intersection-over-union (UIoU)*, a joint performance metric for valid and invalid regions which generalizes standard IoU, reducing to the latter when no invalid prediction exists. UIoU rewards predictions with confidence that is *consistent* to human annotators, i.e. with higher confidence on valid regions than invalid ones, meeting the aforementioned requirement.

Finally, we present *Dark Zurich*, a dataset of real images which contains corresponding images of the same driving scenes at daytime, twilight and nighttime. We use this dataset to feed real data to GCMA and to create a benchmark with 151 nighttime images for our uncertainty-aware evaluation. Our dataset and code are publicly available[1].

2. Related Work

Vision at Nighttime. Nighttime has attracted a lot of attention in the literature due to its ubiquitous nature. Several works pertain to human detection at nighttime, using FIR cameras [10, 37], visible light cameras [15], or a combination of both [4]. In driving scenarios, a few methods have been proposed to detect cars [17] and vehicles' rear lights [30]. Contrary to these domain-specific methods, previous work also includes both methods designed for robustness to illumination changes, by employing domain-invariant representations [1, 25] or fusing information from complementary modalities and spectra [33], and datasets with adverse illumination [21, 29] for localization benchmarking. A recent work [8] on semantic nighttime segmentation shows that images captured at twilight are helpful for supervision transfer from daytime to nighttime. Our work is partially inspired by [8] and extends it by proposing a guided curriculum adaptation framework which learns jointly from stylized images and unlabeled real images of increasing darkness and exploits scene correspondences.

Domain Adaptation. Performance of semantic segmentation on daytime scenes has increased rapidly in recent years. As a consequence, attention is now turning to adaptation to adverse conditions [3, 11, 33, 35]. A case in point are recent efforts to adapt clear-weather models to fog [7, 26, 27], by using both labeled synthetic images and unlabeled real images of increasing fog density. This work instead focuses on the nighttime domain, which poses very different and—as we would claim—greater challenges than the foggy domain (e.g. artificial light sources casting very different illumination patterns at night). A major class of adaptation approaches, including [5, 12, 13, 19, 28, 31, 34, 36, 43, 46], involves adversarial confusion or feature alignment between domains. The general concept of curriculum learning has been applied to domain adaptation by ordering tasks [42] or target-domain pixels [47], while we order domains. Cross-domain correspondences as guidance have only been used very recently in [18], which requires pixel-level matches, while we use more generic image-level correspondences.

Semantic Segmentation Evaluation. Semantic segmentation evaluation is commonly performed with the IoU metric [9]. Cityscapes [6] introduced an instance-level IoU (iIoU) to remove the large-instance bias, as well as mean average precision for the task of instance segmentation. The two tasks have recently been unified into panoptic segmentation [16], with a respective panoptic quality metric. The most closely related work to ours in this regard is WildDash [40], which uses standard IoU together with a fine-grained evaluation to measure the impact of visual hazards on performance. In contrast, we introduce UIoU, a new semantic segmentation metric that handles images with regions of uncertain semantic content and is suited for adverse conditions. Our uncertainty-aware evaluation is *complementary* to uncertainty-aware methods such as [14] that explicitly incorporate uncertainty in their model formulation and aims to promote the development of such methods, as UIoU rewards models that accurately capture heteroscedastic aleatoric uncertainty [14] in the input images through the different treatment of invalid and valid regions.

[1] https://trace.ethz.ch/projects/adverse/GCMA_UIoU

3. Guided Curriculum Model Adaptation

3.1. Problem Formulation

GCMA involves a source domain $\mathcal{S}$, an ultimate target domain $\mathcal{T}$, and an intermediate target domain $\dot{\mathcal{T}}$. In this work, $\mathcal{S}$ is daytime, $\mathcal{T}$ is nighttime, and $\dot{\mathcal{T}}$ is twilight time with an intermediate level of darkness between $\mathcal{S}$ and $\mathcal{T}$. GCMA adapts semantic segmentation models through this sequence of domains $(\mathcal{S}, \dot{\mathcal{T}}, \mathcal{T})$, which is sorted in ascending order with respect to level of darkness. The approach proceeds progressively and adapts the model from one domain in the sequence to the next. The knowledge is transferred through the domain sequence via this gradual adaptation process. The transfer is performed using two coupled branches: 1) learning from labeled synthetic stylized images and 2) learning from real data without annotations, to jointly leverage the assets of both. Stylized images inherit the human annotations of their original counterparts but contain unrealistic artifacts, whereas real images have less reliable pseudo-labels but are characterized by artifact-free textures.

Let us use $z \in \{1, 2, 3\}$ as the index in $(\mathcal{S}, \dot{\mathcal{T}}, \mathcal{T})$. Once the model for the current domain z is trained, its knowledge can be distilled on unlabeled real data from z, and then used, along with a new version of synthetic data from the next domain $z+1$ to adapt the current model to $z+1$.

Before diving into the details, we first define all datasets used. The inputs for GCMA consist of: 1) a labeled daytime set with M real images $\mathcal{D}^1_{lr} = \{(I^1_m, Y^1_m)\}^M_{m=1}$, e.g. Cityscapes [6], where $Y^1_m(i,j) \in \mathcal{C} = \{1, ..., C\}$ is the ground-truth label of pixel (i,j) of I^1_m; 2) an unlabeled daytime set of N_1 images $\mathcal{D}^1_{ur} = \{I^1_n\}^{N_1}_{n=1}$; 3) an unlabeled twilight set of N_2 images $\mathcal{D}^2_{ur} = \{I^2_n\}^{N_2}_{n=1}$; and 4) an unlabeled nighttime set of N_3 images $\mathcal{D}^3_{ur} = \{I^3_n\}^{N_3}_{n=1}$. In order to perform knowledge transfer with annotated data, $\mathcal{D}^1_{lr}$ is rendered in the style of $\mathcal{D}^2_{ur}$ and $\mathcal{D}^3_{ur}$. We use CycleGAN [45] to perform this style transfer, leading to two more sets: $\mathcal{D}^2_{ls} = \{(\bar{I}^2_m, Y^1_m)\}^M_{m=1}$ and $\mathcal{D}^3_{ls} = \{(\bar{I}^3_m, Y^1_m)\}^M_{m=1}$, where $\bar{I}^2_m$ and $\bar{I}^3_m$ are the stylized twilight and nighttime version of I^1_m respectively, and labels are copied. For $z = 1$, the semantic segmentation model ϕ^1 is trained directly on $\mathcal{D}^1_{lr}$. In order to perform knowledge transfer with unlabeled data, pseudo-labels for all three unlabeled real datasets need to be generated. The pseudo-labels for $\mathcal{D}^1_{ur}$ are generated using the model ϕ^1 via $\hat{Y}^1_n = \phi^1(I^1_n)$. For $z > 1$, training ϕ^z and generating $\hat{Y}^z_m$ is performed progressively as GCMA proceeds, as is detailed in Sec. 3.1.1. All six datasets are summarized in Table 1.

Table 1. The training sets used in GCMA. I indicates an image and Y its label map; $\bar{I}$ is a synthetic image and $\hat{Y}$ a pseudo-label map. See the text for details.

	Labeled		Unlabeled
	Real	Synthetic	Real
1. Daytime	$\{(I^1_m, Y^1_m)\}^M_{m=1}$		$\{(I^1_n, \hat{Y}^1_n)\}^{N_1}_{n=1}$
2. Twilight time		$\{(\bar{I}^2_m, Y^1_m)\}^M_{m=1}$	$\{(I^2_n, \hat{Y}^2_n)\}^{N_2}_{n=1}$
3. Nighttime		$\{(\bar{I}^3_m, Y^1_m)\}^M_{m=1}$	$\{(I^3_n, \hat{Y}^3_n)\}^{N_3}_{n=1}$

3.1.1 Guided Curriculum Model Adaptation

Since the method proceeds in an iterative manner, we present the algorithmic details only for a single adaptation step from $z-1$ to z. The presented algorithm is straightforward to generalize to multiple intermediate target domains. In order to adapt the semantic segmentation model ϕ^{z-1} from the previous domain $z-1$ to the current domain z, we generate synthetic stylized data in domain z: $\mathcal{D}^z_{ls}$.

For real unlabeled images, since no human annotations are available, we rely on a strategy of self-learning or curriculum learning. Our motivating assumption is that objects are generally easier to recognize in lighter conditions, so the tasks are solved in ascending order with respect to the level of darkness and the easier, solved tasks are used to retrain the model to further solve the harder tasks. This is in line with the concept of curriculum learning [2]. In particular, the model ϕ^{z-1} for domain $z-1$ can be applied to the unlabeled real images of domain $z-1$ to generate supervisory labels for training ϕ^z. Specifically, the dataset of real images with pseudo-labels for adaptation to domain z is $\mathcal{D}^{z-1}_{ur} = \{(I^{z-1}_n, \hat{Y}^{z-1}_n)\}^{N_{z-1}}_{n=1}$, where $\hat{Y}^{z-1}_n$ denotes the predicted labels of image I^{z-1}_n. A simple way to get these labels is by directly feeding I^{z-1}_n to ϕ^{z-1}, similar to the approach of [7, 26] for the case of fog. This choice, however, suffers from accumulation of substantial errors in the prediction of ϕ^{z-1} into the subsequent training step if domain $z-1$ is not the daytime domain. We instead propose a method to refine these errors by using *guidance* from the semantics of a daytime image I^1_n that *corresponds* to I^{z-1}_n, i.e. depicts roughly the same scene as I^{z-1}_n (the difference in the camera pose is small):

$$\hat{Y}^{z-1}_n = G\left(\phi^{z-1}(I^{z-1}_n), I^{z-1}_n, \phi^1(I^1_{A_{z-1\to1}(n)})\right), \quad (1)$$

where G is a guidance function which will be defined in Sec. 3.2 and $z-1 > 1$. $A_{z-1\to1}(n)$ is the correspondence function giving the index of the daytime image that corresponds to I^{z-1}_n.

Once we have the two training sets $\mathcal{D}^{z-1}_{ur}$ (with labels inferred through (1)) and $\mathcal{D}^z_{ls}$, learning ϕ^z is performed by optimizing a loss function that involves both datasets:

$$\min_{\phi^z}\left(\sum_{\substack{(I,Y)\\ \in\mathcal{D}^z_{ls}}} L(\phi^z(I), Y) + \mu \sum_{\substack{(I,\hat{Y})\\ \in\mathcal{D}^{z-1}_{ur}}} L(\phi^z(I), \hat{Y})\right), \quad (2)$$

where $L(.,.)$ is the cross entropy loss and μ is a hyper-parameter balancing the contribution of the two datasets.

In order to leverage the *place* prior at large scale to improve predictions through the guided label refinement defined in (1), specific aligned datasets need to be compiled. With this aim, we collected the *Dark Zurich* dataset by driving several laps in disjoint areas of Zurich; each lap was driven multiple times during the same day, starting from daytime through twilight to nighttime. The recordings include GPS readings and are split into three sets: daytime, twilight and nighttime (cf. Sec. 5). Since different drives of the same lap correspond to the same route, the camera orientation at a certain point of the lap is similar across all drives. We implement the correspondence function $A_{z\to 1}$ that assigns to each image in domain z its daytime counterpart using a GPS-based nearest neighbor assignment. The method presented in Sec. 3.2 carefully handles the effects of misalignment and dynamic objects in paired images.

3.2. Guided Segmentation Refinement

In the following presentation of our guided segmentation refinement for dark images using corresponding daytime images, we drop for brevity the subscript which was used to indicate this correspondence. The guidance function G which models our refinement approach and was introduced in a general form in (1) can be written more specifically as

$$G\left(\phi^z(I^z), I^z, \phi^1(I^1)\right) = R\left(\phi^z(I^z),\, B(\phi^1(I^1), I^z)\right), \tag{3}$$

i.e. as the composition of a cross bilateral filter B on the daytime predictions, which aligns them to the dark image, with a fusion function R, which adaptively combines the aligned daytime predictions with the initial dark image predictions to refine the latter.

3.2.1 Cross Bilateral Filter for Prediction Alignment

The correspondences between real images that are used in GCMA are not perfect, in the sense that they are not aligned at a pixel-accurate level. Therefore, to leverage the prediction for the daytime image I^1 as *guidance* for refining the respective prediction for the dark image I^z, it is necessary to first align the former prediction to I^z. To this end, we operate on *soft* predictions and define a cross bilateral filter on the initial soft prediction map $\mathbf{S}^1 = \phi^1(I^1)$ which uses the color of the dark image I^z as reference:

$$\tilde{\mathbf{S}}^1(\mathbf{p}) = \frac{\sum\limits_{\mathbf{q}\in\mathcal{N}(\mathbf{p})} G_{\sigma_s}(\|\mathbf{q}-\mathbf{p}\|)G_{\sigma_r}(\|I^z(\mathbf{q})-I^z(\mathbf{p})\|)\mathbf{S}^1(\mathbf{q})}{\sum\limits_{\mathbf{q}\in\mathcal{N}(\mathbf{p})} G_{\sigma_s}(\|\mathbf{q}-\mathbf{p}\|)G_{\sigma_r}(\|I^z(\mathbf{q})-I^z(\mathbf{p})\|)}. \tag{4}$$

In (4), $\mathbf{p}$ and $\mathbf{q}$ denote pixel positions, $\mathcal{N}(\mathbf{p})$ is the neighborhood of $\mathbf{p}$, G_{σ_s} is the spatial-domain Gaussian kernel and G_{σ_r} is the color-domain kernel. The definition of the filter implies that only pixels $\mathbf{q}$ with similar color to the examined pixel $\mathbf{p}$ in the dark image I^z contribute to the output $\tilde{\mathbf{S}}^1(\mathbf{p})$, which shifts salient edges in the initial daytime prediction to their correct position in the dark image. For the color-domain kernel, we use the CIELAB version of I^z, as it is more appropriate for measuring color similarity [24]. We set the spatial parameter σ_s to 80 to account for large misalignment, and σ_r to 10 following [24, 26].

3.2.2 Confidence-Adaptive Prediction Fusion

The final step in our refinement approach is to fuse the aligned prediction $\tilde{\mathbf{S}}^1$ for I^1 with the initial prediction $\mathbf{S}^z = \phi^z(I^z)$ for I^z in order to obtain the refined prediction $\hat{\mathbf{S}}^z$, the hard version of which is subsequently used in training. We propose an adaptive fusion scheme, which uses the *confidence* associated with the two predictions at each pixel to weigh their contribution in the output and addresses disagreements due to dynamic content by properly adjusting the fusion weights. Let us denote the confidence of the aligned prediction $\tilde{\mathbf{S}}^1$ for I^1 at pixel $\mathbf{p}$ by $F^1(\mathbf{p}) = \max_{c\in\mathcal{C}} \tilde{S}^1_c(\mathbf{p})$ and respectively the confidence of the initial prediction $\mathbf{S}^z$ for I^z by $F^z(\mathbf{p})$. Our confidence-adaptive fusion is then defined as

$$\hat{\mathbf{S}}^z = \frac{F^z}{F^z + \alpha F^1}\mathbf{S}^z + \frac{\alpha F^1}{F^z + \alpha F^1}\tilde{\mathbf{S}}^1, \tag{5}$$

where $0 < \alpha = \alpha(\mathbf{p}) \leq 1$ may vary and we have completely dropped the pixel argument $\mathbf{p}$ for brevity. In this way, we allow the daytime image prediction to have a greater effect on the output at regions of the dark image which were not easy for model ϕ^z to classify, while preserving the initial prediction $\mathbf{S}^z$ at lighter regions of the dark image where $\mathbf{S}^z$ is more reliable.

Our fusion distinguishes between dynamic and static scene content by regulating α. In particular, α downweights $\tilde{\mathbf{S}}^1$ to induce a preference towards $\mathbf{S}^z$ when both predictions have high confidence. However, apart from imperfect alignment, the two scenes also differ due to dynamic content. Intuitively, the prediction of a dynamic object in the daytime image should be assigned an even lower weight in case the corresponding prediction in the dark image does not agree, since this object might only be present in the former scene. More formally, we denote the subset of $\mathcal{C}$ that includes dynamic classes by $\mathcal{C}_d$ and define

$$\alpha(\mathbf{p}) = \begin{cases} \alpha_l, & \text{if } c_1 = \arg\max\limits_{c\in\mathcal{C}} \tilde{S}^1_c(\mathbf{p}) \in \mathcal{C}_d \text{ and } S^z_{c_1}(\mathbf{p}) \leq \eta \\ & \text{or } c_2 = \arg\max\limits_{c\in\mathcal{C}} S^z_c(\mathbf{p}) \in \mathcal{C}_d \text{ and } \tilde{S}^1_{c_2}(\mathbf{p}) \leq \eta, \\ \alpha_h & \text{otherwise.} \end{cases} \tag{6}$$

(a) Dark image I^z (b) Daytime image I^1 (c) Initial prediction $\mathbf{S}^z$ for I^z (d) Our refined prediction $\hat{\mathbf{S}}^z$ for I^z

Figure 1. Example pair of corresponding images from *Dark Zurich*, initial prediction for the dark image and our refined prediction.

In our experiments, we manually tune $\alpha_l = 0.3$, $\alpha_h = 0.6$ and $\eta = 0.2$ on a couple of training images (no grid search). A result of our guided refinement is shown in Fig. 1.

4. Uncertainty-Aware Evaluation

Images taken under adverse conditions such as nighttime contain invalid regions, i.e. regions with indiscernible semantic content. Invalid regions are closely related to the concept of negative test cases which was considered in [40]. However, invalid regions constitute intra-image entities and can co-exist with valid regions in the same image, whereas a negative test case refers to an entire image that should be treated as invalid. We build upon the evaluation of [40] for negative test cases and generalize it to be applied uniformly to all images in the evaluation set, whether they contain invalid regions or not. Our annotation and evaluation framework includes invalid regions in the set of evaluated pixels, but treats them differently from valid regions to account for the high uncertainty of their content. In the following, we elaborate on the generation of ground-truth annotations using privileged information through the day-night correspondences of our dataset and present our UIoU metric.

4.1. Annotation with Privileged Information

For each image I, the annotation process involves two steps: 1) creation of the ground-truth invalid mask J, and 2) creation of the ground-truth semantic labeling H.

For the semantic labels, we consider a predefined set $\mathcal{C}$ of C classes, which is equal to the set of Cityscapes [6] evaluation classes ($C = 19$). The annotator is first presented only with I and is asked to mark the valid regions in it as the regions which she can unquestionably assign to one of the C classes or declare as not belonging to any of them. The result of this step is the invalid mask J, which is set to 0 at valid pixels and 1 at invalid pixels.

Secondly, the annotator is asked to mark the semantic labels of I, only that this time she also has access to an *auxiliary* image I'. This latter image has been captured with roughly the same 6D camera pose as I but under more favorable conditions. In our dataset, I' is captured at daytime whereas I is captured at nighttime. The large overlap of static scene content between the two images allows the annotator to label certain regions in H with a legitimate semantic label from $\mathcal{C}$, even though the same regions have been annotated as invalid (and are kept as such) in J. This allows joint evaluation on valid and invalid regions, as it creates regions which can accept both the *invalid* label and the ground-truth label from $\mathcal{C}$ as correct predictions. Due to the imperfect match of the camera poses for I and I', the labeling of invalid regions in H is done conservatively, marking a coarse boundary which may leave unlabeled zones around the true semantic boundaries in I, so that no pixel is assigned a wrong label. The parts of I which remain indiscernible even after inspection of I' are left unlabeled in H. These parts as well as instances of classes outside $\mathcal{C}$ are not considered during evaluation. We illustrate a visual example of our annotation inputs and outputs in Fig. 2.

4.2. Uncertainty-Aware Predictions

The semantic segmentation prediction that is fed to our evaluation is expected to include pixels labeled as *invalid*. Instead of defining a separate, explicit *invalid* class, which would potentially require the creation of new training data to incorporate this class, we allow a more flexible approach for soft predictions with the original set of semantic classes by using a *confidence threshold*, which affords an evaluation curve for our UIoU metric by varying this threshold.

In particular, we assume that the evaluated method outputs an intermediate soft prediction $\mathbf{S}(\mathbf{p})$ at each pixel $\mathbf{p}$ as a probability distribution among the C classes, which is subsequently converted to a hard assignment by outputting the class $\tilde{H}(\mathbf{p}) = \arg\max_{c \in \mathcal{C}} \{S_c(\mathbf{p})\}$ with the highest probability. In this case, $S_{\tilde{H}(\mathbf{p})}(\mathbf{p}) \in [1/C, 1]$ is the effective confidence associated with the prediction. This assumption is not very restrictive, as most recent semantic segmentation methods are based on CNNs with a softmax layer that outputs such soft predictions.

The final evaluated output $\hat{H}$ is computed based on a free parameter $\theta \in [1/C, 1]$ which acts as a confidence threshold by invalidating those pixels where the confidence of the prediction is lower than θ, i.e. $\hat{H}(\mathbf{p}) = \tilde{H}(\mathbf{p})$ if $S_{\tilde{H}(\mathbf{p})}(\mathbf{p}) \geq \theta$ and *invalid* otherwise. Increasing θ results in more pixels being predicted as *invalid*. This approach is motivated by the fact that ground-truth invalid regions are identified during annotation by the uncertainty of their se-

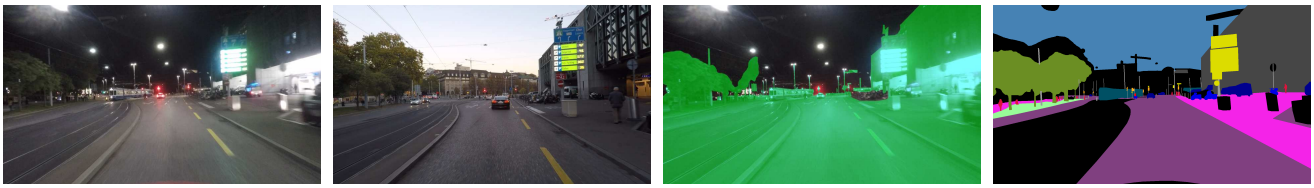

(a) Input image I (b) Auxiliary image I' (c) GT invalid mask J (d) GT semantic labeling H

Figure 2. Example input images from *Dark Zurich-test* and output annotations with our protocol. Valid pixels in J are marked green.

mantic content, which implies that a model should ideally place lower confidence (equivalently higher uncertainty) in predictions on invalid regions than on valid ones, so that the former get invalidated for lower values of θ than the latter. The formulation of our UIoU metric rewards this behavior as we shall see next. Note that our evaluation does not strictly require soft predictions, as UIoU can be normally computed for fixed, hard predictions $\hat{H}$.

4.3. UIoU

We propose UIoU as a generalization of the standard IoU metric for evaluation of semantic segmentation predictions which may contain pixels labeled as *invalid*. UIoU reduces to standard IoU if no pixel is predicted to be invalid, e.g. when $\theta = 1/C$.

The calculation of UIoU for class c involves five sets of pixels, which are listed along with their symbols: true positives (TP), false positives (FP), false negatives (FN), true invalids (TI), and false invalids (FI). Based on the ground-truth invalid masks J, the ground-truth semantic labelings H and the predicted labels $\hat{H}$ for the set of evaluation images, these five sets are defined as follows:

$$\text{TP} = \{\mathbf{p} : H(\mathbf{p}) = \hat{H}(\mathbf{p}) = c\}, \tag{7}$$

$$\text{FP} = \{\mathbf{p} : H(\mathbf{p}) \neq c \text{ and } \hat{H}(\mathbf{p}) = c\}, \tag{8}$$

$$\text{FN} = \{\mathbf{p} : H(\mathbf{p}) = c \text{ and } \hat{H}(\mathbf{p}) \notin \{c, \textit{invalid}\}\}, \tag{9}$$

$$\text{TI} = \{\mathbf{p} : H(\mathbf{p}) = c \text{ and } \hat{H}(\mathbf{p}) = \textit{invalid} \text{ and } J(\mathbf{p}) = 1\}, \tag{10}$$

$$\text{FI} = \{\mathbf{p} : H(\mathbf{p}) = c \text{ and } \hat{H}(\mathbf{p}) = \textit{invalid} \text{ and } J(\mathbf{p}) = 0\}. \tag{11}$$

UIoU for class c is then defined as

$$\text{UIoU} = \frac{|\text{TP}| + |\text{TI}|}{|\text{TP}| + |\text{TI}| + |\text{FP}| + |\text{FN}| + |\text{FI}|}. \tag{12}$$

Note that a true invalid prediction results in equal reward to predicting the correct semantic label of the pixel. Moreover, an invalid prediction does not come at no cost: it incurs the same penalty on valid pixels as predicting an incorrect label.

When dealing with multiple classes, we modify our notation to $\text{UIoU}^{(c)}$ (similarly for the five sets of pixels related to class c), which we avoided in the previous definitions to reduce clutter. The overall semantic segmentation performance on the evaluation set is reported as the mean UIoU over all C classes. By varying the confidence threshold θ and using the respective output, we obtain a parametric expression $\text{UIoU}(\theta)$. When $\theta = 1/C$, no pixel is predicted as invalid and thus $\text{UIoU}(1/C) = \text{IoU}$.

We motivate the usage of UIoU instead of standard IoU in case the test set includes ground-truth invalid masks by showing in Th. 1 that UIoU is guaranteed to be larger than IoU for some $\theta > 1/C$ under the assumption that predictions on invalid regions are associated with lower confidence than those on valid regions, which lies in the heart of our evaluation framework. The proof is in the supplement.

Table 2. Comparison of *Dark Zurich* against related datasets with nighttime semantic annotations. "Night annot.": annotated nighttime images, "Invalid": can invalid regions get legitimate labels?

Dataset	Night annot.	Classes	Reliable GT	Fine GT	Invalid
WildDash [40]	13	19	✓	✓	×
Raincouver [32]	95	3	✓	×	×
BDD100K [39]	345	19	×	✓	×
Nighttime Driving [8]	50	19	✓	×	×
Dark Zurich	151	19	✓	✓	✓

Theorem 1. *Assume that there exist* θ_1, θ_2 *such that* $\theta_1 < \theta_2$, $\forall p : J(p) = 1 \Rightarrow S_{\tilde{H}(p)}(p) \leq \theta_1$ *and* $J(p) = 0 \Rightarrow S_{\tilde{H}(p)}(p) \geq \theta_2$. *If we additionally assume that* $\exists p \in \text{FN}^{(c)}(1/C) \cup \text{FP}^{(c)}(1/C) : J(p) = 1$, *then* $\text{IoU}^{(c)} < \text{UIoU}^{(c)}(\theta_1)$.

5. The Dark Zurich Dataset

Dark Zurich was recorded in Zurich using a 1080p GoPro Hero 5 camera, mounted on top of the front windshield of a car. The collection protocol with multiple drives of several laps to establish correspondences is detailed in Sec. 3.

We split *Dark Zurich* and reserve one lap for testing. The rest of the laps remain unlabeled and are used for training. They comprise 3041 daytime, 2920 twilight and 2416 nighttime images extracted at 1 fps, which are named *Dark Zurich-*{*day*, *twilight*, *night*} respectively and correspond to the three sets in the rightmost column of Table 1. From the testing night lap, we extract one image every 50m or 20s, whichever comes first, and assign to it the corresponding

daytime image to serve as the auxiliary image I' in our annotation (cf. Sec. 4.1). We annotate 151 nighttime images with fine pixel-level Cityscapes labels and invalid masks following our protocol and name this set *Dark Zurich-test*. In total, 272.2M pixels have been annotated with semantic labels and 56.7M of these pixels are marked as invalid. We validate the quality of our annotations by having 20 images annotated twice by different subjects and measuring consistency. 93.5% of the labeled pixels are consistent in the semantic annotations and respectively 95% in the invalid masks. We compare to existing annotated nighttime sets in Table 2, noting that most large-scale sets for road scene parsing, such as Cityscapes [6] and Mapillary Vistas [23], contain few or no nighttime scenes. Nighttime Driving [8] and Raincouver [32] only include *coarse* annotations. *Dark Zurich-test* contains ten times more nighttime images than WildDash [40]—the only other dataset with *reliable* fine nighttime annotations. Detailed inspection showed that ∼70% of the 345 densely annotated nighttime images of BDD100K [39] contain severe labeling errors which render them unsuitable for evaluation, especially in dark regions we treat as invalid (e.g. *sky* is often mislabeled as *building*). Our annotation protocol helps avoid such errors by properly defining invalid regions and using daytime images to aid annotation, and *Dark Zurich-test* is an initial high-quality benchmark to promote our uncertainty-aware evaluation.

6. Results

Our architecture of choice for implementing GCMA is RefineNet [20]. We use the publicly available *RefineNet-res101-Cityscapes* model, trained on Cityscapes, as the baseline model to be adapted to nighttime. Throughout our experiments, we train this model with a constant learning rate of 5×10^{-5} on mini-batches of size 1.

Comparison to Other Adaptation Methods. Our first experiment compares GCMA to state-of-the-art approaches for adaptation of semantic segmentation models to nighttime. To obtain the synthetic labeled datasets for GCMA, we stylize Cityscapes to twilight using a CycleGAN model that is trained to translate Cityscapes to *Dark Zurich-twilight* (respectively to nighttime with *Dark Zurich-night*). The real training datasets for GCMA are *Dark Zurich-day*, instantiating $\mathcal{D}^1_{ur}$, and *Dark Zurich-twilight*, instantiating $\mathcal{D}^2_{ur}$. Each adaptation step comprises 30k SGD iterations and uses $\mu = 1$. For the second step, we apply our guided refinement to the labels of *Dark Zurich-twilight* that are predicted by model ϕ^2 fine-tuned in the first step, using the correspondences of *Dark Zurich-twilight* to *Dark Zurich-day*.

We evaluate GCMA on *Dark Zurich-test* against the state-of-the-art adaptation approaches AdaptSegNet [31] and DMAda [8] and report standard IoU performance in Table 3, including *invalid* pixels which are assigned a legitimate semantic label in the evaluation. We have trained AdaptSegNet to adapt from Cityscapes to *Dark Zurich-night*. For fair comparison, we also report the performance of the respective baseline Cityscapes models for each method. RefineNet is the common baseline of GCMA and DMAda. GCMA significantly outperforms the other methods for most classes and achieves a substantial 10% improvement in the overall mIoU score against the second-best method. The improvement with GCMA is pronounced for classes which usually appear dark at nighttime, such as *sky*, *vegetation*, *terrain* and *person*, indicating that our method successfully handles large domain shifts from its source daytime domain. These findings are supported by visually assessing the predictions of the compared methods, as in the examples of Fig. 3. We repeat the above comparison on Nighttime Driving [8] in Table 4 and show that GCMA generalizes very well to different datasets.

Ablation Study for GCMA. We measure the individual effect of the main components of GCMA in Table 5 by evaluating its ablated versions on *Dark Zurich-test*. Direct adaptation to nighttime in a single step using only Cityscapes images stylized as nighttime with CycleGAN is a strong baseline, due to the reliable ground-truth labels that accompany the stylized Cityscapes, its high diversity and the limited artifacts of CycleGAN-based translation. Adding our real images to the training algorithm and applying our two-stage curriculum significantly improves upon this baseline. Finally, our guided segmentation refinement in the second step of GCMA brings an additional 2.6% benefit, as it corrects a lot of errors in the pseudo-labels of the real twilight images, which helps compute more reliable gradients from the corrected loss during the subsequent training.

Comparisons with UIoU. In Fig. 4, we use our novel UIoU metric to evaluate GCMA against DMAda and our baseline RefineNet model on *Dark Zurich-test* for varying confidence threshold θ and plot the resulting mean UIoU(θ) curves. Note that standard mean IoU can be read out from the leftmost point of each curve. First, our expectation based on Th. 1 is confirmed for all methods, i.e. maximum UIoU values over the range of θ are larger than IoU by ca. 3%. This implies that on *Dark Zurich-test*, these models generally have lower confidence on invalid regions than valid ones. Second, the comparative performance of the methods is the same across all values of θ —GCMA substantially outperforms the other two—which shows that UIoU is generally consistent with standard IoU and is a suitable substitute of the latter in adverse settings where declaring the input as invalid is relevant.

7. Conclusion

In this paper, we have introduced GCMA, a method to gradually adapt semantic segmentation models from day-

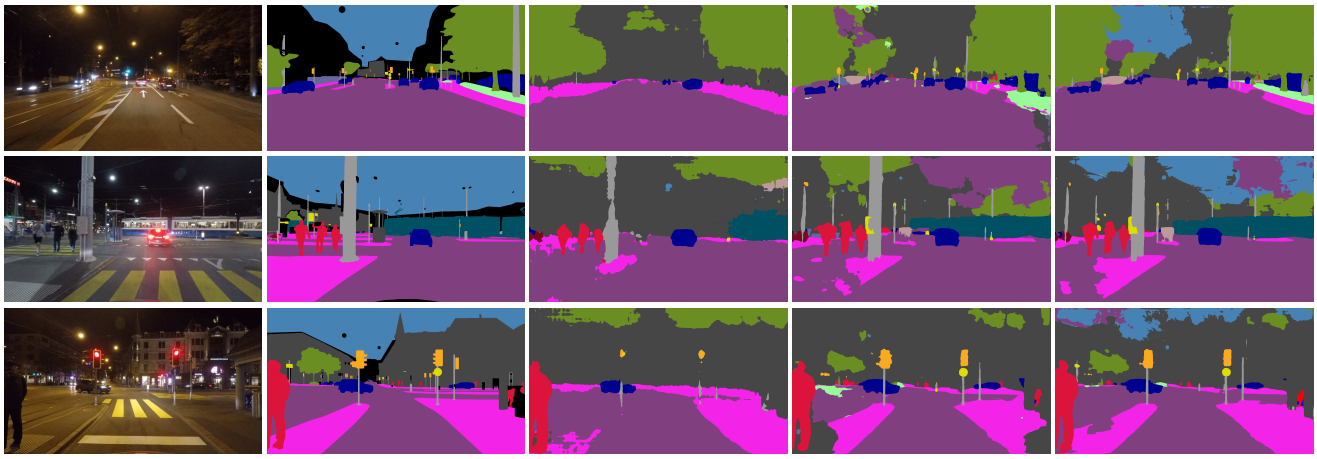

(a) Image (b) Semantic GT (c) AdaptSegNet [31] (d) DMAda [8] (e) GCMA (Ours)

Figure 3. Qualitative semantic segmentation results on *Dark Zurich-test*. "AdaptSegNet" adapts from Cityscapes to *Dark Zurich-night*.

Table 3. Comparison on *Dark Zurich-test*. AdaptSegNet-Cityscapes→*DZ-night* denotes adaptation from Cityscapes to *Dark Zurich-night*.

Method	road	sidew.	build.	wall	fence	pole	light	sign	veget.	terrain	sky	person	rider	car	truck	bus	train	motorc.	bicycle	mIoU
RefineNet [20]	68.8	23.2	46.8	20.8	12.6	29.8	30.4	26.9	43.1	14.3	0.3	36.9	49.7	63.6	6.8	**0.2**	24.0	33.6	9.3	28.5
AdaptSegNet-Cityscapes [31]	79.0	21.8	53.0	13.3	11.2	22.5	20.2	22.1	43.5	10.4	18.0	37.4	33.8	64.1	6.4	0.0	**52.3**	30.4	7.4	28.8
AdaptSegNet-Cityscapes→*DZ-night* [31]	**86.1**	44.2	55.1	**22.2**	4.8	21.1	5.6	16.7	37.2	8.4	1.2	35.9	26.7	68.2	**45.1**	0.0	50.1	33.9	15.6	30.4
DMAda [8]	75.5	29.1	48.6	21.3	14.3	34.3	36.8	29.9	49.4	13.8	0.4	43.3	**50.2**	69.4	18.4	0.0	27.6	**34.9**	11.9	32.1
Ours: GCMA	81.7	**46.9**	**58.8**	22.0	**20.0**	**41.2**	**40.5**	**41.6**	**64.8**	**31.0**	**32.1**	**53.5**	47.5	**75.5**	39.2	0.0	49.6	30.7	**21.0**	**42.0**

Table 4. Comparison on Nighttime Driving [8]. Read as Table 3.

Method	mIoU (%)
RefineNet [20]	31.5
AdaptSegNet-Cityscapes [31]	32.6
AdaptSegNet-Cityscapes→*DZ-night* [31]	34.5
DMAda [8]	36.1
Ours: GCMA	**45.6**

Table 5. Ablations of GCMA on *Dark Zurich-test*, reporting mIoU.

Daytime baseline: RefineNet [20]	28.5%
+direct CycleGAN adapt. (w/o real, w/o curriculum)	37.1%
+GCMA w/o guided refinement	39.4%
+GCMA w/ guided refinement	**42.0%**

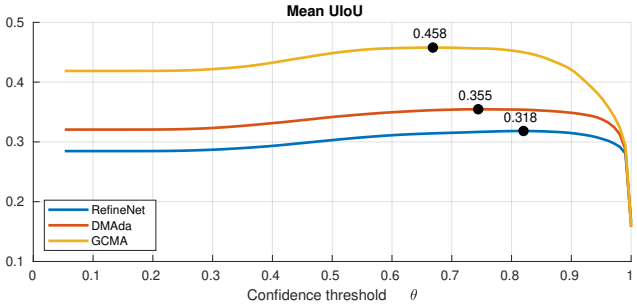

Figure 4. Uncertainty-aware evaluation of RefineNet [20], DMAda [8] and GCMA on *Dark Zurich-test*. We evaluate mean UIoU across the entire range $[1/C, 1]$ of confidence threshold θ. For each method, the point at which mean UIoU is maximized is marked black and labeled with this maximum mean UIoU value.

time to nighttime with stylized data and unlabeled real data of increasing darkness, as well as UIoU, a novel evaluation metric for semantic segmentation designed for images with indiscernible content. We have also presented *Dark Zurich*, a large-scale dataset of real scenes captured at multiple times of day with cross-time-of-day correspondences, and annotated 151 nighttime scenes of it with a new protocol which enables our evaluation. Detailed evaluation with standard IoU on real nighttime sets demonstrates the merit of GCMA, which substantially improves upon competing state-of-the-art methods. Finally, evaluation on our benchmark with UIoU shows that invalidating predictions is useful when the input includes ambiguous content.

Acknowledgements. This work is funded by Toyota Motor Europe via the research project TRACE-Zürich. We thank Simon Hecker for his advice on decoding GoPro GPS data.

References

[1] Jose M. Alvarez and Antonio M. Lopez. Road detection based on illuminant invariance. *IEEE Transactions on Intelligent Transportation Systems*, 12(1):184–193, 2011. 2

[2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *International Conference on Machine Learning*, 2009. 3

[3] Mario Bijelic, Tobias Gruber, and Werner Ritter. Benchmarking image sensors under adverse weather conditions for autonomous driving. In *IEEE Intelligent Vehicles Symposium (IV)*, 2018. 2

[4] Yuxi Chen and Chongzhao Han. Night-time pedestrian detection by visual-infrared video fusion. In *World Congress on Intelligent Control and Automation*, 2008. 2

[5] Yuhua Chen, Wen Li, and Luc Van Gool. ROAD: Reality oriented adaptation for semantic segmentation of urban scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2

[6] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 3, 5, 7

[7] Dengxin Dai, Christos Sakaridis, Simon Hecker, and Luc Van Gool. Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *International Journal of Computer Vision*, 2019. 2, 3

[8] Dengxin Dai and Luc Van Gool. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *IEEE International Conference on Intelligent Transportation Systems*, 2018. 1, 2, 6, 7, 8

[9] Mark Everingham, Luc Van Gool, Christopher K.I. Williams, John Winn, and Andrew Zisserman. The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 1, 2

[10] Junfeng Ge, Yupin Luo, and Gyomei Tei. Real-time pedestrian detection and tracking at nighttime for driver-assistance systems. *IEEE Transactions on Intelligent Transportation Systems*, 10(2):283–298, 2009. 2

[11] Judy Hoffman, Trevor Darrell, and Kate Saenko. Continuous manifold based adaptation for evolving visual domains. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2

[12] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, 2018. 2

[13] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. Conditional generative adversarial network for structured domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2

[14] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, 2017. 2

[15] Jong Hyun Kim, Hyung Gil Hong, and Kang Ryoung Park. Convolutional neural network-based human detection in nighttime images using visible light camera sensors. *Sensors*, 17(5), 2017. 2

[16] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

[17] Hulin Kuang, Kai-Fu Yang, Long Chen, Yong-Jie Li, Leanne Lai Hang Chan, and Hong Yan. Bayes saliency-based object proposal generator for nighttime traffic images. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):814–825, 2018. 2

[18] Mans Larsson, Erik Stenborg, Lars Hammarstrand, Marc Pollefeys, Torsten Sattler, and Fredrik Kahl. A cross-season correspondence dataset for robust semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

[19] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

[20] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. RefineNet: Multi-path refinement networks with identity mappings for high-resolution semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 7, 8

[21] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The Oxford RobotCar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017. 2

[22] Srinivasa G. Narasimhan and Shree K. Nayar. Vision and the atmosphere. *International Journal of Computer Vision*, 48(3):233–254, 2002. 1

[23] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kontschieder. The Mapillary Vistas dataset for semantic understanding of street scenes. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 7

[24] Sylvain Paris and Frédo Durand. A fast approximation of the bilateral filter using a signal processing approach. *International Journal of Computer Vision*, 81(1):24–52, 2009. 4

[25] German Ros and Jose M. Alvarez. Unsupervised image transformation for outdoor semantic labelling. In *IEEE Intelligent Vehicles Symposium (IV)*, 2015. 2

[26] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *The European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 3, 4

[27] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, 2018. 2

[28] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2

[29] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi

Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DOF outdoor visual localization in changing conditions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2

[30] Ravi Kumar Satzoda and Mohan Manubhai Trivedi. Looking at vehicles in the night: Detection and dynamics of rear lights. *IEEE Transactions on Intelligent Transportation Systems*, 2016. 2

[31] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 7, 8

[32] Frederick Tung, Jianhui Chen, Lili Meng, and James J. Little. The Raincouver scene parsing benchmark for self-driving in adverse weather and at night. *IEEE Robotics and Automation Letters*, 2(4):2188–2193, 2017. 6, 7

[33] Abhinav Valada, Johan Vertens, Ankit Dhall, and Wolfram Burgard. AdapNet: Adaptive semantic segmentation in adverse environmental conditions. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017. 2

[34] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S. Davis. DCAN: Dual channel-wise alignment networks for unsupervised scene adaptation. In *The European Conference on Computer Vision (ECCV)*, 2018. 2

[35] Markus Wulfmeier, Alex Bewley, and Ingmar Posner. Addressing appearance change in outdoor robotics with adversarial domain adaptation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017. 2

[36] Markus Wulfmeier, Alex Bewley, and Ingmar Posner. Incremental adversarial domain adaptation for continually changing environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018. 2

[37] Fengliang Xu, Xia Liu, and Kikuo Fujimura. Pedestrian detection and tracking with night vision. *IEEE Transactions on Intelligent Transportation Systems*, 6(1):63–71, 2005. 2

[38] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations*, 2016. 1

[39] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. BDD100K: a diverse driving video database with scalable annotation tooling. *CoRR*, abs/1805.04687, 2018. 6, 7

[40] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steininger, and Gustavo Fernandez Dominguez. WildDash - creating hazard-aware benchmarks. In *The European Conference on Computer Vision (ECCV)*, 2018. 2, 5, 6, 7

[41] Oliver Zendel, Markus Murschitz, Martin Humenberger, and Wolfgang Herzner. How good is my test data? Introducing safety analysis for computer vision. *International Journal of Computer Vision*, 125(1):95–109, 2017. 1

[42] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 2

[43] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, and Tao Mei. Fully convolutional adaptation networks for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2

[44] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1

[45] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 3

[46] Xinge Zhu, Hui Zhou, Ceyuan Yang, Jianping Shi, and Dahua Lin. Penalizing top performers: Conservative loss for semantic segmentation adaptation. In *The European Conference on Computer Vision (ECCV)*, 2018. 2

[47] Yang Zou, Zhiding Yu, B.V.K. Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *The European Conference on Computer Vision (ECCV)*, 2018. 2