This CVPR paper is the Open Access version, provided by the Computer Vision Foundation.
Except for this watermark, it is identical to the accepted version;
the final published version of the proceedings is available on IEEE Xplore.

Heuristic Self-Paced Learning for Domain Adaptive Semantic Segmentation under Adverse Conditions

Shiqin Wang[1,2,4*] Haoyang Chen[1,2,3,4*] Huaizhou Huang[1,2,4*] Yinkan He[1,2,4] Dongfang Sun[1,2,4]
Xiaoqing Chen[3,5] Xingyu Liu[1,2,4] Zheng Wang[1,2,3,4†] Kaiyan Zhao[4]
[1]National Engineering Research Center for Multimedia Software, Institute of Artificial Intelligence, School of Computer Science, Wuhan University [2]Hubei Key Laboratory of Multimedia and Network Communication Engineering
[3]Zhongguancun Academy, Beijing, China. 100094 [4]School of Computer Science, Wuhan University, Wuhan, China
[5]School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

Abstract

The learning order of semantic classes significantly impacts unsupervised domain adaptation for semantic segmentation, especially under adverse weather conditions. Most existing curricula rely on handcrafted heuristics (e.g., fixed uncertainty metrics) and follow a static schedule, which fails to adapt to a model's evolving, high-dimensional training dynamics, leading to category bias. Inspired by Reinforcement Learning, we cast curriculum learning as a sequential decision problem and propose an autonomous class scheduler. *This scheduler consists of two components: (i) a high-dimensional state encoder that maps the model's training status into a latent space and distills key features indicative of progress, and (ii) a category-fair policy-gradient objective that ensures balanced improvement across classes. Coupled with mixed source–target supervision, the learned class rankings direct the network's focus to the most informative classes at each stage, enabling more adaptive and dynamic learning. It is worth noting that our method achieves state-of-the-art performance on three widely used benchmarks (e.g., ACDC, Dark Zurich, and Nighttime Driving) and shows generalization ability in synthetic-to-real semantic segmentation.*

1. Introduction

Semantic segmentation is one of the core technologies for autonomous driving systems to achieve robust environmental perception. However, existing models trained on controlled datasets, when deployed to the real world full of uncertainties, their performance will degrade significantly under real-world adverse weather conditions (such as heavy fog, nighttime, and heavy rain) [4, 25–27, 30]. Sensor data undergoes complex non-linear degradation under these conditions, directly leading to the failure of the model's ability to perceive safe-critical obstacles [18, 21, 36].

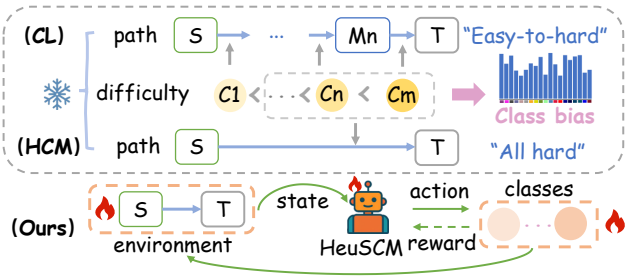

Figure 1. From design curriculum to self-paced learning. Traditional Curriculum Learning (CL) and Hard Class Mining (HCM) both rely on fixed handcrafted priors and fixed learning paths. CL adopts a static, easy-to-hard curriculum, while HCM focuses solely on difficult classes. Both strategies induce class bias, manifesting as significant accuracy discrepancies across semantic classes. Differently, we adaptively mine the most informative classes via dynamically perceived models' evolving (S: Source domain, T: Target domain, M_n: the n-th intermediate domain, Cn: the n-th Semantic Class).

Unsupervised Domain Adaptation (UDA) is a mainstream paradigm to address domain shift issues and reduce reliance on target domain annotations. However, extreme weather poses an extraordinary compound challenge to UDA. This challenge consists of two intertwined problems: (1) Complex non-uniform domain shift: for example, fog concentration and rain intensity are continuously changing, resulting in feature shifts that are highly variable and non-linear; (2) Exacerbated class imbalance: In severe weather, the observability of inherently rare safety-critical semantic classes is further reduced, sharply amplifying the long-tail property of data distribution.

Prevailing UDA methods attempt to tackle this com-

*These authors contributed equally to this work.
†Corresponding author.

pound challenge using separate, often decoupled, mechanisms. They typically employ style transfer [13, 14, 22, 35] to mitigate the domain shift, while concurrently employing optimization strategies like Curriculum Learning (CL) and Hard Class Mining (HCS) to address the exacerbated class imbalance. However, existing CL and HCS methods [8] are generally limited by a common, fundamental paradigm flaw: they rely on a prior human definitions. Specifically, (1) the "difficulty" of the curriculum is quantified by fixed artificially designed metrics (such as prediction uncertainty or confidence); (2) the "path" of the curriculum is driven by artificially designed rules (such as "from easy to hard" or "all hard"). We believe that this "prescriptive paradigm" is fundamentally unreasonable. The "cognitive state" of a model during training is a high-dimensional, dynamic, and non-monotonic evolutionary process. Trying to statically plan this learning path with a fixed, one-dimensional, artificially defined scalar is suboptimal. The rigidity of this strategy makes it unable to adapt to the model's own ever-changing internal state, leading to unavoidable local optima when facing "compound" problems like extreme weather, such as underfitting to noise or overfitting to the majority class, as shown in Figure 1.

Inspired by Reinforcement Learning (RL), we propose a paradigm shift from "designing a curriculum" to "learning a curriculum". We argue that the optimal learning trajectory should not be specified by human prior assumptions but should be autonomously discovered by the model based on its own learning state. In this paper, we groundbreakingly re-envision the training process of UDA as a "sequential decision problem". We propose Heuristic Semantic Class Mining (HeuSCM), a framework that includes an autonomous scheduler (Agent). The efficacy of this method is built upon two core technical designs: *i*. Autonomous State Perception: The decision-making of the Agent no longer relies on a single human-defined metric. We design a High-dimensional Semantic State Extraction (HSSE) network that enables it to comprehensively characterize the current learning progress of the model from high-dimensional state vectors. *ii*. Dynamic Policy Optimization: The optimization objective of the Agent is not merely maximizing the summation of their individual value functions. We propose Categorical α-Fairness for Policy Gradients (CαPG), which optimizes the policy via our designed global fairness objective function for an equitable reward distribution across all semantic classes.

Our agent autonomously discovers a dynamic curriculum, adjusting in real-time according to the model's state rather than following a pre-set path. This RL-inspired approach allows the model to adapt to complex, real-world scenarios, overcoming the limitations of traditional curriculum learning. Our main contributions are as follows:

- We are the first to redefine UDA curriculum learning from a "human-defined heuristic" problem to an "autonomously learned strategy" problem, enabling the model to dynamically adjust its learning path based on its evolving internal state.
- We propose HeuSCM, a novel framework with the core being our HSSE and CαPG design. HeuSCM realizes the faithful perception of high-dimensional semantic states and dynamic policy optimization that ensures categorical reward equity.
- Extensive experiments on highly challenging extreme weather semantic segmentation benchmarks verify the effectiveness of our method, particularly achieving state-of-the-art performance 72.9 mIoU [%] on the ACDC test. Moreover, our Heuristic Class Sampling Policy (HCSP) exhibits superior generalization capability on the synthetic-to-real segmentation benchmark.

2. Related Work

2.1. Unsupervised Domain Adaptation Semantic Segmentation Under Adverse Weather

To bridge the large domain gap between the source domain (clear weather) and the target domain (adverse weather), some researchers generated the intermediate mixed domain via cross-domain mixed sampling and performed the domain adaptation from the source domain to the mixed domain [3, 16]. Other research efforts introduce style transfer [13, 14, 22, 35] or image generation [23] networks to mitigate discrepancies in visual appearance, or alternatively focus on reducing domain-specific discrepancies at the feature level [12]. Later, Bruggemann *et al.* [2] designed Contrastive Model Adaptation (CMA) to learn domain-invariant features via aligning the features of target paired image pairs. However, these methods uniformly adapt source knowledge to adverse target weather, neglecting the inherent variations among challenging scenarios and causing the model to generate hallucinations (erroneous class predictions). To address this, Gong *et al.* [8] introduced intermediate domains and performed the first-easy-then-hard domain adaptation (CoDA). Conversely, Chen *et al.* [6] focuses on learning hard classes that are visually similar within the target domain (AMSC). Regardless of the specific mechanism, whether easy-to-hard curriculum learning or hard class mining, it is strongly validated that the learning sequence of different samples or special attention to specific classes significantly promotes the field. However, the difficulty assessment in these methods relies on a single, manually designed metric, and the learning path is fixed, often leading to insufficient model learning.

2.2. Class Curriculum Learning

In cross-domain adaptation, methods concerning class curriculum learning primarily fall into Curriculum Learning

(CL), sequencing learning from easy to hard, and Hard Class Mining (HCM), focusing on hard sample learning. These methods involve two main steps: difficulty assessment and curriculum scheduling. Regarding difficulty criteria, one line of research measures class difficulty based on target domain predictions, such as class frequency [1, 15, 25], prediction uncertainty [29], or confidence [37]. Another line further incorporates domain discrepancy [25] or feature similarity [28]. Regarding the learning scheduling, mainstream strategies can be categorized into three types. The first assigns selected classes higher sampling probabilities when performing cross-domain mixing sampling [25, 37]. The second assigns them higher loss weights to intensify the model's focus [17, 29, 33]. The third directly utilizes these identified hard classes to retrain the model [1, 15].

Despite significant progress, the difficulty assessment criteria are manually fixed based on prior knowledge, and the class curriculum learning remains static. This prevents the model from adaptively learning the most informative classes according to its learning state, consequently resulting in insufficient learning of semantic classes.

2.3. Reinforcement Learning in Unsupervised Domain Adaptation

Reinforcement Learning (RL) has demonstrated its effectiveness in learning complex policies by interacting with the environment. Some researchers [7, 11, 24, 32] are exploring RL paradigms for unsupervised domain adaptation. Zhang *et al.* [32] selected the most relevant cross-domain features via RL, and then applied adversarial learning to minimize the domain shift. This method focused on image classification. In parallel, other studies [7, 11, 24] have explored RL paradigms for Unsupervised Domain Adaptation in Semantic Segmentation. Dong *et al.* [7] maximized the transfer gain under reinforcement learning manner. Usmani *et al.* [24] utilized RL techniques to realize cross-domain feature-level alignment. Judge *et al.* [11] introduced RL for 2D + time echocardiography segmentation. Despite these efforts, most methods focus on feature alignment, with no attention to class sampling.

3. Method

In this paper, we propose a Heuristic Semantic Class Mining (HeuSCM) self-paced curriculum framework, conceptually inspired by reinforcement learning, that dynamically perceives the learning progress of the semantic segmentation model from a high-dimensional, evolving state space. Instead of relying on fixed, predefined rules, HeuSCM learns a policy-based class scheduler that autonomously adjusts the learning strategy based on a reward signal, rather than a static curriculum. As illustrated in Figure 2, we first obtain the learning status of the current semantic segmentation model from the high-dimensional segmentation state (Sec . 3.1). Then we continuously optimize ClassGen to generate ranked semantic classes (sorted in descending order of informativeness, Sec . 3.2), which guide the cross-domain mixed sampling and in turn update the segmentation model via the SegLoss.

Segmentation Model: UDA-SS under adverse weather aims to adapt the semantic segmentation knowledge from the source domain $D_S = \left\{(x_s^i, y_s^i)\right\}_{i=1}^{N_s}$ to the target domain $D_T = \left\{(x_t^j)\right\}_{j=1}^{N_t}$. To bridge the large domain gap, cross-domain mixed sampling between the D_S and D_T is performed to generate the mixed domain $D_M = \{(\mathcal{X}_{Tmix}^t, \mathcal{Y}_{Rmix}^t)\}_{t=1}^{N_m}$ as the bridge, and then $D_S \rightarrow D_M$ domain adaptation is performed. N_s, N_t, and N_m denote the source, target domain, and mixed domain sample sizes.

The t-th mixed image ($\mathcal{X}_{Tmix}^t$) and label ($\mathcal{Y}_{Rmix}^t$) is generated as follows:

$$\begin{cases} \mathcal{X}_{Tmix}^t = \mathcal{H} \odot x_s^i + (1-\mathcal{H}) \odot x_t^j, \\ \mathcal{Y}_{Rmix}^t = \mathcal{H} \odot y_s^i + (1-\mathcal{H}) \odot g_\theta(x_t^j)^r, \\ \mathcal{H}_{ijk} = \mathbb{I}(L_S(m,n) = k) \cdot \mathbb{I}(k \in \mathcal{C}_{\text{low}}), \\ \mathcal{C}_{\text{low}} = \{\mathcal{R}_S^i(k) \mid \lfloor N_S^i/2 \rfloor + 1 \leq k \leq N_S^i\}, \\ \mathcal{R}_S^i = (c_k \in \mathcal{R} \mid c_k \in C_S^i)), \end{cases} \tag{1}$$

where g_θ, x_s^i, x_t^j, y_s^i denote the segmentation model, the i-th source image, j-th target image, and i-th source label, respectively. $\mathcal{H} \in \{0,1\}^{h \times w \times c}$, $\mathbb{I}$ are the binary mask and an indicator function, respectively. (m, n) denotes the spatial coordinates of the pixel, and k is the semantic class index. C_S^i and N_S^i denote the semantic classes and classes number of the i-th source label. $\mathcal{R}$ is the "Ranked Classes". The core mechanism of HeuSCM is that this list $\mathcal{R}$ is not fixed; it is the action output by our autonomous agent, dynamically generated based on the model's current state.

After that, the segmentation model (g_θ) is optimized via the following training loss function (*i.e.,* SegLoss):

$$\mathcal{L}_{seg} = \lambda_1 \mathcal{L}_{CE}\left(g_\theta(x_s), y_s\right) + \lambda_2 \mathcal{L}_{CE}\left(g_\theta(\mathcal{X}_{Tmix}), \mathcal{Y}_{Rmix}\right), \tag{2}$$

where $\mathcal{L}_{CE}$ is cross-entropy loss. λ_1, λ_2 are the coefficient.

3.1. High-dimensional Semantic State Extraction

3.1.1. Low-dimensional State Representation Learning

The state space reflecting the semantic segmentation model's learning status is high-dimensional and redundant, comprising coupled features and complex interactions that jointly capture the subtle nuances of the domain adaptation process. Direct learning within such a space presents significant challenges for policy network optimization. To tackle this, our goal is to construct a low-dimensional latent state representation that effectively captures the intricate relationships among these high-dimensional features.

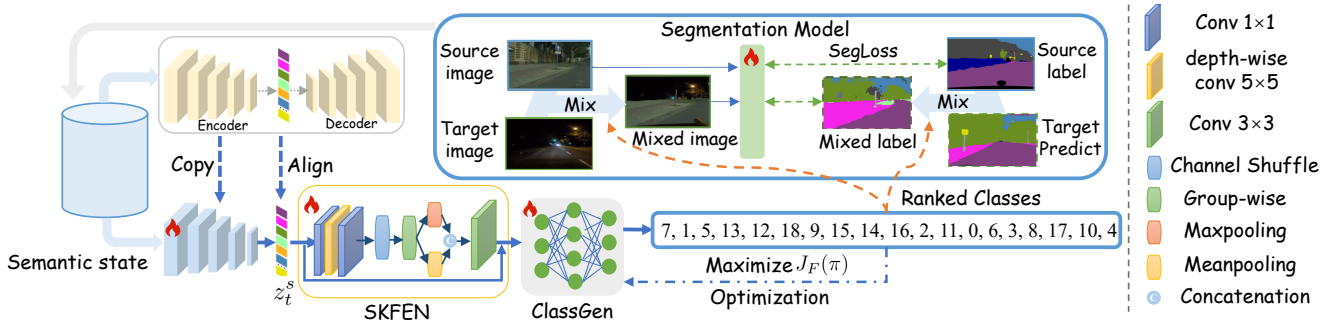

Figure 2. The framework of our designed Heuristic Semantic Class Mining (HeuSCM). First, a Gaussian Mixture VAE (GM-VAE) encodes high-dimensional semantic states into latent features z_t^s. Our SKFEN processes z_t^s to distill key features reflecting the model's learning status. Conditioned on these features, ClassGen outputs ranked classes (sorted in descending order of informational value). These rankings guide the generation of mixed image and mixed label pairs, which optimize the segmentation model via SegLoss. Concurrently, we maximize the objective $J_F(\pi)$ to jointly optimize the copied GM-VAE encoder, SKFEN, and ClassGen.

To model the multi-modal nature of the semantic segmentation model's learning, we employ a Gaussian Mixture Variational Autoencoder (GM-VAE) to encode the high-dimensional learning features into a compact latent space, where each dimension represents a distinct aspect of the model's domain adaptation status. During training, given a high-dimensional state $\mathbf{s} \in \mathcal{S}$, the GM-VAE encoder ($\text{Enc}_\psi$) learns to infer the posterior over both the discrete component, *i.e.*, categorical distribution $q_\psi(c|\mathbf{s}) = \text{Cat}(c; \text{Enc}_\psi(\mathbf{s}))$ (representing the learning mode), and the continuous latent state posterior $q_\psi(\mathbf{z}|\mathbf{s}, c) = \mathcal{N}(\mathbf{z}; \mu_\psi(\mathbf{s}, c), \Sigma_\psi(\mathbf{s}, c))$. The model is trained by maximizing the following variational lower bound:

$$\begin{aligned}\mathcal{L}_{\text{GM-VAE}}(\mathbf{s}) = \mathbb{E}_{q_\psi(c|\mathbf{s})}\left[\mathbb{E}_{q_\psi(\mathbf{z}|\mathbf{s},c)}\left[\log p_\theta(\mathbf{s}|\mathbf{z})\right]\right.\\ \left.-\text{KL}\left(q_\psi(\mathbf{z}|\mathbf{s},c)\|p(\mathbf{z}|c)\right)\right] - \text{KL}\left(q_\psi(c|\mathbf{s})\|p(c)\right),\end{aligned} \tag{3}$$

where p_θ is the probabilistic decoder which reconstructs the original state $\mathbf{s}$ from the latent variable $\mathbf{z}$. q_ψ represents the approximate posteriors and $\text{KL}(\cdot\|\cdot)$ denotes the Kullback–Leibler divergence. $p(\mathbf{z}|c)$ and $p(c)$ are the priors.

After unsupervised pre-training, we freeze the GM-VAE decoder and retain the encoder to perform the mapping from the high-dimensional state space to the low-dimensional latent space. To accommodate task-specific requirements, we jointly fine-tune the encoder together with the policy network optimization. As the policy adjusts based on the reward signal, the encoder is co-updated to align the latent learning state with the reward-relevant semantic class outputs. At the same time, we regularize the encoder so that the learned latent representation still preserves the structure of the original state manifold, via the reconstruction loss as follows:

$$\mathcal{L}_{\text{recon}} = \mathbb{E}_{\mathbf{s}_t \sim \mathcal{D}_{\text{state}}}\left[\left\|\mathbf{s}_t - p_\theta\left(\text{Enc}_\psi(\mathbf{s}_t)\right)\right\|^2\right], \tag{4}$$

where $\mathbf{s}_t$ denotes the high-dimensional segmentation state at step t, Enc_ψ is the encoder of the copied GM-VAE, and p_θ is the frozen decoder.

3.1.2. Semantic Key Feature Extraction Network

To address the large redundant information within the low-dimensional latent space and distill the key features that truly reflect the model's learning status, we propose the Semantic Key Feature Extraction Network (SKFEN). Inspired by traditional spatial pooling and grouped convolution, SKFEN does not operate on spatial compression; rather, we introduce a novel feature refinement mechanism designed innovatively from the perspective of reducing feature redundancy among different channels.

Our SKFEN contains two main phases: (1) Initial Transformation and Grouping: the low-dimensional state z_t^s first undergoes initial feature fusion and interaction modeling. Subsequently, its channel dimension is expanded and permuted to prepare for grouped processing as follows:

$$f_{\text{group}} = \text{Group}(\text{Shuffle}(\text{Conv}_{1\times1}^{\text{expand}}(\text{Conv}_{5\times5}^{\text{ds}}(\text{Conv}_{1\times1}^{\text{fuse}}(z_t^s))))) = \bigcup_{g=1}^{G} f_g, \tag{5}$$

where z_t^s is the input low-dimensional latent state, $\text{Conv}_{1\times1}^{\text{fuse}}(\cdot)$ denotes an initial 1×1 convolution for feature fusion, $\text{Conv}_{5\times5}^{\text{ds}}(\cdot)$ represents a 5×5 depth-wise separable convolution for spatial interaction modeling, and $\text{Conv}_{1\times1}^{\text{expand}}(\cdot)$ is a 1×1 convolution that expands the channel dimension to n. $\text{Shuffle}(\cdot)$ is the channel shuffle operation, $\text{Group}(\cdot)$ splits the feature map into G groups along the channel dimension, resulting in feature maps f_g for each group $g \in \{1, \ldots, G\}$.

(2) Group-wise Feature Aggregation: These grouped features f_g are then processed in parallel to distill salient and statistical information. These two distinct representations are concatenated and fused by a final convolution to obtain the refined features. This entire aggregation and fu-

Algorithm 1 Heuristic Class Mining.

Initialize: Segmentation Model g_θ, GM-VAE encoder Enc_ψ and decoder p_θ, SKFEN, policy network (ClassGen) π, Replay Buffer $\mathcal{D}$, Value (Critic) Networks $V_{c,\xi}$ (for $c = 1..C$)
for $t \leftarrow 1$ to T **do**
  Observe high-dim state $\mathbf{s}$ from g_θ
  Encode low-dim state $z_t^s = Enc_\psi(\mathbf{s})$
  Distill key features z_{key} ▷ Eq. (5), (6)
  Generate Ranked Classes $\mathcal{R}_t = \pi_\theta(z_{key})$
  Generate mixed data $(\mathcal{X}_{Tmix}, \mathcal{Y}_{Rmix})$ ▷ Eq.(1)
  Update g_θ by minimizing $\mathcal{L}_{seg}$ ▷ Eq. (2)
  Compute reward $\vec{r}_t = [r_1(t), ..., r_C(t)]$ ▷ Eq. (8)
  Observe next state $z_{key,t+1}$
  Record transition $(z_{key,t}, \mathcal{R}_t, \vec{r}_t, z_{key,t+1})$ in $\mathcal{D}$
  Record state $\mathbf{s}_t$ in $\mathcal{D}_{\text{state}}$
end for
  for t = 1 to $T_{\text{agent_step}}$ **do**
    Sample a batch of transitions from $\mathcal{D}$
    Update π_θ, SKFEN, Enc_ψ via policy gradient $\nabla_\theta J_F(\pi_\theta)$ ▷ Eq. (10)
    Sample a batch of states $\mathbf{s}_t$ from $\mathcal{D}_{\text{state}}$
    Update Enc_ψ by minimizing $\mathcal{L}_{recon}$ ▷ Eq. (4)

  end for
end for

sion process is formulated as:

$$z_{\text{out}} = \text{Conv}_{3\times3}\left(\text{Concat}\left(\bigcup_{g=1}^{G}\left[\max_{c=1}^{C_g} f_g(c,\cdot,\cdot)\right], \bigcup_{g=1}^{G}\left[\frac{1}{C_g}\sum_{c=1}^{C_g} f_g(c,\cdot,\cdot)\right]\right)\right) + z_t^s, \tag{6}$$

where z_{out} is the final refined latent state, and z_t^s is the original input state from Eq. (5) used for the residual connection. $\text{Conv}_{3\times3}$ denotes the final fusion convolution. For each feature group f_g (with $C_g = n/G$ channels) obtained from f_{group}, the term $\left[\max_{c=1}^{C_g}\{f_g(c,\cdot,\cdot)\}\right]$ represents the channel-wise max pooling operation, which computes the maximum value across all channels c for each spatial location. Similarly, $\left[\frac{1}{C_g}\sum_{c=1}^{C_g} f_g(c,\cdot,\cdot)\right]$ represents the channel-wise average pooling operation. The $\bigcup_{g=1}^{G}$ operator signifies concatenation along the channel dimension across all G groups, and $\text{Concat}(\cdot)$ concatenates the resulting max-pooled and average-pooled feature maps.

3.2. Categorical α-Fairness for Policy Gradients

Due to inherent training bias, the model naturally favors certain semantic classes. However, traditional policy gradient calculation solely focuses on reward maximization, thereby neglecting the concept of fairness (*i.e.*, the need for the agent to treat every semantic class equally). This issue is paramount in adverse weather autonomous driving scenarios, where robust and balanced segmentation performance across all classes is essential. To achieve fairness among the learning progress of different semantic classes, we adapt the concept of multi-agent fairness to our single-agent, multi-class problem.

We first define the value function for the c-th semantic class, $V_c^\pi(s)$, as the expected total discounted return that category c can obtain, starting from state s and following policy π:

$$V_c^\pi(s) := \mathbb{E}_\pi\left[\sum_{k=0}^{\infty}\gamma^k r_c(t+k) \mid S_t = s\right], \tag{7}$$

where $r_c(t)$ is the reward signal computed specifically for class c at time t.

Inspired by the insight that models with enhanced transferability and discriminability yield superior target-domain performance, we propose a reward mechanism to explicitly quantify these two factors. Given the absence of target labels, we therefore compute an unsupervised, composite reward that synthesizes both, calculated as follows:

$$\begin{cases} r_c(t) = \underbrace{\dfrac{A_c^S(t)\cdot A_c^T(t)}{\|A_c^S(t)\|\cdot\|A_c^T(t)\|}}_{\text{Transferability}} + \lambda\cdot\underbrace{\sum_{k\in\mathcal{C},k\neq c}\left(1-\dfrac{A_c^T(t)\cdot A_k^T(t)}{\|A_c^T(t)\|\cdot\|A_k^T(t)\|}\right)}_{\text{Discriminability}}, \\ A_c^S(t) = \dfrac{1}{|\Lambda_c^S(t)|}\sum_{(x_s,y_s)\in D_S}\mathbb{I}\{y_s = c\}\cdot f(x_s)|_c, \\ A_c^T(t) = \dfrac{1}{|\Lambda_c^T(t)|}\sum_{(x_t,g_\theta(x_t))\in D_T}\mathbb{I}\{g_\theta(x_t) = c\}\cdot f(x_t)|_c, \end{cases} \tag{8}$$

where $|\Lambda_c^S(t)|$ and $|\Lambda_c^T(t)|$ is the number of pixels belongs to category c in the source domain and target domain, respectively. $\mathbb{I}$ is an indicator function. $f(x)|_c$ denotes the feature output of our segmentation network g_θ for c. λ is a hyperparameter balancing the two objectives.

To achieve fairness among the returns of all classes, we no longer optimize the standard policy objective $J^{\text{sum}}(\pi)$ (which sums the values of all classes and is thus susceptible to preference bias), but instead optimize a different global fairness objective $J_F(\pi)$:

$$J_F(\pi) := F(V_1^\pi(s), V_2^\pi(s), \ldots, V_C^\pi(s)) = \sum_{c=1}^{C}\frac{1}{1-\alpha}\left(V_c^\pi(s)\right)^{1-\alpha}, \tag{9}$$

where F is the function of each class's return, chosen to enforce fairness. To mitigate the model fairness issue exacerbated by the greedy nature of traditional RL reward mechanisms, the α-fairness objective is employed as the objective function F.

When the $J_F(\pi)$ is applied to the policy gradient calculation, it corresponds to a weighted advantage function.

Table 1. Comparison of the state-of-the-art in Cityscapes→ACDC domain adaptation on the ACDC test set. The best results are presented in bold.

Method	Backbone	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	**mIoU**
DeepLab-v2 [5]	DeepLab-v2	71.9	26.2	51.1	18.8	22.5	19.7	33.0	27.7	67.9	28.6	44.2	43.1	22.1	71.2	29.8	33.3	48.4	26.2	35.8	38.0
Refign [3]	DeepLab-v2	49.5	56.7	79.8	31.2	25.7	34.1	48.0	48.7	76.2	42.5	38.5	48.3	24.7	75.8	46.5	43.9	64.3	34.1	43.6	48.0
CMA [2]	DeepLab-v2	83.1	52.7	65.4	18.7	30.5	**44.5**	56.3	53.9	76.7	39.7	79.0	54.2	31.2	76.7	40.2	39.3	47.4	29.8	38.6	50.4
CompUDA [35]	DeepLab-v2	52.4	54.5	75.6	30.6	26.8	35.6	44.7	47.8	74.5	40.5	39.1	45.1	20.6	76.3	47.2	40.5	64.9	36.2	36.2	47.0
VBLC [14]	DeepLab-v2	49.6	39.3	79.4	35.8	29.5	42.6	57.2	57.5	69.1	42.7	39.8	**54.5**	29.3	77.8	43.0	36.2	32.7	38.7	**53.4**	47.8
ATP [29]	DeepLab-v2	76.2	47.3	71.4	**42.7**	31.4	44.2	55.4	**62.0**	**89.0**	34.7	79.1	49.9	16.6	77.5	30.0	19.7	47.7	**44.0**	39.4	50.5
CISS [22]	DeepLab-v2	70.5	36.7	67.0	29.4	30.2	31.6	45.6	48.9	70.4	24.7	65.5	48.2	31.1	76.6	45.7	47.0	62.8	26.8	38.9	47.2
HeuSCM (Ours)	DeepLabv2	**91.4**	**66.6**	**84.3**	40.6	**32.3**	38.0	**57.9**	54.8	82.8	**50.4**	**94.2**	53.8	**33.0**	**80.1**	**50.8**	**52.4**	**72.2**	29.3	50.7	**58.7**
DAFormer [9]	DAFormer	56.9	45.4	84.7	44.7	35.1	48.6	44.8	57.4	69.5	52.9	45.8	57.1	28.2	82.8	57.2	63.9	84.0	40.2	50.5	55.3
Gaussian [31]	DAFormer	62.8	51.6	83.0	34.7	35.0	52.1	30.1	56.4	73.0	55.9	60.9	62.7	33.8	80.2	59.5	58.5	81.8	47.5	52.3	56.4
Refign [3]	DAFormer	89.5	63.4	87.3	43.6	34.3	52.3	63.2	61.4	86.9	58.5	95.7	62.1	39.3	84.1	65.7	71.3	85.4	47.9	52.8	65.5
VBLC [14]	DAFormer	89.2	59.8	85.9	44.0	37.2	53.5	**64.5**	63.2	72.4	56.3	84.1	65.5	37.7	85.1	60.1	71.8	85.2	47.7	56.3	64.2
CoPT [19]	DAFormer	49.1	**70.3**	83.6	**59.4**	**42.4**	**58.5**	48.3	**67.2**	73.5	**60.7**	45.0	**69.3**	**45.2**	83.4	**76.3**	**74.5**	**88.2**	**54.4**	**61.4**	63.7
Instance-Warp [34]	DAFormer	83.0	53.2	85.5	47.4	38.3	46.0	51.4	57.8	73.9	56.2	82.1	61.3	32.3	**85.5**	69.0	68.9	82.5	46.7	52.0	61.7
HeuSCM (Ours)	DAFormer	**90.2**	64.4	**87.3**	43.5	34.9	51.9	63.6	61.6	**86.9**	59.1	**95.8**	62.1	39.2	84.3	65.4	71.3	85.4	48.3	52.2	**65.7**
HRDA [10]	HRDA	88.3	57.9	88.1	55.2	36.7	56.3	62.9	65.3	74.2	57.7	85.9	68.8	45.7	88.5	76.4	82.4	87.7	52.7	60.4	68.0
Refign [3]	HRDA	93.8	**75.7**	90.0	57.9	43.3	55.6	67.4	68.2	88.2	61.8	96.1	67.5	50.8	88.8	75.2	83.4	**89.6**	54.5	61.2	72.1
CompUDA [35]	HRDA	92.7	71.5	89.5	**61.6**	39.8	51.0	**72.0**	67.2	82.8	58.7	92.9	67.0	46.4	89.3	75.3	81.2	88.7	56.3	62.4	71.1
VBLC [14]	HRDA	90.2	63.9	87.8	44.1	42.3	54.1	67.0	65.5	74.4	58.9	85.9	66.4	43.8	87.5	72.0	83.9	84.2	48.5	57.1	67.2
CoDA [8]	HRDA	93.1	72.7	**90.7**	57.3	**47.4**	56.8	69.9	70.0	87.3	59.8	95.4	**71.4**	47.6	**90.3**	**77.1**	83.8	89.1	54.7	**64.1**	72.6
ACSegFormer [16]	HRDA	**94.3**	75.3	90.2	57.9	42.1	55.7	71.4	**71.6**	87.8	61.5	96.0	69.8	51.2	88.7	75.5	85.8	89.5	53.8	62.3	72.7
CISS [22]	HRDA	92.0	69.6	89.2	57.3	40.5	55.8	67.1	67.3	75.3	59.7	86.4	70.0	47.5	88.9	73.1	77.5	87.0	55.6	61.7	69.6
HeuSCM (Ours)	HRDA	93.6	74.6	90.5	59.9	42.7	**57.3**	71.1	70.3	**88.7**	**62.4**	**96.4**	70.9	**52.3**	88.7	74.9	**86.3**	87.0	**57.3**	61.0	**72.9**

Table 2. Comparison of the state-of-the-art in Cityscapes→ACDC domain adaptation on the ACDC val set. The best results are presented in bold.

Method	Backbone	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	**mIoU**
DeepLab-v2 [5]	DeepLab-v2	69.5	15.3	52.4	11.0	12.0	25.4	48.2	39.6	68.3	24.2	69.0	40.1	15.2	64.7	14.2	19.8	28.8	21.8	17.2	34.6
Refign [3]	DeepLab-v2	91.5	67.5	**81.4**	43.8	32.5	43.0	66.4	50.1	**82.1**	35.0	93.3	47.7	27.4	78.7	52.4	51.8	**64.1**	13.9	**37.7**	55.8
CMA [2]	DeepLab-v2	83.9	51.3	60.1	25.3	28.2	45.6	**69.1**	52.3	79.0	30.1	79.2	51.5	23.5	77.4	43.4	37.0	36.2	24.9	28.6	48.8
VBLC [14]	DeepLab-v2	50.0	42.5	73.6	34.5	26.3	**48.0**	65.6	**55.1**	67.8	33.3	39.1	**56.3**	**39.6**	**80.8**	45.7	28.2	29.9	**31.8**	25.6	46.0
HeuSCM (Ours)	DeepLabv2	**92.7**	**69.1**	81.0	**43.9**	**34.0**	42.8	65.9	48.8	81.6	**35.6**	**93.3**	47.3	24.9	79.4	**55.5**	**55.5**	62.0	14.4	37.5	**56.1**
DAFormer [9]	DAFormer	71.4	50.0	79.1	42.1	33.1	55.5	40.0	50.0	72.2	35.4	68.2	54.3	17.5	83.0	71.2	74.8	80.9	41.2	31.5	55.3
Refign [3]	DAFormer	89.4	62.4	**85.5**	48.6	36.6	57.7	71.0	55.0	85.3	41.0	95.1	57.3	33.1	82.9	73.6	82.5	86.0	43.9	48.1	65.0
VBLC [14]	DAFormer	88.5	57.6	81.9	41.2	35.2	58.0	**72.8**	57.5	71.7	39.3	82.1	62.2	36.2	**87.1**	82.6	86.6	84.1	41.6	44.9	63.7
CoPT [19]	DAFormer	55.4	**71.7**	79.1	**57.7**	**47.2**	**62.9**	60.5	**65.2**	73.0	40.3	51.0	**68.9**	**46.6**	83.1	**82.7**	**90.8**	**88.2**	**47.8**	**58.6**	64.8
Instance-Warp [34]	DAFormer	82.9	56.1	79.8	44.6	40.3	52.7	60.8	52.5	72.0	38.4	78.0	56.6	30.5	84.9	80.2	86.9	86.4	44.5	45.8	61.8
HeuSCM (Ours)	DAFormer	**90.8**	65.6	85.2	49.3	38.0	58.6	70.8	55.7	**85.4**	**42.4**	**95.1**	58.0	33.4	83.1	73.2	81.0	87.7	44.4	45.5	**65.4**
HRDA [10]	HRDA	86.5	52.5	83.7	50.6	34.8	61.4	72.7	60.7	72.3	39.7	81.9	65.5	45.5	88.1	84.7	82.8	74.3	48.4	53.7	65.2
Refign [3]	HRDA	94.7	76.5	87.1	52.9	43.4	62.1	77.4	65.8	86.4	44.0	95.3	64.5	42.4	87.9	85.5	90.4	89.8	47.7	56.3	71.1
VBLC [14]	HRDA	89.5	61.8	84.9	46.3	**48.5**	61.3	74.8	60.1	72.3	39.7	82.1	62.9	41.9	87.8	85.7	76.0	87.4	48.2	52.6	66.5
CoDA [8]	HRDA	93.3	72.6	**88.3**	**59.6**	48.2	**63.3**	76.6	66.2	85.9	40.2	94.4	**71.3**	48.5	**90.5**	83.6	91.5	91.7	53.6	56.0	72.4
ACSegFormer [16]	HRDA	**95.5**	77.1	87.5	56.9	45.8	59.5	79.4	66.2	86.3	42.9	**95.3**	68.0	49.9	87.3	**86.5**	**94.1**	91.7	**55.3**	54.2	72.6
CISS [22]	HRDA	92.3	69.9	85.2	50.5	41.3	60.7	76.0	61.2	73.1	41.2	82.4	67.5	45.0	89.6	83.8	88.3	89.1	53.2	54.5	68.7
HeuSCM (Ours)	HRDA	95.3	**77.9**	87.0	57.2	44.1	60.3	**79.2**	**66.7**	**86.5**	**44.7**	95.2	66.5	**50.7**	87.7	86.4	93.4	**91.8**	50.8	**60.9**	**72.7**

The gradient of the policy π_θ is computed using a fairness-weighted aggregate advantage, $\tilde{A}_\alpha$ as follows:

$$\begin{aligned}\nabla_\theta J_F(\pi_\theta) &= \mathbb{E}_\pi\left[\nabla_\theta \log \pi_\theta(a_t \mid s_t) \cdot \tilde{A}_\alpha(s_t, a_t)\right] \\ &= \mathbb{E}_\pi\left[\nabla_\theta \log \pi_\theta(a_t \mid s_t) \cdot \sum_{c=1}^{C} w_c(s_t) \cdot A_c(s_t, a_t)\right],\end{aligned} \tag{10}$$

where $\tilde{A}_\alpha(s_t, a_t)$ is our α-fair advantage function, A_c is the advantage function for semantic class c. And $w_c(s_t)$ denotes the fairness weight that is controlled by α and is inversely proportional to the α-th power of the category's current value: $w_c(s_t) := V_c^\pi(s_t)^{-\alpha}$. By maximizing $J_F(\pi_\theta)$, we can compel the policy π to pursue a fairer and more balanced learning trajectory. The full training procedure is

outlined in Algorithm 1.

4. Experiments

4.1. Datasets and evaluation metrics

The mean Intersection-over-Union (mIoU) is adopted as the evaluation metric, where a higher value indicates better performance. We validate the effectiveness of our method on 1) unsupervised domain adaptation semantic segmentation (UDA-SS) under adverse weather: Cityscapes→ACDC domain adaptation; and 2) UDA night semantic segmentation: Cityscapes→DarkZurich domain adaptation. Additionally, we evaluate the generalization ability of our method in the synthetic-to-real semantic segmentation, *i.e.,* GTA5 → Cityscapes.

4.2. Experimental settings

Our proposed framework is implemented using the PyTorch framework on an NVIDIA A800 GPU. DeepLab-v2 [5], DAFormer [9], and HRDA [10] as the backbone. λ_1 and λ_2 in Eq. 2 are both set to 1.0. And λ in Eq. 8 is set to 1.0. Training is conducted for 60k iterations using 1024×1024 random crops from the Cityscapes and ACDC datasets. We train our model using the AdamW optimizer, setting the weight decay to 1e-4. For mixed image generation, we first employ our HeuSCM method, followed by standard augmentations including Color Jittering and Gaussian Blurring. Additionally, we adopt the rare class sampling strategy from [9] to mitigate the source domain's long-tail distribution, setting the α parameter to 0.999.

4.3. Comparison with State-of-the-art Methods

4.3.1. Comparison on ACDC

We present comparisons to several kinds of semantic segmentation methods, including 1) *backbones*: DeepLab-v2 [5], DAFormer [9], and HRDA [10]; 2) *UDA-SS methods under Adverse Weather*: Refign [3], CMA [2], VBLC [14], CompUDA [35], CoDA [8] and ACSegFormer [16]; and 3) *general UDA-SS methods*: ATP [29], CISS [22], Gaussian [31], CoPT [19] and Instance-Warp [34]. The quantitative results of mIoU performances on the ACDC test set and the ACDC val set are reported in Table 1 and Table 2, respectively. We observe that our method consistently outperforms existing approaches across different backbones on both the ACDC test and validation sets. Notably, when built upon the HRDA backbone, our method achieves mIoU scores of 72.9 and 72.7 mIoU [%] on the ACDC test and val sets, respectively, establishing state-of-the-art performance for Cityscapes → ACDC domain adaptation.

The qualitative comparison of our method with existing UDA on the ACDC val set is shown in Figure 3. Under the same backbone, our results are visually closer to the ground truth than existing UDA-SS under Adverse Weather methods. Furthermore, our HeuSCM, built upon the HRDA backbone, yields the best semantic segmentation performance, further validating its effectiveness.

Table 3. Comparison with state-of-the-art methods on the Dark Zurich-val set and Nighttime Driving test set.

Method	Backbone	Dark Zurich-val	Nighttime Driving
		mIoU	mIoU
DAFormer [9]	DAFormer	37.1	54.0
InforMS [25]	DAFormer	45.1	56.0
HeuSCM (Ours)	DAFormer	**45.5**	**56.7**
HRDA [10]	HRDA	42.1	54.1
InforMS [25]	HRDA	52.5	58.5
HeuSCM (Ours)	HRDA	**52.8**	**59.3**

Table 4. Ablation Study on several model variants of our method on the ACDC val. LSRL, SKFEN, CαPG are Low-dimensional State Representation Learning in 3.1.1, Semantic Key Feature Extraction Network in 3.1.2, Categorical α-Fairness for Policy Gradients in 3.2, respectively.

LSRL		✓			✓	✓		✓
SKFEN			✓		✓		✓	✓
CαPG				✓		✓	✓	✓
mIoU	71.1 (+0.0)	72.2 (+1.1)	71.7 (+0.6)	71.6 (+0.5)	72.3 (+1.2)	72.0 (+0.9)	72.2 (+1.1)	**72.7 (+1.6)**

4.3.2. Comparison on Dark Zurich

To verify the effectiveness of our method on night scenes, we further conduct comparative experiments on the Dark Zurich-val dataset, and the results are shown in Table 3. Notably, built upon HRDA, our method achieves state-of-the-art performance on Dark Zurich-val with 52.8 [%] mIoU. The visualization results are presented in Figure 4, which verify the effectiveness of our method.

Table 5. Comparison of semantic segmentation performance on GTA5 → Cityscapes on the Cityscapes validation set. Best results are in **bold**.

Method	Backbone	GTA5 → Cityscapes
		mIoU
IAST [20]	Deeplab-v2	52.2
IAST + HCSP (Ours)	Deeplab-v2	52.4
HIAST [37]	Deeplab-v2	56.3
HIAST + HCSP (Ours)	Deeplab-v2	**56.5**

4.3.3. Comparison on Nighttime Driving

To show our method's generalization on night scenes, we also evaluate our approach on the Nighttime Driving test set in Table 3 when performing cityscapes→DarkZurich domain adaptation, with sample visualization results presented in Figure 4. With HRDA as the backbone, our method consistently achieves the highest performance on this dataset, reaching 59.3 [%] mIoU. These results confirm the strong generalization capability of our method on the Nighttime Driving dataset.

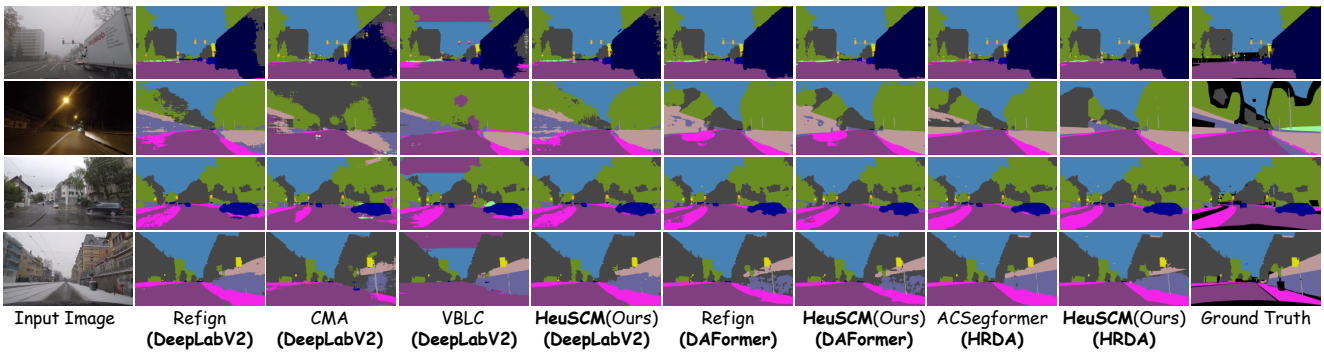

Figure 3. The qualitative comparison between our method and existing state-of-the-art methods based on DeepLabV2, DAFormer, and HRDA on the ACDC val. Compared with the existing state-of-the-art UDA method (*e.g.,* Refign, CMA, VBLC, ACSegFormer), our method achieves better performance under the same backbone. Importantly, our method, based on HRDA, achieves the best performance, with predictions closely matching the ground truth, validating the effectiveness of our method.

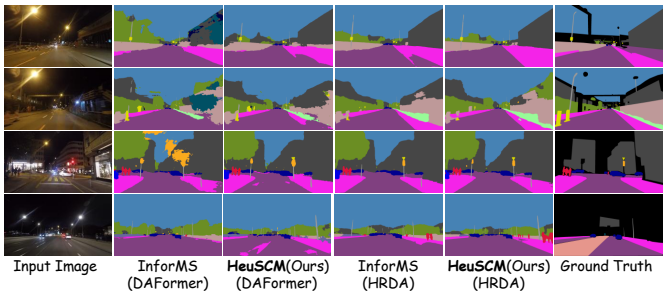

Figure 4. A qualitative comparison between our method and existing state-of-the-art approaches based on DAFormer and HRDA is conducted on the Dark Zurich val (Top two rows) and Nighttime Driving test (Bottom two rows).

4.4. Ablation Study

In this section, we validate the effectiveness of our three core innovations. For the Cityscapes $\rightarrow$ ACDC domain adaptation, we trained several model variants of our HeuSCM (HRDA) and evaluated their performance on the ACDC val set, as shown in Table 4. We adopt Refign [3] as our baseline and incrementally incorporate our proposed components. The results show that adding LSRL, SKFEN, and CαPG individually yields performance gains of 1.1, 0.6, and 0.5 mIoU, respectively. Furthermore, combining two components (LSRL and SKFEN, or LSRL and CαPG, or SKFEN and CαPG) leads to improvements of 1.2, 0.9, and 1.1 mIoU. And full integration of all designs achieves a top performance of 72.7 [%] mIoU, confirming the collective effectiveness and indispensability of each module.

4.5. Generalization Study

To verify the generalization of our proposed Heuristic Class Sampling Policy (HCSP), we replace the sampling strategy of existing hard-class mining methods with our HCSP for synthetic-to-real semantic segmentation. Experimental results are summarized in Table 5. Notably, our HCSP consistently yields significant performance gains, confirming its superior generalization capability.

5. Conclusion

In this paper, we introduce HeuSCM, a novel reinforcement learning framework for unsupervised domain adaptation semantic segmentation. Different from existing class curriculum learning that relies on predefined, human-designed heuristics, we proposed a paradigm shift from "designing" curricula to "learning" them. Our framework employs an autonomous agent guided by two key technical innovations: (1) a High-dimensional Semantic State Extraction that perceives the learning status of the semantic segmentation model, and (2) a Categorical α-Fairness for Policy Gradients that achieves equitable rewards across semantic classes. Extensive experiments demonstrated that our method achieves state-of-the-art performance on challenging adverse weather segmentation benchmarks. Furthermore, its superior results on synthetic-to-real semantic segmentation validate the strong generalization capability of our approach. We believe our work opens a promising new avenue for unsupervised domain adaptation, demonstrating that learned, strategic, multi-objective scheduling policies can significantly outperform traditional fixed-heuristic methods in complex adaptation scenarios.

Acknowledgement

This work was funded by the National Natural Science Foundation of China (Grant No. 62571379) and the Hubei Provincial Key Research and Development Program (Grant No. 2024BAB050). The numerical calculations in this paper have been done on the supercomputing system in the Supercomputing Center of Wuhan University.

References

[1] Liu Bo, Qiulei Dong, and Zhanyi Hu. Hardness sampling for self-training based transductive zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16499–16508, 2021. 3

[2] David Brüggemann, Christos Sakaridis, Tim Brödermann, and Luc Van Gool. Contrastive model adaptation for cross-condition robustness in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11378–11387, Paris, France, 2023. IEEE. 2, 6, 7

[3] David Bruggemann, Christos Sakaridis, Prune Truong, and Luc Van Gool. Refign: Align and refine for adaptation of semantic segmentation to adverse conditions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1–17, Waikoloa, HI, USA, 2023. IEEE. 2, 6, 7, 8

[4] Haoyang Chen, Dongfang Sun, Caoyuan Ma, Shiqin Wang, Kewei Zhang, Zheng Wang, and Zhixiang Wang. Subjective camera 1.0: Bridging human cognition and visual reconstruction through sequence-aware sketch-guided diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17838–17847, 2025. 1

[5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 6, 7

[6] Yang Chen, Liang Shan, Yi Qu, Weilong Zhang, and Lu Chang. Amsc: Adaptive masking and structure-constraint learning for domain adaptive semantic segmentation under adverse conditions. *IEEE Signal Processing Letters*, 31:181–185, 2023. 2

[7] Jiahua Dong, Yang Cong, Gan Sun, Yuyang Liu, and Xiaowei Xu. Cscl: Critical semantic-consistent learning for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 745–762. Springer, 2020. 3

[8] Ziyang Gong, Fuhao Li, Yupeng Deng, Deblina Bhattacharjee, Xianzheng Ma, Xiangwei Zhu, and Zhenming Ji. Coda: Instructive chain-of-domain adaptation with severity-aware visual prompt tuning. In *European Conference on Computer Vision*, pages 130–148. Springer, 2024. 2, 6, 7

[9] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9924–9935, 2022. 6, 7

[10] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In *European conference on computer vision*, pages 372–391. Springer, 2022. 6, 7

[11] Arnaud Judge, Nicolas Duchateau, Thierry Judge, Roman A Sandler, Joseph Z Sokol, Christian Desrosiers, Olivier Bernard, and Pierre-Marc Jodoin. Reinforcement learning for unsupervised domain adaptation in spatio-temporal echocardiography segmentation. *arXiv preprint arXiv:2510.14244*, 2025. 3

[12] Sohyun Lee, Namyup Kim, Sungyeon Kim, and Suha Kwak. Frest: Feature restoration for semantic segmentation under multiple adverse conditions. In *European Conference on Computer Vision*, pages 1–18. Springer, 2025. 2

[13] Fuhao Li, Ziyang Gong, Yupeng Deng, Xianzheng Ma, Renrui Zhang, Zhenming Ji, Xiangwei Zhu, and Hong Zhang. Parsing all adverse scenes: Severity-aware semantic segmentation with mask-enhanced cross-domain consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 13483–13491, 2024. 2

[14] Mingjia Li, Binhui Xie, Shuang Li, Chi Harold Liu, and Xinjing Cheng. Vblc: visibility boosting and logit-constraint learning for domain adaptive semantic segmentation under adverse conditions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8605–8613, Washington, DC, USA, 2023. 2, 6, 7

[15] Bo Liu, Lihua Hu, Zhanyi Hu, and Qiulei Dong. Hardboost: boosting zero-shot learning with hard classes. *arXiv preprint arXiv:2201.05479*, 2022. 3

[16] Wenyu Liu, Song Wang, Jianke Zhu, Xuansong Xie, and Lei Zhang. Domain adaptation transformer for unsupervised driving-scene segmentation in adverse conditions. *IEEE Transactions on Intelligent Transportation Systems*, 2024. 2, 6, 7

[17] Yahao Liu, Jinhong Deng, Xinchen Gao, Wen Li, and Lixin Duan. Bapa-net: Boundary adaptation and prototype alignment for cross-domain semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8801–8811, 2021. 3

[18] Xianzheng Ma, Zhixiang Wang, Yacheng Zhan, Yinqiang Zheng, Zheng Wang, Dengxin Dai, and Chia-Wen Lin. Both style and fog matter: Cumulative domain adaptation for semantic foggy scene understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18922–18931, 2022. 1

[19] Cristina Mata, Kanchana Ranasinghe, and Michael S Ryoo. Copt: Unsupervised domain adaptive segmentation using domain-agnostic text embeddings. In *European Conference on Computer Vision*, pages 424–440. Springer, 2025. 6, 7

[20] Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. In *European conference on computer vision*, pages 415–430. Springer, 2020. 7

[21] Junhan Peng, Jia Su, Yongqing Sun, Zheng Wang, and Chia-wen Lin. Semantic nighttime image segmentation via illumination and position aware domain adaptation. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1034–1038. IEEE, 2021. 1

[22] Christos Sakaridis, David Bruggemann, Fisher Yu, and Luc Van Gool. Condition-invariant semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2, 6, 7

[23] Fengyi Shen, Li Zhou, Kagan Kuecuekaytekin, George Basem Fouad Eskandar, Ziyuan Liu, He Wang, and Alois Knoll. W-controluda: Weather-controllable diffusion-

assisted unsupervised domain adaptation for semantic segmentation. *IEEE Robotics and Automation Letters*, 2025. 2

[24] Usman Ahmad Usmani, Jafreezal Jaafar, Izzatdin Abdul Azis, Junzo Watada, and Abdullah Yousuf Usmani. Reinforced domain adaptation for medical semantic segmentation: Style adaptation and boundary enhancement. *Available at SSRN 4896025*, 2023. 3

[25] Shiqin Wang, Xin Xu, Xianzheng Ma, Kui Jiang, and Zheng Wang. Informative classes matter: Towards unsupervised domain adaptive nighttime semantic segmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 163–172, 2023. 1, 3, 7

[26] Shiqin Wang, Xin Xu, Haoyang Chen, Kui Jiang, Zheng Wang, and Ke Tang. Low-light salient object detection meets the small size. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 9(4):2754–2766, 2024.

[27] Shiqin Wang, Xin Xu, Haoyang Chen, Kui Jiang, and Zheng Wang. The parables of the mustard seed and the yeast: Extremely low-budget, high-performance nighttime semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7853–7861, 2025. 1

[28] Yuxi Wang, Jian Liang, Jun Xiao, Shuqi Mei, Yuran Yang, and Zhaoxiang Zhang. Informative data mining for one-shot cross-domain semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1064–1074, 2023. 3

[29] Yuxi Wang, Jian Liang, and Zhaoxiang Zhang. A curriculum-style self-training approach for source-free semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 3, 6, 7

[30] Xin Xu, Shiqin Wang, Zheng Wang, Xiaolong Zhang, and Ruimin Hu. Exploring image enhancement for salient object detection in low light images. *ACM transactions on multimedia computing, communications, and applications (TOMM)*, 17(1s):1–19, 2021. 1

[31] Jun Yu, Guochen Xie, Quansheng Liu, Zhen Kan, Lei Wang, Tianyu Liu, Qiang Ling, Wei Xu, and Fang Gao. Contrastive learning with multiple prototypes for unsupervised domain adaptive semantic segmentation. *IEEE Transactions on Multimedia*, 2025. 6, 7

[32] Youshan Zhang, Hui Ye, and Brian D Davison. Adversarial reinforcement learning for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 635–644, 2021. 3

[33] Yukun Zhang, Xiansheng Guo, Henry Leung, and Lin Li. Cross-task and cross-domain sar target recognition: A meta-transfer learning approach. *Pattern Recognition*, 138: 109402, 2023. 3

[34] Shen Zheng, Anurag Ghosh, and Srinivasa G Narasimhan. Instance-warp: Saliency guided image warping for unsupervised domain adaptation. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 8197–8206. IEEE, 2025. 6, 7

[35] Ziqiang Zheng, Yingshu Chen, Binh-Son Hua, and Sai-Kit Yeung. Compuda: Compositional unsupervised domain adaptation for semantic segmentation under adverse conditions. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7675–7681, Detroit, MI, USA, 2023. IEEE. 2, 6, 7

[36] Xian Zhong, Shidong Tu, Xianzheng Ma, Kui Jiang, Wenxin Huang, and Zheng Wang. Rainy wcity: A real rainfall dataset with diverse conditions for semantic driving scene understanding. In *Ijcai*, pages 1743–1749, 2022. 1

[37] Chuang Zhu, Kebin Liu, Wenqi Tang, Ke Mei, Jiaqi Zou, and Tiejun Huang. Hard-aware instance adaptive self-training for unsupervised cross-domain semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 3, 7