

Self-Aligned Concave Curve: Illumination Enhancement for Unsupervised Adaptation

Wenjing Wang

Peking University, Beijing, China
daoshee@pku.edu.cn

Haofeng Huang

Peking University, Beijing, China
hhf@pku.edu.cn

Zhengbo Xu

Peking University, Beijing, China
icey.x@pku.edu.cn

Jiaying Liu*

Peking University, Beijing, China
liujiaying@pku.edu.cn

ABSTRACT

Low light conditions not only degrade human visual experience, but also reduce the performance of downstream machine analytics. Although many works have been designed for low-light enhancement or domain adaptive machine analytics, the former considers less on high-level vision, while the latter neglects the potential of image-level signal adjustment. How to restore underexposed images/videos from the perspective of machine vision has long been overlooked. In this paper, we are the first to propose a learnable illumination enhancement model for high-level vision. Inspired by real camera response functions, we assume that the illumination enhancement function should be a concave curve, and propose to satisfy this concavity through discrete integral. With the intention of adapting illumination from the perspective of machine vision without task-specific annotated data, we design an asymmetric cross-domain self-supervised training strategy. Our model architecture and training designs mutually benefit each other, forming a powerful unsupervised normal-to-low light adaptation framework. Comprehensive experiments demonstrate that our method surpasses existing low-light enhancement and adaptation methods and shows superior generalization on various low-light vision tasks, including classification, detection, action recognition, and optical flow estimation. All of our data, code, and results will be available online upon publication of the paper.

CCS CONCEPTS

• **Computing methodologies** → **Image processing; Computer vision problems.**

KEYWORDS

Low-light, domain adaptation, low-level vision, high-level vision

*Corresponding Author. This work was supported by the National Natural Science Foundation of China under Contract No.62172020, and a research achievement of Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Technology).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '22, October 10–14, 2022, Lisboa, Portugal

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9203-7/22/10...\$15.00

<https://doi.org/10.1145/3503161.3547991>

ACM Reference Format:

Wenjing Wang, Zhengbo Xu, Haofeng Huang, and Jiaying Liu. 2022. Self-Aligned Concave Curve: Illumination Enhancement for Unsupervised Adaptation. In *Proceedings of the 30th ACM International Conference on Multimedia (MM '22)*, October 10–14, 2022, Lisboa, Portugal. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3503161.3547991>

1 INTRODUCTION

Insufficient lighting is a common kind of image degradation, caused by dark environments, corrupted equipment, or improper shooting settings. It may degrade the visual quality of images, causing low visibility, lost details, and aesthetic distortion. Besides, with the rise of machine learning, visual analytics has been playing an increasingly critical role in many applications. Low light can also pose threats to machine analytics, presenting challenges to high-level vision tasks in the low-light condition, such as nighttime autonomous driving and surveillance video analysis.

The research of restoring low-light images has attracted wide attention. From early manually designed algorithms [28] to recent data-driven deep models [39], a great number of works have effectively improved human visual quality of low-light images. However, most existing low-light enhancement methods do not take machine vision into account. Some methods introduce noise removal [60] or detail reconstruction [64], which enhances visual quality but might mislead high-level analytics models. Some other methods take semantic perception into consideration [40], but they still focus on human vision and have unsatisfactory performance on downstream high-level vision tasks.

In recent years, to promote the development of machine analytics in low-light scenarios, many datasets have been built [42, 67], giving birth to a series of machine vision models tailored for underlit environments. With a large amount of training data, many methods improve normal light frameworks by incorporating the properties of insufficient lighting conditions [6]. Some works further explore the task of adapting normal light models to low-light without task-specific annotation, which is more universal and application-driven [32, 58, 61]. However, existing normal-to-low light adaptation algorithms either rely on multiple sources [61], adopt troublesome multi-stage and multi-level processes [58], or fail in darker cases [32]. Moreover, most adaptive methods concentrate on the high-dimensional features of machine analytics models, but neglect the characteristics of input images themselves.

Different from existing normal-to-low light adaptation methods, we make full use of the potential of illumination adjustment. We first build an illumination enhancement model which can maximize

the model’s abilities while being easy to learn. Inspired by camera response functions, we design a new model paradigm: *deep concave curve*, which can determine the new pixel value in the enhanced result with a high degree of freedom. To effectively satisfy concavity, we propose to first predict a non-positive second derivative, then apply discrete integral implemented by convolutions. To train this model towards unsupervised adaptation, we design *asymmetric self-supervised alignment*. On the normal light side, we learn decision heads with a self-supervised pretext task. Then on the low-light side, we fix the decision heads and let our model improve the pretext task performance through enhancing the input image. In this way, even without annotated data, our model can learn how to make the machine analytics model better perceive the enhanced low-light image. To make full use of image information and provide good guidance for illumination enhancement, we propose a new rotated jigsaw permutation task. Experiments show that our model architecture and training design are compatible with each other. On one hand, our self-learned strategy can better restore illumination compared with other feature adaptation strategies; on the other hand, our deep concave curve can best maximize the potential of self-learned illumination alignment.

The proposed illumination enhancement model, self-aligned concave curve (SACC), can serve as a powerful tool for unsupervised low-light adaptation. Although SACC does not require normal or low-light annotations and does not even adjust the downstream model, it achieves superior performance on a variety of low-light vision tasks. To further deal with noises and semantic domain gaps, we propose to adapt downstream analytics models by pseudo labeling. Finally, we build an adaptation framework SACC+, which is concise and easy to implement but can outperform existing low-light enhancement and adaptation methods by a large margin.

In summary, our contributions are threefold:

- We are the first to propose a learnable pure illumination enhancement model for high-level vision. Inspired by camera response functions, we design a deep concave curve. Through discrete integral, the concavity constraint can be satisfied through the model architecture itself.
- Towards unsupervised normal-to-low light adaptation, we design an asymmetric cross-domain self-supervised training strategy. Guided by the rotated jigsaw permutation pretext task, our curve can adjust illumination from the perspective of machine vision.
- To verify the effectiveness of our method, we explore various high-level vision tasks, including classification, detection, action recognition, and optical flow estimation. Experiments demonstrate our superiority over both low-light enhancement and adaptation state-of-the-art methods.

2 RELATED WORKS

Low-light Enhancement. Early methods manually design illumination models and enhancement strategies. In the Retinex theory [31], images are decomposed into reflectance (albedo) and shading (illumination). On this basis, Retinex-based methods [10, 15] first decompose images and then either separately or simultaneously process the two components. Histogram equalization and its variants [48] instead redistribute the intensities on the histogram.

Recent methods are mainly based on deep learning. Some mod-

els mimic the Retinex decomposition process [60, 70]. RUAS [39] unrolls the optimization process of Retinex-inspired models and searches desired network architectures. EnlightenGAN [27] introduces adversarial learning. Zero-DCE [14] designs a curve-based low-light enhancement model and learns in a zero-reference way. Some methods also target RAW images [5], videos [4, 25], and introduce extra light sources [62, 63]. Interested readers may refer to [37] and [33] for comprehensive surveys.

Existing low-light enhancement methods disregard downstream machine learning tasks. In comparison, our model targets high-level vision and greatly benefits machine vision in low-light scenarios.

High-Level Vision in Low-light scenarios. With an increasing demand for autonomous driving and surveillance analysis, low-light high-level vision has attracted ever-higher attention in recent years. For dark object detection, Sasagawa *et al.* [50] merged pretrained models in different domains with glue layers and generative knowledge distillation. MAET [6] learns through jointly decoding degrading transformations and detection predictions. HLA-Face [58] adopts a joint pixel-level and feature-level adaptation framework. For nighttime semantic segmentation, DANNet [61] employs adversarial training to adapt models in one stage without additional day-night image transferring. For general tasks, CConv [32] designed a color invariant representation. Some works also focus on tasks of image retrieval [24], depth estimation [57], and matching [52] in low-light conditions.

Despite all these progress on high-level vision in low-light scenarios, many methods rely on low-light annotations, which are neither robust nor flexible enough. Existing unsupervised adaptation methods focus on feature migration and ignore the importance of pixel-level adjustment. Based on illumination enhancement, we propose a new method for low-light adaptation that outperforms existing methods by a wide margin.

3 DEEP CONCAVE CURVE

In this section, we introduce the motivation and detailed architecture of our illumination enhancement model.

3.1 From CRF to Concave Curve

Digital photographic cameras use camera response functions (CRFs) when mapping irradiance to intensities. Although scene illumination changes linearly on the irradiance level, to fit the logarithmic perception of human vision, cameras employ non-linear CRFs, making illumination adjustment complicated on the intensity level. To utilize the linearity of irradiance, some low-light enhancement methods [19, 20] transform intensities to irradiance, adjust the irradiance, and then map irradiance back to intensities. However, back-and-forth irradiance \leftrightarrow intensity mapping is inconvenient and difficult to introduce high-level machine vision guidance.

We propose to simplify the above complex pipeline into one single intensity-level adjustment, which is denoted by g . We first analyze what form g should take. Ignoring spatial variations like lens fall-off [1], vignetting, and signal-dependent noise, CRF can be assumed to be the same for each pixel in an image. Accordingly, we set g to be spatially shared. Second, to follow the numerical range of pixels and preserve order, g should pass $(0, 0)$, $(1, 1)$, and increase monotonically. Additionally, although pixels are discrete, we want g to appear roughly continuous, *i.e.*, like a curve. Despite the




Figure 1: Left: real camera CRFs from the DoRF [13] dataset. Right: the heat map of second-order derivatives in DoRF. Each column represents a CRF.

above constraints, the solution space of g is still too large, making optimization difficult.

DoRF [13] is a dataset of real camera CRFs. As shown in Fig. 1, we find that most CRFs are concave. Statistically, 89.5% of the discrete second derivatives are negative. This is because human perception of brightness is more sensitive to relative differences between darker tones than differences between lighter tones. Accordingly, enhancing illumination on irradiance results in concave transformation on intensities. Therefore, we assume that g should be concave.

With the concavity constraint narrowing the solution space, in Sec. 4 we will develop a learning strategy, through which an appropriate g can be automatically found.

3.2 Form and Implementation

Based on the above analysis, our illumination enhancement adjustment function should be *monotonically increasing* and *concave*. The model is called: deep concave curve.

Now we introduce the detailed design of our model. We use a neural network C to predict g from the input low-light image I_L . Our g is a mapping from original pixel values to new pixel values. If the color space has P numeric values (e.g., 256 for 8-bit images), then g will be an \mathbb{R}^P vector. For original pixels of value $(p-1)/P$ on the input image I_L , where $p \in \{1, 2, \dots, P\}$, their new pixel values would be the p -th element of g . For ease of understanding, we provide an example in the supplementary.

Intuitively, the conditions of monotonic increasing and concavity can be fulfilled by penalty terms on the discrete first and second derivatives. However, this naïve strategy will introduce new losses, adding the burden of balancing multiple learning objectives.

We propose to satisfy these two constraints through the model architecture itself. Instead of directly estimating g , we first predict a minus second derivative $-\nabla^2 c = C(I_L)$. By setting the last layer of network C to ReLU, we ensure that $-\nabla^2 c \geq 0$. Then, we integrate $-\nabla^2 c$ into c . Finally, to fit the range of $[0, 1]$, we normalize c into $g = \text{Norm}(c)$. Since $-\nabla^2 c \geq 0$, g is guaranteed to be concave.

The discrete integral from $-\nabla^2 c$ to ∇c can be implemented as subsequence summation. Denote matrix-vector multiplication $\nabla c = A \times (-\nabla^2 c)$, we set $A = [a_{ij}]$ as an upper triangular matrix:

$$a_{ij} = \begin{cases} 1, & i \leq j, \\ 0, & i > j. \end{cases} \quad (1)$$

Since $-\nabla^2 c \geq 0$ and $a_{ij} \geq 0$, we can derive that $\nabla c \geq 0$, thus g is guaranteed to increase monotonically. Similarly, denote $c =$

$B \times (\nabla c)$, we set $B = [b_{ij}]$ as a strictly lower triangular matrix:

$$b_{ij} = \begin{cases} 0, & i \leq j, \\ 1, & i > j. \end{cases} \quad (2)$$

For simplicity, we merge the two integral into one step $c = D \times (-\nabla^2 c)$, where $D = B \times A$ is pre-calculated. Since the first element of g is known to be zero, we only need to predict the other $P-1$ values, i.e., $-\nabla^2 c \in \mathbb{R}^{P-1}$ and $c \in \mathbb{R}^P$. Finally, for normalization, we divide the vector by its maximum value: $g = \text{Norm}(c) = c / \|c\|_\infty$.

3.3 Network Architecture

We predict an independent g for each color channel. For RGB images, we have g_R, g_G, g_B . Since different color channels share the same summation matrix, we put g_R, g_G , and g_B side-by-side to form a 1D tensor, and implement integral as a 1D convolution with 256 kernels of size 1. The convolution weight is fixed to $D = B \times A$.

The architecture is illustrated in the red part of Fig. 2. To increase receptive field and acceleration, we first down-sample the input image to 16×16 resolution. Curve predictor C consists of a shallow U-net [49], two convolutions, and one fully connected layer. For 8-bit RGB, the output dimension is 765. Then, we reshape the output into 255×3 , and feed it to the integral 1D convolution. Finally, we apply g on the image I_L , which is efficiently implemented with the `index_select` function in PyTorch [47].

4 ASYMMETRIC SELF-SUPERVISED ALIGNMENT

In the previous section, we have proposed the deep concave curve for illumination enhancement. In this section, we will introduce how to train it towards machine vision.

4.1 How to Guide Illumination Enhancement

Our goal is to align the brightness of low-light images to normal light and improve the performance of downstream tasks in low-light conditions. The main idea is to employ high-level vision models as guidance. We first pretrain downstream models on normal light images, then fix their backbones and let our deep concave curve bridge the feature-level gap between low and normal light.

The challenge is that our deep concave curve can only adjust illumination, namely signal restoration, but low and normal light also differs in semantics. To narrow feature-level gaps, existing methods adopt discrepancy metrics [54] and adversarial learning [11]. However, discrepancy metrics and adversarial discriminators bring extra semantic supervision, which can mislead our model and make training hard to converge.

To disentangle signal restoration and semantics, we adopt cross-domain self-supervised pretext tasks. The framework is shown in Fig. 2. We first train normal light self-supervised learning with a fixed-weight feature backbone and a trainable multilayer perceptrons (MLP) head. Although enhanced low-light data and normal light data may contain different objects and scenes, they all follow the distribution of natural images. Therefore, the normal light pretrained model can also accomplish pretext tasks on enhanced low-light images. Using this property, we fix both the backbone and the head, and train our deep concave curve by the self-supervised loss on enhanced low-light images. Since semantics are taken care of by the MLP head, our deep concave curve will not be disturbed

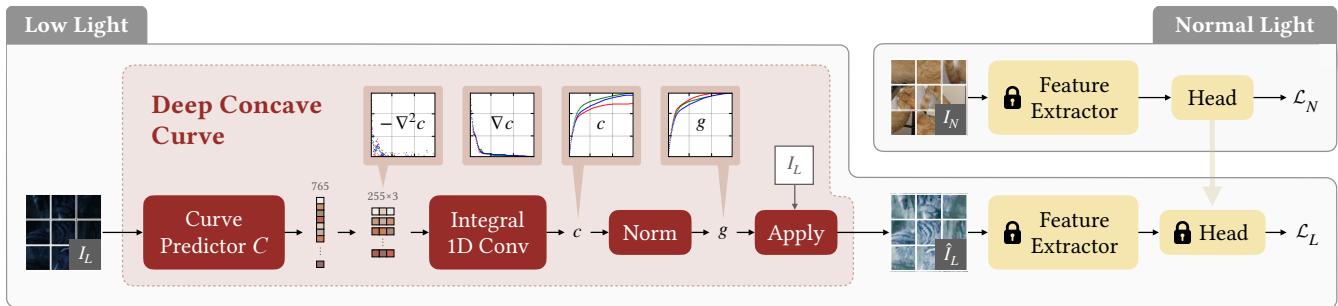


Figure 2: The framework of our model. Our model first predicts a non-negative minus second derivative $-\nabla^2 c$, and then integrates and normalizes it into a concave curve g , which controls the illumination of the enhanced result. The model is trained in an asymmetric self-supervised way. Based on a pretrained and fixed-weight feature extractor, we first train a pretext head on normal light images, and then train our model with the fixed-weight pretext head on low-light images.

by it. Through training to improve pretext task performance, our model learns how to make the feature extractor better perceive the enhanced low-light image, therefore implicitly benefiting downstream tasks in low-light conditions.

4.2 Pretext Task Design

Now we need to find a pretext task that is able to provide sufficient supervision for illumination adaptation. Contrastive learning [16] is one of the most popular self-supervised learning methods nowadays. However, it adopts signal-related augmentations such as color jittering, which may destroy the model’s perception of illumination. Moreover, contrastive learning maximizes the feature distance between different samples, which may mislead the enhancement model into generating diverse but abnormal results. Rotation prediction [12] and jigsaw puzzling [46] do not interfere with the enhancement process, but they are not powerful enough, providing limited supervision for illumination adaptation.

To effectively guide the enhancement model, we propose rotated jigsaw puzzles. We first rotate the input image by random angles, then apply 3×3 jigsaw shuffling and ask the network to recognize the permutation. Compared with traditional jigsaw puzzling, rotation augmentation can increase the difficulty of the pretext task, which forces the MLP heads to better understand semantics, thus introduces more supervision for adaptation.

The learning objective is as follows:

$$\begin{aligned} \mathcal{L}_L &= \mathcal{L}_c(\hat{p}_L^{r-jig}, p_L^{r-jig}), \\ \mathcal{L}_N &= \mathcal{L}_c(\hat{p}_N^{r-jig}, p_N^{r-jig}), \end{aligned} \quad (3)$$

where \mathcal{L}_c is the cross-entropy loss, \hat{p}^{r-jig} is the predicted jigsaw permutation of the randomly rotated and jigsaw shuffled image, p^{r-jig} is the ground truth permutation, L stands for low light and N stands for normal light.

5 FRAMEWORK DESIGN AND ANALYSIS

5.1 Unsupervised Low-light Adaptation

With the above designs, we propose a new unsupervised low-light adaptation framework: self-aligned concave curve (SACC). Given a pretrained normal-light downstream model, we use its feature backbone to train our deep concave curve with Eq. (3). While testing, we first enhance the input image and then apply the downstream

Table 1: Effects of different learning strategies for guiding our illumination enhancement model.

Category	Strategy	mAP (%)
Baseline	-	16.09
Discrepancy Metrics	CMD [68]	28.90
	MMD [2]	33.76
Adversarial Learning	LSGAN [43]	38.05
Self-Supervised Learning	MoCo [16]	38.46
	Rotation [12]	39.59
	Jigsaw [46]	41.01
	Rotated Jigsaw (proposed)	41.31

model. Different from existing low-light adaptation methods, SACC is blind to both normal and low-light annotations and does not even need to adjust the downstream model. Although only performing illumination enhancement, SACC achieves superior performance by taking full advantage of image data and machine features.

Besides insufficient light, there are other low-light factors, such as color bias and noise. For the former, SACC can handle color bias. First, SACC has independent curves for each color channel. Second, downstream feature extractors are usually trained on high-quality normal-light images without color bias, thus encouraging our model to restore balanced color. For the latter, real camera noise is quite hard to remove, but improving the model’s robustness to noise is much easier. On this basis, we adopt pseudo labeling to adapt the downstream model. We predict labels of the enhanced low-light data, remove predictions of low confidence, and then fine-tune the downstream model on normal light data with labels and low-light data with pseudo labels. Our strategy can not only implicitly handle noises but also narrow other kinds of domain gaps, such as scene differences. This advanced version is called SACC+.

5.2 Analysis: Why Self-Supervised Alignment?

In the following, we verify our self-supervised alignment strategy through experiments. The analysis is based on DARK FACE [67] and WIDER FACE [67]. The goal is to improve the performance of WIDER-FACE-pretrained DSFD [35] on DARK FACE by enhancing input images. We do not use multiscale testing for simplicity.

Table 2: Effects of different low-light enhancement backbones under our asymmetric self-supervised strategy.

Category	Method	mAP (%)
Baseline	-	16.09
Other	EnlightenGAN [27]	18.42
	Zero-DCE [14]	25.36
	Gamma Correction x^γ	38.25
Ours	No Constraint	12.44
	$\nabla g \geq 0$	34.40
	$\nabla g \geq 0$ and $\nabla^2 g \leq 0$ (proposed)	41.31
	$\nabla g \geq 0$, $\nabla^2 g \leq 0$, and $\nabla^3 g \geq 0$	40.74

Table 3: Comparison with other low-light enhancement networks under different training strategies.

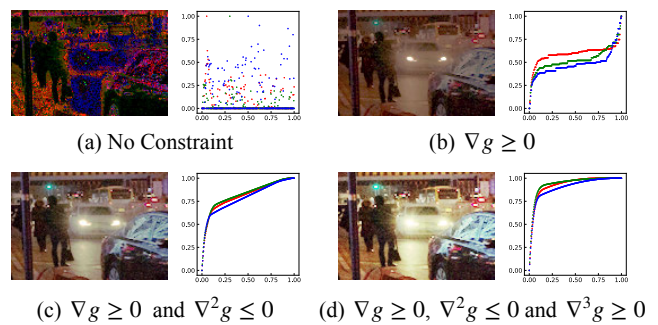
Method	Version	mAP (%)
Baseline	-	16.09
EnlightenGAN [27]	Original	28.85
	Retrained	20.08
	Original + Eq. (3)	32.88
Zero-DCE [14]	Original	38.32
	Retrained	32.53
	Original + Eq. (3)	40.54
SACC (Ours)	-	41.31

Adapting from DARK FACE to WIDER FACE is especially hard since WIDER FACE is collected from the Internet while DARK FACE is captured at nighttime streets. Due to this huge scene gap, discrepancy metrics do not perform well as shown in Table 1. Adversarial learning leads to a better result, but the training is very unstable. In experiments, the balance between the low-light enhancement model and the discriminator often crashes. In comparison, self-supervised learning is more effective and easier to train. The contrastive learning method MoCo [16] has been proved powerful for training networks from scratch. But for adapting illumination, MoCo is less effective than rotation [12] and jigsaw [46]. Finally, our rotated jigsaw strategy achieves the best result, which is in line with our motivation to increase the difficulty of the pretext task.

5.3 Analysis: Why Deep Concave Curve?

Next, we raise another question: under the proposed training strategy, is deep concave curve superior to other network architectures?

We first explore two representative models: EnlightenGAN [27] and Zero-DCE [14]. EnlightenGAN adopts a widely-used U-net [49] like image-to-image architecture. Zero-DCE adjusts images by element-wise second-order curves. We retrain these two networks with Eq. (3). Other experimental settings are the same as Sec. 5.2. As shown in Table 2, their mAP scores are less than 30. In Fig. 3(b) and (c), we can see that there are weird line artifacts. As far as we are concerned, these lines are trivial solutions for the rotated jigsaw permutation problem. In other words, EnlightenGAN and Zero-DCE get overfitted to Eq. (3), therefore cannot generalize to the downstream high-level vision task. In comparison, our deep concave curve obtains better visual results and mAP scores. This is

**Figure 3: Effects of asymmetric self-supervised learning with different low-light enhancement model backbones.****Figure 4: Low-light enhancement results (left) and curve shapes (right) under different curve form constraints.**

because our spatially sharing, monotonically increasing and concavity constraints act as regularization and prevent the model from carving cheating symbols or hints on images.


Naïve gamma correction x^γ ($0 < \gamma < 1$) has been commonly used for low-light enhancement, but the adaptation performance is not as good as ours as shown in Table 2. This is because, for each color channel, x^γ only has one learnable parameter γ . In contrast, our curve has 255 parameters. Compared with x^γ , our curve has a higher degree of freedom and thus is more powerful.

We further explore the necessity of each constraint in Table 2 and Fig. 4. Without any constraint, the enhanced results are far from natural images. With only $\nabla g \geq 0$, the curve does not look continuous and leads to low contrast and dull color as shown in Fig. 4(b). With $\nabla g \geq 0$ and $\nabla^2 g \leq 0$, the model achieves the best result. We also try $\nabla^3 g \geq 0$, but the mAP score gets slightly worse and the image gets partially over-exposed. The reason might be that too strong constraints limit the expression ability of the model. Also, integrating too many times makes the weights of the integral matrix extremely large, resulting in small gradients and increasing the difficulty of neural network training.

Finally, we discuss: if the key lies in preventing overfitting to Eq. (3), will other ‘regularizations’ also work? We explore the constraints of adversarial learning in EnlightenGAN and a group of spatial, exposure, color, and smoothness non-reference loss functions in Zero-DCE. As shown in Table 3, when we retrain EnlightenGAN and Zero-DCE on DARK FACE and WIDER FACE, their performances get degraded, because their learning strategies are

Table 4: Effects of denoising enhanced low-light images and the proposed advanced version SACC+.

Method	Version	mAP (%)
Baseline	SACC	41.31
Denoising	SACC + BM3D [7]	24.51
	SACC + Neighbor2Neighbor [21]	40.54
Proposed	SACC+	45.51

**Figure 5: Example result for classification. The original image is wrongly categorized as “Chair” with a confidence of 0.98, after enhancement, it can be correctly recognized as “Cat”.**

specially designed for their original datasets. Then, we try to keep their loss functions on their original datasets, but add Eq. (3) on DARK FACE and WIDER FACE. In Table 3, we can see that the mAP scores are higher, demonstrating the effectiveness of Eq. (3) on other models. This also indicates that adversarial learning and non-reference loss functions are effective ‘regularizations’ as well. But even when using Eq. (3), EnlightenGAN and Zero-DCE still perform worse than our deep concave curve. This is because adversarial and non-reference learning introduces extra conflicts against high-level vision, while our deep concavity is a purer regularization, and thus can maximize the potential of Eq. (3).

5.4 Analysis: Denoising or SACC+?

We further discuss the necessity of denoising. We evaluate two representative methods: the widely-used BM3D [7] and the learning-based Neighbor2Neighbor [21]. These two methods can remove noise from the perspective of human vision. However, denoising also blurs details, which impedes machine vision. As a result, the accuracy of face detection degrades as shown in Table 4. In comparison, fine-tuning with low-light pseudo labels greatly improves the performance, supporting our design of SACC+.

6 EXPERIMENTS

In this section, we provide implementation details, benchmarking results, and more experimental analysis and applications.

6.1 Implementation Details

Our deep concave curve can be applied to various models and vision tasks. In this paper, we show the results of ResNet [17], DSFD [35] (VGG [51] backbone), I3D [3] (3D-CNN), and PWC-Net [55] (feature pyramid backbone), involving four different visual tasks.

For SACC, we first pretrain the pretext task head on normal light images with the feature extractor fixed. Then, we fix the pretext task and train our deep concave curve on low light images. For SACC+, we first fix the pretrained SACC and then fine-tune the downstream model on normal-light and SACC-enhanced low-light data. Due to page limit, detailed experimental settings for each task

Table 5: Adaptive low-light classification comparison results. † denotes that the low-light enhancement model is retrained.

Category	Method	Top-1 (%)
Baseline	ResNet-18 [17]	55.04
Fully Supervised	Finetuned ResNet-18 [17]	71.52
Low-Light Enhancement	RetinexNet [60]	44.72
	Zero-DCE† [14]	50.80
	EnlightenGAN [27]	57.76
	Zero-DCE++ [34]	58.56
	MF [10]	59.20
	KinD [70]	59.28
	LIME [15]	59.44
	Zero-DCE [14]	59.44
	RUAS† [39]	59.84
	EnlightenGAN† [27]	60.24
LLFlow [59]	60.72	
	SACC (Ours)	61.44
Unsupervised Domain Adaptation	CMD [68]	55.92
	AdaBN [36]	59.68
	DANN [11]	59.76
	CICov [32]	60.96
	SACC+ (Ours)	63.92

are provided in the supplementary material.

6.2 Low-light Object Classification

We first evaluate our approach by one of the most fundamental vision tasks: image classification. CODaN [32] is a 10-class dataset for low-light adaptation. It contains 13k daytime and 2.5k nighttime images. We use its official daytime setting and split 1.25k nighttime images for training/validation and the others for testing. The task is to adapt ResNet-18 [17] from daytime to nighttime. Our method is compared with nine low-light enhancement and four unsupervised domain adaptation methods. Enhancement methods only adjust input images, while adaptation methods also adjust the classifier. They all cannot obtain low-light labels. We additionally show the result of fully supervised learning as a reference of performance upper bound, where the baseline is trained with low-light labels.

The results are shown in Table 5. Low-light enhancement methods restore input low-light images from the perspective of human visual quality and neglect machine vision, thus resulting in limited performance on CODaN. Unsupervised domain adaptation methods adapt the classifier to nighttime. CMD [68], AdaBN [36], and DANN [11] are designed for normal light, but the day-night domain gap is more complex. Accordingly, their accuracies are all below 60%. CICov [32] uses a color invariant layer. Benefitting from its robust feature, CICov is superior to other low-light enhancement and unsupervised domain adaptation methods. But compared with our methods, CICov is less effective. Through adjusting illumination from the perspective of machine vision, our SACC best adapts the daytime classifier to nighttime. With the help of pseudo label fine-tuning to adjust the downstream classifier, the accuracy of our SACC+ is even higher. Fig. 5 shows an example result. $-\nabla^2c$ is very discrete, but after integral g appears like a continuous curve. As the input image suffers from red color bias, our model brightens more on the green channel, which adjusts the color balance.




Figure 6: Subjective comparison results for dark face detection. † denotes that the low-light enhancement model is retrained. The color of the bounding boxes represents the confidence of recognition, with yellow indicating higher confidence.

Table 6: Adaptive dark face detection comparison results. † denotes that the low-light enhancement model is retrained.

Category	Method	mAP (%)
Baseline	DSFD [35]	16.09
Fully Supervised	Finetuned DSFD [35]	45.99
Low-Light Enhancement	RetinexNet [60]	12.04
	KinD [70]	15.84
	EnlightenGAN† [27]	20.77
	EnlightenGAN [27]	31.31
	Zero-DCE † [14]	37.30
	LLFlow [59]	37.41
	RUAS† [39]	38.36
	LIME [15]	40.71
	Zero-DCE++ [34]	40.90
	Zero-DCE [14]	41.27
MF [10]	41.43	
	SACC (Ours)	44.57
Unsupervised Domain Adaptation	CICov [32]	4.40
	OSHOT [8]	25.38
	Progressive DA [18]	28.47
	Pseudo Labeling [23]	35.07
	HLA-Face [58]	44.44
	SACC+ (Ours)	46.91

6.3 Dark Face Detection

We further conduct a benchmark on unsupervised dark face detection. WIDER FACE [66] contains 32k images of various events and scenes. DARK FACE [67] contains 10k nighttime street scene images. We use their official splittings. The baseline is DSFD [35].

Results are shown in Table 6. Although RetinexNet [60] and KinD [70] consider detail reconstruction and noise reduction, they introduce extra artifacts for machine vision and thus degrade the performance of face detection. Other low-light enhancement methods can benefit DSFD, but their mAP scores are lower than ours. These methods may make enhanced images look unnatural in machine vision, which can mislead the face detector as shown in Fig. 6.

Unsupervised domain adaptation methods OSHOT [8], Progressive DA [18], and Pseudo Labeling [23] cannot effectively handle the normal-low light gap. Although CICov [32] is effective on CO-DaN, due to the much lower illumination of nighttime street scenes, the color invariant layer of CICov fails on DARK FACE. As visualized in Fig. 6(i), the representation of CICov suffers from severe noise. HLA-Face [58] is the state-of-the-art for WIDER-DARK unsupervised adaptation. It adopts a complex low-level and high-level joint adaptation strategy. In comparison, our SACC only adjusts illumination but still outperforms HLA-Face. Our advanced version SACC+ even outperforms the fully supervised upper bound, demonstrating the great potential of our method.

Now we analyze the generalization of DSFD-trained SACC to other face detectors: PyramidBox [56] and MogFace [41]. As shown in Table 7, our performance is consistently good even for unseen downstream models, indicating that high-level vision knowledge is universal among different machine models.

6.4 Low-light Action Recognition

Although initially designed for images, our method can also generalize to video tasks. Here we evaluate our SACC with 11-class low-light video action recognition. We gather about 800 low-light video clips from ARID [65]. Normal light data consists of 2.6k normal

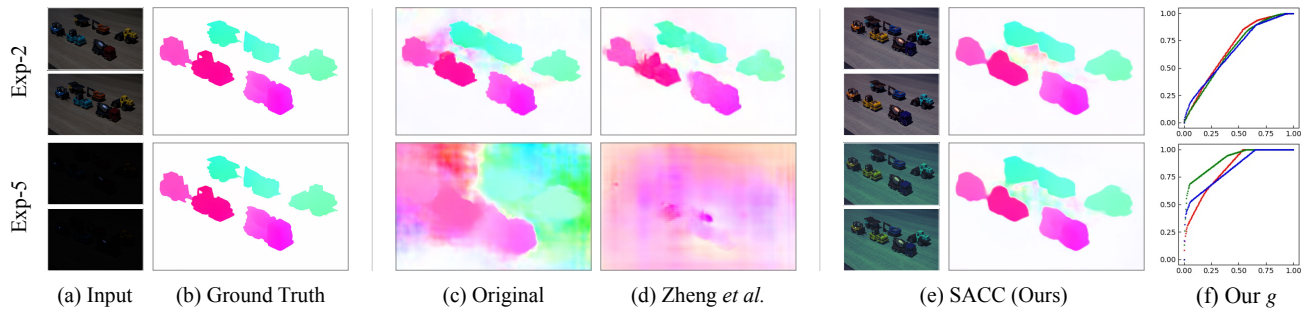


Figure 7: Optical flow estimation results of the same scene under different illumination levels.

Table 7: More results on DARK FACE [67]. Red denotes the best and blue denotes the second best performance.

	Detection mAP (%)				NR-IQA	
	[56]	[35]	[41]	AVG	[44]	[38]
Original	13.99	16.09	16.36	15.48	8.5	24.1
RetinexNet [60]	11.42	12.04	14.56	12.67	7.7	21.9
KinD [70]	15.61	15.84	21.27	17.57	9.8	42.2
EnlightenGAN [†] [27]	19.54	20.77	24.02	21.44	9.0	22.6
EnlightenGAN [27]	28.45	31.31	35.79	31.85	9.7	14.1
Zero-DCE [†] [14]	33.41	37.30	37.75	36.15	6.7	18.4
LLFlow [59]	32.84	37.41	41.08	37.11	8.4	10.3
RUAS [†] [39]	32.77	38.36	40.71	37.28	6.2	4.4
LIME [15]	35.69	40.71	42.82	39.74	6.5	5.7
Zero-DCE++ [34]	35.56	40.90	43.45	39.97	6.4	6.7
Zero-DCE [14]	35.95	41.27	43.62	40.28	6.4	7.9
MF [10]	37.49	41.43	43.87	40.93	6.5	8.2
SACC (Ours)	39.20	44.57	46.45	43.41	6.3	3.2

Table 8: Video action recognition comparison results.

Original	StableLLVE [69]	SMOID [25]	SACC (Ours)
47.14%	49.70%	50.18%	52.13%

light video clips from HMDB51 [30], UCF101 [53], Kinetics-600 [29], and Moments in Time [45]. The action recognizer is I3D [3] based on 3D-ResNet [9]. We report the top-1 accuracy.

We extend SACC to videos by merging all frames into a big image, predicting a unified g for the whole video clip, then applying g to each frame. Since g is temporally shared, temporal consistency can be naturally maintained.

In Table 8, video enhancement methods StableLLVE [69] and SMOID [25] can benefit action recognition but attain limited performance gain. Our SACC instead increases the accuracy by about 5%, showing the superiority of our method for videos.

6.5 Optical Flow Estimation in the Dark

Last but not least, we explore a task that differs a lot from classification and detection: optical flow estimation. The VBOF [71] dataset contains 10k image pairs under various brightness levels. We select Exp-2 to Exp-5 subsets as the adaptation target domain, where Exp-5 is the darkest. The ground truth flow fields are originally

Table 9: Adaptive low-light optical flow estimation results.

	Exp-2	Exp-3	Exp-4	Exp-5	AVG
Original	10.21	11.90	14.36	17.82	13.57
Zheng <i>et al.</i> [71]	9.04	8.81	8.78	9.31	8.99
SACC (Ours)	6.70	7.03	7.57	8.47	7.44

estimated by FlowNet 2.0 [22] on the sharp and normal light Exp-1 subset, we re-calculate them with a more state-of-the-art approach GMA [26]. The baseline to adapt is PWC-Net [55]. Performance is evaluated by end-point error (EPE).

Zheng *et al.* [71] proposed to adapt optical flow models by simulating the noise of dark images and synthesizing a low-light training dataset. However, the synthesis process only considers signal distribution and neglects machine vision, and thus the performance gain is worse than ours as shown in Table 9. Furthermore, as shown in Fig. 7, our SACC is robust to different input illumination levels.

6.6 Subjective Human Vision

We further analyze the relationship between human and machine vision. We benchmark two non-reference image quality assessment (NR-IQA) scores: NIQE [44] and SSEQ [38] in Table 7. Although our SACC only considers high-level tasks, our NR-IQA scores still rank high among all methods. As shown in Fig. 6, our enhanced result has less artifacts and better color balance.

The correlation coefficient of NIQE/SSEQ and the average mAP of all detectors is -0.69/-0.84. This experiment indicates that there is indeed a correlation between human and machine vision, but good visual quality does not always guarantee good high-level performance, which validates our motivation for designing enhancement models tailored for machine vision.

7 CONCLUSION

Towards high-level vision, we propose deep concave curve, which migrates illumination from low to normal light through a self-supervised strategy. Our model architecture and learning strategy cooperate with each other, which form a powerful normal-to-low light adaptation framework. Experimental results demonstrate the superiority of our framework.

REFERENCES

- [1] Naoki Asada, Akira Amano, and Masashi Baba. 1996. Photometric Calibration of Zoom Lens Systems. In *ICPR*.
- [2] Karsten M. Borgwardt, Arthur Gretton, Malte J. Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alexander J. Smola. 2006. Integrating Structured Biological Data by Kernel Maximum Mean Discrepancy. In *ISMB*.
- [3] João Carreira and Andrew Zisserman. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*.
- [4] Chen Chen, Qifeng Chen, Minh N. Do, and Vladlen Koltun. 2019. Seeing Motion in the Dark. In *ICCV*.
- [5] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. 2018. Learning to See in the Dark. In *CVPR*.
- [6] Ziteng Cui, Guo-Jun Qi, Lin Gu, Shaodi You, Zenghui Zhang, and Tatsuya Harada. 2021. Multitask AET With Orthogonal Tangent Regularity for Dark Object Detection. In *ICCV*.
- [7] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. 2007. Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering. *IEEE TIP* 16, 8 (2007), 2080–2095.
- [8] Antonio D’Innocente, Francesco Cappio Borlino, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. 2020. One-Shot Unsupervised Cross-Domain Detection. In *ECCV*.
- [9] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slow-Fast Networks for Video Recognition. In *ICCV*.
- [10] Xueyang Fu, Delu Zeng, Yue Huang, Yinghao Liao, Xinghao Ding, and John W. Paisley. 2016. A Fusion-Based Enhancing Method for Weakly Illuminated Images. *Signal Processing* 129 (2016), 82–96.
- [11] Yaroslav Ganin and Victor S. Lempitsky. 2015. Unsupervised Domain Adaptation by Backpropagation. In *ICML*.
- [12] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised Representation Learning by Predicting Image Rotations. In *ICLR*.
- [13] Michael D. Grossberg and Shree K. Nayar. 2004. Modeling the Space of Camera Response Functions. *IEEE TPAMI* 26, 10 (2004), 1272–1282.
- [14] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. 2020. Zero-Reference Deep Curve Estimation for Low-Light Image Enhancement. In *CVPR*.
- [15] Xiaojie Guo, Yu Li, and Haibin Ling. 2017. LIME: Low-Light Image Enhancement via Illumination Map Estimation. *IEEE TIP* 26, 2 (2017), 982–993.
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *CVPR*.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*.
- [18] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Kumar Singh, and Ming-Hsuan Yang. 2020. Progressive Domain Adaptation for Object Detection. In *WACV*.
- [19] Haofeng Huang, Wenhan Yang, Yueyu Hu, and Jiaying Liu. 2020. Raw-Guided Enhancing Reprocess of Low-Light Image via Deep Exposure Adjustment. In *ACCV*.
- [20] Haofeng Huang, Wenhan Yang, Yueyu Hu, Jiaying Liu, and Ling-Yu Duan. 2022. Towards Low Light Enhancement With RAW Images. *IEEE TIP* 31 (2022), 1391–1405.
- [21] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. 2021. Neighbor2Neighbor: Self-Supervised Denoising From Single Noisy Images. In *CVPR*.
- [22] Eddy Ilg, Nikolaus Mayer, Tomoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. 2017. FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *CVPR*.
- [23] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2018. Cross-Domain Weakly-Supervised Object Detection Through Progressive Domain Adaptation. In *CVPR*.
- [24] Tomás Jeníček and Ondrej Chum. 2019. No Fear of the Dark: Image Retrieval Under Varying Illumination Conditions. In *ICCV*.
- [25] Haiyang Jiang and Yinqiang Zheng. 2019. Learning to See Moving Objects in the Dark. In *ICCV*.
- [26] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard I. Hartley. 2021. Learning to Estimate Hidden Motions with Global Motion Aggregation. In *ICCV*.
- [27] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. 2021. EnlightenGAN: Deep Light Enhancement Without Paired Supervision. *IEEE TIP* 30 (2021), 2340–2349.
- [28] Daniel J. Jobson, Zia-ur Rahman, and Glenn A. Woodell. 1997. A Multiscale Retinex for Bridging the Gap between Color Images and the Human Observation of Scenes. *IEEE TIP* 6, 7 (1997), 965–976.
- [29] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. 2017. The Kinetics Human Action Video Dataset. *arXiv* (2017).
- [30] Hildegard Kuehne, Hueihan Jhuang, Estibaliz Garrote, Tomaso A. Poggio, and Thomas Serre. 2011. HMDB: A Large Video Database for Human Motion Recognition. In *ICCV*.
- [31] Edwin Herbert Land. 1977. The Retinex Theory of Color Vision. *Scientific American* 237 6 (1977), 108–28.
- [32] Attila Lengyel, Sourav Garg, Michael Milford, and Jan C. van Gemert. 2021. Zero-Shot Domain Adaptation with a Physics Prior. In *ICCV*.
- [33] Chongyi Li, Chunle Guo, Ling-Hao Han, Jun Jiang, Ming-Ming Cheng, Jinwei Gu, and Chen Change Loy. 2021. Low-Light Image and Video Enhancement Using Deep Learning: A Survey. *IEEE TPAMI* (2021).
- [34] Chongyi Li, Chunle Guo, and Chen Change Loy. 2021. Learning to Enhance Low-Light Image via Zero-Reference Deep Curve Estimation. *IEEE TPAMI* (2021).
- [35] Jian Li, Yabiao Wang, Changan Wang, Ying Tai, Jianjun Qian, Jian Yang, Chengjie Wang, Jilin Li, and Feiyue Huang. 2019. DSFD: Dual Shot Face Detector. In *CVPR*.
- [36] Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. 2017. Revisiting Batch Normalization For Practical Domain Adaptation. In *ICLR Workshops*.
- [37] Jiaying Liu, DeJia Xu, Wenhan Yang, Minhao Fan, and Haofeng Huang. 2021. Benchmarking Low-Light Image Enhancement and Beyond. *IJCV* 129, 4 (2021), 1153–1184.
- [38] Lixiong Liu, Bao Liu, Hua Huang, and Alan Conrad Bovik. 2014. No-Reference Image Quality Assessment Based on Spatial and Spectral Entropies. *Signal Processing: Image Communication* 29, 8 (2014), 856–863.
- [39] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. 2021. Retinex-Inspired Unrolling With Cooperative Prior Architecture Search for Low-Light Image Enhancement. In *CVPR*.
- [40] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. 2020. Single-Image HDR Reconstruction by Learning to Reverse the Camera Pipeline. In *CVPR*.
- [41] Yang Liu, Fei Wang, Baigui Sun, and Hao Li. 2021. MogFace: Rethinking Scale Augmentation on the Face Detector. *arXiv* (2021).
- [42] Yuen Peng Loh and Chee Seng Chan. 2019. Getting to Know Low-Light Images with the Exclusively Dark Dataset. *CVIU* 178 (2019), 30–42.
- [43] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. 2017. Least Squares Generative Adversarial Networks. In *ICCV*.
- [44] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. 2013. Making a "Completely Blind" Image Quality Analyzer. *IEEE SPL* 20, 3 (2013), 209–212.
- [45] Mathew Monfort, Carl Vondrick, Aude Oliva, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa M. Brown, Quanfu Fan, and Dan Gutfreund. 2020. Moments in Time Dataset: One Million Videos for Event Understanding. *IEEE TPAMI* 42, 2 (2020), 502–508.
- [46] Mehdi Noroozi and Paolo Favaro. 2016. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In *ECCV*.
- [47] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*.
- [48] S. M. Pizer, R. E. Johnston, J. P. Ericksen, B. C. Yankaskas, and K. E. Muller. 1990. Contrast-Limited Adaptive Histogram Equalization: Speed and Effectiveness. In *VBC*.
- [49] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*.
- [50] Yukihiro Sasagawa and Hajime Nagahara. 2020. YOLO in the Dark - Domain Adaptation Method for Merging Multiple Models. In *ECCV*.
- [51] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*.
- [52] Wenzheng Song, Masanori Suganuma, Xing Liu, Noriyuki Shimobayashi, Daisuke Maruta, and Takayuki Okatani. 2021. Matching in the Dark: A Dataset for Matching Image Pairs of Low-light Scenes. In *ICCV*.
- [53] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv* (2012).
- [54] Baochen Sun and Kate Saenko. 2016. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In *ECCV Workshops*.
- [55] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. 2018. PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. In *CVPR*.
- [56] Xu Tang, Daniel K. Du, Zeqiang He, and Jingtuo Liu. 2018. PyramidBox: A Context-Assisted Single Shot Face Detector. In *ECCV*.
- [57] Kun Wang, Zhenyu Zhang, Zhiqiang Yan, Xiang Li, Baobei Xu, Jun Li, and Jian Yang. 2021. Regularizing Nighttime Weirdness: Efficient Self-supervised Monocular Depth Estimation in the Dark. In *ICCV*.
- [58] Wenjing Wang, Wenhan Yang, and Jiaying Liu. 2021. HLA-Face: Joint High-Low Adaptation for Low Light Face Detection. In *CVPR*.
- [59] Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex C. Kot. 2022. Low-Light Image Enhancement with Normalizing Flow. In *AAAI*.
- [60] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. 2018. Deep Retinex Decomposition for Low-Light Enhancement. In *BMVC*.
- [61] Xinyi Wu, Zhenyao Wu, Hao Guo, Lili Ju, and Song Wang. 2021. DANNet: A One-Stage Domain Adaptation Network for Unsupervised Nighttime Semantic Segmentation. In *CVPR*.

- [62] Zhihao Xia, Michael Gharbi, Federico Perazzi, Kalyan Sunkavalli, and Ayan Chakrabarti. 2021. Deep Denoising of Flash and No-Flash Pairs for Photography in Low-Light Environments. In *CVPR*.
- [63] Jinhui Xiong, Jian Wang, Wolfgang Heidrich, and Shree K. Nayar. 2021. Seeing in Extra Darkness Using a Deep-Red Flash. In *CVPR*.
- [64] Ke Xu, Xin Yang, Baocai Yin, and Rynson W. H. Lau. 2020. Learning to Restore Low-Light Images via Decomposition-and-Enhancement. In *CVPR*.
- [65] Yuecong Xu, Jianfei Yang, Haozhi Cao, Kezhi Mao, Jianxiong Yin, and Simon See. 2020. ARID: A New Dataset for Recognizing Action in the Dark. *arXiv* (2020).
- [66] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2016. WIDER FACE: A Face Detection Benchmark. In *CVPR*.
- [67] Wenhan Yang and et al. Ye Yuan. 2020. Advancing Image Understanding in Poor Visibility Environments: A Collective Benchmark Study. *IEEE TIP* 29 (2020), 5737–5752.
- [68] Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. 2017. Central Moment Discrepancy (CMD) for Domain-Invariant Representation Learning. In *ICLR*.
- [69] Fan Zhang, Yu Li, Shaodi You, and Ying Fu. 2021. Learning Temporal Consistency for Low Light Video Enhancement From Single Images. In *CVPR*.
- [70] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. 2019. Kindling the Darkness: A Practical Low-light Image Enhancer. In *ACM MM*.
- [71] Yinqiang Zheng, Mingfang Zhang, and Feng Lu. 2020. Optical Flow in the Dark. In *CVPR*.