

Unsupervised Image Prior via Prompt Learning and CLIP Semantic Guidance for Low-Light Image Enhancement

Igor Morawski¹ Kai He³ Shusil Dangi³ Winston H. Hsu^{1,2}

¹National Taiwan University

²Mobile Drive Technology

³Qualcomm Inc.

Abstract

Currently, low-light conditions present a significant challenge for machine cognition. In this paper, rather than optimizing models by assuming that human and machine cognition are correlated, we use zero-reference low-light enhancement to improve the performance of downstream task models. We propose to improve the zero-reference low-light enhancement method by leveraging the rich visual-linguistic CLIP prior without any need for paired or unpaired normal-light data, which is laborious and difficult to collect. We propose a simple but effective strategy to learn prompts that help guide the enhancement method and experimentally show that the prompts learned without any need for normal-light data improve image contrast, reduce over-enhancement, and reduce noise over-amplification. Next, we propose to reuse the CLIP model for semantic guidance via zero-shot open vocabulary classification to optimize low-light enhancement for task-based performance rather than human visual perception. We conduct extensive experimental results showing that the proposed method leads to consistent improvements across various datasets regarding task-based performance and compare our method against state-of-the-art methods, showing favorable results across various low-light datasets.

1. Introduction

Low-light conditions present a significant challenge for image quality, adversely impacting the performance of high-level computer vision models. In addition to high-level noise caused by photon-limited imaging, low-light images often contain out-of-focus and motion blur, unnatural appearance caused by camera flash and image signal processor failures such as incorrect white-balancing or tone-mapping, all affecting the image quality negatively. With the increased integration of high-level computer vision methods into daily life, addressing the robustness of machine cognition in low-light conditions becomes ever more important.

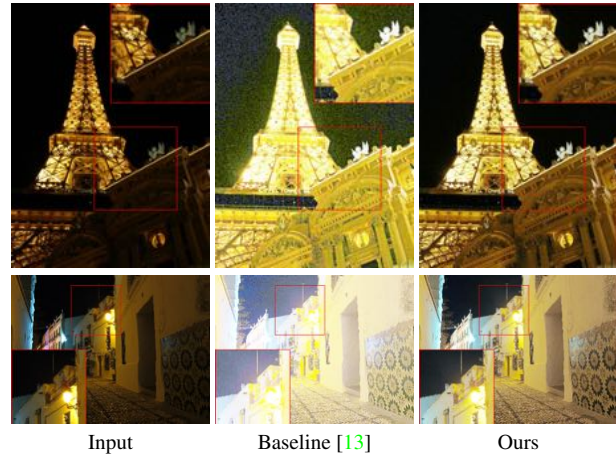


Figure 1. Our proposed method leverages the CLIP [40] model for unsupervised image prior via prompt learning and open-vocabulary semantic guidance. Our proposed method improves the over-all image hue, reduces over-enhancement and reduces noise over-amplification. Further, we conduct extensive experiments to show that our proposed method significantly improves machine cognition as measured by task-based performance of down-stream tasks models, without incurring any additional computation costs on the light-weight enhancement baseline model [13].

While most works try to improve machine cognition by correlating the human visual perception with the image quality, we study machine cognition-oriented low-light image enhancement in this paper.

Many works [9, 51, 53–56, 59, 63–66] propose to supervise the training using paired low- and normal-light data. However, the collection of paired low- and normal-light data is laborious because it usually requires capturing the same scene with different sensor exposure time. Because of the collection costs of low- and normal-light data, some works [13, 20, 27, 32, 58] propose to learn image enhancement in an unsupervised fashion. While some methods [20, 32, 58] still require the selection of unpaired low- and normal-light datasets, zero-reference methods such as [13, 27, 67] only require low-light data for training. Moti-

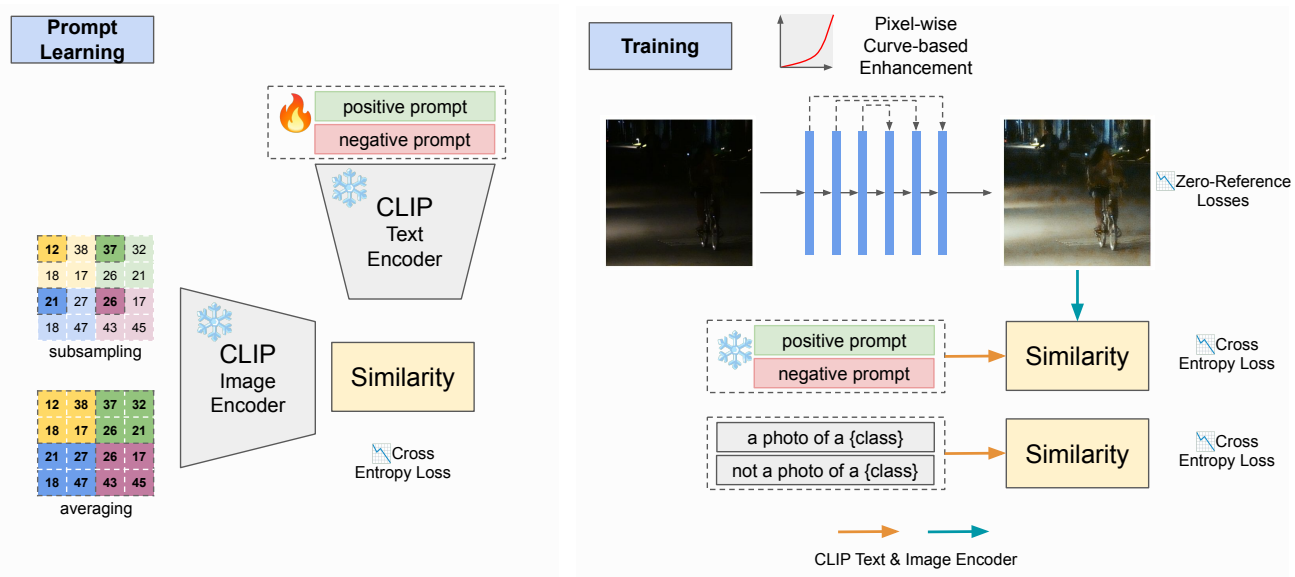


Figure 2. Our proposed two-stage training process leverages the pre-trained CLIP model that can capture lighting conditions and quality of images. We propose to use the CLIP model to learn the positive and negative image priors with a simple data augmentation strategy without any need for paired or unpaired normal-light data via prompt learning, and use them for guiding the image enhancement model. During the training, we use the learned prompts and reuse the CLIP model for semantic guidance to improve the quality of the enhanced images. Our proposed method uses open-vocabulary classification, so it can be easily extended to any dataset, without limiting object categories, with annotated bounding boxes or any type of annotation that can be used to extract patches with an object category, as well as to paired low- and normal-light datasets, increasing the variety of the data in the training.

vated by difficulties in collecting paired normal- and low-light data, we focus on zero-reference image enhancement in this work. While zero-reference methods [13, 27, 67] solve the problems associated with collection costs, they do not integrate semantic knowledge during the training process.

Motivated by the zero-shot capabilities of open vocabulary image understanding [7, 24, 62, 68] enabled by the CLIP [40] model, we incorporate the rich CLIP prior that can capture the lighting conditions and image quality [32, 44] into the zero-reference training process.

Our contributions are as follows:

- We propose to learn a low-light image prior via prompt learning, using a simple data augmentation strategy, without any need for paired or unpaired normal-light data. The proposed prompts help to guide the image enhancement model by improving the contrast, reducing loss of information caused by over-enhancement of bright regions, and reducing over-amplification of noise.
- We reuse the CLIP model for semantic guidance via open vocabulary image classification. While our strategy is simple, it scales favorably to include datasets with any object categories. In addition to that, paired low- and normal-light datasets can also be included in

training by using high-scoring detections inferred by a generic object detector on normal-light data.

- We conduct extensive ablation and comparison experiments to show that our proposed method leads to consistent improvements in task-based performance across many high-level low-light datasets.

2. Related Work

2.1. Low-Light Enhancement

Traditional methods. Traditional approaches to low-light enhancement include histogram-based and Retinex-based methods. Histogram-based methods employ a mapping of intensity values based on histograms at a global [12, 45] or local [23, 43] level. While simple and efficient, traditional histogram-based methods typically discount structural information, do not address the problem of denoising and are prone to producing unnatural artifacts including color shift. On the other hand, methods based on the Retinex theory [25] aim to decompose a low-light image into reflectance and illumination components and treat the estimated reflectance or its modification as the final output. Seminal works single-scale [22] and multi-scale [21] Retinex propose to estimate image illumination with single- and multi-scale Gaussian filters. Over the years, more

complex Retinex-based methods such as [11, 15, 46] have been proposed, including joint enhancement and denoising methods [28, 29], to produce more natural-looking results. Because traditional Retinex-based methods rely on hand-crafted features and assumptions about images, they require careful parameter tuning, might not generalize well to the variety of real-life images, and often lack robustness against image degradation.

Learning-based methods. In recent years, motivated by impressive results demonstrated by deep learning methods in computer vision, many deep learning low-light image enhancement methods have been proposed. Existing methods can be generally divided into end-to-end methods [8, 9, 20, 53–56, 66] that directly enhance the input image and Retinex-based methods [10, 32, 34, 37, 50–52, 59, 63–65] that typically decompose the image into reflectance and illumination components. Many of these methods require large amounts of paired low- and normal-light images to train [9, 51, 53–56, 59, 63–66]. While synthetic data is easy to generate, models trained on it have difficulties generalizing to real-life images. On the other hand, real data requires significant collection and annotation efforts because the paired data is collected either by post-processing images by trained experts [3] or varying camera and lighting settings [4–6, 16, 51]. Although it is possible to use a beam splitter to capture the same scene with two different sensors [19, 26], the majority of the existing datasets are collected by capturing the same scene multiple times with varying exposure settings, and are thus limited to static scenes. Motivated by the costs and limitations of real paired datasets, some works [20, 32, 58] propose to use unpaired low- and normal-light data instead. However, careful selection of normal-light data is still necessary. Consequently, Guo *et al.* [13, 27] proposed a zero-reference method which does not require any paired or unpaired normal-light data. In [13, 27, 67] Guo *et al.* formulated image enhancement as a curve estimation problem and proposed a set of zero-reference losses based on assumptions about natural images, such as average image brightness and gray world hypothesis.

2.2. Low-Light Image Understanding

Loh and Chan [35] have experimentally demonstrated that convolutional neural networks trained on low- and normal-light data for image recognition do not learn to normalize deep image features. That is, Loh and Chan [35] showed that extracted low- and normal-light features belong to distinct feature clusters. Since then, low-light image understanding has further received the interest of researchers in high-level tasks such as face detection [31, 47, 60], object detection [18, 35, 38, 39], pose estimation [26] and action recognition [57]. Another line of works focuses on improving machine cognition by introducing low-light en-

hancement to the cognition framework and learning the enhancement under supervision of a high-level task model [17, 31, 39, 49, 67]. In contrast with previous works in image enhancement and restoration that use semantic information for guidance at a loss [1, 33] or feature level [30, 48], task-oriented enhancement models [17, 31, 39, 49, 67] optimize for machine perception rather than assuming correlation between human perception and downstream task performances [17, 38, 39, 41, 49, 61]. Moreover, as an enhanced image can be reused by many different generic downstream-task models, this approach might be attractive in resource-limited scenarios.

2.3. CLIP in Vision


Following the remarkable success of Contrastive Language-Image Pre-Training (CLIP) [40], the CLIP visual-linguistic prior has been leveraged to generalize to open-vocabulary problems in high-level vision [7, 24, 62, 68]. Recently, Wang *et al.* [44] showed that CLIP can be used to capture the quality and abstract *feel* of the images, and Liang *et al.* [32] showed that the CLIP prior can differentiate between various lighting conditions in natural images. Motivated by this, Liang *et al.* [32] proposed to leverage the CLIP prior for image enhancement by iteratively learning prompts to discriminate between normal-light and back-lit images.

As we focus on low-light machine cognition, we make minimal assumptions about correlating the human and machine visual perception and employ task-based low-light performance to evaluate our proposed enhancement method.

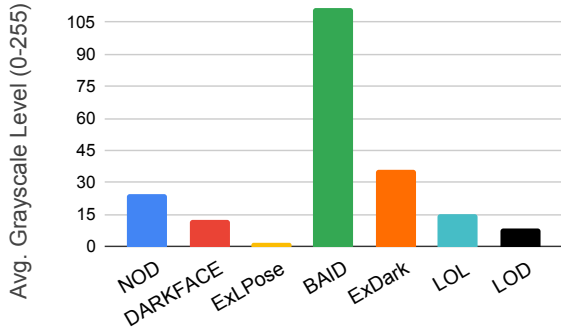
3. Proposed Method

We propose a two-stage training process that leverages the pre-trained CLIP model for semantic guidance and learning an image prior from low-light data without any additional paired or unpaired normal-light data.

During the pre-training stage, we propose a simple but effective strategy to learn the pair of positive and negative image prompts to help to guide the enhancement. Next, we use the learned prompts and additionally reuse the CLIP model for semantic guidance to improve the enhancement of low-light images without any need for paired or unpaired low-light data. Because our proposed semantic guidance is realized by open-vocabulary CLIP classification, the proposed method can be easily extended to any low-light dataset containing bounding-box annotations or any type of annotation that can be used to extract patches with an object category.



(a) Proportions of samples in the dataset.



(b) Average image brightness statistics.

Figure 3. Statistics of the datasets used for the ablation study.




Figure 4. Our strategy to learn positive and negative image prompts: 1 : 4 subsampled (left) negative and 4 × 4 averaged (right) positive image. Subsampling preserves the noise in the image while averaging acts as a fast and simply proxy for denoising.

3.1. Unsupervised Image Prior via Prompt Learning

Motivated by the observation that the CLIP image prior can capture diverse lighting conditions and image quality, we propose to leverage CLIP to learn an image prior from low-light data without the need for paired or unpaired normal-light data. During the first stage of training, we use the pre-trained CLIP model to learn positive and negative prompts to discriminate between averaged and subsampled low-light data, respectively.

Given a low-light image $I \in \mathbb{R}^{H \times W \times C}$ and a pair of

randomly initialized positive and negative prompts, $P_p \in \mathbb{R}^{N \times 512}$ and $P_n \in \mathbb{R}^{N \times 512}$, where N denotes the prompt length, we first augment the image using random photometric augmentation into I' to avoid learning prompts that would overly constraint the overall image brightness. Next, we generate a pair of positive and negative images.

We generate a positive image $I_p = avg_{m \times m}(I')$ by applying $m \times m$ average pooling to the augmented image I' as a fast and simple proxy for denoising. We generate a negative image $I_n = sub_{1:m}(I')$ by applying 1 : m subsampling to the augmented image I' , keeping the original noise levels. The process is illustrated in Fig. 2 and the effect of averaging and subsampling is showed in Fig. 4.

We use the binary cross-entropy loss to learn the prompt pair to capture the image prior while differentiating between the image quality.

$$L_{prompt \text{ init.}} = -(y \log \hat{y} + (1 - y) \log (1 - \hat{y})), \quad (1)$$

where y is the image label (0 for a positive, 4×4 averaged and 1 for a negative, 1 : 4 subsampled image) and \hat{y} is based on the softmax of cosine similarity between each of the prompts and each of the images $I \in \{I_p, I_n\}$:

$$\hat{y} = \frac{e^{\cos(\Phi_{img}(I), \Phi_{txt}(P_p))}}{\sum_{i \in \{p, n\}} e^{\cos(\Phi_{img}(I), \Phi_{txt}(P_i))}}, \quad (2)$$

where Φ_{img} is the CLIP image encoder and Φ_{txt} is the CLIP text encoder.

3.2. Unsupervised Low-Light Enhancement

Motivated by the success of curve-based enhancement [13] for zero-reference low-light enhancement, we adopt DCE-Net [13] as our lightweight baseline model that formulates enhancement as an pixel-wise curve prediction applied to the input image I as below:

$$LE_n(x) = LE_{n-1}(x) + A_n(x)LE_{n-1}(x)(1 - LE_{n-1}(x)), \quad (3)$$

where A is a set of N pixel-wise parameter maps $\alpha \in [-1, 1]$ applied iteratively to the input image I .

3.2.1 Zero-Reference Image Losses

We follow Guo *et al.* [13] and adopt a set of zero-reference loss functions for low-light image enhancement. For clarity, we briefly discuss the employed exposure control L_{exp} , spatial consistency L_{spa} , color constancy L_{RGB} and illumination smoothness L_{TV_A} below.

Exposure control loss L_{exp} encourages exposure correction of under- and over-exposed regions by setting an expected average region intensity E :

$$L_{exp} = \frac{1}{M} \sum_{i=1}^M \sum_{i=1}^M |\hat{I}_i - E|, \quad (4)$$

where \hat{I}_k is the k -th 16×16 patch of the enhanced image, M is the number of non-overlapping 16×16 image patches and E is the expected average intensity. E is empirically set to 0.6

Spatial consistency loss L_{spa} ensures spatial coherence of the enhanced image by preserving relative differences between adjacent image regions:

$$L_{spa} = \frac{1}{K} \sum_{i=1}^K \sum_{j \in \Omega(i)} (|\hat{I}_i - \hat{I}_j| - |I_i - I_j|)^2, \quad (5)$$

where I and \hat{I} are the original and enhanced images, respectively, K is the number of local regions, and $\Omega(i)$ denotes the 4-connected neighborhood centered at the 4×4 region i .

Color loss L_{RGB} follows the Gray-World hypothesis [2] to constraint differences between color channels:

$$L_{RGB} = (\hat{I}_{R,\mu} - \hat{I}_{G,\mu})^2 + (\hat{I}_{R,\mu} - \hat{I}_{B,\mu})^2 + (\hat{I}_{G,\mu} - \hat{I}_{B,\mu})^2 \quad (6)$$

where $\hat{I}_{c,\mu}$ is the average intensity of color channel $c \in \{R, G, B\}$ of the enhanced image \hat{I} .

Illumination smoothness loss L_{TV_A} encourages monotonicity relation between neighboring image pixels and is applied to intermediate curve parameter map A :

$$L_{TV_A} = \sum_{c \in \{R, G, B\}} (|\nabla_x A_c| + |\nabla_y A_c|)^2, \quad (7)$$

where A_c is the intermediate channel-wise curve parameter used to enhance the image \hat{I} , and ∇ denotes the gradient operation.

Additional zero-reference losses. We have experimented with additional zero-reference losses, such as total-variation smoothness loss applied to the enhanced image \hat{I} , total-variation smoothness loss applied to μ -tonemapped enhanced image \hat{I} , comparing differently sampled enhanced image \hat{I} and input image \hat{I} . While the additional constraints we tried led to improved performance for some datasets, we found that the impact is not consistent across many datasets and sometimes lead to significant decrease in performance.

3.2.2 Leveraging CLIP for Image Enhancement

We propose to leverage the pre-trained CLIP model which can capture lighting conditions, image quality and semantic information for image enhancement at the loss level.

Learned CLIP Image Priors. To further constrain the enhanced image \hat{I} and improve the quality of the enhanced images, we use the learned image prompts introduced in subsection 3.1. We found that the image prior learned via prompt learning using our methods reduces over-amplification of the image noise and overexposure of bright regions.

To this end, we first encode the enhanced image \hat{I} using the CLIP image encoder Φ_{img} and the learned prompt pair using the CLIP text encoder Φ_{txt} . Next, we compute the cosine similarity between the embedded enhanced image $\Phi_{img}(\hat{I})$ and the learned prompt pair $\Phi_{txt}(P)$, where $P = P_p, P_n$ is a pair of the positive and negative learned prompts.

$$\hat{y}_{prompt} = \frac{e^{\cos(\Phi_{img}(\hat{I}), \Phi_{txt}(P_p))}}{\sum_{i \in \{p, n\}} e^{\cos(\Phi_{img}(\hat{I}), \Phi_{txt}(P_i))}}, \quad (8)$$

Finally, we compute the cross-entropy loss, assuming that the enhanced image should match the positive prompt P_p .

Semantic guidance. We propose to leverage the zero-shot capabilities of the CLIP model to introduce semantic guidance during training in a straightforward way. Thanks to the simplicity, the proposed semantic guidance method can be used with any low-light dataset containing bounding-box annotations or any type of annotation that can be used to extract patches with an object category, which is advantageous given the sparsity and high collection costs of low-light datasets. During the training, we use annotation to extract image patches counting objects of a class cls . Because we use CLIP to perform zero-shot classification, the label set is not pre-fixed or limited and can be adjusted for a single batch dynamically during the training. For each instance, we use a pair of positive and negative class prompts, $T = \{“a photo of a cls”, “not a photo of a cls”\}$ and experimentally show that this results in improved task-based performance of our method.

Given an enhanced low-light image \hat{I} and its class label cls , we perform the classification based on the cosine similarity between the embedded enhanced image $\Phi_{img}(\hat{I})$ and a pair of antonym prompts $\Phi_{txt}(T)$, where $T = \{“a photo of a cls”, “not a photo of a cls”\}$ is a pair of positive and negative prompts.

$$\hat{y}_{cls} = \frac{e^{\cos(\Phi_{img}(\hat{I}), \Phi_{txt}(P_{cls}))}}{\sum_{i \in \{cls, \neg cls\}} e^{\cos(\Phi_{img}(\hat{I}), \Phi_{txt}(P_i))}}, \quad (9)$$

where P_{cls} and $P_{\neg cls}$ is an antonym pair of positive and negative class prompts, “a photo of a cls ” and “not a photo of a cls ”.

We use binary cross-entropy loss to guide the enhancement process in a way that facilitates high-level machine cognition.

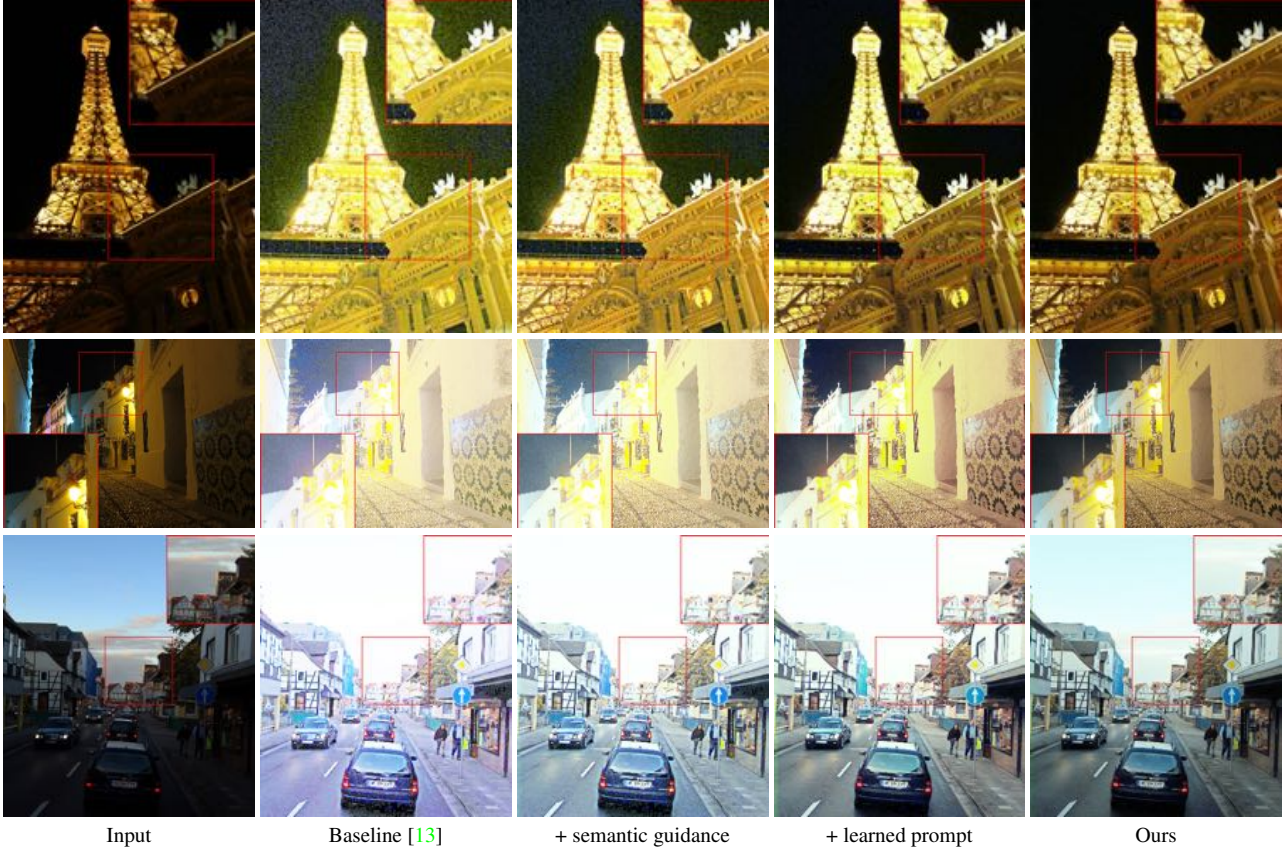


Figure 5. Ablation study of our proposed method. Semantic guidance improves the color distribution of the images, while the learned prompt improves image contrast, reduces overexposure and reduces over-amplification of the noise.

Zero-ref. baseline [13]	Open vocab. class.	Learned img. prior	NOD [39] mAP ↑	NOD SE [39] mAP ↑	LOD [18] mAP ↑	ExDark [35] mAP ↑	ExDark [35] class. acc. ↑	DarkFace [60] mAP@.5 ↑
	(no enhancement)		42.1%	22.8%	37.8%	40.3%	22.8%	33.3%
✓			40.4%	22.0%	38.6%	33.3%	40.9%	33.8%
✓	✓		41.1%	22.5%	42.5%	34.5%	41.2%	28.6%
✓		✓	42.4%	23.3%	47.1%	38.0%	42.7%	35.5%
✓	✓	✓	43.2%	26.5%	47.4%	39.7%	43.6%	38.0%

Legend: **best**, **second best** result.

Table 1. Qualitative results of the ablation study of our proposed method in terms of task-based performance. Our proposed improvements leads to consistent improvement over the baseline method.

In contrast with methods that perform semantic guidance using pre-trained normal-light high-level models, our method can be applied to datasets without paired normal-light ground-truth data. Moreover, our method leverages the zero-shot capabilities of CLIP and is scalable to any dataset with at least bounding box-level annotation via open-vocabulary classification. Additionally, the method is easily extendable to unannotated paired low- and normal-light datasets by extracting high-confidence detections using a normal-light object detector.

Total loss The total training loss is expressed as:

$$L_{total} = \lambda_{exp}L_{exp} + \lambda_{spa}L_{spa} + \lambda_{RGB}L_{RGB} + \lambda_{TV_A}L_{TV_A} + \lambda_{cls}L_{cls} + \lambda_{prompt}L_{prompt}, \quad (10)$$

where L_{exp} , L_{spa} , L_{RGB} , L_{TV_A} are the zero-reference losses described in subsection 3.2.1, L_{cls} and L_{prompt} are the proposed losses described in subsections 3.2.2, and λ are the corresponding loss weights.



Figure 6. Qualitative comparison with related methods, all trained using the LOL [51] dataset.

Method	Requires paired data	NOD [39] mAP ↑	NOD SE [39] mAP ↑	LOD [18] mAP ↑	ExDark [35] mAP ↑	ExDark [35] class. acc. ↑	DarkFace [60] mAP@.5 ↑	FLOPS MACs ↓	Params. ↓
(no enhancement)	N/A	42.1%	22.8%	37.8%	40.3%	22.8%	33.3%	-	-
Histogram Equalization	✗	39.8%	18.4%	38.9%	33.3%	39.9%	32.9%	-	-
DRBN [58]+SKF [53] (CVPR20+CVPR23)	✓	43.1%	25.1%	46.5%	39.5%	38.5%	40.7%	560G	887M
SNR [55]+SKF [53] (CVPR22+CVPR23)	✓	41.6%	20.1%	41.8%	39.0%	39.6%	37.2%	887G	105M
RUAS [34] (CVPR21)	✗	38.6%	21.3%	41.1%	36.7%	34.1%	21.6%	4G	3.4K
SGZ [67] (WACV22)	✗	41.3%	23.9%	42.6%	34.7%	41.7%	37.0%	11G	11K
Zero-DCE [13] (CVPR20)	✗	40.4%	20.5%	38.8%	32.8%	43.1%	33.3%	83G	79K
SCI [37] (CVPR22)	✗	42.9%	26.1%	46.9%	38.7%	43.8%	39.9%	0.37G	0.36K
Ours	✗	42.8%	26.1%	47.0%	38.5%	47.86%	37.3%	83G	79K

Legend: **best**, **second best** result.

Table 2. Qualitative comparisons of the proposed and related works in terms of task-based performance. All methods were trained on the LOL [51] dataset that comprises 500 paired low- and normal-light images.

4. Experiments

In our work, we focus on improving low-light cognition by enhancing images before processing with downstream models. Similar to related work [17, 38, 39, 41, 49, 61], rather than assuming correlation between human perception and downstream task performances, we focus on task-based evaluation. Such evaluation has been used previously in [17, 38, 39, 49].

Implementation details¹. We use Zero-DCE [13] as our baseline model and CLIP for image classification and the learned prompt. For training the prompt, we learn the positive and negative prompts of length 16 using brightness, contrast and hue augmentation, and down-sample image crops by the scale factor of 4 for positive samples and down-sample by nearest-neighbor downsampling by the factor of 4 for negative samples. For training the enhancement model, we extract image patches containing objects based on dataset bounding box annotation and include a portion of the image around the bounding box if the object is very small (e.g. in the DARKFACE [60] dataset). For paired low- and normal-light datasets without annotated object in-

¹The experiments were completed by authors at National Taiwan University.

stances, we ran an open-vocabulary detector YOLO-World [7] on paired normal-light data using 365 labels in the Objects365 [42] dataset and extract patches based on predictions with a confidence score over 0.3. We train our enhancement network on 224×224 image patches with a batch size 8, for 105K steps, using the Adam optimizer with a learning rate 0.0001, weight decay 0.0001 and gradient norm clipping set to 0.1. We use zero-reference loss terms proposed by Guo *et al.* [14] and set the the weights of color loss term to 5, spatial consistency term to 1, exposure control term to 10, and TV term 200.

4.1. Ablation Study

We conduct an extensive ablation study on multiple low-light datasets to show the contribution of each of the proposed loss terms. Following [13] recommendations, we train the model using a variety of different lighting condition. To this end, we use a collection of data extracted from the NOD [39], ExDark [35], DarkFace [60], ExLPose [26], LOL [51] and BAID [36] datasets to train all of the variants in Tab. 1. To further increase the diversity of data, we use paired low- and normal-light datasets by using an open-vocabulary detector on normal-light data. For the LOL [51] and BAID [36] datasets which do not contain any instance

annotation, we extract object instances from low-light images based on paired detections on normal-light data with a confidence score over 0.3 using YOLO-World [7], using 365 labels in the Objects365 [42] dataset. In addition to a large variety of illumination condition, our collection has a variety of different labels categories, and our method leverages zero-shot capabilities of the CLIP [40] model by adjusting the class prompts for each batch on-line. We show the basics dataset statistics in Fig. 3.

Task-based performance. As seen in the ablation study results presented in Tab. 1, the proposed open-vocabulary classification loss leads to consistent improvements over the baseline method [13], except for the DarkFace [60] dataset. Similarly, introducing the proposed learned prompt leads to consistent improvement over both the baseline method [13] and detection with unenhanced images. The two proposed improvements together lead to consistent improvements over the baseline method [13] except the ExDark dataset [35] where the differences between the methods in the study are relatively small and well above detection performance using unenhanced images. Possible explanations for the observed difference for the ExDark [35] dataset include small proportion of the ExDark data during training or relatively higher average brightness levels in the dataset as seen in Fig. 3.

Qualitative results. Qualitative results of the ablation study are shown in Fig. 5. Semantic guidance improves the color distribution of the images, removing the color hue shift visible in images produced by the baseline method. The learned prompt improves the color distribution of the images and contributes to improving the contrast in the images. Compared with the baseline method, the learned prompt reduces the overexposure that leads to information loss from the original input image. Moreover, the learned prompt reduces over-amplification of the noise in the images.

4.2. Task-based Comparison with Related Methods

We show task-based performance comparison with related methods in Tab. 2, all trained on the LOL [51]. We first compare our supervised methods DRBN [58] and SNR-LLNet [55] with semantic guidance [53]. Our method is on-par with DRBN+SKF [53, 58] and outperforms SNR-LLNet [55] in terms of task-based performance across the tested datasets. However, if compared to the computational complexity of the DRBN+SKF [53, 58] our method relies on a much more lightweight architecture [13]. Similarly to the ablation study, our method shows improvement in term of task-based performance across all tested datasets except the ExDark [35] dataset. Again, possible explanations for the observed difference for the ExDark [35] dataset include small proportion of the ExDark data during training or relatively higher average brightness levels in the dataset as seen

in Fig. 3.

5. Conclusion

We proposed to leverage the rich CLIP prior and CLIP's zero-shot capabilities at the training stage to improve zero-reference low-light image enhancement. We proposed to first pre-train a pair of prompts that capture enhanced low-light image prior via prompt learning with a simple data augmentation strategy without any need for paired or unpaired normal-light data. We experimentally showed that the learned prompt helps guiding the enhancement by improving the image contrast, reducing over-enhancement and reducing over-amplification of noise. Next, we proposed to further reuse the CLIP model during the training process using a straightforward yet effective and scalable semantic segmentation via zero-shot open vocabulary classification. We conducted extensive experiments that show consistent improvements over the baseline method across various low-light datasets in terms of task-based performance.

Acknowledgements

This work was supported in part by National Science and Technology Council, Taiwan, under Grant NSTC 111-2634-F-002-022 and by Qualcomm through a Taiwan University Research Collaboration Project.

References

- [1] Andreas Aakerberg, Anders S Johansen, Kamal Nasrollahi, and Thomas B Moeslund. Semantic segmentation guided real-world super-resolution. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 449–458, 2022. 3
- [2] Gershon Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin institute*, 310(1):1–26, 1980. 5
- [3] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input/output image pairs. In *CVPR 2011*, pages 97–104. IEEE, 2011. 3
- [4] Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018. 3
- [5] Chen Chen, Qifeng Chen, Minh N Do, and Vladlen Koltun. Seeing motion in the dark. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 3185–3194, 2019. 3
- [6] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3291–3300, 2018. 3
- [7] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xinggang Wang, and Ying Shan. Yolo-world: Real-

- time open-vocabulary object detection. *arXiv preprint arXiv:2401.17270*, 2024. 2, 3, 7, 8
- [8] Xingbo Dong, Wanyan Xu, Zhihui Miao, Lan Ma, Chao Zhang, Jiewen Yang, Zhe Jin, Andrew Beng Jin Teoh, and Jiajun Shen. Abandoning the bayer-filter to see in the dark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17431–17440, 2022. 3
- [9] Chi-Mao Fan, Tsung-Jung Liu, and Kuan-Hsien Liu. Half wavelet attention on m-net+ for low-light image enhancement. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 3878–3882. IEEE, 2022. 1, 3
- [10] Huiyuan Fu, Wenkai Zheng, Xiangyu Meng, Xin Wang, Chuanming Wang, and Huadong Ma. You do not need additional priors or regularizers in retinex-based low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18125–18134, 2023. 3
- [11] Xueyang Fu, Delu Zeng, Yue Huang, Yinghao Liao, Xinghao Ding, and John Paisley. A fusion-based enhancing method for weakly illuminated images. *Signal Processing*, 129:82–96, 2016. 3
- [12] Rafael C Gonzalez. *Digital image processing*. Pearson education india, 2009. 2
- [13] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1780–1789, 2020. 1, 2, 3, 4, 6, 7, 8
- [14] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 7
- [15] Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on image processing*, 26(2):982–993, 2016. 3
- [16] Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Fengzhu Zou, Fang Lin, and Songchen Han. R2rnet: Low-light image enhancement via real-low to real-normal network. *Journal of Visual Communication and Image Representation*, 90:103712, 2023. 3
- [17] Khurram Azeem Hashmi, Goutham Kallempudi, Didier Stricker, and Muhammad Zeshan Afzal. Featenhancer: Enhancing hierarchical features for object detection and beyond under low-light vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6725–6735, 2023. 3, 7
- [18] Yang Hong, Kaixuan Wei, Linwei Chen, and Ying Fu. Crafting object detection in very low light. In *BMVC*, volume 1, page 3, 2021. 3, 6, 7
- [19] Haiyang Jiang and Yinqiang Zheng. Learning to see moving objects in the dark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7324–7333, 2019. 3
- [20] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE transactions on image processing*, 30:2340–2349, 2021. 1, 3
- [21] Daniel J Jobson, Zia-ur Rahman, and Glenn A Woodell. A multiscale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image processing*, 6(7):965–976, 1997. 2
- [22] Daniel J Jobson, Zia-ur Rahman, and Glenn A Woodell. Properties and performance of a center/surround retinex. *IEEE transactions on image processing*, 6(3):451–462, 1997. 2
- [23] Joung-Youn Kim, Lee-Sup Kim, and Seung-Ho Hwang. An advanced contrast enhancement using partially overlapped sub-block histogram equalization. *IEEE transactions on circuits and systems for video technology*, 11(4):475–484, 2001. 2
- [24] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. F-vlm: Open-vocabulary object detection upon frozen vision and language models. *arXiv preprint arXiv:2209.15639*, 2022. 2, 3
- [25] Edwin H Land. The retinex theory of color vision. *Scientific american*, 237(6):108–129, 1977. 2
- [26] Sohyun Lee, Jaesung Rim, Boseung Jeong, Geonu Kim, ByungJu Woo, Haechan Lee, and Suha Kwak Sunghyun Cho. Human pose estimation in extremely low-light conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3, 7
- [27] Chongyi Li, Chunle Guo, and Chen Change Loy. Learning to enhance low-light image via zero-reference deep curve estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4225–4238, 2021. 1, 2, 3
- [28] Mading Li, Jiaying Liu, Wenhan Yang, and Zongming Guo. Joint denoising and enhancement for low-light images via retinex model. In *International Forum on Digital TV and Wireless Multimedia Communications*, pages 91–99. Springer, 2017. 3
- [29] Mading Li, Jiaying Liu, Wenhan Yang, Xiaoyan Sun, and Zongming Guo. Structure-revealing low-light image enhancement via robust retinex model. *IEEE Transactions on Image Processing*, 27(6):2828–2841, 2018. 3
- [30] Xiaoming Li, Chaofeng Chen, Shangchen Zhou, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. Blind face restoration via deep multi-scale component dictionaries. In *European conference on computer vision*, pages 399–415. Springer, 2020. 3
- [31] Jinxiu Liang, Jingwen Wang, Yuhui Quan, Tianyi Chen, Jiaying Liu, Haibin Ling, and Yong Xu. Recurrent exposure generation for low-light face detection. *IEEE Transactions on Multimedia*, 24:1609–1621, 2021. 3
- [32] Zhexin Liang, Chongyi Li, Shangchen Zhou, Ruicheng Feng, and Chen Change Loy. Iterative prompt learning for unsupervised backlit image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8094–8103, 2023. 1, 2, 3
- [33] Ding Liu, Bihan Wen, Xianming Liu, Zhangyang Wang, and Thomas S Huang. When image denoising meets high-

- level vision tasks: A deep learning approach. *arXiv preprint arXiv:1706.04284*, 2017. 3
- [34] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10561–10570, 2021. 3, 7
- [35] Yuen Peng Loh and Chee Seng Chan. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding*, 178:30–42, 2019. 3, 6, 7, 8
- [36] Xiaoqian Lv, Shengping Zhang, Qinglin Liu, Haozhe Xie, Bineng Zhong, and Huiyu Zhou. Backlitnet: A dataset and network for backlit image enhancement. *Computer Vision and Image Understanding*, 218:103403, 2022. 7
- [37] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5637–5646, 2022. 3, 7
- [38] Igor Morawski, Yu-An Chen, Yu-Sheng Lin, Shushil Dangi, Kai He, and Winston H Hsu. Genisp: neural isp for low-light machine cognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 630–639, 2022. 3, 7
- [39] Igor Morawski, Yu-An Chen, Yu-Sheng Lin, and Winston H Hsu. Nod: Taking a closer look at detection under extreme low-light conditions with night object detection dataset. *arXiv preprint arXiv:2110.10364*, 2021. 3, 6, 7
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 2, 3, 8
- [41] Nicolas Robidoux, Luis E Garcia Capel, Dong-eun Seo, Avinash Sharma, Federico Ariza, and Felix Heide. End-to-end high dynamic range camera pipeline optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6297–6307, 2021. 3, 7
- [42] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. 7, 8
- [43] J Alex Stark. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transactions on image processing*, 9(5):889–896, 2000. 2
- [44] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 2555–2563, 2023. 2, 3
- [45] Qing Wang and Rabab K Ward. Fast image/video contrast enhancement based on weighted thresholded histogram equalization. *IEEE transactions on Consumer Electronics*, 53(2):757–764, 2007. 2
- [46] Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE transactions on image processing*, 22(9):3538–3548, 2013. 3
- [47] Wenjing Wang, Wenhan Yang, and Jiaying Liu. Hla-face: Joint high-low adaptation for low light face detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16195–16204, 2021. 3
- [48] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 606–615, 2018. 3
- [49] Yudong Wang, Jichang Guo, Ruining Wang, Wanru He, and Chongyi Li. Tienet: task-oriented image enhancement network for degraded object detection. *Signal, Image and Video Processing*, pages 1–8, 2023. 3, 7
- [50] Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex Kot. Low-light image enhancement with normalizing flow. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 2604–2612, 2022. 3
- [51] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018. 1, 3, 7, 8
- [52] Wenhui Wu, Jian Weng, Pingping Zhang, Xu Wang, Wenhan Yang, and Jianmin Jiang. Uretinex-net: Retinex-based deep unfolding network for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5901–5910, 2022. 3
- [53] Yuhui Wu, Chen Pan, Guoqing Wang, Yang Yang, Jiwei Wei, Chongyi Li, and Heng Tao Shen. Learning semantic-aware knowledge guidance for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1662–1671, June 2023. 1, 3, 7, 8
- [54] Ke Xu, Xin Yang, Baocai Yin, and Rynson WH Lau. Learning to restore low-light images via decomposition-and-enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2281–2290, 2020. 1, 3
- [55] Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17693–17703, 2022. 1, 3, 7, 8
- [56] Xiaogang Xu, Ruixing Wang, and Jiangbo Lu. Low-light image enhancement via structure modeling and guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9893–9903, June 2023. 1, 3
- [57] Yuecong Xu, Jianfei Yang, Haozhi Cao, Kezhi Mao, Jianxiong Yin, and Simon See. Arid: A new dataset for recognizing action in the dark. In *Deep Learning for Human Activity Recognition: Second International Workshop, DL-HAR 2020, Held in Conjunction with IJCAI-PRICAI 2020, Kyoto, Japan, January 8, 2021, Proceedings 2*, pages 70–84. Springer, 2021. 3
- [58] Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, and Jiaying Liu. From fidelity to perceptual quality: A semi-

- supervised approach for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3063–3072, 2020. 1, 3, 7, 8
- [59] Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *IEEE Transactions on Image Processing*, 30:2072–2086, 2021. 1, 3
- [60] Wenhan Yang, Ye Yuan, Wenqi Ren, Jiaying Liu, Walter J. Scheirer, Zhangyang Wang, Zhang, and et al. Advancing image understanding in poor visibility environments: A collective benchmark study. *IEEE Transactions on Image Processing*, 29:5737–5752, 2020. 3, 6, 7, 8
- [61] Masakazu Yoshimura, Junji Otsuka, Atsushi Irie, and Takeshi Ohashi. Dynamicisp: dynamically controlled image signal processor for image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12866–12876, 2023. 3, 7
- [62] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary detr with conditional matching. In *European Conference on Computer Vision*, pages 106–122. Springer, 2022. 2, 3
- [63] Yonghua Zhang, Xiaojie Guo, Jiayi Ma, Wei Liu, and Jiawan Zhang. Beyond brightening low-light images. *International Journal of Computer Vision*, 129:1013–1037, 2021. 1, 3
- [64] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *Proceedings of the 27th ACM international conference on multimedia*, pages 1632–1640, 2019. 1, 3
- [65] Zhao Zhang, Huan Zheng, Richang Hong, Mingliang Xu, Shuicheng Yan, and Meng Wang. Deep color consistent network for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1899–1908, 2022. 1, 3
- [66] Chuanjun Zheng, Daming Shi, and Wentian Shi. Adaptive unfolding total variation network for low-light image enhancement. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4439–4448, 2021. 1, 3
- [67] Shen Zheng and Gaurav Gupta. Semantic-guided zero-shot learning for low-light image/video enhancement. In *Proceedings of the IEEE/CVF Winter conference on applications of computer vision*, pages 581–590, 2022. 1, 2, 3, 7
- [68] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, pages 696–712. Springer, 2022. 2, 3