

# CLIP-DQA: Blindly Evaluating Dehazed Images from Global and Local Perspectives Using CLIP

Yirui Zeng<sup>1,\*</sup>, Jun Fu<sup>1,\*</sup>, Hadi Amirpour<sup>2</sup>, Huasheng Wang<sup>3</sup>, Guanghui Yue<sup>4</sup>, Hantao Liu<sup>1</sup>, Ying Chen<sup>3</sup>, Wei Zhou<sup>1,†</sup>

<sup>1</sup>School of Computer Science and Informatics, Cardiff University, United Kingdom

<sup>2</sup>Christian Doppler Laboratory ATHENA, Alpen-Adria-Universität, Klagenfurt, Austria

<sup>3</sup>Alibaba Group

<sup>4</sup>School of Biomedical Engineering, Shenzhen University, China

Email: zengyr5@mail2.sysu.edu.cn

fujun@mail.ustc.edu.cn

zhouw26@cardiff.ac.uk

**Abstract**—Blind dehazed image quality assessment (BDQA), which aims to accurately predict the visual quality of dehazed images without any reference information, is essential for the evaluation, comparison, and optimization of image dehazing algorithms. Existing learning-based BDQA methods have achieved remarkable success, while the small scale of DQA datasets limits their performance. To address this issue, in this paper, we propose to adapt Contrastive Language-Image Pre-Training (CLIP), pre-trained on large-scale image-text pairs, to the BDQA task. Specifically, inspired by the fact that the human visual system understands images based on hierarchical features, we take global and local information of the dehazed image as the input of CLIP. To accurately map the input hierarchical information of dehazed images into the quality score, we tune both the vision branch and language branch of CLIP with prompt learning. Experimental results on two authentic DQA datasets demonstrate that our proposed approach, named CLIP-DQA, achieves more accurate quality predictions over existing BDQA methods. The code is available at <https://github.com/JunFu1995/CLIP-DQA>.

**Index Terms**—Dehazed image quality assessment, Prompt learning, CLIP.

## I. INTRODUCTION

Haze is a common natural phenomenon that significantly reduces visibility in scenes, causing many computer vision algorithms, such as object detection [1], [2] and image recognition [3], to experience severe performance degradation. To alleviate this issue, considerable image dehazing algorithms (DHAs) [4]–[10] have been proposed. However, before deploying these DHAs at scale, it is necessary to evaluate their effectiveness, *i.e.*, assess the quality of dehazed images they generate.

The most accurate way to measure the quality of dehazed images is subjective quality evaluation [11], where the quality of dehazed images is directly evaluated by a certain number of human subjects. However, subjective quality evaluation has a narrow range of applications since it needs to conduct time-consuming and labor-intensive subjective experiments. As a result, objective quality evaluation is proposed, aiming to automatically assess the quality of dehazed images without human involvement.

In general, objective dehazed quality measures can be divided into three categories: full-reference dehazed image

quality assessment [12]–[14], reduced-reference dehazed image quality assessment [15]–[20], and no-reference dehazed image quality assessment [11], [21]–[24]. When evaluating the quality of dehazed images, the former two categories require reference images, while the last category only takes dehazed images as input. In real-world scenarios, reference images are typically unavailable. Therefore, NR DQA, also known as blind dehazed image quality assessment (BDQA), has received considerable attention in recent years.

Existing BDQA approaches are mainly composed of traditional methods [11], [21]–[23] and learning-based methods [19], [24], [25]. Traditional methods usually manually design some haze-related features for quality evaluation. Since handcrafted features have limited capabilities in representing distortion and content of dehazed images, the performance of traditional BDQA methods is often unsatisfactory. To this end, learning-based methods employ deep neural networks to automatically extract representative features from dehazed images. Compared to traditional BDQA methods, learning-based BDQA methods achieve more accurate quality prediction. However, their performance is still limited by the small size of the DQA dataset.

To alleviate this issue, we resort to Contrastive Language-Image Pre-Training (CLIP) [26], which has shown good generalization ability across various image recognition tasks [27], [28]. In the BDQA task, the perceptual characteristics of the human visual system need to be considered [11], that is, image understanding based on hierarchical features. Therefore, we feed patches that preserve the local details of the dehazed image and the resized image that preserves the global structure of the dehazed image into CLIP. Then, we tune CLIP by layerwisely inserting learnable prompts into its vision and language branches to accurately map the input hierarchical information of dehazed images into the quality score.

The contributions of this paper are mainly two-fold:

- We present the first CLIP-based BDQA method, CLIP-DQA, which evaluates dehazed images from global and local perspectives.
- We conduct extensive experiments, including ablation studies and visualizations, to verify the efficacy of the proposed method.

\* denotes equal contribution, and † indicates the corresponding author.

## II. PROPOSED METHOD

### A. Problem Formulation

Given a dehazed image  $I$ , blind dehazed image quality assessment aims to estimate the visual quality score of the dehazed image without any reference information. In existing works, a patch-based evaluation framework is widely used for BDQA. In such a framework, we first crop  $N$  patches  $\{I_p^i\}_{i=0}^{N-1}$  from the input dehazed image  $I$ , then estimate the quality score for each patch, and finally use the average score as the quality score of the whole picture:

$$\hat{Q} = \frac{1}{N} \sum_{i=0}^{N-1} f(I_p^i), \quad (1)$$

where  $f(\cdot)$  denotes the BDQA method, and  $\hat{Q}$  is the estimated result.

Inspired by the human visual system using hierarchical features for image understanding, we evaluate the dehazed image from global and local perspectives:

$$\hat{Q} = \frac{1}{N} \sum_{i=0}^{N-1} f(I_p^i, I_s), \quad (2)$$

where  $I_p^i$  keeps the local details and  $I_s$ , a resized version of  $I$ , maintains the global structure. Moreover, we model  $f(\cdot)$  using CLIP, which learns well-generalized knowledge from millions of image-text pairs. Next, we will detail how to evaluate dehazed images with CLIP.

### B. Zero-shot BDQA

Since CLIP can judge the similarity between images and natural language descriptions, we can directly use CLIP to evaluate the dehazed image without training, *i.e.*, zero-shot BDQA.

Let  $T_p$  and  $T_n$  be a pair of antonym text prompts, *e.g.*, "Good photo." and "Bad photo.". The estimated visual quality,  $f(I_p^i, I_s)$  in Equation. 2, can be calculated as follows:

$$f(I_p^i, I_s) = \frac{e^{\text{sim}(t_p, c_i)}}{e^{\text{sim}(t_p, c_i)} + e^{\text{sim}(t_n, c_i)}}, \quad (3)$$

where  $t_p$  and  $t_s$  are the textual representation extracted by the language branch of CLIP from  $T_p$  and  $T_n$ , respectively.  $c_i$  is the visual representation extracted by the vision branch of CLIP from the pair of  $I_p^i$  and  $I_s$ . Here,  $\text{sim}(\cdot, \cdot)$  calculates the cosine distance between textual and visual representations.

However, the performance of the zero-shot BDQA is typically far from satisfactory. There are two main reasons for this result. First, the handcrafted antonym text prompts are often sub-optimal for BDQA since designing effective antonym text prompts requires considerable expertise. Second, the visual representation extracted by CLIP may not be discriminative for BDQA as dehazed images differ from images used for training CLIP in terms of distortion type and appearance.

### C. Multi-modal Prompt Tuning

In order to better adapt CLIP to BDQA, it is necessary to fine-tune CLIP using the DQA dataset. Motivated by the success of prompt learning [30]–[33], we tune CLIP with multi-modal prompts, including textual and visual prompt tuning, as shown in Fig. 1. The core idea of multi-modal prompt tuning is to use learnable textual prompts for automatically mining useful soft antonym text prompts from DQA datasets, and use learnable visual prompts for mitigating the domain gap between dehazed images and natural images.

**Textual Prompt Tuning.** The language branch of CLIP consists of  $K$  transformer layers, and the  $i$ -th transformer layer can be defined as:

$$[W_i] = \mathcal{L}_i(W_{i-1}), \quad i = 1, \dots, K, \quad (4)$$

where  $W_{i-1}$  and  $W_i$  are the input and output of the  $i$ -th transformer layer  $\mathcal{L}_i$ , respectively. The input of the first transform layer,  $W_0$ , corresponds to word embeddings of handcrafted antonym text prompts. To tune the language branch of CLIP, we layerwisely insert learnable prompts:

$$[\_, W_i] = \mathcal{L}_i([F_{i-1}(P_{i-1}), W_{i-1}]), \quad (5)$$

where we map the set of learnable prompts  $P_{i-1}$  into the same space as  $W_{i-1}$  through a fully-connected layer  $F_{i-1}$ . Notably, the output of  $P_{i-1}$  is discarded after the  $i$ -th transformer layer  $\mathcal{L}_i$ .

**Visual Prompt Tuning.** In this paper, the vision branch of CLIP is transformer-based and also contains  $K$  transformer layers. The formulation of each transformer layer is defined as follows:

$$[c_i, E_i^l, E_i^g] = \mathcal{V}_i([c_{i-1}, E_{i-1}^l, E_{i-1}^g]), \quad i = 1, \dots, K, \quad (6)$$

where  $c_i$ ,  $E_i^l$ , and  $E_i^g$  are the class token, the token set of the input patch, and the token set of the input resized dehazed image, respectively. To tune the vision branch of CLIP, we also layerwisely insert learnable prompts:

$$[c_i, E_i^l, E_i^g, \_] = \mathcal{V}_i([c_{i-1}, E_{i-1}^l, E_{i-1}^g, \hat{P}_{i-1}(\hat{P}_{i-1})]), \quad (7)$$

where we map the set of learnable prompts  $\hat{P}_{i-1}$  into the same space as  $E_{i-1}$  through a fully-connected layer  $\hat{F}_{i-1}$ . Similar to textual prompt tuning, we discard the output of the learnable prompts  $\hat{P}_{i-1}$  after the transformer layer  $\mathcal{V}_i$ .

### D. Loss Function

For the DQA task, we use Mean Square Error (MSE) as the training objective:

$$L_{MSE} = \frac{1}{B} \sum_{j=1}^B \|\hat{Q}_j - Q_j\|_2^2, \quad (8)$$

where  $B$  is the batch size,  $\hat{Q}_j$  and  $Q_j$  are the predicted quality score and the mean opinion score of the  $j$ -th dehazed image.

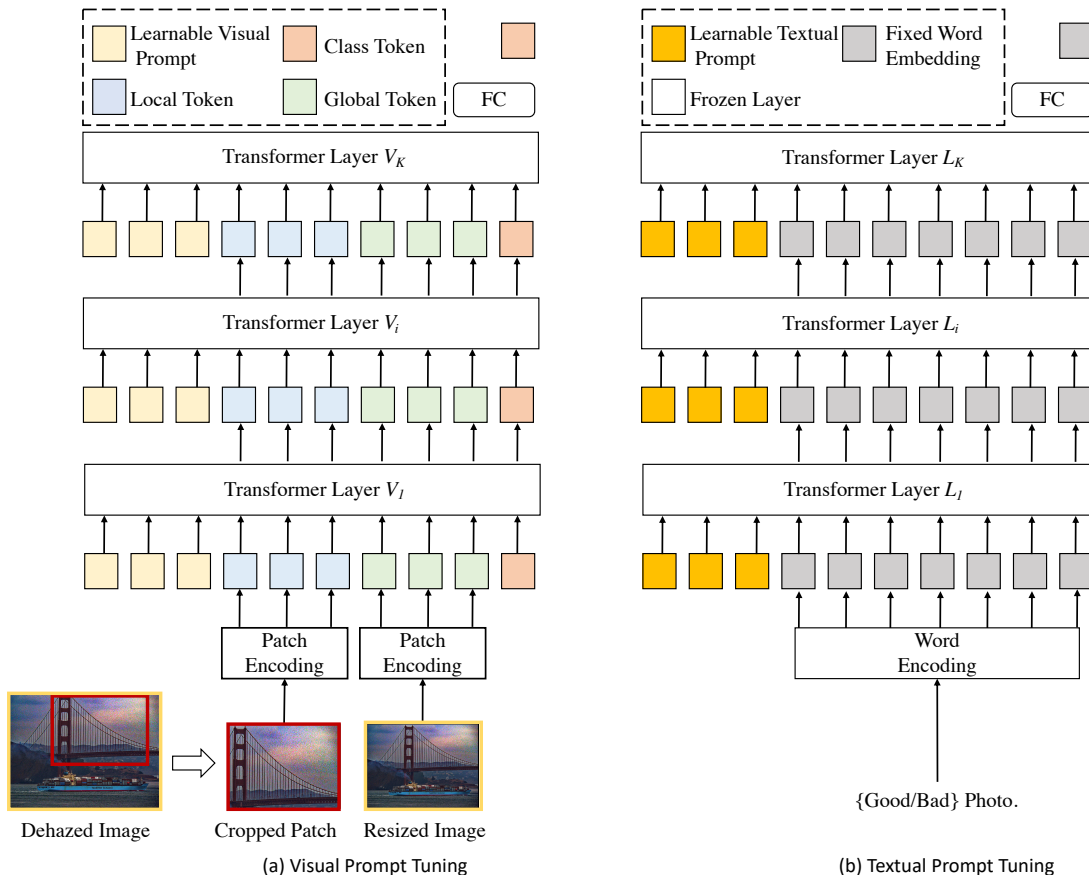


Fig. 1. Illustration of multi-modal prompt tuning for blind dehazed image quality assessment. “FC” means the fully connected layer. We use “Good photo.” and “Bad photo.” as antonym text prompts following the CLIP-based general image quality assessment method [29].

### III. EXPERIMENTS

#### A. Experimental Protocols

To verify the effectiveness of our proposed method, we conduct experiments on two authentic dehazed image quality databases:

- DHQ database [17]: It consists of 250 hazy images and 1,750 dehazed images generated by 7 image dehazing algorithms. Each dehazed image is labeled by a mean opinion score (MOS) ranging from 0 to 100.
- exBeDDE database [13]: It contains 12 haze-free images, 167 hazy images, and 1,670 dehazed images produced by 10 image dehazing algorithms. Each dehazed image is annotated by a MOS ranging from 0 to 1.

As suggested by the video quality expert group (VQEG) [34], we employ Spearman rank order correlation coefficient (SRCC), Pearson linear correlation coefficient (PLCC), and Kendall rank order correlation coefficient (KRCC) for performance comparisons. All three evaluation criteria ranged from 0 to 1, and the higher the value, the better the performance.

For fair comparison, we evaluate each image quality assessment method 10 times, and report the average results. At each

time, we randomly split the dataset into two parts based on the content of the image, 80% for training and 20% for testing.

We build the proposed method based on ViT-B/32 CLIP, and the length of learnable prompts at each transformer layer is set to 8. During training, we keep CLIP frozen and optimize the remaining parts using the Adam optimizer [35] with a learning rate of  $1e-4$ . The total training epoch and the batch size are 50 and 64, respectively. All experiments are run on a single NVIDIA RTX 4090 GPU.

#### B. Performance Comparisons

To validate the effectiveness of the proposed method, we compare it with 13 representative methods, which include three full-reference general image quality assessment (FR GIQA) methods, two FR DQA metrics, three no-reference general image quality assessment (NR GIQA) methods, three conventional NR DQA methods, and two CLIP-based NR DQA methods. All NR methods, except for FADE [21] and HazDetNet [25], are retrained with the same settings as ours. The results are reported in Table I. From this table, we have the following observations. First, five FR IQA methods fail to produce results on the DHQ dataset due to the lack of haze-free images, and underperform on the exBeDDE dataset. Second,

TABLE I  
THE RESULTS OF DIFFERENT METHODS ON THE DHQ AND exBeDDE DATASETS.

Types	Method	DHQ			exBeDDE		
		SRCC	PLCC	KRCC	SRCC	PLCC	KRCC
FR GIQA	PSNR	-	-	-	0.5375	0.5129	0.3758
	SSIM [36]	-	-	-	0.5730	0.5584	0.3990
	VSI [37]	-	-	-	0.5895	0.5789	0.4145
FR DQA	VI [13]	-	-	-	0.5089	0.4909	0.3545
	RI [13]	-	-	-	0.5013	0.4969	0.3515
NR GIQA	CNNIQA [38]	0.6756	0.6937	0.4911	0.8435	0.8142	0.6482
	TReS [39]	0.8395	0.8420	0.6537	0.9090	0.9188	0.7380
	HyperIQA [40]	0.8646	0.8684	0.6848	0.9168	0.9269	0.7497
Conventional NR DQA	FADE [21]	0.2502	0.1845	0.1712	0.7283	0.6659	0.5319
	HazDesNet [25]	0.3175	0.2999	0.2158	0.6848	0.6875	0.4847
	BDQM [24]	0.7733	0.7858	0.5818	0.6465	0.5938	0.4738
CLIP-based NR DQA	CLIPQA [29]	0.3630	0.3545	0.2467	0.3142	0.3161	0.2168
	CLIPQA+ [29]	0.8392	0.8510	0.6523	0.9111	0.9315	0.7394
	CLIP-DQA (Ours)	<b>0.9179</b>	<b>0.9232</b>	<b>0.7568</b>	<b>0.9227</b>	<b>0.9376</b>	<b>0.7589</b>

The performance of FR metrics on the DHQ dataset is not reported due to a lack of reference images.

TABLE II  
ABLATION STUDY ON EACH COMPONENT OF THE PROPOSED METHOD.

Components \ Method		Method			
		$M_1$	$M_2$	$M_3$	Ours
Handcrafted Text Prompts		✓			
Textual Prompt Tuning			✓		✓
Visual Prompt Tuning				✓	✓
Global Information					✓
Local Information		✓	✓	✓	✓
DHQ	SRCC	0.3630	0.8392	0.8925	<b>0.9179</b>
	PLCC	0.3545	0.8510	0.9002	<b>0.9232</b>
	KRCC	0.2467	0.6523	0.7200	<b>0.7568</b>
exBeDDE	SRCC	0.3142	0.8984	0.9111	<b>0.9227</b>
	PLCC	0.3161	0.9219	0.9315	<b>0.9376</b>
	KRCC	0.2168	0.7182	0.7394	<b>0.7589</b>

three NR GIQA methods exhibit competitive performance against three NR DQA methods on the two datasets. This is mainly because FADE and HazDesNet are designed for haze density prediction rather than quality evaluation, while BDQM is limited by its shallow network architecture. Third, the CLIP-based NR DQA method, CLIPQA+, is comparable to HyperIQA, which achieves the second-best performance on DHQ and exBeDDE datasets. This indicates the promising potential of CLIP models for NR DQA. Fourth, our proposed method, *i.e.*, CLIP-DQA, outperforms HyperIQA by a noticeable margin, especially on the DHQ dataset. This confirms the effectiveness of our proposed method.

### C. Ablation Study

**Effectiveness of Textual Prompt Tuning.** We explore two variants of our proposed method, *i.e.*,  $M_1$  and  $M_2$ . Both methods disable visual prompt tuning and conduct the quality evaluation based on local information of dehazed images. However,  $M_1$  uses handcrafted text prompts, while  $M_2$  employs textual prompt tuning. The comparison results of these two methods are presented in Table II. As we can see,  $M_2$  significantly outperforms  $M_1$ . For example, on the DHQ dataset, the SRCC of  $M_2$  is more than double the SRCC of  $M_1$ . Therefore, we can draw the conclusion that textual prompt tuning is effective.




Fig. 2. Illustration of the average attention map for the last visual transformer layer. The first and third rows are resized versions and patches of six dehazed images, respectively. The second and fourth rows visualize attention on resized images and patches, respectively. Each column corresponds to a dehazed image.

**Effectiveness of Visual Prompt Tuning.** We explore another variant of our proposed method, *i.e.*,  $M_3$ . Compared to  $M_2$  with only textual prompt tuning,  $M_3$  turns on both textual and visual prompt tuning. As we can see,  $M_3$  shows clear advantages over  $M_2$  in all three evaluation metrics, especially on the DHQ dataset. As a result, we can conclude that visual prompt tuning is useful.

**Effectives of Hierarchical Information Perception.** Compared to  $M_3$ , our proposed method exploits both local and global information of dehazed images for quality prediction. As shown in Table II, the proposed method achieves better performance than  $M_3$  on both datasets. This confirms the efficacy of hierarchical information perception.

### D. Case Study

To understand how the proposed method perceives dehazed images, we visualize the average attention map of the last visual transformer layer. In the average attention map, spatial regions with brighter colors contribute more to quality predictions. According to Fig. 2, we have the following findings:

First, for the resized version of dehazed images, the proposed method tends to spread attention to the whole image. Second, for patches cropped from dehazed images, the proposed method mainly focuses on hazy regions. In summary, the proposed method employs global and local information on dehazed images for quality prediction. This is highly consistent with the human visual system.

## IV. CONCLUSION

In this paper, we present the first preliminary study on introducing CLIP to BDQA, dubbed CLIP-DQA. Motivated by characteristics of the human vision system, CLIP-DQA blindly evaluates dehazed images from global and local perspectives. Furthermore, CLIP-DQA employs learnable multi-modal prompts to tune CLIP for accurate quality prediction. Comprehensive experiments demonstrate that the proposed method achieves state-of-the-art performance.

## REFERENCES

- [1] Z. Chen, C. Yang, Q. Li, F. Zhao, Z.-J. Zha, and F. Wu, "Disentangle your dense object detector," in *Proceedings of the 29th ACM international conference on multimedia*, 2021, pp. 4939–4948.
- [2] X. Xie, C. Lang, S. Miao, G. Cheng, K. Li, and J. Han, "Mutual-assistance learning for object detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [3] S. Wang, Z. Wang, H. Li, and W. Ouyang, "Category-specific semantic coherency learning for fine-grained image recognition," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 174–183.
- [4] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 12, pp. 2341–2353, 2010.
- [5] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao, "Dehazenet: An end-to-end system for single image haze removal," *IEEE Transactions on Image Processing*, vol. 25, no. 11, pp. 5187–5198, 2016.
- [6] X. Liu, H. Li, and C. Zhu, "Joint contrast enhancement and exposure fusion for real-world image dehazing," *IEEE Transactions on Multimedia*, vol. 24, pp. 3934–3946, 2021.
- [7] Y. Li, H. Chen, Q. Miao, D. Ge, S. Liang, Z. Ma, and B. Zhao, "Image hazing and dehazing: From the viewpoint of two-way image translation with a weakly supervised framework," *IEEE Transactions on Multimedia*, 2022.
- [8] Y. Song, Z. He, H. Qian, and X. Du, "Vision transformers for single image dehazing," *IEEE Transactions on Image Processing*, vol. 32, pp. 1927–1941, 2023.
- [9] D. Cheng, Y. Li, D. Zhang, N. Wang, J. Sun, and X. Gao, "Progressive negative enhancing contrastive learning for image dehazing and beyond," *IEEE Transactions on Multimedia*, 2024.
- [10] R. Wang, W. Li, X. Liu, C. Li, Z. Zhang, X. Min, and G. Zhai, "Hazeclip: Towards language guided real-world image dehazing," *arXiv preprint arXiv:2407.13719*, 2024.
- [11] W. Zhou, R. Zhang, L. Li, G. Yue, J. Gong, H. Chen, and H. Liu, "Dehazed image quality evaluation: from partial discrepancy to blind perception," *IEEE Transactions on Intelligent Vehicles*, 2024.
- [12] X. Min, G. Zhai, K. Gu, Y. Zhu, J. Zhou, G. Guo, X. Yang, X. Guan, and W. Zhang, "Quality evaluation of image dehazing methods using synthetic hazy images," *IEEE Transactions on Multimedia*, vol. 21, no. 9, pp. 2319–2333, 2019.
- [13] S. Zhao, L. Zhang, S. Huang, Y. Shen, and S. Zhao, "Dehazing evaluation: Real-world benchmark datasets, criteria, and baselines," *IEEE Transactions on Image Processing*, vol. 29, pp. 6947–6962, 2020.
- [14] W. Liu, F. Zhou, T. Lu, J. Duan, and G. Qiu, "Image defogging quality assessment: Real-world database and method," *IEEE Transactions on Image Processing*, vol. 30, pp. 176–190, 2020.
- [15] S. Fang, J. Yang, J. Zhan, H. Yuan, and R. Rao, "Image quality assessment on image haze removal," in *Chinese Control and Decision Conference*. IEEE, 2011, pp. 610–614.
- [16] Y. Song, H. Luo, R. Lu, and J. Ma, "Dehazed image quality assessment by haze-line theory," in *Journal of Physics: Conference Series*, vol. 844, no. 1. IOP Publishing, 2017, p. 012045.
- [17] X. Min, G. Zhai, K. Gu, X. Yang, and X. Guan, "Objective quality evaluation of dehazed images," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 8, pp. 2879–2892, 2018.
- [18] C. Wang, H. Fan, H. Zhang, and Z. Li, "Pixel-level dehazed image quality assessment based on dark channel prior and depth," in *IEEE International Conference on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCloud/SocialCom/SustainCom)*. IEEE, 2019, pp. 1545–1549.
- [19] T. Zeng, L. Zhang, W. Zou, X. Li, and S. Tian, "Trg-dqa: Texture residual-guided dehazed image quality assessment," in *IEEE International Conference on Image Processing*. IEEE, 2023, pp. 3075–3079.
- [20] T. Guan, C. Li, Y. Zheng, X. Wu, and A. C. Bovik, "Dual-stream complex-valued convolutional network for authentic dehazed image quality assessment," *IEEE Transactions on Image Processing*, 2023.
- [21] L. K. Choi, J. You, and A. C. Bovik, "Referenceless prediction of perceptual fog density and perceptual image defogging," *IEEE Transactions on Image Processing*, vol. 24, no. 11, pp. 3888–3901, 2015.
- [22] W. Shen, S. Hao, J. Qian, and L. Li, "Blind quality assessment of dehazed images by analyzing information, contrast, and luminance," *J. Neww. Intell.*, vol. 2, no. 1, pp. 139–146, 2017.
- [23] T. Guan, C. Li, K. Gu, H. Liu, Y. Zheng, and X.-J. Wu, "Visibility and distortion measurement for no-reference dehazed image quality assessment via complex contourlet transform," *IEEE Transactions on Multimedia*, 2022.
- [24] X. Lv, T. Xiang, Y. Yang, and H. Liu, "Blind dehazed image quality assessment: a deep cnn-based approach," *IEEE Transactions on Multimedia*, 2023.
- [25] J. Zhang, X. Min, Y. Zhu, G. Zhai, J. Zhou, X. Yang, and W. Zhang, "Hazdesnet: an end-to-end network for haze density prediction," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 4, pp. 3087–3102, 2020.
- [26] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [27] X. Li, D. Lian, Z. Lu, J. Bai, Z. Chen, and X. Wang, "Graphadapter: Tuning vision-language models with dual knowledge graph," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [28] R. Zhang, Z. Guo, W. Zhang, K. Li, X. Miao, B. Cui, Y. Qiao, P. Gao, and H. Li, "Pointclip: Point cloud understanding by clip," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8552–8562.
- [29] J. Wang, K. C. Chan, and C. C. Loy, "Exploring clip for assessing the look and feel of images," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 2, 2023, pp. 2555–2563.
- [30] K. Zhou, J. Yang, C. C. Loy, and Z. Liu, "Conditional prompt learning for vision-language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16816–16825.
- [31] M. Jia, L. Tang, B.-C. Chen, C. Cardie, S. Belongie, B. Hariharan, and S.-N. Lim, "Visual prompt tuning," in *European Conference on Computer Vision*. Springer, 2022, pp. 709–727.
- [32] K. Zhou, J. Yang, C. C. Loy, and Z. Liu, "Learning to prompt for vision-language models," *International Journal of Computer Vision*, vol. 130, no. 9, pp. 2337–2348, 2022.
- [33] M. U. Khattak, H. Rasheed, M. Maaz, S. Khan, and F. S. Khan, "Maple: Multi-modal prompt learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 19113–19122.
- [34] J. Antkowiak, T. J. Baina, F. V. Baroncini, N. Chateau, F. FranceT-elecom, A. C. F. Pessoa, F. S. Colonnese, I. L. Contin, J. Caviedes, and F. Philips, "Final report from the video quality experts group on the validation of objective models of video quality assessment march 2000," *Final report from the video quality experts group on the validation of objective models of video quality assessment march*, vol. 2000, 2000.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [36] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [37] L. Zhang, Y. Shen, and H. Li, "Vsi: A visual saliency-induced index for perceptual image quality assessment," *IEEE Transactions on Image Processing*, vol. 23, no. 10, pp. 4270–4281, 2014.
- [38] L. Kang, P. Ye, Y. Li, and D. Doermann, "Convolutional neural networks for no-reference image quality assessment," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1733–1740.
- [39] S. A. Golestaneh, S. Dadsetan, and K. M. Kitani, "No-reference image quality assessment via transformers, relative ranking, and self-consistency," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 1220–1230.
- [40] S. Su, Q. Yan, Y. Zhu, C. Zhang, X. Ge, J. Sun, and Y. Zhang, "Blindly assess image quality in the wild guided by a self-adaptive hyper network," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3667–3676.