

Adaptive Language-Aware Image Reflection Removal Network

Siyan Fang¹, Yuntao Wang¹, Jinpu Zhang², Ziwen Li¹ and Yuehuan Wang^{1*}

¹Huazhong University of Science and Technology

²National University of Defense Technology

fangsiyanfsy@163.com, yuntaowang@hust.edu.cn, zhangjinpu@nudt.edu.cn, {D201980722, yuehwang}@hust.edu.cn

Abstract

Existing image reflection removal methods struggle to handle complex reflections. Accurate language descriptions can help the model understand the image content to remove complex reflections. However, due to blurred and distorted interferences in reflected images, machine-generated language descriptions of the image content are often inaccurate, which harms the performance of language-guided reflection removal. To address this, we propose the Adaptive Language-Aware Network (ALANet) to remove reflections even with inaccurate language inputs. Specifically, ALANet integrates both filtering and optimization strategies. The filtering strategy reduces the negative effects of language while preserving its benefits, whereas the optimization strategy enhances the alignment between language and visual features. ALANet also utilizes language cues to decouple specific layer content from feature maps, improving its ability to handle complex reflections. To evaluate the model’s performance under complex reflections and varying levels of language accuracy, we introduce the Complex Reflection and Language Accuracy Variance (CRLAV) dataset. Experimental results demonstrate that ALANet surpasses state-of-the-art methods for image reflection removal. The code and dataset are available at <https://github.com/fashion/ALANet>.

1 Introduction

When capturing images through glass, reflections diminish image quality by obscuring details and distorting colors, thereby impairing the image’s usability and hindering downstream computer vision tasks [Xie *et al.*, 2021; Zhang *et al.*, 2023]. The widely used reflection model [Wan *et al.*, 2018; Yang *et al.*, 2018] considers the reflected image I as a linear combination of the transmission layer T and the reflection layer R , where T represents the content transmitted through the glass, and R represents the reflected content.

To obtain a clear T from I , multi-image reflection removal methods [Li and Brown, 2013; Liu *et al.*, 2020] leverage

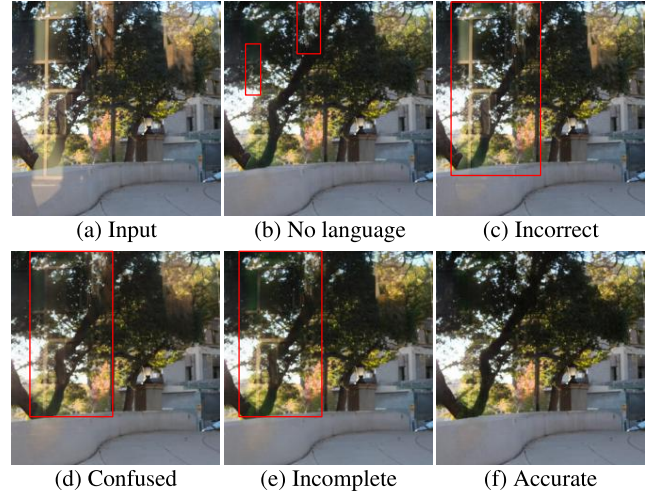


Figure 1: The impact of language-guided reflection removal with different types of language inputs. Inaccurate language inputs result in worse outcomes than having no language. The specific language inputs for each subfigure are provided in the supplementary material.

the distinct motions between T and R from different viewpoints to distinguish them. However, these methods rely on controlled environments and multi-angle observations, limiting their applicability in real-world scenarios. Earlier single-image reflection removal methods [Levin *et al.*, 2004; Chung *et al.*, 2009; Shih *et al.*, 2015] primarily rely on handcrafted priors, but are only suitable for simple reflection scenarios. With the rise of deep learning, many methods [Zhang *et al.*, 2018; Wei *et al.*, 2019; Li *et al.*, 2020; Hu and Guo, 2023] use neural networks to model the different layers in reflected images. However, these methods struggle to remove complex reflections due to the limited information available in a single image.

Language has shown great effectiveness in various visual tasks [Radford *et al.*, 2021; Li *et al.*, 2022; Wang *et al.*, 2022]. Language descriptions can provide additional information about objects within a scene, enhancing the understanding and processing capabilities of networks. For example, with complex reflections in an image, language descriptions can indicate which areas belong to the reflection layer and which belong to the transmission layer. Zhong *et al.*

*Corresponding author.

[Zhong *et al.*, 2024] introduced language into the image reflection removal field to provide contextual clues, which facilitate the separation of different layers. However, this method requires that the language descriptions accurately match the image content; mismatches can harm performance, as illustrated in Figure 1. The L-DIFFER [Hong *et al.*, 2025] employed a language-based diffusion model to remove reflections, but does not account for cases where the language descriptions are inaccurate. Since manually annotating images is time-consuming and labor-intensive, the most convenient approach is to use language models to automatically generate descriptions. However, reflections interfere with the language model’s ability to understand the image content. For example, on the reflection removal datasets Nature, Real, and SIR², using BLIP [Li *et al.*, 2022] to generate captions for reflected images versus reflection-free images, the average BLEU, METEOR, CIDEr, and ROUGE-L scores drop by 50.95%, 34.11%, 55.07%, and 31.45%, respectively, indicating that reflections mislead the language model into generating inaccurate descriptions. These inaccuracies fall into three categories: (1) Incorrect: Describing non-existent content in the image, misleading the model to process fictitious elements and adversely affecting its handling of actual content. (2) Confused: Mixing up of parts of the transmission and reflection layers, preventing the model from distinguishing the features of each layer. (3) Incomplete: Omitting details, especially those obscured by reflections, which prevents the model from focusing on the critical areas that need processing.

To mitigate the negative impact of inaccurate language descriptions on the performance of image reflection removal, we propose the Adaptive Language-Aware Network (ALANet) through two well-designed strategies: filtering and optimization. The filtering strategy aims to filter out the negative effects of inaccurate language while retaining its positive effects. For this, we propose the Language-Aware Competition Attention Module (LCAM), which enables language-guided attention and visually-driven attention to compete with each other, dynamically adjusting their influence. The optimization strategy seeks to refine language features so they can align with the content of the corresponding layers. For this, we propose the Adaptive Language Calibration Module (ALCM), which uses visual features to fine-tune language features. Furthermore, to effectively utilize language features for separating specific information from images, we design the Language-Guided Spatial-Channel Cross Attention (LSCA) mechanism. This mechanism adjusts the spatial and channel structures of the feature map by language, accurately extracting different layers from intertwined scenes.

To evaluate model performance under complex reflections and varying levels of language accuracy, we introduce the real-world Complex Reflection and Language Accuracy Variance (CRLAV) dataset. This dataset includes multiple types of complex reflections, with each image paired with language descriptions of varying accuracies. It not only allows for assessing a model’s ability to remove complex reflections, but it also evaluates the model’s robustness under conditions of varying language accuracy guidance.

In summary, the main contributions are as follows:

- We propose ALANet to improve the model’s reflection removal performance with inaccurate language descriptions through filtering and optimization strategies.
- We introduce the real-world CRLAV dataset to evaluate the performance of models under complex reflections and varying levels of language accuracy.
- Experiments demonstrate that the proposed ALANet surpasses state-of-the-art (SOTA) methods and achieves solid performance even with inaccurate language inputs.

2 Related Work

2.1 Image Reflection Removal

The existing methods for reflection removal can be divided into multi-image-based and single-image-based methods. While multi-image-based methods leverage information from multiple images to remove reflections, their applicability is limited by data acquisition constraints. This paper focuses on the more broadly applicable single-image reflection removal. Early single-image reflection removal methods primarily relied on priors such as gradient sparsity [Levin *et al.*, 2004; Levin and Weiss, 2007], layer smoothness [Chung *et al.*, 2009; Yan *et al.*, 2014], and ghosting cues [Shih *et al.*, 2015]. However, these methods often perform poorly in real-world scenes due to their heavy dependence on prior assumptions.

With the rise of deep learning, CEILNet [Fan *et al.*, 2017] was the first to apply deep learning to the task of single-image reflection removal, using edge features as auxiliary information to eliminate reflections. Zhu *et al.* [Zhu *et al.*, 2024] designed the max-min reflection filter to characterize the reflection locations in paired images. However, these methods often face bottlenecks in complex scenes due to the lack of additional information. Given that language conveys human prior knowledge about the world [Deng *et al.*, 2023], Zhong *et al.* [Zhong *et al.*, 2024] facilitated layer separation by establishing correspondences between language descriptions and image layers. However, when the language descriptions do not accurately match the corresponding layer content, they can have negative effects. L-DIFFER [Hong *et al.*, 2025] used language as a condition for the diffusion model to remove reflections, but it did not consider the harm caused by inaccurate language descriptions. To address this issue, we propose the ALANet to reduce the negative impact of inaccurate language descriptions on reflection removal.

2.2 Applications of Language in Image Processing

The ability of CLIP [Radford *et al.*, 2021] to understand language and interpret images has been leveraged in various tasks. BLIP [Li *et al.*, 2022] proposed using bootstrapping to synthesize captions, obtaining higher quality data. CLIP-LIT [Liang *et al.*, 2023] learned initial prompt pairs by constraining the text-image similarity between prompts and corresponding images in the CLIP latent space. NeRCO [Yang *et al.*, 2023] employed the priors of a pre-trained vision-language model to provide perceptually guided instruction for learning lighting conditions. DA-CLIP [Luo *et al.*, 2023] predicted the degraded embeddings of low-quality images

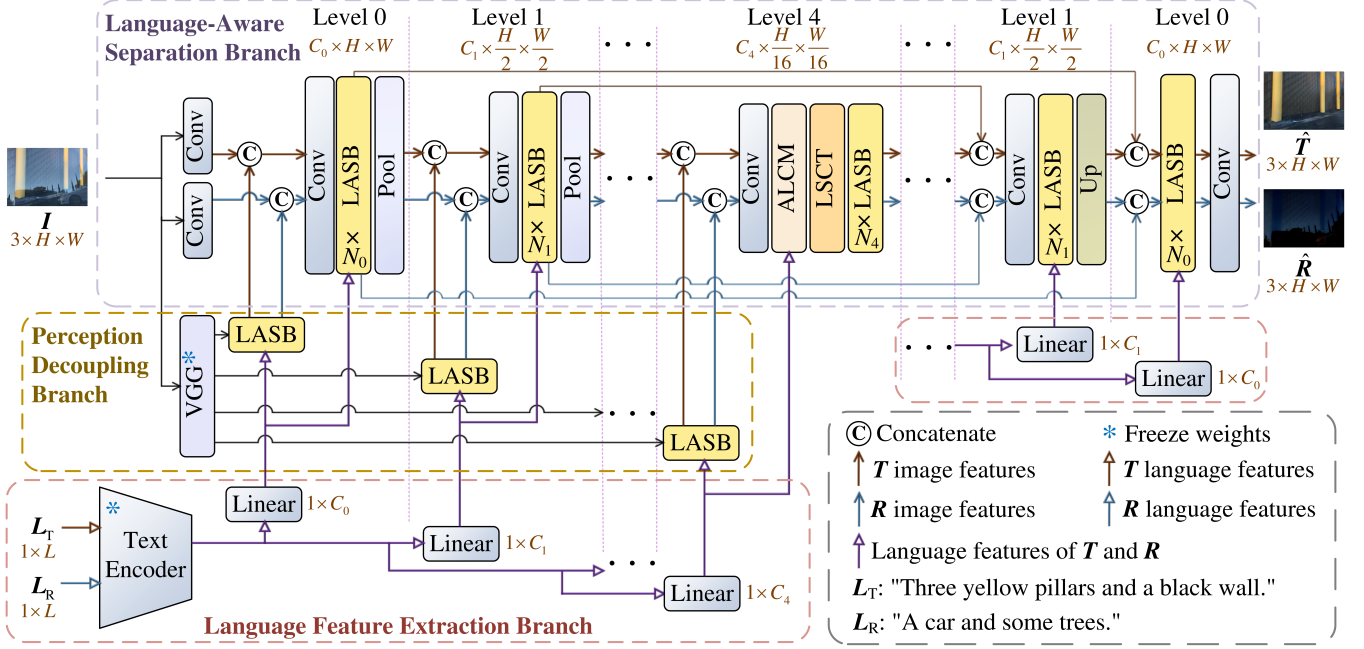


Figure 2: Overview of the proposed ALANet, which comprises various modules that use language adaptively to remove reflections. T and R represent the transmission and reflection layers, respectively.

through a controller, and directed the CLIP image encoder to output high-quality content embeddings. Zhong et al. [Zhong et al., 2024] introduced language descriptions to provide layer content for addressing the ill-posed problem of reflection separation. However, since images with reflections are difficult to be accurately described by language models, how to separate specific image layers under imperfect language conditions remains an unresolved challenge.

3 Proposed Method

3.1 Adaptive Language-Aware Network

The Adaptive Language-Aware Network (ALANet) consists of the Language-Aware Separation Branch (LSBranch), the Perception Decoupling Branch (PDBranch), and the Language Feature Extraction Branch (LEBranch), as shown in Figure 2. The LEBranch is responsible for encoding the input language, and adjusting the channel dimensions to fit the network structure. The PDBranch uses a pre-trained Visual Geometry Group (VGG) model [Simonyan and Zisserman, 2014] to extract high-level visual features. It then decouples the corresponding features through language to facilitate the separation process in the LSBranch. In the LSBranch, the Language-Aware Separation Block (LASB) adjusts the influence of language-guided attention according to the accuracy of the language description, thereby preventing issues of misguidance caused by inaccurate language.

The structure of the LASB is shown in Figure 3. Its purpose is to use semantic information from language that matches with the visual content, guiding the separation of the transmission and reflection layers. The LASB primarily employs the LCAM and the Multi-Receptive Field Decoupling




Figure 3: Structure of the LASB. As the core of LASB, LCAM utilizes language features from different layers to facilitate the separation of those layers.

Module (MFDM) to separate the different layers. Details of MFDM are provided in the supplementary material.

3.2 Language-Aware Competition Attention Module

The Language-Aware Competition Attention Module (LCAM) is shown in Figure 4. It competitively adjusts the weight distribution of language-guided attention and channel attention, based on the matching degree between language cues and the visual content, aiming to retain the positive effects of accurate language while suppressing the negative effects of inaccurate language. Given the image features $F_I \in \mathbb{R}^{C \times H \times W}$ and language features $F_L \in \mathbb{R}^{1 \times C}$, where H and W are the height and width of the image, and C is the number of channels. F_I and F_L both undergo processing and through matrix multiplication, the language-image similarity matrix $M_L \in \mathbb{R}^{C \times C}$ is obtained. Each element (i, j) in M_L represents the similarity between the i -th channel of language features and the j -th channel of image features.

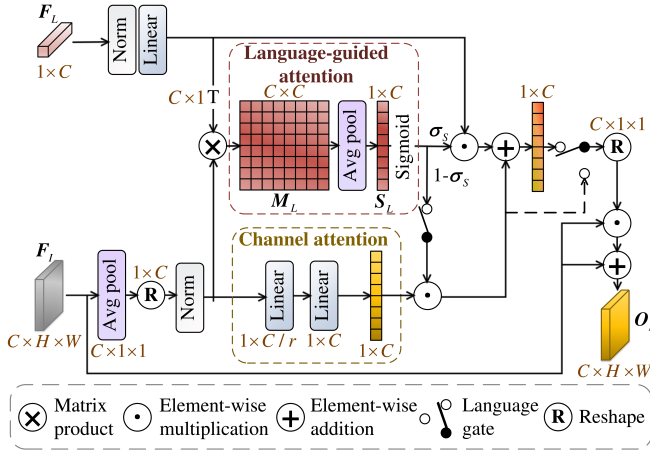


Figure 4: Structure of the LCAM. The dashed lines indicate the scenario without language input.

Subsequently, M_L is average-pooled along the language dimension to obtain the language-image similarity scores $S_L \in \mathbb{R}^{1 \times C}$, where each element indicates the similarity of each image channel to the overall language information. After passing through a sigmoid function, the adjustment vector σ_s is used to modify the weight of language-guided attention, and its complement $1 - \sigma_s$ is used to adjust the weight of channel attention. The channel attention differentiates between the two layers by considering their physical characteristics, specifically the structural continuity in the transmission layer and the specular sparsity in the reflection layer, thereby complementing language-guided attention from a visual perspective.

3.3 Adaptive Language Calibration Module

The Adaptive Language Calibration Module (ALCM), as shown in Figure 5, adjusts and optimizes language features by leveraging visual features to enhance the consistency between language and visual content. In the ALCM, image features F_I and language features F_L are processed and then concatenated. Subsequently, they pass through a linear layer and a sigmoid function to generate an adjustment vector $\sigma_c \in \mathbb{R}^{1 \times C}$, which dynamically controls the fusion ratio of language and image features. In this process, the linear layer acts as a mediator for feature fusion, optimizing the combination points between the two types of information.

3.4 Language-Guided Spatial-Channel Cross Transformer

The Language-Guided Spatial-Channel Cross Transformer (LSCT), as shown in Figure 6, features the Language-Guided Spatial-Channel Cross Attention (LSCA) as its core. The LSCA utilizes the semantic information from language to interact with the spatial and channel dimensions of the feature map, leveraging language-driven attention mechanisms to decouple specific information. As illustrated, the LSCA first applies spatial and channel pooling to the input image features F_I , obtaining $F_S \in \mathbb{R}^{1 \times C}$ and $F_C \in \mathbb{R}^{HW \times 1}$, where $HW = H \times W$. The language features F_L interact with F_S

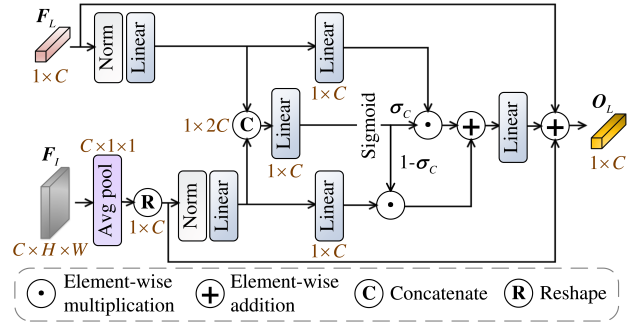


Figure 5: Structure of the ALCM. The ALCM enhances the consistency between language features and visual content.

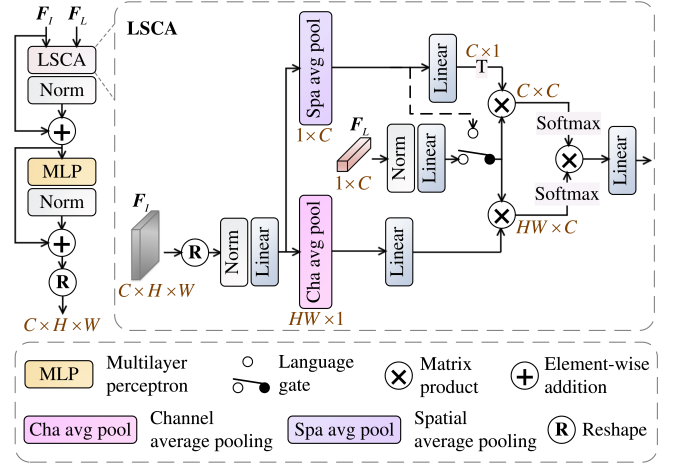


Figure 6: Structure of the LSCT and its core component, LSCA. The dashed lines in LSCA indicate the scenario without language input.

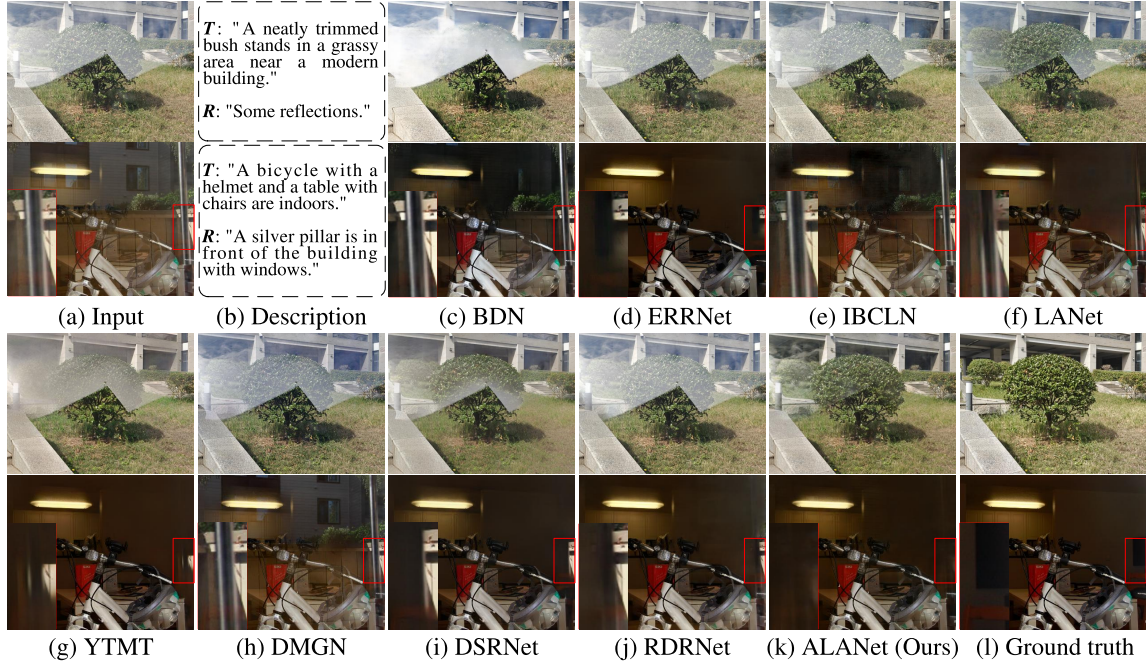
to generate $M_{SL} \in \mathbb{R}^{C \times C}$. Each element (i, j) in M_{SL} represents the correlation between the i -th channel of the image features and the j -th channel of the language features, indicating a global language-image interaction effect. Simultaneously, the language features F_L interact with F_C to produce $M_{LC} \in \mathbb{R}^{HW \times C}$. Each element (h, c) in M_{LC} indicates the correlation between the h -th spatial position of the image and the c -th channel of the language features, reflecting the relevance of the language description to specific local regions in the image. Finally, both M_{LC} and M_{SL} undergo softmax and are multiplied to obtain $M_{LCSL} \in \mathbb{R}^{HW \times C}$. This matrix combines global language guidance with local image features, enhancing the model's ability to interpret language descriptions, and focus on specific regions of the image that correspond to these descriptions.

4 Complex Reflection and Language Accuracy Variance Dataset

To evaluate the model's performance under complex reflections and varying levels of language accuracy, we propose the Complex Reflection and Language Accuracy Variance (CRLAV) dataset. The dataset encompasses real-world indoor and outdoor scenes, comprising a total of 600 image

Methods	Nature (20)		Real (20)		Wild (55)		Postcard (199)		Solid (200)		Average	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
BDN (ECCV'18)	18.83	0.737	18.68	0.728	22.02	0.822	20.54	0.857	22.68	0.856	20.55	0.800
ERRNet (CVPR'19)	20.39	0.766	22.28	0.796	25.14	0.873	21.53	0.877	23.55	0.880	22.58	0.838
IBCLN (CVPR'20)	23.78	0.784	21.59	0.764	24.46	0.885	22.95	0.875	<u>24.74</u>	0.893	23.50	0.840
LANet (ICCV'21)	23.55	0.811	22.51	0.815	<u>26.06</u>	0.900	24.14	0.907	24.30	0.898	24.11	<u>0.866</u>
YTMT (NIPS'21)	24.08	0.814	22.68	0.798	25.24	0.888	21.86	0.880	23.79	0.887	23.53	0.853
DMGN (TIP'21)	20.63	0.764	20.28	0.763	21.34	0.774	22.65	0.879	23.27	0.872	21.63	0.810
DSRNet (ICCV'23)	24.84	0.823	22.09	0.790	26.00	<u>0.902</u>	20.05	0.883	23.96	0.887	23.39	0.857
RDRNet (CVPR'24)	<u>25.33</u>	0.835	<u>22.76</u>	0.804	26.97	0.905	22.11	0.881	24.42	0.892	<u>24.32</u>	0.863
ALANet (Ours)	25.56	<u>0.829</u>	23.89	<u>0.812</u>	25.93	0.900	<u>23.47</u>	<u>0.897</u>	24.85	0.900	24.74	0.868

Table 1: Quantitative comparison with SOTA methods on public datasets. Bold and underline indicate top 1st and 2nd rank, respectively.


 Figure 7: Qualitative comparison with SOTA methods on the CRLAV dataset (top) and the Real dataset (bottom). ALANet excels in identifying and removing reflections in complex environments. *T* and *R* represent the transmission and reflection layers, respectively.

pairs. Using a tripod-mounted smartphone, we capture images with reflection artifacts by placing glass and acrylic sheets of varying thicknesses as obstructions. Ground truth images without reflections are obtained by removing these obstructions. To achieve the three key characteristics of complex reflections—high intensity, large coverage, and indistinguishability—we employ several strategies during data collection. By adjusting the tilt angles of the obstructions, we control the area and intensity of the reflections, allowing them to cover larger regions and increase the overlap between the reflected light and the transmission layer content, thereby creating high-intensity and large-area reflection artifacts. Additionally, we select complex scenes with rich textures or multiple objects to enhance the confusion and complexity of the reflected content, making it more challenging to distinguish from the actual transmission content.

To assess model performance under language conditions of varying accuracy, each image is annotated with both accurate and inaccurate language descriptions. Inaccurate de-

scriptions are categorized into three types: incorrect, confused, and incomplete. Each type is further divided into four levels—slightly, moderately, severely, and entirely inaccurate—corresponding to adjustments of 25%, 50%, 75%, and 100% of the label content, respectively. This setup simulates the impact of language errors on reflection removal models.

5 Experiments

5.1 Implementation Details

To balance performance and parameter count, the channel numbers from level 0 to level 4 of the network are set to $C_0, C_1, C_2, C_3, C_4 = [64, 128, 128, 160, 160]$, and the number of LASBs at each level (N_0 to N_4) is set to 2. The initial learning rate is 10^{-4} , with a batch size of 1, a patch size of 224×224 , and random flipping applied for data augmentation. The model is trained for 70 epochs using the Adam optimizer [Kingma and Ba, 2014] with a single RTX 3080 Ti GPU. The learning rate decreases to 10^{-5} at 50 epochs.

Methods	CRLAV (600)		Param (M)	FLOPs (G)
	PSNR	SSIM		
BDN (ECCV'18)	17.46	0.686	75.16	12.70
ERRNet (CVPR'19)	18.93	0.702	18.95	116.72
IBCLN (CVPR'20)	18.81	0.701	21.61	98.16
LANet (ICCV'21)	19.28	<u>0.709</u>	10.93	83.81
YTMT (NIPS'21)	18.92	0.697	76.90	110.98
DMGN (TIP'21)	18.49	0.698	45.49	116.85
DSRNet (ICCV'23)	18.58	0.693	9.87	32.34
RDRNet (CVPR'24)	<u>19.51</u>	0.706	29.09	5.14
ALANet (Ours)	19.68	0.719	8.69	32.92

Table 2: Quantitative comparison with SOTA methods on the CRLAV dataset, including parameters and FLOPs (computed for a 128×128 RGB image).

5.2 Dataset and Evaluation Metrics

We employ both synthetic and real-world images to train our model. For synthetic images, we generate data using the popular image captioning dataset Flickr8k [Hodosh *et al.*, 2013], which contains 8,091 images, each with five different language descriptions. We randomly select images from the Flickr8k dataset to serve as the transmission and reflection layers. These are combined through linear blending [Zhang *et al.*, 2018] to generate synthetic images. For the real-world training data, following prior works [Zhong *et al.*, 2024; Hu and Guo, 2023; Zhu *et al.*, 2023], we train our model using 200 image pairs from the Nature dataset [Li *et al.*, 2020] and 90 image pairs from the Real dataset [Zhang *et al.*, 2018].

We use the remaining images from the Nature and Real datasets, along with the three subsets—Wild, Postcard, and Solid—from the SIR² dataset [Wan *et al.*, 2017] as public test sets. The CRLAV dataset is also included as a test set. Following prior works [Wei *et al.*, 2019; Hu and Guo, 2021], to prevent memory overload, we resize the images in the Real test set by setting the longer side length to 420 while preserving the original aspect ratio. We use the commonly employed peak signal-to-noise ratio (PSNR) [Huynh-Thu and Ghanbari, 2008] and structural similarity index measure (SSIM) [Wang *et al.*, 2004] as evaluation metrics, which are calculated in the RGB color space. Higher values indicate better performance.

5.3 Comparison Results

To assess the performance of the proposed ALANet, we compared it with eight SOTA methods: BDN [Yang *et al.*, 2018], ERRNet [Wei *et al.*, 2019], IBCLN [Li *et al.*, 2020], LANet [Dong *et al.*, 2021], YTMT [Hu and Guo, 2021], DMGN [Feng *et al.*, 2021], DSRNet [Hu and Guo, 2023], and RDRNet [Zhu *et al.*, 2024]. To ensure a fair comparison, for methods with publicly available training code, we fine-tuned their models on our training datasets and selected the version that performed best. Quantitative comparison results across public datasets are shown in Table 1. It can be observed that our ALANet achieves the best or second-best results across multiple datasets, ultimately delivering the best average performance. This demonstrates ALANet’s advantage in image reflection removal, confirming its effectiveness and reliability in various reflection scenarios.

To compare the capability of removing complex reflections, Table 2 presents a quantitative comparison between

Language type for training			Language type for testing			Average	
Correct	Random	None	Correct	Random	None	PSNR	SSIM
✓	×	×	✓	×	×	24.74	0.868
✓	×	×	×	✓	×	24.09	0.861
✓	×	×	×	×	✓	24.02	0.856
×	✓	×	✓	×	×	24.27	0.864
×	✓	×	×	✓	×	24.11	0.860
×	✓	×	×	×	✓	23.83	0.852
×	×	✓	×	×	✓	24.02	0.857

Table 3: Ablation experiments for different types of language inputs during training and testing on public datasets.




Figure 8: Visual effects of our ALANet with different types of language inputs. The specific language inputs for each subfigure are provided in the supplementary material.

ALANet and other methods on the CRLAV dataset, where ALANet achieves the best performance. Table 2 also shows the parameter count and floating point operations (FLOPs), demonstrating that ALANet maintains a relatively balanced parameter count and FLOPs compared to SOTA methods.

Figure 7 presents a qualitative comparison of ALANet with SOTA methods across various scenes. In the complex outdoor scenario from the CRLAV dataset, only ALANet achieves the most thorough removal of reflections. In the complex indoor scene from the Real dataset, although ERRNet, DSRNet, and RDRNet can remove reflections from the silver column, they fail to eliminate reflections from the light. Only ALANet successfully removes reflections from both the light and the silver column. This demonstrates that ALANet is more effective at removing complex reflections compared to other methods.

5.4 Ablation Studies

In this section, we delve into the impact of different components of ALANet by conducting various ablation studies.

Impact of language input. To explore the impact of language and its accuracy on model performance, we conducted experiments using different types of language inputs during the training and testing phases, with the results shown in Table 3. It can be observed that when correct language is used for training, the performance progressively decreases when tested with correct, random, and no language inputs, indicating that correct language input is most beneficial for enhancing model performance. Simultaneously, even with the lower accuracy caused by random language inputs, benefiting from our filtering and optimization strategies, the performance with random language remains better than with no lan-

	Degree	Incorrect		Confused		Incomplete	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Public datasets (494)	Slightly (25%)	24.55	0.865	24.48	0.866	24.36	0.864
	Moderately (50%)	24.35	0.862	24.29	0.861	24.22	0.861
	Severely (75%)	24.08	0.859	24.01	0.859	24.03	0.857
	Entirely (100%)	22.70	0.838	24.01	0.854	24.02	0.856
CRLAV (600)	Slightly (25%)	19.43	0.713	19.65	0.714	19.52	0.718
	Moderately (50%)	19.25	0.707	19.59	0.712	19.33	0.716
	Severely (75%)	19.23	0.705	19.51	0.711	19.24	0.714
	Entirely (100%)	19.06	0.703	19.41	0.708	19.13	0.697

Table 4: Ablation experiments on the effects of varying degrees of language inaccuracy.

ALCMLSCT	LCAM		Average		Param FLOPs	
	Language-guided attention	Channel attention	PSNR	SSIM	(M)	(G)
✓	✓	✓	24.74	0.868	8.69	32.92
✓	✓	×	24.52	0.865	8.46	32.91
✓	✓	×	24.05	0.859	8.24	32.92
✓	✓	×	24.11	0.858	8.01	32.90
×	✓	✓	24.41	0.867	8.51	32.92
×	×	✓	24.19	0.863	8.20	32.89

Table 5: Ablation experiments for different modules on public datasets. FLOPs for a 128x128 RGB image.

guage input. When training with random language, testing with both correct and random language still achieved commendable performance, surpassing that of RDRNet and DSRNet respectively on the SSIM metric. This demonstrates that random language during training did not mislead the model into confusion, showing that the model can still perform well even under conditions of low language accuracy.

Figure 8 illustrates the visual effects of our ALANet in removing complex reflections with three types of inaccurate descriptions: incorrect, confused, and incomplete. It can be seen that ALANet can utilize the portions of inaccurate language inputs that match the semantics of the corresponding layers, resulting in de-reflection outcomes that are superior to those without any language inputs.

Table 4 illustrates the performance of our ALANet under varying degrees of language accuracy. It can be observed that the model’s performance decreases as the degree of accuracy declines on both public datasets and the CRLAV dataset. Notably, even with severely incorrect, confused, or incomplete input, the model’s performance remains superior to that with no language input. This demonstrates ALANet’s robustness to severely inaccurate language.

Ablation study on LCAM. Table 5 illustrates the contributions of language-guided attention and channel attention to performance within LCAM. It can be observed that the experiments combining both types of attention achieved the best performance, demonstrating the effectiveness of LCAM’s competitive attention mechanism. Figure 9 further shows the feature maps within LCAM and the corresponding language-guided attention weights. It can be observed that the closer the feature map content is to the language description, the higher the language-guided attention weight becomes. This helps to emphasize the regions in the image related to the language description, thereby decoupling the described object.

Ablation study on ALCM. Table 5 showcases the perfor-




Figure 9: Feature maps in LCAM and corresponding language-guided attention weights, with the language description: “A star-shaped toy on the tabletop.”




Figure 10: Comparison of Pre-ALCM and Post-ALCM similarities between language and corresponding image features across different datasets.

mance of models with and without ALCM. It is clear that the models equipped with ALCM outperform those without, demonstrating that ALCM enhances the role of language in image reflection removal. To further validate that ALCM can improve the alignment between language and image features, Figure 10 displays the cosine similarities between language features and image features before and after processing through ALCM. The similarities post-ALCM are consistently higher across various datasets, indicating an improved alignment between language and image features following the application of ALCM.

Ablation study on LSCT. To demonstrate the effectiveness of LSCT, we conducted experiments by further removing LSCT on top of removing ALCM, with the results shown in Table 5. It can be observed that the performance without LSCT is lower than that with LSCT. This indicates that LSCT plays a positive role in enhancing reflection removal performance, especially when paired with ALCM, which amplifies the effects of LSCT.

6 Conclusion

In this paper, we propose the ALANet to remove complex reflections with low dependence on language accuracy. Specifically, ALANet employs filtering and optimization strategies to mitigate the effects of inaccurate language and enhance the alignment between language and visual features, while leveraging language cues to decouple layer content from feature maps. Additionally, we introduce the CRLAV dataset to evaluate the model’s performance under complex reflections and varying levels of language accuracy. Experimental results demonstrate the effectiveness of the ALANet and its superiority over SOTA methods.

References

- [Chung *et al.*, 2009] Yun-Chung Chung, Shyang-Lih Chang, Jung-Ming Wang, and Sei-Wang Chen. Interference reflection separation from a single image. In *2009 Workshop on Applications of Computer Vision (WACV)*, pages 1–6. IEEE, 2009.
- [Deng *et al.*, 2023] Congyue Deng, Chiyu Jiang, Charles R Qi, Xinchun Yan, Yin Zhou, Leonidas Guibas, Dragomir Anguelov, et al. Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20637–20647, 2023.
- [Dong *et al.*, 2021] Zheng Dong, Ke Xu, Yin Yang, Hujun Bao, Weiwei Xu, and Rynson WH Lau. Location-aware single image reflection removal. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5017–5026, 2021.
- [Fan *et al.*, 2017] Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, and David Wipf. A generic deep architecture for single image reflection removal and image smoothing. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3238–3247, 2017.
- [Feng *et al.*, 2021] Xin Feng, Wenjie Pei, Zihui Jia, Fanglin Chen, David Zhang, and Guangming Lu. Deep-masking generative network: A unified framework for background restoration from superimposed images. *IEEE Transactions on Image Processing*, 30:4867–4882, 2021.
- [Hodosh *et al.*, 2013] Micah Hodosh, Peter Young, and Julia Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899, 2013.
- [Hong *et al.*, 2025] Yuchen Hong, Haofeng Zhong, Shuchen Weng, Jinxiu Liang, and Boxin Shi. L-differ: Single image reflection removal with language-based diffusion model. In *European Conference on Computer Vision*, pages 58–76. Springer, 2025.
- [Hu and Guo, 2021] Qiming Hu and Xiaojie Guo. Trash or treasure? an interactive dual-stream strategy for single image reflection separation. *Advances in Neural Information Processing Systems*, 34:24683–24694, 2021.
- [Hu and Guo, 2023] Qiming Hu and Xiaojie Guo. Single image reflection separation via component synergy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13138–13147, 2023.
- [Huynh-Thu and Ghanbari, 2008] Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of psnr in image/video quality assessment. *Electronics letters*, 44(13):800–801, 2008.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Levin and Weiss, 2007] Anat Levin and Yair Weiss. User assisted separation of reflections from a single image using a sparsity prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(9):1647–1654, 2007.
- [Levin *et al.*, 2004] Anat Levin, Assaf Zomet, and Yair Weiss. Separating reflections from a single image using local features. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004.
- [Li and Brown, 2013] Yu Li and Michael S Brown. Exploiting reflection change for automatic reflection removal. In *Proceedings of the IEEE international conference on computer vision*, pages 2432–2439, 2013.
- [Li *et al.*, 2020] Chao Li, Yixiao Yang, Kun He, Stephen Lin, and John E Hopcroft. Single image reflection removal through cascaded refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3565–3574, 2020.
- [Li *et al.*, 2022] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022.
- [Liang *et al.*, 2023] Zhixin 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.
- [Liu *et al.*, 2020] Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Learning to see through obstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14215–14224, 2020.
- [Luo *et al.*, 2023] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models for universal image restoration. *arXiv preprint arXiv:2310.01018*, 2023.
- [Radford *et al.*, 2021] 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.
- [Shih *et al.*, 2015] YiChang Shih, Dilip Krishnan, Fredo Durand, and William T Freeman. Reflection removal using ghosting cues. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3193–3201, 2015.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Wan *et al.*, 2017] Renjie Wan, Boxin Shi, Ling-Yu Duan, Ah-Hwee Tan, and Alex C Kot. Benchmarking single-image reflection removal algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3922–3930, 2017.

- [Wan *et al.*, 2018] Renjie Wan, Boxin Shi, Ling-Yu Duan, Ah-Hwee Tan, and Alex C Kot. Crnr: Multi-scale guided concurrent reflection removal network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4777–4785, 2018.
- [Wang *et al.*, 2004] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [Wang *et al.*, 2022] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11686–11695, 2022.
- [Wei *et al.*, 2019] Kaixuan Wei, Jiaolong Yang, Ying Fu, David Wipf, and Hua Huang. Single image reflection removal exploiting misaligned training data and network enhancements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8178–8187, 2019.
- [Xie *et al.*, 2021] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34:12077–12090, 2021.
- [Yan *et al.*, 2014] Qing Yan, Yi Xu, Xiaokang Yang, and Truong Nguyen. Separation of weak reflection from a single superimposed image. *IEEE Signal Processing Letters*, 21(10):1173–1176, 2014.
- [Yang *et al.*, 2018] Jie Yang, Dong Gong, Lingqiao Liu, and Qinfeng Shi. Seeing deeply and bidirectionally: A deep learning approach for single image reflection removal. In *Proceedings of the european conference on computer vision (ECCV)*, pages 654–669, 2018.
- [Yang *et al.*, 2023] Shuzhou Yang, Moxuan Ding, Yanmin Wu, Zihan Li, and Jian Zhang. Implicit neural representation for cooperative low-light image enhancement. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12918–12927, 2023.
- [Zhang *et al.*, 2018] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4786–4794, 2018.
- [Zhang *et al.*, 2023] Jinpu Zhang, Ziwen Li, Ruonan Wei, and Yuehuan Wang. Progressive domain-style translation for nighttime tracking. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7324–7334, 2023.
- [Zhong *et al.*, 2024] Haofeng Zhong, Yuchen Hong, Shuchen Weng, Jinxiu Liang, and Boxin Shi. Language-guided image reflection separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24913–24922, 2024.
- [Zhu *et al.*, 2023] Yurui Zhu, Xueyang Fu, Zheyu Zhang, Aiping Liu, Zhiwei Xiong, and Zheng-Jun Zha. Hue guidance network for single image reflection removal. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [Zhu *et al.*, 2024] Yurui Zhu, Xueyang Fu, Peng-Tao Jiang, Hao Zhang, Qibin Sun, Jinwei Chen, Zheng-Jun Zha, and Bo Li. Revisiting single image reflection removal in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25468–25478, 2024.