

# Progressive Image Restoration via Text-Conditioned Video Generation

Peng Kang<sup>\*1</sup>, Xijun Wang<sup>†1</sup>, Yu Yuan<sup>†</sup>

<sup>\*</sup>Department of Computer Science, University of Illinois Springfield, Springfield, IL, USA

<sup>†</sup>School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA

**Abstract**—Recent text-to-video models have demonstrated strong temporal generation capabilities, yet their potential for image restoration remains underexplored. In this work, we repurpose CogVideo for progressive visual restoration tasks by fine-tuning it to generate restoration trajectories rather than natural video motion. Specifically, we construct synthetic datasets for super-resolution, deblurring, and low-light enhancement, where each sample depicts a gradual transition from degraded to clean frames. Two prompting strategies are compared: a uniform text prompt shared across all samples, and a scene-specific prompting scheme generated via LLaVA multi-modal LLM and refined with ChatGPT. Our fine-tuned model learns to associate temporal progression with restoration quality, producing sequences that improve perceptual metrics such as PSNR, SSIM, and LPIPS across frames. Extensive experiments show that CogVideo effectively restores spatial detail and illumination consistency while maintaining temporal coherence. Moreover, the model generalizes to real-world scenarios on the ReLoBlur dataset without additional training, demonstrating strong zero-shot robustness and interpretability through temporal restoration.

**Index Terms**—Image Restoration, Text-to-Video Generation, Temporal Enhancement, CogVideo

## I. INTRODUCTION

Recent advances in text-to-video diffusion models [1], [2], [3], [4], [5], [6], [7] have enabled the generation of visually coherent and temporally consistent video content directly from textual prompts. Among these models, CogVideo [8] has demonstrated remarkable capability in synthesizing high-quality videos that capture realistic motion dynamics and scene details. While such models have primarily been used for creative video generation, their potential for visual restoration and enhancement tasks remains largely unexplored.

In this work, we propose a novel paradigm that leverages CogVideo as a generative bridge for visual enhancement tasks, including image super-resolution, deblurring, and low-light enhancement. Instead of treating these tasks as independent image-to-image translation problems, we view them through the lens of video progression learning: that is, generating a sequence of temporally ordered frames that evolve from degraded inputs (e.g., low-resolution, dark, or blurred) toward their enhanced versions. By fine-tuning CogVideo on specialized datasets that encode these degradations as temporal transitions, the model learns to capture the intrinsic relationships between visual degradation and recovery in a dynamic, context-aware manner.

Our key insight is that temporal consistency and motion priors learned by CogVideo can be exploited to improve

spatial detail recovery and stability in single-frame restoration tasks. Traditional restoration methods [9], [10], [11], [12], [13], including CNN- and transformer-based architectures, often focus on spatial fidelity without modeling progressive transitions. In contrast, our approach harnesses the temporal diffusion process of CogVideo to implicitly model the gradual transformation from degraded to clean states, yielding results that are both perceptually smooth and detail-preserving.

To evaluate this idea, we construct and fine-tune CogVideo on three newly formalized datasets:

- Resolution Progression Dataset – videos simulating transitions from low to high resolution;
- Illumination Progression Dataset – videos progressing from dark to well-lit conditions;
- Blur-to-Sharp Dataset – videos depicting gradual recovery from motion or defocus blur.

Comprehensive experiments show that our fine-tuned CogVideo not only produces visually appealing videos but also serves as a powerful generator for restoration tasks when sampled frame-by-frame. Furthermore, to assess the generalization of our approach, we apply the fine-tuned CogVideo model to the ReLoBlur dataset [14], a challenging real-world video deblurring benchmark. Despite not being explicitly trained on real-world motion blur, our model demonstrates superior visual restoration and temporal stability, producing realistic and coherent enhancement sequences.

In summary, the main contributions of this work are three-fold:

- We propose a novel framework that repurposes CogVideo for image enhancement tasks by modeling restoration as a video progression process.
- We construct and release three benchmark datasets representing progressive visual degradation-to-restoration transitions.
- We demonstrate that fine-tuned CogVideo achieves competitive performance in super-resolution, deblurring, and low-light enhancement, offering a new perspective on leveraging text-to-video diffusion models for image restoration.

## II. METHODOLOGY

Building upon the motivation introduced earlier, our goal is to repurpose CogVideo, a large-scale text-to-video diffusion model, for image restoration and enhancement tasks. Rather than designing a new network architecture from scratch, we exploit the generative prior and temporal reasoning capabilities already embedded in CogVideo. By fine-tuning it

<sup>1</sup>Both authors contributed equally to this work. Peng Kang is the corresponding author. Email: pkang8@uis.edu

on carefully constructed progression datasets, we guide the model to interpret “visual enhancement” as a gradual temporal transformation.

Our method proceeds in three stages. First, we formalize the restoration problem as a temporal video generation task, where a degraded image progressively evolves into its clean counterpart. Second, we construct specialized datasets for super-resolution, deblurring, and low-light correction, providing both uniform and scene-adaptive text conditions for comparison. Finally, we fine-tune CogVideo using a framework that conditions the diffusion process on textual descriptions of the restoration. The overall pipeline is designed to be modular and data-driven, allowing CogVideo to generalize across multiple degradation types.

### A. Problem Definition

Given a degraded image  $I_d$  (e.g., low-resolution, blurred, or low-light) and its corresponding clean target  $I_c$ , conventional restoration methods aim to learn a direct mapping

$$f_\theta : I_d \rightarrow I_c, \quad (1)$$

where  $f_\theta$  is typically a CNN- or Transformer-based model optimized for pixel-level fidelity or perceptual similarity.

In this work, we reinterpret image restoration as a **temporal video generation** problem. Instead of predicting a single restored frame, we model the progressive enhancement process as a temporally ordered sequence

$$\mathcal{V} = \{I_1, I_2, \dots, I_T\}, \quad (2)$$

where  $I_1 = I_d$  is the degraded input,  $I_T = I_c$  is the restored target, and the intermediate frames  $\{I_t\}_{t=2}^{T-1}$  represent gradual transitions from degraded to enhanced states.

Let  $G_\phi$  denote the CogVideo generator parameterized by  $\phi$ . During inference, the model takes as input both a textual prompt  $p$  that semantically describes the desired restoration process and the initial degraded frame  $I_d$ . The generation process can thus be formulated as

$$\mathcal{V} = G_\phi(p, I_d), \quad (3)$$

where  $\mathcal{V}$  is a short video depicting a smooth enhancement trajectory from  $I_d$  to  $I_T$ . The final frame  $I_T$  is extracted as the restored image, while preceding frames visualize interpretable intermediate restoration stages.

This formulation offers two key advantages. First, by leveraging the temporal diffusion dynamics of CogVideo, the model captures the gradual, continuous nature of visual enhancement. Second, the text prompt  $p$  provides an additional conditioning signal that aligns the generation process. Through this joint conditioning, the model learns to couple textual guidance with spatiotemporal visual transformations, enabling a unified framework for super-resolution, deblurring, and low-light enhancement.

### B. Dataset Construction

To enable CogVideo to learn visual restoration as a progressive video generation task, we construct three domain-specific datasets derived from the DIV2K high-resolution image corpus [15]. Each dataset contains short video clips, each composed of  $T = 9$  frames showing a smooth transition from a degraded state to a clean target. For every image in the DIV2K dataset, we generate two distinct videos using random degradation parameters to increase diversity. All generated videos are stored at  $1360 \times 768$  resolution and 5 fps.

**1) Resolution Progression Dataset.** We simulate low-to-high resolution transitions by progressively upscaling bicubic-downsampled images. Each video is generated by resizing the image with scale factors  $\{s_1, s_2, \dots, s_9\}$  linearly increasing from a random minimum  $s_1 \in [0.05, 0.25]$  to  $s_9 = 1.0$ . JPEG compression artifacts are added at lower resolutions to mimic real-world degradation. The corresponding Python implementation (simplified) is:

$$I_t = \text{Resize}(\text{Resize}(I_{\text{HR}}, s_t), 1/s_t), \quad t = 1, \dots, 9. \quad (4)$$

**2) Blur-to-Sharp Dataset.** To emulate motion blur recovery, we apply directional motion kernels with gradually decreasing blur intensity. For each image, a random blur angle  $\alpha \in [0, 360]$  and a maximum kernel length  $k_{\text{max}} \in [40, 200]$  are sampled. Each frame  $I_t$  is obtained by applying a kernel of size  $k_t = k_{\text{max}} \cdot (1 - (t - 1)/8)$ . The last frame ( $t = 9$ ) is left unblurred.

$$I_t = \text{MotionBlur}(I_{\text{HR}}, k_t, \alpha), \quad t = 1, \dots, 9. \quad (5)$$

**3) Low-Light Progression Dataset.** For low-light enhancement, we generate realistic darkening effects using exposure roll-off, white-balance drift, Gaussian noise, and mild blur. Each video depicts a gradual illumination increase from dark to bright:

$$I_t = \mathcal{L}(I_{\text{HR}}, s_t, n_t), \quad (6)$$

where  $\mathcal{L}(\cdot)$  represents the low-light transformation parameterized by strength  $s_t = 1 - (t - 1)/8$  and noise level  $n_t \in [0.02, 0.08]$ . This formulation produces visually convincing nighttime conditions with sensor noise and color casts.

**Data Organization.** Each dataset directory contains the generated videos, along with two text files: `videos.txt` listing relative paths to all clips, and `prompt.txt` containing the corresponding textual prompts. For every image, two videos are generated with random degradation strengths, resulting in twice the number of samples as the source dataset.

**Uniform vs. Scene-Adaptive Prompts.** We prepare two prompt versions for fine-tuning:

- **Uniform Text Version:** All videos within the same task share a fixed prompt, e.g., “*The image becomes sharper and higher in resolution. Nothing moves. Static image.*” for the resolution dataset.
- **Scene-Adaptive Text Version:** Each video uses a distinct prompt generated automatically by LLaVA [16] and refined with ChatGPT [17], [18], [19], [20], incorporating

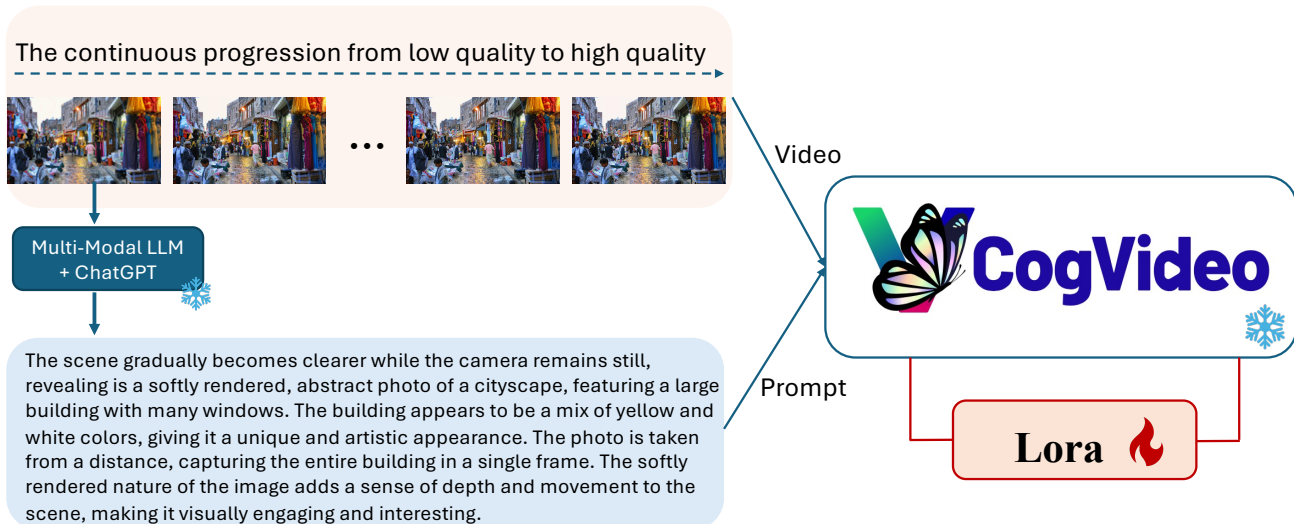


Fig. 1. Illustration of the CogVideo fine-tuning process for image restoration. The model receives a video showing progressive visual enhancement (e.g., from low-resolution to super-resolution) and a corresponding textual prompt describing the scene and restoration dynamics. Using LoRA fine-tuning, CogVideo learns to align temporal visual improvements with text semantics, enabling it to generate restoration sequences conditioned on both the input frame and the prompt.




Fig. 2. Frame-wise restoration performance across three enhancement tasks: (a) super-resolution, (b) deblurring, and (c) low-light enhancement. For each task, PSNR and SSIM generally increase while LPIPS decreases as frames progress, confirming that the model learns temporal enhancement dynamics. The various-prompt version consistently provides slightly higher perceptual quality and smoother progression, highlighting the benefit of scene-aware textual conditioning.

semantic details from the scene while maintaining the same restoration theme (e.g., “A night street gradually brightens under lamplight.” or “A blurred portrait becomes focused.”).

### C. Fine-Tuning and Inference

We fine-tune the pretrained CogVideo model to learn progressive restoration from degraded to clean images. The overall pipeline is illustrated in Fig. 1. Given a textual prompt  $p$  describing the enhancement process and an initial degraded image  $I_d$ , the model generates a video sequence  $\mathcal{V} = \{I_1, I_2, \dots, I_T\}$  that depicts a smooth visual transition toward restoration. The last frame  $I_T$  serves as the final restored output.

We employ lightweight LoRA fine-tuning to adapt CogVideo efficiently while keeping most parameters frozen.

**Inference.** At inference time, CogVideo takes a degraded input image  $I_d$  and a prompt  $p$  (either uniform or scene-adaptive) and generates a video  $\mathcal{V} = G_\phi(p, I_d)$ . The resulting frames  $\{I_t\}$  visualize the restoration trajectory, and the last frame  $I_T$

is used as the enhanced output image. This process enables unified inference for super-resolution, deblurring, and low-light enhancement within a single generative framework.

## III. EXPERIMENTS

We conduct a series of experiments to evaluate the effectiveness of our fine-tuned CogVideo framework on three synthetic datasets (super-resolution, deblurring, and low-light enhancement) and one real-world dataset for motion blur restoration. We report both quantitative metrics and qualitative visual comparisons. All experiments are performed on one NVIDIA H100 GPU using the LoRA fine-tuned CogVideoX-5B-I2V model.

### A. Quantitative Analysis

To comprehensively evaluate the restoration dynamics across different degradation types, we summarize the frame-wise quantitative trends of the fine-tuned CogVideo on the super-resolution, deblurring, and low-light enhancement

TABLE I  
FRAME-WISE LPIPS VALUES ON THE ReLoBLUR TEST SET. LOWER LPIPS INDICATES BETTER PERFORMANCE.

Frame	1	2	3	4	5	6	7	8	9
LPIPS	0.335	0.330	0.323	0.319	0.316	0.318	0.316	0.317	0.317

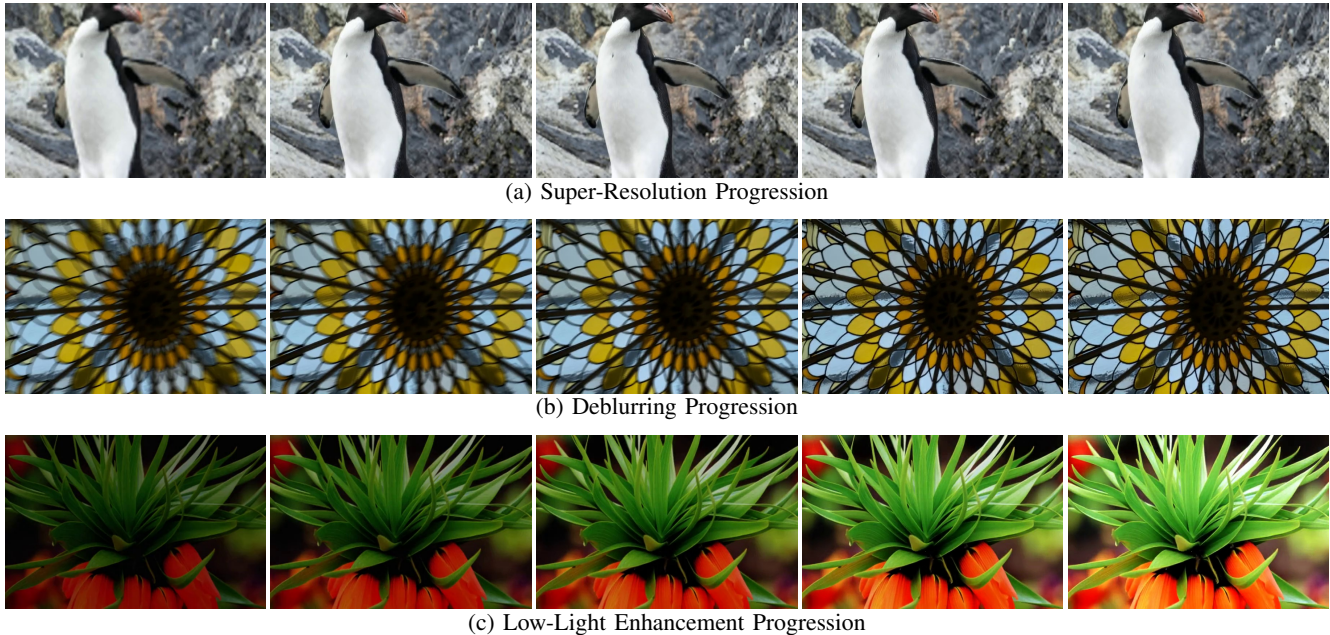


Fig. 3. Qualitative restoration results across three enhancement tasks using the various-prompt fine-tuned CogVideo. Each row visualizes frames 1, 3, 5, 7, and 9 to illustrate the temporal restoration trajectory. The model demonstrates smooth progression, gradually recovering fine details, edges, and illumination while preserving global structure and perceptual consistency.

datasets. Figure 2 visualizes PSNR, SSIM, and LPIPS progression curves for both the uniform- and scene-adaptive various-prompt configurations across all three tasks.

Across all tasks, the fine-tuned CogVideo exhibits a clear temporal enhancement trend: initial frames correspond to heavily degraded conditions, while later frames progressively restore spatial details and perceptual fidelity. In the super-resolution setting, performance improves rapidly during early frames and then stabilizes, indicating fast convergence toward the clean state. For deblurring, the progression is smoother and more gradual, suggesting that temporal modeling helps capture the evolution of motion sharpness. In low-light enhancement, the improvement is monotonic across all frames, reflecting the model’s ability to simulate illumination recovery over time.

Comparing prompt strategies, the various-prompt version consistently achieves higher PSNR and SSIM and lower LPIPS across all three datasets, demonstrating better generalization and semantic-text alignment. This confirms that integrating diverse, scene-specific prompts during fine-tuning encourages the model to more effectively couple textual context with visual enhancement, leading to both quantitative and perceptual gains.

### B. Qualitative Results

To further demonstrate the restoration capability and temporal coherence of the proposed framework, we visualize

frame-wise qualitative results from the various-prompt fine-tuned CogVideo model across all three enhancement tasks. Figure 3 shows the generated progression sequences for super-resolution, deblurring, and low-light enhancement, respectively. Each row represents a single task, and frames 1, 3, 5, 7, and 9 illustrate the restoration trajectory over time.

Across all tasks, the qualitative results confirm that the fine-tuned CogVideo successfully learns a temporally consistent enhancement process. For super-resolution, edges and textures become increasingly sharper across frames. For deblurring, motion-induced artifacts are gradually removed while maintaining realistic contours. In the low-light case, brightness and contrast improve steadily with minimal color distortion. Overall, the visual results illustrate that the model not only restores perceptual quality but also encodes the restoration process itself as a smooth temporal evolution, aligning well with our quantitative observations in Section 3.1.

### C. Evaluation on the Real-World Dataset

To evaluate the generalization capability of our fine-tuned CogVideo model beyond synthetic training conditions, we conduct a zero-shot test on the **ReLoBlur** dataset, a real-world benchmark that captures natural motion blur across diverse lighting and motion scenarios. No additional fine-tuning or adaptation is performed for this experiment.

As shown in Table I, the LPIPS metric decreases during the early frames and stabilizes after frame 5, indicating that




Fig. 4. Qualitative example from the **ReLoBlur** dataset focusing on a moving football scene. The fine-tuned CogVideo progressively restores the sharp structure and texture of the football while reducing surrounding motion streaks and background smear. From frame 1 to frame 9, the football becomes increasingly clear and well-defined, demonstrating the model’s ability to recover localized high-frequency motion details and maintain temporal consistency without over-sharpening artifacts.

the model progressively enhances perceptual similarity before convergence. Despite being trained solely on synthetic video progressions, CogVideo generalizes well to real-world motion blur, capturing the temporal evolution from heavily blurred to sharp regions.

Overall, the quantitative and qualitative results on ReLoBlur demonstrate that the proposed fine-tuning strategy enables CogVideo to learn a robust restoration prior that transfers effectively from synthetic to real-world degradations. The ability to model restoration as a temporal generation process appears to promote generalization across diverse motion patterns and lighting conditions.

#### IV. CONCLUSION

We presented a novel approach to image restoration using **CogVideo** fine-tuned for progressive enhancement. By framing restoration as a temporal generation process, the model learns to transform degraded inputs into gradually improved sequences, covering super-resolution, deblurring, and low-light enhancement. Our results show consistent improvement in perceptual metrics and clear restoration trajectories, with strong generalization to real-world motion blur in the **ReLoBlur** dataset. This study highlights the potential of text-conditioned video generation as an interpretable and extensible paradigm for unified visual restoration tasks.

#### REFERENCES

- [1] Shenghai Yuan, Jinfa Huang, Xianyi He, Yongyang Ge, Yujun Shi, Liuhan Chen, Jiebo Luo, and Li Yuan, “Identity-preserving text-to-video generation by frequency decomposition,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 12978–12988.
- [2] Penghui Ruan, Pichao Wang, Divya Saxena, Jiannong Cao, and Yuhui Shi, “Enhancing motion in text-to-video generation with decomposed encoding and conditioning,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 70101–70129, 2024.
- [3] Taegyeong Lee, Soyeong Kwon, and Taehwan Kim, “Grid diffusion models for text-to-video generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 8734–8743.
- [4] Shenghai Yuan, Jinfa Huang, Yujun Shi, Yongqi Xu, Ruijie Zhu, Bin Lin, Xinhua Cheng, Li Yuan, and Jiebo Luo, “Magictime: Time-lapse video generation models as metamorphic simulators,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [5] Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan, “Videocrafter2: Overcoming data limitations for high-quality video diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 7310–7320.
- [6] Xiang Wang, Shiwei Zhang, Hangjie Yuan, Zhiwu Qing, Biao Gong, Yingya Zhang, Yujun Shen, Changxin Gao, and Nong Sang, “A recipe for scaling up text-to-video generation with text-free videos,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 6572–6582.
- [7] Hangjie Yuan, Shiwei Zhang, Xiang Wang, Yujie Wei, Tao Feng, Yining Pan, Yingya Zhang, Ziwei Liu, Samuel Albanie, and Dong Ni, “Instructvideo: Instructing video diffusion models with human feedback,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 6463–6474.
- [8] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al., “Cogvideox: Text-to-video diffusion models with an expert transformer,” *arXiv preprint arXiv:2408.06072*, 2024.
- [9] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang, “Restormer: Efficient transformer for high-resolution image restoration,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5728–5739.
- [10] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee, “Accurate image super-resolution using very deep convolutional networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1646–1654.
- [11] Ying Tai, Jian Yang, Xiaoming Liu, and Chunyan Xu, “Memnet: A persistent memory network for image restoration,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 4539–4547.
- [12] Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, and Thomas S Huang, “Non-local recurrent network for image restoration,” *Advances in neural information processing systems*, vol. 31, 2018.
- [13] Stamatios Lefkimmiatis, “Universal denoising networks: a novel cnn architecture for image denoising,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3204–3213.
- [14] Haoying Li, Ziran Zhang, Tingting Jiang, Peng Luo, Huajun Feng, and Zhihai Xu, “Real-world deep local motion deblurring,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023, vol. 37, pp. 1314–1322.
- [15] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, Lei Zhang, Bee Lim, et al., “Ntire 2017 challenge on single image super-resolution: Methods and results,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- [16] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee, “Visual instruction tuning,” *Advances in neural information processing systems*, vol. 36, pp. 34892–34916, 2023.
- [17] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al., “Improving language understanding by generative pre-training,” 2018.
- [18] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al., “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, pp. 9, 2019.
- [19] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al., “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [20] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al., “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.