Plug-and-Play Image Dehazing

Zherui Zhang

Abstract—Single-image dehazing is an ill-posed inverse problem where recovering the clean scene radiance requires estimating both the transmission map and atmospheric light. While classical methods like the Dark Channel Prior (DCP) rely on handcrafted assumptions that often fail in bright regions, recent end-to-end deep learning approaches suffer from poor generalization to real-world scenes and require extensive training data. This work proposes a unified, training-free Plug-and-Play (PnP) optimization framework that combines physical modeling with off-the-shelf learned image priors. We systematically evaluate three distinct deep priors—DnCNN, DRUNet, and diffusion-based models—within an ADMM solver without performing any task-specific fine-tuning. Our comparative study reveals that the choice of prior critically determines stability and performance. We demonstrate that simpler CNN priors achieve superior trade-offs in restoration quality, computational efficiency, and convergence stability compared to more complex architectures. Experiments on the RESIDE-SOTS benchmark show consistent improvements over classical baselines, with our best configuration achieving significant gains while maintaining practical runtime.

Index Terms—Image dehazing, Plug-and-Play priors, ADMM, Image restoration, Deep learning priors, Denoiser architectures

1 Introduction

Outdoor photography in adverse weather conditions is often degraded by atmospheric haze, which reduces visibility and contrast. The physical process of haze formation is commonly modeled by the atmospheric scattering equation [1]:

$$\mathbf{I}(\mathbf{x}) = \mathbf{J}(\mathbf{x})t(\mathbf{x}) + \mathbf{A}(1 - t(\mathbf{x})), \tag{1}$$

where $\mathbf{I}(\mathbf{x}) \in \mathbb{R}^3$ is the observed hazy image at pixel $\mathbf{x}$, $\mathbf{J}(\mathbf{x}) \in \mathbb{R}^3$ is the clean scene radiance, $t(\mathbf{x}) \in [0, 1]$ is the transmission map, and $\mathbf{A} \in \mathbb{R}^3$ is the global atmospheric light. Recovering $\mathbf{J}$ from a single image is highly ill-posed because both t and $\mathbf{A}$ are unknown.

Classical approaches estimate these quantities using physical priors. The Dark Channel Prior (DCP), widely adopted due to its efficiency, assumes that haze-free outdoor images contain pixels with very low intensity in at least one color channel [1]. By estimating the transmission map through this prior and refining it with a guided filter [2], DCP has become a standard baseline due to its computational efficiency and interpretability. However, DCP breaks down in bright regions (e.g., sky, white surfaces) and often produces color shifts and halo artifacts.

More recently, end-to-end deep learning methods such as AOD-Net and DehazeNet have achieved strong performance on synthetic benchmarks by directly learning the mapping from I to J [3], [4]. Benchmarks like RESIDE-SOTS provide large-scale paired hazy/clean datasets to train and evaluate such models [5]. However, these models often rely heavily on the training distribution and can generalize poorly to real-world haze with different atmospheric conditions or camera responses. A recent comparative study on real hazy images finds that many data-driven dehazing networks trained on synthetic datasets degrade significantly under dense, real-world haze, while hybrid or physics-aware methods are more robust [6].

Plug-and-Play (PnP) priors offer a middle ground: they integrate pre-trained image denoisers into classical optimization frameworks, combining physical constraints with learned image priors without task-specific retraining [7]. Instead of explicitly defining a regularizer, PnP methods replace the proximal operator in iterative algorithms (such as ADMM or HQS) with an off-the-shelf denoiser. Crucially, this approach is training-free: it leverages the robust feature representations of generic denoisers without requiring hazy-clean image pairs for fine-tuning. For dehazing, joint optimization frameworks that couple transmission estimation with CNN-based priors have shown that combining physics-based models with learned representations can improve stability and visual quality [8]. However, systematic comparisons across modern denoiser architectures remain limited.

In this work, we develop a unified PnP-ADMM framework and evaluate three representative deep priors: a feed-forward CNN (DnCNN) [9], a residual U-Net (DRUNet) [10], and a diffusion-based prior (DiffPIR) [11], [12]. We provide a comparative analysis revealing how denoiser architecture and hyperparameters affect optimization stability and final restoration quality. Our key findings are:

1) Architectural complexity does not guarantee better performance in iterative optimization.
2) Different denoisers have distinct optimal iteration counts and convergence properties.
3) Proper implementation of diffusion-based priors is critical for success in inverse problems.

We evaluate our methods on the RESIDE-SOTS benchmark [5] and discuss trade-offs between quality, runtime,

• *Zherui Zhang is an undergraduate student with the Department of Computer Science, University of Toronto, Toronto, ON, Canada*
E-mail: zherui.zhang@mail.utoronto.ca

and stability.

2 Related Work

2.1 Classical single-image dehazing

Early dehazing methods rely on the atmospheric scattering model and handcrafted priors to recover the clean image. The Dark Channel Prior (DCP), proposed by He et al., assumes that in most haze-free outdoor patches, at least one color channel has very low intensity [1]. By estimating the transmission map through this prior and refining it with a guided filter [2], DCP has become a standard baseline due to its computational efficiency and interpretability. However, DCP often fails in bright regions (sky, white surfaces) where the dark channel assumption breaks down, leading to color shifts and halo artifacts.

More recent classical methods employ other priors, such as color lines and color constancy, but similarly suffer from limitations when these priors are violated. These handcrafted constraints motivate the search for more flexible, learnable approaches.

2.2 End-to-end deep learning for dehazing

End-to-end CNN-based methods, such as AOD-Net and DehazeNet, directly learn the mapping from hazy to clean images [3], [4]. Trained on synthetic datasets like RESIDE-SOTS [5], these approaches achieve high performance on in-distribution benchmarks. However, they are often treated as "black boxes" that ignore the underlying physical model, and their performance can degrade substantially on real-world hazy images. A recent comparative analysis on real hazy scenes reports that many purely data-driven methods exhibit color distortions or loss of detail in dense haze, whereas physically grounded or hybrid approaches are more robust across varying conditions [6]. Additionally, these end-to-end models lack interpretability, making it difficult to understand when and why they fail.

2.3 Plug-and-play image restoration

The Plug-and-Play (PnP) framework offers a middle ground by combining classical optimization with learned denoisers. Rather than end-to-end training, PnP methods replace the proximal operator in iterative algorithms (such as ADMM or HQS) with a pre-trained image denoiser [7]. This approach has been successfully applied to image deblurring and super-resolution, demonstrating that off-the-shelf denoisers can serve as powerful implicit priors without task-specific training.

For dehazing and related inverse problems, Zhang et al. proposed joint optimization frameworks that couple transmission estimation with CNN priors [8]. These works show that combining physics-based models with learned representations improves stability and reconstruction quality compared to purely handcrafted or purely data-driven approaches. Building on this direction, the current work extends PnP to systematically compare multiple denoiser architectures within a unified ADMM framework, investigating how denoiser choice affects convergence behavior and restoration quality.

2.4 Deep denoisers: CNN architectures

DnCNN, a 20-layer feed-forward convolutional network trained on additive Gaussian noise, has become a widely adopted denoiser in many restoration tasks due to its simplicity and effectiveness [9]. More recently, residual U-Net architectures like DRUNet employ encoder–decoder structures with skip connections and deeper feature hierarchies to model complex degradations, and have been used as powerful priors within PnP-style image restoration frameworks [10]. While DRUNet achieves state-of-the-art results in traditional denoising benchmarks, its stability and applicability within iterative optimization frameworks like PnP remains an open question. This motivates a direct comparison between simpler and more complex denoisers in the context of single-image dehazing.

2.5 Diffusion models as priors for inverse problems

Diffusion models have emerged as powerful generative priors for inverse problems. Recent works demonstrate that score-based diffusion models can be incorporated into iterative solvers to reconstruct high-quality images while enforcing data consistency [11]. Methods such as DiffPIR apply diffusion priors to image restoration by iteratively refining estimates through learned reverse processes and carefully designed projection operators [12]. However, the success of diffusion-based methods depends critically on proper score matching, accurate noise scheduling, and well-posed data consistency formulations. A naïve wrapper of a standard denoiser as a "diffusion prior" may lack the learned reverse process and refined projections necessary for complex inverse problems such as dehazing, an issue we investigate empirically in this work.

2.6 Positioning of this work

This work provides a systematic comparison of three representative deep learning priors (DnCNN, DRUNet, and diffusion-based models) within a unified PnP-ADMM framework for single-image dehazing. Unlike prior works that evaluate a single method or architecture [3], [4], [8], [12], we investigate how architectural complexity, iteration dynamics, and implementation details affect both optimization stability and restoration quality. Our findings reveal that simpler architectures can outperform more complex denoisers in this setting, and that diffusion priors require proper pretraining and careful implementation to succeed in dehazing. These insights have implications for the broader use of PnP methods and deep generative priors in imaging inverse problems [6], [7], [10], [11].

3 Proposed Method

We investigate three deep learning priors—DnCNN [9], DRUNet [10], and diffusion models [12]—within a unified Plug-and-Play (PnP) ADMM framework. Our approach decouples the physical haze formation model from the image prior, allowing us to systematically compare how different denoisers affect restoration quality and optimization stability.

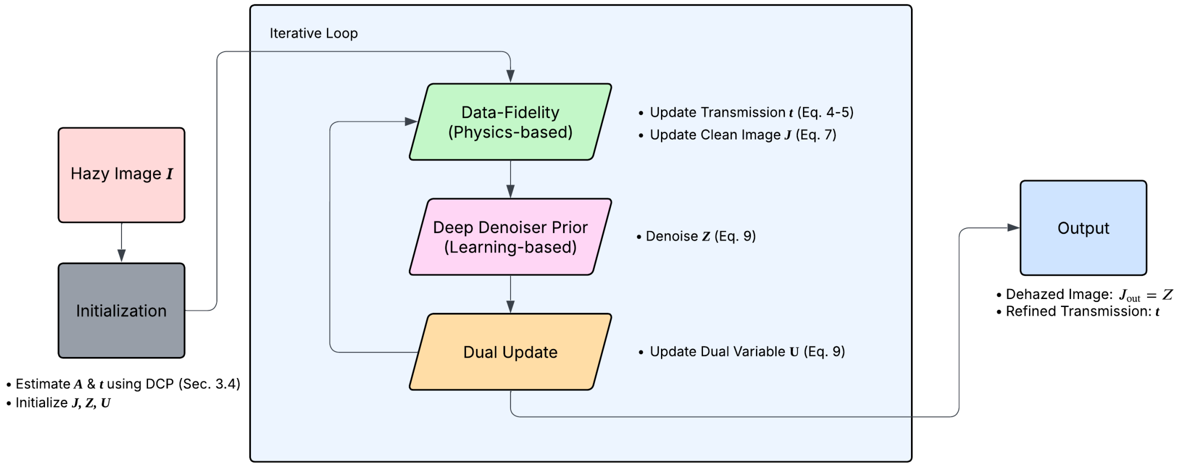

Fig. 1. Overview of our Plug-and-Play ADMM Dehazing Framework. The method alternates between physics-based updates (transmission and clean image estimation) and a learning-based prior step (denoising). The clean image update ensures consistency with the atmospheric scattering model, while the denoiser enforces natural image statistics.

3.1 Problem Formulation

3.1.1 Optimization Objective

Based on the atmospheric scattering model defined in Eq. (1), we formulate single-image dehazing as a regularized inverse problem. We seek to jointly estimate the clean image $\mathbf{J}$ and transmission map t by minimizing:

$$\min_{\mathbf{J},t} \underbrace{\|\mathbf{I} - (\mathbf{J} \odot t + \mathbf{A} \odot (1-t))\|_2^2}_{\text{Data Fidelity}} + \underbrace{\lambda \mathcal{R}(\mathbf{J})}_{\text{Image Prior}}, \tag{2}$$

where $\odot$ denotes element-wise multiplication (broadcasting the scalar t to the 3 color channels), and $\mathcal{R}(\mathbf{J})$ is an implicit prior enforced by a learned denoiser. The atmospheric light $\mathbf{A}$ is estimated during initialization and fixed throughout the optimization (see Section 3.4).

3.2 PnP-ADMM Framework

We employ the Plug-and-Play ADMM algorithm [7] to decouple data fidelity from the image prior, enabling the use of off-the-shelf denoisers. We introduce an auxiliary variable $\mathbf{Z}$ and the constraint $\mathbf{J} = \mathbf{Z}$ to reformulate the problem as:

$$\min_{\mathbf{J},\mathbf{Z},t} \|\mathbf{I} - (\mathbf{J} \odot t + \mathbf{A} \odot (1-t))\|_2^2 + \frac{\rho}{2} \|\mathbf{J} - \mathbf{Z} + \mathbf{U}\|_2^2, \tag{3}$$

where $\mathbf{U} \in \mathbb{R}^{H \times W \times 3}$ is the scaled dual variable, and $\rho > 0$ is the penalty parameter controlling the trade-off between data fidelity and prior strength.

The optimization alternates between the following update steps (summarized in Fig. 1):

3.2.1 Transmission Update (t-step)

Given the current estimate $\mathbf{J}^{(k)}$, we estimate the raw transmission by inverting the haze model per pixel:

$$t_{\text{raw}}^{(k+1)}(\mathbf{x}) = \text{median}_{c \in \{R,G,B\}} \left(\frac{I_c(\mathbf{x}) - A_c}{J_c^{(k)}(\mathbf{x}) - A_c} \right), \tag{4}$$

clipped to $[0.01, 1]$ to prevent division by zero. To suppress noise and artifacts, we refine this raw map using Guided Image Filtering [2]:

$$t^{(k+1)} = \text{GuidedFilter}(t_{\text{raw}}^{(k+1)}, \text{guide} = \mathbf{I}, r = 40, \epsilon = 10^{-3}). \tag{5}$$

This refinement step is performed at every iteration.

3.2.2 Clean Image Update ($\mathbf{J}$-step)

Fixing t and $\mathbf{Z}$, the subproblem for $\mathbf{J}$ is a pixel-wise quadratic minimization:

$$\mathbf{J}^{(k+1)} = \arg\min_{\mathbf{J}} \left\| \mathbf{I} - \left(\mathbf{J} \odot t^{(k+1)} + \mathbf{A}' \right) \right\|_2^2 + \frac{\rho}{2} \left\| \mathbf{J} - \mathbf{B}^{(k)} \right\|_2^2, \tag{6}$$

where $\mathbf{A}' = \mathbf{A} \odot (1 - t^{(k+1)})$ and $\mathbf{B}^{(k)} = \mathbf{Z}^{(k)} - \mathbf{U}^{(k)}$. The closed-form solution for each channel c at pixel $\mathbf{x}$ is:

$$J_c^{(k+1)}(\mathbf{x}) = \frac{1}{t(\mathbf{x})^2 + \frac{\rho}{2}} \Big[t(\mathbf{x}) \left(I_c(\mathbf{x}) - A_c(1 - t(\mathbf{x})) \right) + \frac{\rho}{2} (Z_c^{(k)}(\mathbf{x}) - U_c^{(k)}(\mathbf{x})) \Big]. \tag{7}$$

The result is clipped to $[0, 1]$.

3.2.3 Prior Step ($\mathbf{Z}$-step)

We enforce the image prior by applying a pre-trained denoiser $\mathcal{D}_\sigma$:

$$\mathbf{Z}^{(k+1)} = \mathcal{D}_\sigma \left(\mathbf{J}^{(k+1)} + \mathbf{U}^{(k)} \right). \tag{8}$$

We compare two denoiser architectures:

1) PnP-DnCNN: Uses a 20-layer feed-forward CNN trained on Gaussian noise [9].
2) PnP-DRUNet: Uses a deep residual U-Net with skip connections [10], which has a much larger receptive field and parameter count (32.65M vs 0.67M).

3.2.4 Dual Variable Update ($\mathbf{U}$-step)

$$\mathbf{U}^{(k+1)} = \mathbf{U}^{(k)} + \left(\mathbf{J}^{(k+1)} - \mathbf{Z}^{(k+1)} \right). \tag{9}$$

3.3 Diffusion-Based Method (DiffPIR)

To evaluate generative diffusion priors, we implement a variant of DiffPIR [12]. Unlike the PnP-ADMM loop, this method uses a sampling-based approach. We iterate backwards from noise $\mathbf{x}_T$ to $\mathbf{x}_0$ using a diffusion model $\mathcal{M}$ to predict noise, followed by a data-consistency projection.

At each timestep i, the estimated clean image $\mathbf{x}_{0|i}$ is projected onto the solution space of Eq. (1):

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \|\mathbf{x} \odot t + \mathbf{A} \odot (1-t) - \mathbf{I}\|_2^2 + \gamma \left\|\mathbf{x} - \mathbf{x}_{0|i}\right\|_2^2, \tag{10}$$

where γ controls the strength of the data consistency. This projection uses a closed-form solution similar to Eq. (7). We use a SimpleDiffusionModel wrapper around DRUNet for the noise prediction step, consistent with recent diffusion-for-inverse-problems literature.

3.4 Implementation Details

3.4.1 Initialization

We estimate the atmospheric light $\mathbf{A}$ using the top 0.1% brightest pixels in the dark channel [1]. This estimate is fixed throughout the optimization to ensure stability. The initial transmission map $t^{(0)}$ is also derived from the Dark Channel Prior and refined via guided filtering.

3.4.2 Hyperparameters

For PnP-ADMM, we set the penalty parameter $\rho = 0.01$ to prioritize the denoiser's influence. We run both DnCNN and DRUNet for 100 and 200 iterations. For DiffPIR, we use $\gamma = 10.0$ and 30 diffusion steps. All guided filtering operations use radius $r = 40$ and regularization $\epsilon = 10^{-3}$.

4 Experimental Results

4.1 Experimental Setup

4.1.1 Dataset

We evaluate our methods on the RESIDE-SOTS (Synthetic Objective Testing Set) benchmark [5], which provides ground truth clean images paired with synthetic haze. Due to the computational cost of iterative optimization (PnP-ADMM), we select a representative subset of 5 scenes (3 indoor, 2 outdoor) covering diverse challenges: dense haze, depth discontinuities, and bright sky regions.

4.1.2 Baselines and Metrics

We compare our PnP frameworks against the Dark Channel Prior (DCP) [1], a standard physics-based baseline. We report:

1) PSNR (dB) and SSIM: To measure restoration accuracy against the ground truth.
2) Runtime (s): Average inference time per image to evaluate efficiency.

We analyze three PnP configurations: PnP-DnCNN, PnP-DRUNet, and the diffusion-based DiffPIR.

TABLE 1
Quantitative comparison on RESIDE-SOTS subset (average over 5 images).

Method	PSNR	SSIM	Time (s)
DCP (Baseline)	18.08	0.873	**0.5**
PnP-DnCNN (100 iter)	22.40	0.890	8.3
PnP-DnCNN (200 iter)	**23.34**	0.889	16.6
PnP-DRUNet (100 iter)	22.40	0.887	15.3
PnP-DRUNet (200 iter)	15.97	0.797	30.6
DiffPIR (30 steps)	9.48	0.600	25.0

4.2 Quantitative Comparison

Table 1 summarizes the average performance of all methods.

Performance Gains: Our proposed PnP-DnCNN (200 iter) achieves the highest restoration quality, improving over the DCP baseline by +5.25 dB on average. This demonstrates that integrating a learned CNN prior significantly helps in recovering details and correcting color shifts that classical priors miss.

Stability vs. Complexity: Interestingly, the simpler DnCNN prior outperforms the more complex DRUNet prior in terms of both performance and stability. PnP-DnCNN continues to improve as iterations increase from 100 to 200. In contrast, PnP-DRUNet diverges at 200 iterations on multiple scenes (average PSNR drops from 22.40 dB to 15.97 dB), suggesting that complex skip connections in U-Net architectures may accumulate errors in the ADMM loop without proper regularization.

4.3 Visual and Quantitative Results

Fig. 2 presents the comprehensive visual and quantitative comparison across all methods and test images.

4.3.1 Success Cases (PnP-DnCNN and PnP-DRUNet at 100 iterations)

On all 5 test images, PnP-DnCNN consistently outperforms the DCP baseline:

1) Low-texture surfaces (Set2, Indoor): DCP produces a very dark, low-contrast result on low-texture regions like walls and tables (11.20 dB). PnP-DnCNN at 100 iters recovers structure and detail on these challenging surfaces (+5.43 dB), and further improves at 200 iters (+8.11 dB).
2) Outdoor Scenes (Sets 4–5): DCP violates the dark channel assumption in bright sky regions, producing artifacts. PnP-DnCNN maintains natural colors and recovers fine details (+5.6 dB and +4.5 dB respectively).
3) Detail Recovery: Across all scenes, PnP-DnCNN adds sharpness and texture clarity without overshooting.

PnP-DRUNet at 100 iterations provides similar results.

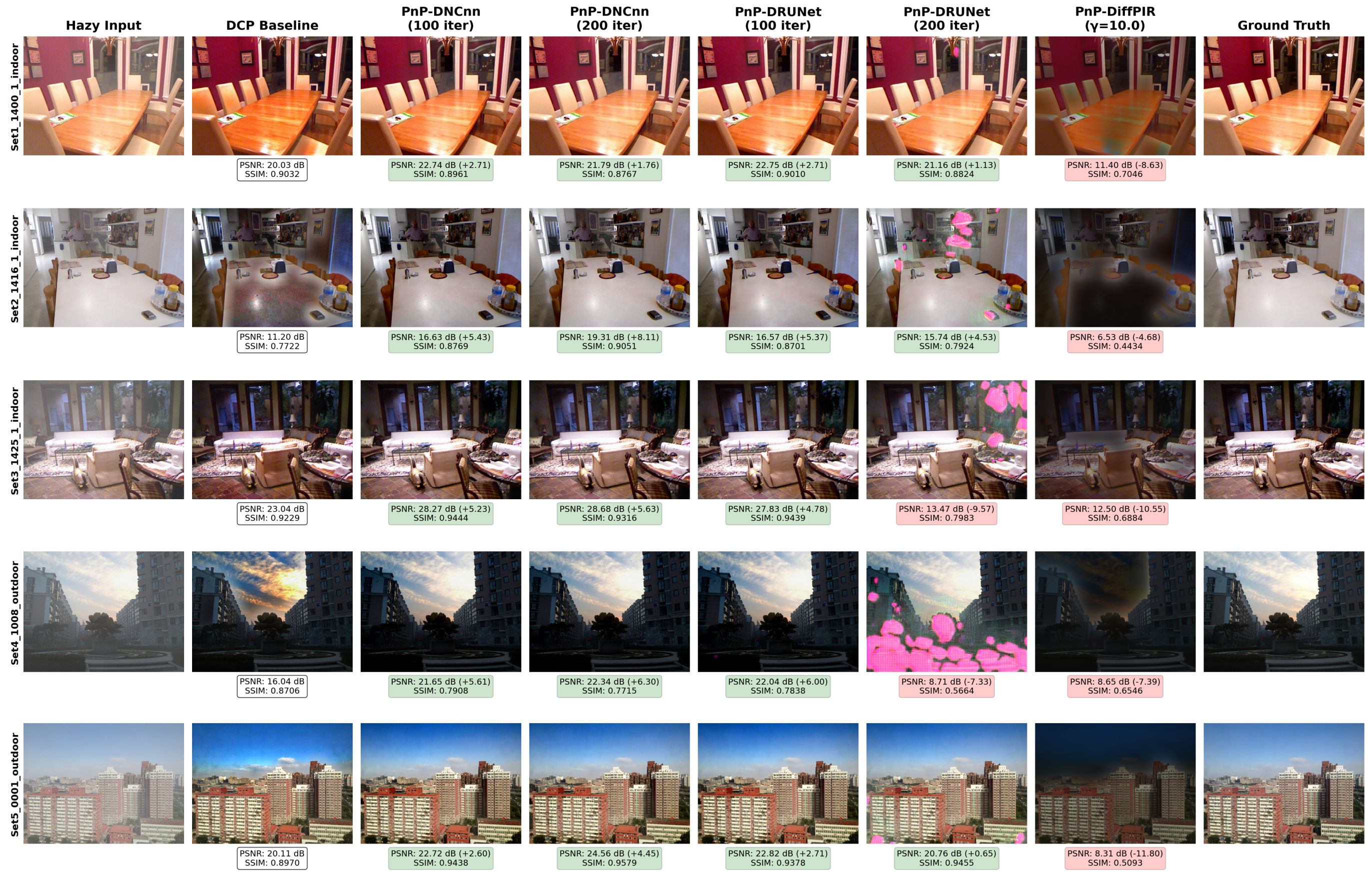

Fig. 2. Comprehensive comparison of all dehazing methods on 5 test scenes (3 indoor, 2 outdoor). Each row shows one scene; columns show hazy input, DCP baseline, and five PnP variants, plus ground truth. PSNR/SSIM and improvements over DCP (green=gain, red=loss) are displayed beneath each result. Note the magenta artifacts visible in the DRUNet-200 column (all rows), indicating optimization divergence.

4.3.2 *Failure Case 1: PnP-DRUNet Divergence at 200 Iterations*

While PnP-DRUNet performs well at 100 iterations, extending to 200 iterations reveals critical instability across all scenes. Visually, magenta artifacts appear in all 5 test images (Fig. 2, DRUNet-200 column), indicating optimization divergence. The severity varies by scene:

1) Set3: DRUNet-100 achieves excellent quality (27.83 dB, +4.78 dB vs. DCP). At 200 iterations, severe magenta artifacts dominate, and PSNR collapses to 13.47 dB (-9.57 dB).
2) Set4: DRUNet-100 performs well (22.04 dB, +6.00 dB). At 200 iterations, widespread magenta noise is visible, dropping to 8.71 dB (-7.33 dB).
3) All Scenes: On Sets 1, 2, and 5 where PSNR degradation is less severe, magenta color artifacts are visually apparent in the DRUNet-200 results, signaling optimization instability.

These artifacts indicate that the ADMM dual variable $\mathbf{U}$ accumulates errors due to the non-expansiveness of DRUNet's skip connections, causing the algorithm to diverge after extended iterations.

4.3.3 *Failure Case 2: DiffPIR Complete Collapse*

DiffPIR fails on all 5 images with severe degradation (-8.61 dB on average):

1) Image Quality: Results are washed-out and low-contrast, lacking spatial structure entirely.
2) Root Cause: The closed-form projection (Eq. (10)) becomes ill-conditioned when transmission $t \to 0$ (dense haze) or atmospheric light $\mathbf{A} \to 1$ (bright scenes). Without a true learned score function, the algorithm cannot correct these projections, resulting in complete failure.

4.4 Critical Role of Visual Inspection

While PSNR and SSIM provide objective quantitative measures, they alone cannot fully capture dehazing quality, particularly when early-stage divergence occurs. This is exemplified by PnP-DRUNet behavior.

4.4.1 *The Divergence Trap*

On Set2 (indoor, dense haze), PnP-DRUNet at 200 iterations achieves PSNR=15.74 dB and SSIM=0.7924—only marginally worse than the 100-iteration result (16.57 dB, 0.8701)—yet the visual output is severely corrupted with magenta artifacts (see Set2, DRUNet-200 in Fig. 2). Similarly, on Set4, the metrics degrade more severely (8.71 dB), revealing widespread magenta patterns that are visually obvious but may be underrepresented in global pixel-level metrics.

4.4.2 *Why Visual Inspection Matters*

The magenta artifacts visible in DRUNet-200 (all 5 scenes) demonstrate why visual inspection is essential:

1) Metric Blindness to Color Artifacts: PSNR and SSIM compute pixel-level errors averaged across the image. Sparse magenta pixels scattered across the image contribute minimally to these global metrics, yet are immediately obvious to human viewers and indicate severe algorithm failure.
2) Detection of Instability: High-frequency magenta noise and color oscillations are hallmark signs of ADMM divergence. These patterns appear before PSNR drops catastrophically, making visual inspection an early warning system.
3) Practical Quality: In real applications, users care about artifact-free, natural-looking results. Methods can achieve decent PSNR yet fail visually (as DRUNet-200 shows), making visual quality equally important as quantitative gains.

Therefore, restoration quality assessment must combine quantitative metrics (Table 1) with careful visual examination (Fig. 2). In practical applications, stability and artifact-free output are as important as raw PSNR gains.

4.5 Runtime and Efficiency

While PnP methods achieve higher quality, they are computationally heavier than the single-pass DCP. PnP-DnCNN (100 iter) offers the best speed-quality trade-off, achieving +4.3 dB gain over DCP while remaining under 10 seconds per image. PnP-DRUNet is approximately $1.8\times$ slower (15.3s vs. 8.3s) due to its larger parameter count (32.65M vs. 0.67M), without offering performance benefits at 100 iterations and with catastrophic divergence at 200 iterations.

5 Conclusion

This work systematically compares three deep learning priors (DnCNN, DRUNet, and diffusion-based models) within a unified Plug-and-Play ADMM framework for single-image dehazing. A key finding challenges conventional deep learning wisdom: architectural complexity does not guarantee better performance in iterative optimization. The simpler DnCNN prior outperforms the more complex DRUNet and fails completely with diffusion-based approaches, demonstrating that thoughtful algorithm design combined with careful prior selection can outperform blindly adopting state-of-the-art architectures.

Our evaluation is limited to 5 synthetic scenes; full-scale validation on real-world hazy images remains future work. Additionally, we assume fixed atmospheric light and use pre-trained off-the-shelf denoisers without task-specific training. We did not implement end-to-end deep learning baselines (AOD-Net, DehazeNet) for comparison, instead focusing exclusively on the classical DCP and our PnP variants; direct comparison with recent end-to-end methods would strengthen the evaluation. Our DiffPIR implementation uses a simplified wrapper rather than proper score-matching diffusion training. Despite these constraints, the systematic comparison provides valuable insights into stability-performance trade-offs of different priors in PnP frameworks.

Future work should focus on real-world validation, theoretical analysis of ADMM convergence with different denoisers, and investigation of whether findings generalize to other restoration tasks. More broadly, this work demonstrates that balanced hybrid approaches combining classical physics and modern learning are more robust than either alone.

Acknowledgments

The author would like to thank Professor David Lindell for an engaging course and insightful lectures, and Shayan Shekarforoush for mentoring and providing valuable suggestions for this project.

References

[1] 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, 2011.
[2] ——, "Guided image filtering," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 6, pp. 1397–1409, 2013.
[3] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng, "Aod-net: All-in-one dehazing network," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 4770–4778.
[4] 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.
[5] B. Li, W. Ren, D. Fu, D. Tao, D. Feng, W. Zeng, and Z. Wang, "Benchmarking single-image dehazing and beyond," *IEEE Transactions on Image Processing*, vol. 28, no. 1, pp. 492–505, 2019.
[6] C. Zheng, W. Ying, and Q. Hu, "Comparative analysis of dehazing algorithms on real-world hazy images," *Scientific Reports*, vol. 15, p. 10822, 2025.
[7] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proceedings of the IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2013, pp. 945–948.
[8] H. Zhang, V. Sindagi, and V. M. Patel, "Joint transmission map estimation and dehazing using deep networks," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 10, pp. 3611–3623, 2020.
[9] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
[10] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. V. Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 6360–6376, 2022.
[11] V. Graikos, A. Malkin, N. Jojic, and D. Samaras, "Diffusion models as plug-and-play priors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 14 715–14 725.
[12] Y. Zhu, K. Zhang, J. Liang, J. Cao, B. Wen, R. Timofte, and L. V. Gool, "Denoising diffusion models for plug-and-play image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2023, pp. 1219–1229.