Guided-Filter Transmission Refinement for Dark Channel Prior-Based Single Image Dehazing

Dr. Pulkit Dwivedi[1], Prof. (Dr.) Shashi Kant Gupta[1]
[1] Lincoln University College, 47301, Petaling Jaya, Selangor Darul Ehsan, Malaysia
Emails: pdwivedi1990@gmail.com , shashigupta@lincoln.edu.my

Abstract: Haze scatters and attenuates light before it reaches an imaging sensor, degrading the contrast and colour fidelity that outdoor vision systems such as autonomous driving, surveillance, and remote sensing depend on. The classical Dark Channel Prior (DCP) recovers a haze-free image by estimating a per-pixel transmission map, but its original box-filtered transmission blurs across depth discontinuities and introduces halo artefacts near edges. This paper implements and quantitatively evaluates a lightweight fix: replacing the box-filtered transmission with an edge-aware guided-filter refinement, and separately testing a CLAHE-based local-contrast post-enhancement stage. Since real hazy images lack paired ground truth, the three variants -- box-filtered baseline, guided-filter-refined, and guided-filter-plus-CLAHE -- are evaluated on a physically synthesised five-image benchmark with known ground truth, using PSNR and SSIM. Guided-filter refinement improves PSNR on three of five images and SSIM on four of five, giving a small net average gain over the baseline, whereas CLAHE post-enhancement raises perceptual contrast but reduces average fidelity to ground truth. These results quantify a concrete, reproducible edge-aware improvement to DCP and expose a fidelity-versus-perceptual-quality trade-off relevant to choosing a post-processing strategy for latency-constrained applications such as driver-assistance and surveillance systems.

Keywords: Image dehazing; Dark channel prior; Guided image filtering; Transmission estimation; Atmospheric scattering model

Introduction

Haze, fog, and smoke are atmospheric phenomena in which suspended particles scatter and absorb light before it reaches an imaging sensor. This scattering is commonly modelled by the Koschmieder atmospheric scattering model, $I(x) = J(x)t(x) + A(1 - t(x))$, where $I(x)$ is the observed hazy image, $J(x)$ is the haze-free scene radiance to be recovered, A is the global atmospheric light, and $t(x)$ is the medium transmission describing the fraction of light reaching the camera without being scattered [1]. The consequence for machine vision is a systematic loss of contrast, colour fidelity, and edge information, degrading downstream tasks such as object detection in autonomous driving, target recognition in surveillance, and land-cover classification in remote sensing.

The Dark Channel Prior (DCP) remains one of the most widely used single-image dehazing techniques because it requires no training data and runs in real time, but its original formulation estimates the raw transmission map and then smooths it with a simple box filter [2]. Because a box filter is not edge-aware, it blurs the transmission map across depth discontinuities and produces halo artefacts around object boundaries -- a well-documented limitation that motivated He et al.'s edge-preserving guided image filter as a general-purpose replacement for box-filter-style smoothing [3]. This paper implements that substitution specifically for the DCP transmission map, adds an optional local-contrast post-enhancement stage, and quantifies the effect of each change on a controlled, paired synthetic-haze benchmark, rather than relying on qualitative visual comparison alone.

Related work

DCP estimates transmission from the empirical observation that, in most haze-free outdoor patches, at least one colour channel has some pixels with very low intensity; the dark channel of a hazy patch is therefore a strong local cue for transmission [2]. Learning-based methods later replaced this hand-crafted cue: DehazeNet trains a CNN to regress the transmission map directly [4], and AOD-Net re-parameterises the scattering model so a compact network outputs the haze-free image in a single forward pass, small enough to run ahead of an object detector [5]. Such networks are typically trained with pixel-wise reconstruction losses combined with perceptual terms; gradient-based regularisation losses proposed for general deep-learning optimisation have also been suggested as a way to stabilise training and sharpen edge recovery in restoration networks of this kind [9]. Attention-based, GAN-based, and transformer-based architectures have since pushed reported PSNR/SSIM further still, at increasing computational cost; DehazeFormer, for example, adapts the Swin Transformer to low-level restoration and reports the highest published PSNR/SSIM on the RESIDE indoor benchmark among widely cited methods [7], and SFX-GAN combines multi-modal spectral fusion with an explainability-oriented adversarial design aimed at complex, non-uniform haze [8]. Standardised comparison across these families is supported by the RESIDE benchmark, which aggregates synthetic and real-world hazy images under a common protocol [6].

The specific limitation this paper addresses -- the box filter used in the original DCP transmission-refinement step -- was identified by He et al., who proposed the guided image filter as an edge-aware, computationally cheap alternative to both box filtering and the soft-matting refinement used in some DCP variants [3]. Table 1 situates the method evaluated in this paper relative to this prior work.

Table 1. Qualitative positioning of the proposed guided-filter-refined DCP relative to prior work.

Method	Physical Scattering Model	Learning-Based	Explicit Transmission Refinement
DCP, box-filtered [2]	Yes	No	Box filter (not edge-aware)
DehazeNet [4]	Yes	Yes (CNN)	N/A -- learned end to end
AOD-Net [5]	Yes (reparam.)	Yes (CNN)	N/A -- learned end to end

DehazeFormer [7]	No	Yes (Transformer)	N/A -- learned end to end
SFX-GAN [8]	No	Yes (GAN)	N/A -- learned end to end
Proposed (this work)	Yes	No	Guided image filter (edge-aware)

Key Contribution

This paper's contribution is a small, concrete, and fully reproducible modification to the classical DCP pipeline together with an honest quantitative evaluation, rather than a new end-to-end learned architecture. Specifically: (i) the box-filter transmission-smoothing step in DCP is replaced with He et al.'s edge-aware guided image filter, using the greyscale hazy image itself as the guidance image [2],[3]; (ii) an optional CLAHE-based local-contrast post-enhancement stage is added and evaluated separately, so its effect on fidelity can be isolated from the transmission-refinement change; (iii) because real hazy photographs have no paired ground truth, a small controlled benchmark is constructed by synthesising haze with a spatially varying, physically based transmission map over five diverse test images, enabling full-reference PSNR/SSIM evaluation rather than qualitative-only comparison; and (iv) both the improvement and its limits are reported per image, including the one test image on which the guided filter under-performed the naive baseline, rather than only the images that favour the proposed change.

Method, Experiments and Results

All three evaluated variants share the standard DCP front end: the dark channel is computed with a 15x15 minimum filter, the atmospheric light A is estimated as the brightest pixel (in the original hazy image) among the top 0.1% of dark-channel intensities, and a raw transmission map is estimated with the usual omega = 0.95 haze-retention factor [2]. The variants differ only in how this raw transmission map is refined before scene radiance is recovered with t0 = 0.1 as the transmission floor: (a) Baseline -- a 15x15 box filter, matching the original DCP formulation; (b) Proposed (GF) -- He et al.'s guided image filter (radius 40, epsilon = 1e-3), using the greyscale hazy image as guide [3]; (c) Proposed (GF+CLAHE) -- the guided-filter output followed by CLAHE (clip limit 0.8, 8x8 tiles) applied to the luminance channel of the recovered image.

Because paired clean/hazy ground truth does not exist for real photographs, five diverse RGB images (320x320) were haze-synthesised using I = Jt + A(1-t) with A = (0.9, 0.9, 0.9) and a spatially varying transmission t = exp(-beta x d(x)), beta = 1.0, where d(x) is a smooth synthetic depth map generated by Gaussian-blurring low-resolution random noise and rescaling to [0.6, 2.8]; this produces non-uniform, depth-correlated haze rather than a single global transmission value. All three pipeline variants were implemented in Python/OpenCV and run on the same five hazy images, and PSNR and SSIM were computed against the known ground truth.

Table 2 reports the resulting averages. Guided-filter refinement improved PSNR on three of the five test images and SSIM on four of five, for a small net average gain of +0.06 dB PSNR and +0.006 SSIM over the box-filtered baseline; the largest gains occurred on images with sharp depth discontinuities near object silhouettes, consistent with the guided filter's edge-preserving behaviour, while the one image dominated by a large low-texture background showed a small decrease (-0.39 dB), where the box filter's extra smoothing incidentally helped. Adding CLAHE increased PSNR on three of five images but reduced it on the other two by up to 2.6 dB, and reduced SSIM on four of five images overall, so its average fidelity was lower than both the guided-filter-only and the box-filtered baseline variants despite visibly higher local contrast in the output images.

Table 2. PSNR and SSIM averaged over the five-image synthetic-haze benchmark (higher is better).

Method	Avg PSNR (dB)	Avg SSIM
Hazy input (no processing)	8.11	0.490
Baseline DCP (box-filtered)	18.18	0.861
Proposed (guided filter)	18.24	0.867
Proposed (guided filter + CLAHE)	17.85	0.845

Discussions

The results show that an edge-aware transmission filter is a small but genuine, reproducible improvement over the original box-filtered DCP formulation, rather than a dramatic one: the guided filter's main benefit is at depth discontinuities, where it avoids smoothing the transmission map across object boundaries and thereby avoids the halo artefacts visible near edges in the baseline. Because both variants share the same dark-channel computation and atmospheric-light estimate, the guided filter is a drop-in replacement that adds negligible computational overhead and no training requirement, which matters for latency-constrained deployments such as advanced driver-assistance systems; this constraint mirrors observations in other real-time computer-vision pipelines, such as real-time human pose and repetition-counting systems, where inference budgets of only a few milliseconds per frame are required for the system to remain usable [10].

The CLAHE result illustrates a fidelity-versus-perceptual-quality trade-off that recurs throughout the dehazing literature: methods optimised for perceptual realism or local contrast, including adversarially trained approaches, do not always maximise pixel-wise fidelity metrics such as PSNR/SSIM against a reference image [8]. Here, CLAHE visibly increased local contrast but reduced average fidelity to the synthetic ground truth, so its use should depend on whether the downstream task rewards visual contrast (e.g., human inspection) or numerical accuracy (e.g., a downstream detector trained on clean images). A limitation of this study is its scale: five images and one synthetic haze model, rather than a large real-world benchmark such as RESIDE, so the reported averages should be read as an indicative, reproducible comparison rather than a general claim.

Conclusions

• **Problem statement addressed / motivation:** Haze degrades contrast and colour fidelity in outdoor images, and the classical Dark Channel Prior's box-filtered transmission map blurs across depth discontinuities, producing halo artefacts that a training-free, real-time-capable dehazing method should avoid.

• **Method used:** The box-filter transmission-refinement step in DCP was replaced with an edge-aware guided image filter, and an optional CLAHE local-contrast post-enhancement stage was added and evaluated separately, using PSNR/SSIM on a five-image paired synthetic-haze benchmark constructed for this study.

• **Key findings:** Guided-filter refinement gave a small net average improvement over the box-filtered baseline (+0.06 dB PSNR, +0.006 SSIM), with gains concentrated on images with strong depth discontinuities; CLAHE post-enhancement increased perceptual contrast but reduced average fidelity to ground truth, demonstrating a measurable fidelity-versus-perceptual-quality trade-off.

• **Limitations and future work:** The evaluation used only five images and one synthetic haze model rather than a large real-world benchmark; future work should validate the guided-filter refinement on RESIDE and real-world hazy photographs, and should compare it quantitatively against learned methods such as DehazeNet, AOD-Net, and DehazeFormer under the same protocol.

References

[1] S. G. Narasimhan and S. K. Nayar, "Vision and the atmosphere," Int. J. Comput. Vis., vol. 48, no. 3, pp. 233–254, 2002.

[2] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," IEEE Trans. Pattern Anal. Mach. Intell., vol. 33, no. 12, pp. 2341–2353, 2011. https://doi.org/10.1109/TPAMI.2010.168

[3] K. He, J. Sun, and X. Tang, "Guided image filtering," IEEE Trans. Pattern Anal. Mach. Intell., vol. 35, no. 6, pp. 1397–1409, 2013.

[4] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao, "DehazeNet: An end-to-end system for single image haze removal," IEEE Trans. Image Process., vol. 25, no. 11, pp. 5187–5198, 2016. https://doi.org/10.1109/TIP.2016.2598681

[5] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng, "AOD-Net: All-in-one dehazing network," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), 2017, pp. 4770–4778.

[6] B. Li et al., "Benchmarking single-image dehazing and beyond," IEEE Trans. Image Process., vol. 28, no. 1, pp. 492–505, 2019. https://doi.org/10.1109/TIP.2018.2867951

[7] Y. Song, Z. He, H. Qian, and X. Du, "Vision transformers for single image dehazing," IEEE Trans. Image Process., vol. 32, pp. 1927–1941, 2023. https://doi.org/10.1109/TIP.2023.3256763

[8] P. Dwivedi, V. Agarwal, N. A. Wani, M. Kajal, and A. Bennour, "SFX-GAN: Sustainable and explainable multi-modal spectral fusion for image dehazing in complex systems," Complex Intell. Syst., vol. 12, no. 3, art. no. 112, 2026.

[9] P. Dwivedi, B. Islam, and M. Kajal, "Smooth gradient loss: A loss function for gradient regularization in deep learning optimization," J. Supercomput., vol. 81, no. 15, art. no. 1457, 2025.

[10] K. Sen, S. Ranjan, and P. Dwivedi, "Real-time exercise repetition counting and form analysis using MediaPipe and computer vision," in Proc. 5th Int. Conf. Adv. Electr. Comput. Technol. (ICAECT), 2025.