

Article

Fast Model-Free Image Dehazing via Haze-Density-Driven Fusion

Jeonghyeon Son ^{1,†}, Dat Ngo ^{2,†} , Suhun Ahn ¹ and Bongsoon Kang ^{1,*} 

¹ Department of Electronics Engineering, Dong-A University, Busan 49315, Republic of Korea; 1922913@donga.ac.kr (J.S.); yusu@donga.ac.kr (S.A.)

² Department of Computer Engineering, Korea National University of Transportation, Chungbuk 27469, Republic of Korea; datngo@ut.ac.kr

* Correspondence: bongsoon@dau.ac.kr; Tel.: +82-51-200-7703

† These authors contributed equally to this work.

Abstract

This paper presents a fast and model-free image dehazing algorithm based on haze-density-driven image fusion. Instead of relying on explicit physical haze models, the proposed approach restores visibility by fusing the input image with its dehazed estimate using spatially adaptive weights derived from a haze-density map. The dehazed estimate is produced by blending multiple synthetically under-exposed versions of the input, where local fusion weights promote stronger enhancement in dense-haze regions while preserving appearance in mild-haze areas. This model-free formulation avoids the limitations inherent to traditional scattering-based models and ensures robust performance under spatially nonuniform haze conditions. The overall framework is lightweight and suitable for embedded, real-time imaging systems due to its reliance on simple local operations. Experimental evaluations demonstrate that the proposed method achieves competitive results compared to state-of-the-art dehazing algorithms in both visual quality and quantitative metrics. A hardware prototype further shows that the method can process high-resolution imagery at real-time rates, achieving 271.74 megapixels per second, or 30.69 frames per second at DCI 4K (4096×2160) resolution. These results establish haze-density-driven fusion as an effective and efficient model-free solution for real-time image dehazing.

Keywords: autonomous dehazing; model-free dehazing; image fusion; FPGA

MSC: 68U10

1. Introduction

Outdoor image degradation under hazy conditions is largely attributed to light attenuation and airlight scattering, which jointly reduce scene visibility and alter color appearance. As haze becomes denser, contrast diminishes and colors shift to a veiled look.

A large body of prior work attempts to reverse this degradation by estimating scene radiance through physical-model inversion, often guided by handcrafted priors [1,2], statistical cues [3–5], or deep networks [6–8]. Although these methods can produce strong results, their dependence on explicit physical modeling reduces robustness under complex or spatially varying haze conditions. Moreover, they implicitly assume that haze is present; when this assumption fails, over-enhancement, color distortion, and halo artifacts often arise. These limitations underscore the need for dehazing strategies that autonomously assess haze conditions and adjust their behavior without relying on model inversion.

Achieving such adaptability in real time remains difficult. High-quality restoration methods are computationally expensive, while lighter designs may lack robustness across



Academic Editor: Manuel G. Forero

Received: 11 December 2025

Revised: 31 December 2025

Accepted: 4 January 2026

Published: 6 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons](#)

[Attribution \(CC BY\) license](#).

diverse scenes. A practical solution must therefore balance accuracy, adaptiveness, and efficiency to support deployment in embedded or low-latency systems.

To address these challenges, we propose a fast, model-free dehazing framework driven by haze-density-aware image fusion. Instead of estimating transmission or atmospheric light, the method constructs a dehazed estimate by fusing the input image with multiple under-exposed variants, whose complementary visibility characteristics enhance structural clarity. Locally adaptive fusion weights, derived from a referenceless haze-density map, regulate the restoration strength across the scene. This formulation avoids artifacts associated with model inversion and exhibits strong robustness to varying haze levels. Furthermore, the algorithm's localized and lightweight operations make it highly suitable for hardware acceleration, enabling real-time processing at high resolutions.

The main contributions of this work are as follows:

- A model-free, haze-density-driven fusion framework that performs autonomous, region-adaptive dehazing without relying on physical-model inversion.
- A complete hardware accelerator that implements the entire pipeline for real-time embedded deployment.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 details the proposed method and hardware design. Section 4 presents quantitative, qualitative, and hardware evaluation results. Section 5 concludes the paper and discusses future directions.

2. Related Work

Over the past decade, image dehazing research has progressed from physically grounded priors to fusion-based techniques, learned models, and efficient hardware realizations. Existing approaches can be broadly grouped into prior-based, fusion-based, deep learning-based, and hardware-accelerated methods.

2.1. Prior-Based Dehazing

Prior-based approaches estimate transmission and atmospheric light using physical or statistical assumptions about natural images. The Dark Channel Prior (DCP) [1] established a widely used baseline by exploiting low-intensity statistics in haze-free regions. The Color Attenuation Prior (CAP) [2] later introduced a linear depth model to improve efficiency, and Haze-Lines [9] characterized RGB distributions for more reliable radiance recovery.

Recent efforts expanded these ideas with more adaptive priors for complex environments, including the Rank-One Prior [10], Color Ellipsoid Prior [11], and scene- or modality-specific priors for UAV [12] and remote-sensing imagery [13]. The RIDCP framework [14] further demonstrated how high-quality codebook priors can enhance real-image dehazing. Although interpretable and broadly applicable, prior-driven methods often depend on explicit physical modeling, which can limit robustness under spatially nonuniform haze or challenging illumination conditions.

2.2. Model-Free Dehazing

Model-free, fusion-based approaches construct multiple image variants and combine them using spatially adaptive weights, avoiding explicit estimation of transmission or atmospheric light. Artificial Multiple-Exposure Image Fusion [15] demonstrated that exposure manipulation and weighted blending can restore visibility across diverse scenes. Extensions such as [16] incorporated local airlight compensation, further improving performance in both daytime and nighttime conditions.

As these methods rely on localized contrast cues rather than explicit physical models, they exhibit strong robustness across different haze conditions. Their emphasis on simple

operations and spatial locality makes them especially well-suited for real-time deployment. This combination of adaptability and efficiency motivates the haze-density-aware fusion strategy developed in this work. Notably, model-free fusion techniques have also inspired hardware-friendly implementations [17], underscoring their practical relevance.

2.3. Learning-Based Dehazing

Deep learning techniques learn the mapping from hazy to clear images directly from data. Early examples such as DehazeNet [18] focused on transmission estimation, while later CNN-based models incorporated multiscale features and structural constraints [19]. Transformer-based approaches, including ViT-based dehazing [8] and MB-TaylorFormer [20], expanded learning capacity for long-range dependencies.

Recent developments combine learning with physical cues or generative modeling, including Feature Physics Models [21], conditional variational approaches [22], diffusion-based techniques [3,6], and detail-enhanced attention mechanisms [7]. Despite state-of-the-art restoration quality, these models often require substantial computation and memory, which limits their deployment in real-time or embedded systems. As a result, lightweight model-free strategies—especially fusion-based pipelines—are increasingly viewed as complementary or preferable alternatives for hardware-oriented settings.

2.4. Hardware-Oriented Dehazing

Hardware research has primarily focused on accelerating existing dehazing algorithms to meet real-time constraints. Early FPGA demonstrations of DCP-based techniques [23,24] confirmed that substantial speedups are achievable through parallelism. Later refinements reduced resource usage, improved temporal stability, and integrated dehazing into larger embedded vision pipelines [25,26].

More recent efforts target high-resolution video, demonstrating real-time throughput for prior-based methods [27,28] or hybrid dehazing models [29]. Fusion-based accelerators [17] further highlight that localized, model-free fusion strategies align naturally with hardware parallelism. These findings reinforce the suitability of lightweight fusion techniques—particularly haze-density-driven formulations—for real-time embedded dehazing systems. Table 1 summarizes representative FPGA realizations and their trade-offs in precision, resource usage, and frame rate.

Table 1. Overview of real-time image dehazing implementations. NA and fps stand for not available and frames per second, respectively.

Year	Paper	Device	Technique	Design Tool	Key Features and Outcomes
2018	[26]	FPGA (NA)	Filtering	NA	<ul style="list-style-type: none"> - Direct FPGA realization of DCP - Proposing a fast airlight estimation method - Achieving a throughput of 88.7 Mpixels/s
2019	[25]	ASIC (TSMC's 0.13 μ m)	Filtering	Verilog HDL	<ul style="list-style-type: none"> - Adding temporal smoothing between successive frames to eliminate flickering - Implementing a 7-stage pipelined ASIC accelerator - Achieving a throughput of 200 Mpixels/s
2021	[30]	FPGA (Xilinx XC7Z020-3CLG484)	CNN	Verilog HDL	<ul style="list-style-type: none"> - Integrating a lightweight CNN for transmission map estimation from the input image and its dark channel - Achieving a throughput of 200 Mpixels/s
2023	[29]	FPGA (Xilinx XC7Z020-CLG484-1)	Filtering	Verilog HDL	<ul style="list-style-type: none"> - Estimating the pixel-wise transmission map directly from saturation - Estimating the atmospheric light using the downsampled input image - Implementing a 7-stage pipelined FPGA accelerator - Achieving a throughput of 85.2 Mpixels/s
2024	[17]	FPGA (Xilinx XC7K325T-2FFG900C)	Image fusion	NA	<ul style="list-style-type: none"> - Combining high-boost filtering with linear intensity stretching for haze removal - Achieving a throughput of 72.299 Mpixels/s

Table 1. *Cont.*

Year	Paper	Device	Technique	Design Tool	Key Features and Outcomes
2025	[27]	FPGA (Xilinx XC7Z020CLG400)	Filtering	NA	<ul style="list-style-type: none"> - Direct implementation of DCP - Refining the transmission map using a guided filter - Post-processing the dehazed image using gamma correction - Processing HD videos at 60 fps
2025	[28]	FPGA (Xilinx XC7S25CSGA225-2)	Filtering	NA	<ul style="list-style-type: none"> - Improving DCP with custom sky segmentation - Leveraging global saturation to detect haze - Leveraging contrast to estimate haze density - Processing Full HD videos at 60 fps

3. Proposed Method

This section presents the proposed haze-density-driven fusion framework and its corresponding hardware accelerator for real-time deployment. We first describe the model-free dehazing algorithm, followed by the hardware architecture that executes this pipeline efficiently at high resolutions.

As illustrated in Figure 1, the method comprises seven modules. The fusion-based dehazing module generates an initial dehazed estimate of the input image. The haziness-degree evaluator computes pixel-wise haze density and derives both patch-based and average haze values. The interpolation module low-pass filters and upsamples the patch-based densities to produce a smooth spatial haze field. This field is then used by the local-blending-weight module to compute fusion weights. The image-blending module fuses the input with its dehazed estimate using these weights. In parallel, the average haze density feeds the self-calibrating-weight module, which controls the adaptive tone-remapping module. This final module enhances the fused image while preventing artifacts such as undershoot, over-enhancement, and color distortion.

To support real-time operation, the modules are scheduled across video timing intervals: the active-period modules (pale green in Figure 1) process streaming pixels continuously, while the interpolation module (pale blue) operates during blanking periods, leveraging temporal reuse to maintain throughput. This organization matches the model-free design of the proposed framework and facilitates efficient hardware realization.

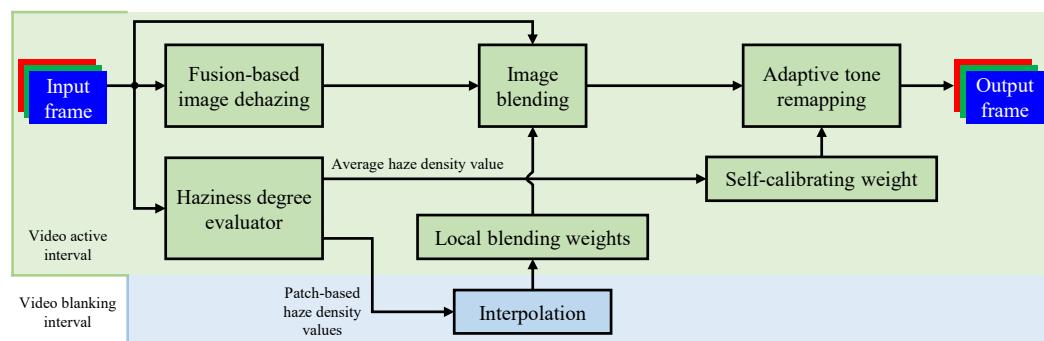


Figure 1. Overview of the proposed haze-density-driven fusion dehazing framework. Green modules operate during the video active interval, whereas the interpolation module (blue) runs during the blanking interval to avoid throughput bottlenecks.

3.1. Fusion-Based Image Dehazing

The first stage produces a dehazed estimate using a multi-variant fusion strategy. As shown in Figure 2, which illustrates a hazy image from the IVC dataset [31], applying different under-exposure levels to a hazy image (for example, gamma values of 1.5 and 2.5) reveals scene structures otherwise obscured by veiling light; objects become increasingly discernible in these darker variants (highlighted in pink, blue, and red rectangles). This

observation motivates the generation of multiple under-exposed versions of the input and their fusion to enhance scene visibility.

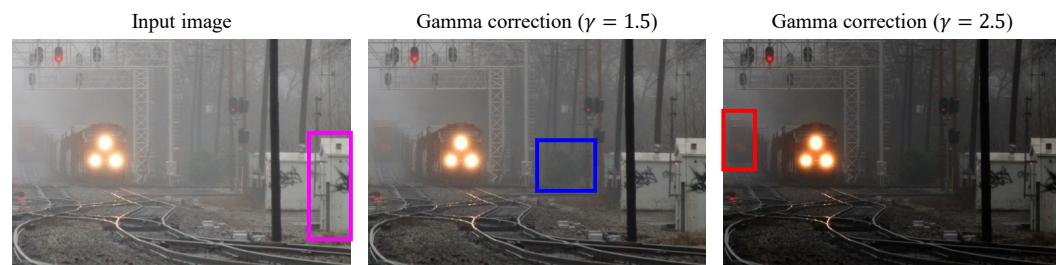


Figure 2. Hazy image and two of its gamma-adjusted variants. Colored rectangles highlight representative regions whose scene details become clearer.

As haze suppresses fine details, a detail-enhancement step is applied prior to gamma correction and fusion (Figure 3). The resulting variants are then combined using dark-channel-based weights, whose strong correlation with haze density makes them both simple and effective. This model-free formulation yields a dehazed estimate that enhances structural details, reduces veiling, and maintains natural color appearance.

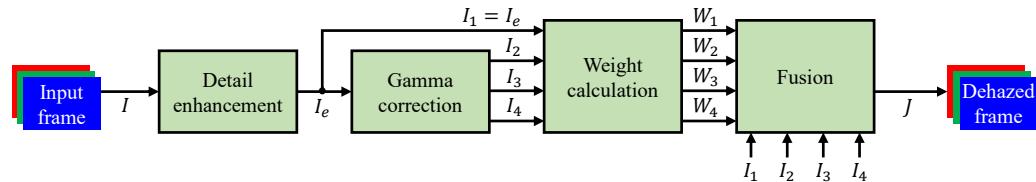


Figure 3. Fusion-based dehazing module. Multiple gamma-adjusted variants are fused using dark-channel-derived weights.

3.1.1. Detail Enhancement

Given an input image (I), the luminance component is first extracted via its conversion from the RGB color space to the YCbCr domain. Detail enhancement is then performed on the luminance channel (Y), where high-frequency information representing local texture and edge content is obtained via convolution with a Laplacian kernel (∇^2). The resulting detail layer (e) is combined with the original luminance through an enhancement weight (ω) that controls the degree of sharpening. This weight is formulated as a piecewise linear function of local variance (v), allowing spatially adaptive amplification of fine details in textured regions while preventing noise over-enhancement in smooth or flat areas. The enhanced luminance (Y_e) is subsequently merged with its Cb and Cr components to reconstruct a refined RGB image (I_e).

This step is mathematically expressed as follows:

$$Y_e = Y + \omega e, \quad (1)$$

$$e = Y \circledast \nabla^2, \quad (2)$$

$$\omega = \begin{cases} \omega_1 & v < v_1 \\ \left(\frac{\omega_2 - \omega_1}{v_2 - v_1} \right) v + \frac{\omega_1 v_2 - \omega_2 v_1}{v_2 - v_1} & v_1 \leq v \leq v_2, \\ \omega_2 & v > v_2 \end{cases} \quad (3)$$

$$v = Y^2 \circledast U - (Y \circledast U)^2, \quad (4)$$

where U denotes the averaging kernel, \circledast represents the convolution operation, and user-defined parameters $\{v_1, v_2, \omega_1, \omega_2\}$ are empirically determined as $\{0.001, 0.01, 2.5, 1\}$.

3.1.2. Gamma Correction

Gamma correction is applied to the detail-enhanced image (I_e) to simulate multiple under-exposure levels. In principle, increasing the number of gamma-adjusted variants can further improve dehazing quality by providing richer visibility cues. However, each additional variant requires an extra look-up table (LUT), dark-channel computation, and fusion branch, causing hardware resource usage to grow exponentially. Considering that dehazing serves as a pre-processing stage in a larger vision system, a balance between algorithmic performance and hardware efficiency is essential.

To this end, three gamma values of 1.9, 1.95, and 2 were empirically selected to yield perceptually distinct under-exposed variants that effectively attenuate haze while keeping resource usage manageable. Together with the uncorrected image, these variants form the set $\{I_1 = I_e, I_2 = (I_e)^{1.9}, I_3 = (I_e)^{1.95}, I_4 = (I_e)^2\}$ used in the subsequent fusion stage.

3.1.3. Weight Calculation and Image Fusion

For each variant and the original image, the complement of the dark channel with respect to one is computed and then L1-normalized, serving as fusion weights that quantify the relative haze concentration at each pixel location. Pixels with lower dark channel intensity—indicative of thinner haze—receive greater weights, indicating that clearer regions contribute more dominantly to the fusion output.

The image fusion step sums over these weighted inputs to generate the final dehazed image (J), combining the enhanced local contrast of the under-exposed variants with the original color fidelity. This process avoids explicit transmission estimation, enabling a lightweight and fully model-free solution well-suited for parallel processing.

$$W_i = 1 - \min_{3 \times 3} \left(\min_{c \in \{R, G, B\}} I_i^c \right), \quad (5)$$

$$\bar{W}_i = \frac{W_i}{\sum_{i=1}^4 W_i}, \quad (6)$$

$$J = \sum_{i=1}^4 \bar{W}_i I_i. \quad (7)$$

3.2. Haze Density Estimation and Interpolation

This section describes the haze-density estimation and interpolation process, corresponding to the haziness-degree-evaluator and interpolation modules in Figure 1. The goal is to derive a spatially coherent haze-density field that subsequently guides the image blending and tone-remapping stages. Section 3.2.1 introduces the referenceless haze-density estimator, and Section 3.2.2 details the smoothing and interpolation procedure used to generate the final haze-density map.

3.2.1. Referenceless Haze-Density Estimation

To achieve autonomous and spatially adaptive dehazing, the system estimates haze density directly from the input image without relying on physical haze parameters or reference data. The resulting density map governs both local blending and global tone remapping. A referenceless evaluator [32] computes the haze-density map ρ_I as

$$\rho_I = 1 - \hat{t}, \quad (8)$$

where \hat{t} is obtained by minimizing a cost function $O(t)$, which is formulated as

$$O(t) = \frac{S(t)V(t)\sigma(t)}{D(t)} + \frac{\lambda}{t}, \quad (9)$$

where $S(t)$, $V(t)$, and $\sigma(t)$ represent the saturation, brightness, and sharpness of the image as functions of the transmission map t , while $D(t)$ denotes the dark channel. The parameter λ provides regularization to prevent overestimation of t . Minimizing $O(t)$ identifies \hat{t} such that the recovered scene exhibits maximum colorfulness, brightness, and detail clarity, while the dark channel is simultaneously suppressed.

To maintain conciseness, detailed derivations are omitted here; interested readers are referred to [32] (Section 3.4 and Appendix A) for the complete mathematical formulation. The final expression for computing the haze-density map is given by

$$\rho_I = I_{m\Psi} + \frac{I_{mc}v}{\lambda} - \sqrt{\frac{I_{mc}v}{\lambda} \left(\frac{I_{mc}v}{\lambda} - 255 + I_{m\Psi} \right)}, \quad (10)$$

where $I_{m\Psi} = \min_{(x,y) \in \Psi} \left[\min_{c \in \{R,G,B\}} I^c(x,y) \right]$ is the minimum intensity within a local patch Ψ , and $I_{mc} = \max_{c \in \{R,G,B\}} I^c - \min_{c \in \{R,G,B\}} I^c$ is the color difference between the maximum and minimum channel intensities. The term v denotes the local luminance variance, defined earlier in Equation (4).

3.2.2. Haze Density Map Interpolation

The estimated haze-density map often exhibits spatial discontinuities caused by abrupt variations in haze levels between neighboring regions. To accommodate spatial heterogeneity while ensuring smooth visual transitions, the input frame is first divided into non-overlapping 8×8 patches. The local haze density ρ_{Ω_i} for the i -th patch Ω_i is then computed as

$$\rho_{\Omega_i} = \max \left(\bar{\rho}_I, \frac{1}{|\Omega_i|} \sum_{(x,y) \in \Omega_i} \rho_I(x,y) \right), \quad (11)$$

where $|\Omega_i|$ denotes the number of pixels in Ω_i and $\bar{\rho}_I$ is the average haze density of the input frame.

Figure 4 illustrates that using raw 8×8 local weights can lead to blocky artifacts in the blended output, especially at patch boundaries where haze densities vary sharply. The left image in Figure 4 visualizes the patch-based haze densities, where the highlighted region shows abrupt transitions from 0.1531 to zero. When these coarse values are directly applied for blending (right image), visible discontinuities appear as the pink-outlined blocks.

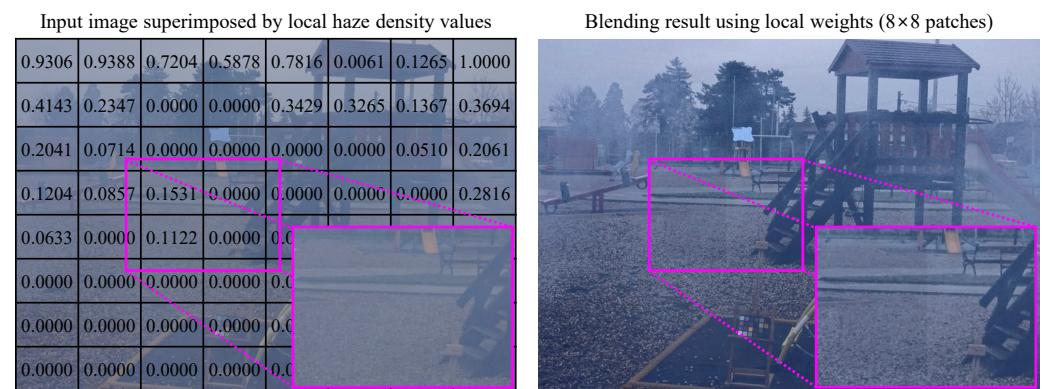


Figure 4. Patch-level haze densities and resulting block artifacts when used directly for blending.

To mitigate these artifacts, a 2×2 low-pass filter is first applied to attenuate abrupt transitions, followed by $4 \times$ bilinear interpolation to upscale the local haze densities from 8×8 patches to an effective resolution of 29×29 . This interpolation process preserves regional haze characteristics while generating a spatially continuous density field suitable for image blending.

As shown in Figure 5, the proposed interpolation strategy significantly reduces blocky artifacts by producing more gradual transitions across adjacent patches. In the same region highlighted in Figure 4, the coarse 2×2 haze densities are smoothed into finer 5×5 variations, yielding a perceptually natural blending behavior in both horizontal and vertical directions. Figures 4 and 5 are adopted from our previous work in [33].

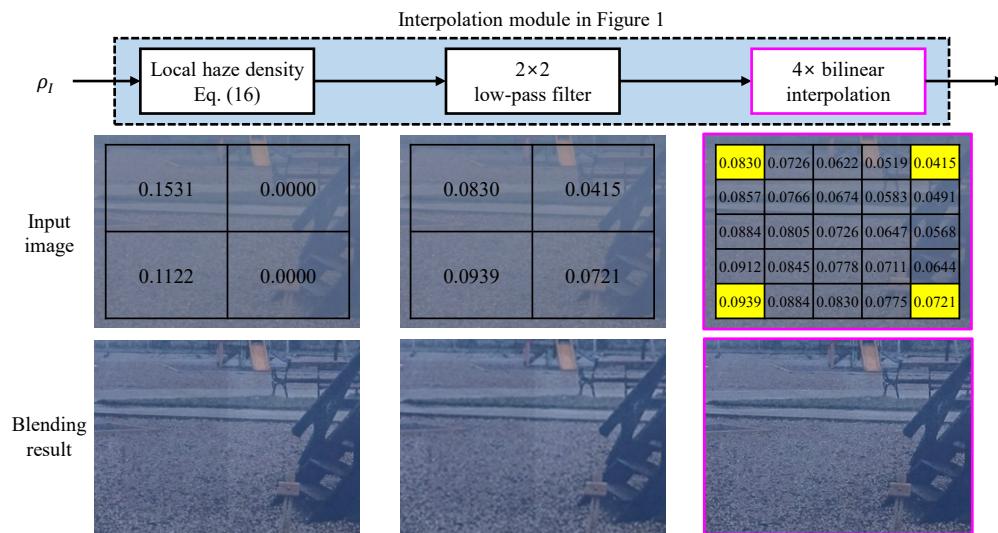


Figure 5. Effect of interpolation. The refined haze-density map yields smoother fusion weights and eliminates blocky artifacts.

3.3. Haze-Aware Image Blending

This section describes the haze-aware image blending process, which corresponds to the local-blending-weight and image-blending modules in Figure 1. Using the interpolated haze-density map, the system computes blending weights and applies them to fuse the input image with its dehazed estimate. Section 3.3.1 details the local blending weight calculation, and Section 3.3.2 presents the image-blending procedure used to attain autonomous dehazing.

3.3.1. Local Blending Weight Calculation

The local blending weight α_i is derived from the interpolated haze-density map ρ_{INT} obtained in Section 3.2.2. Based on perceptual analysis, pixels with higher haze density should rely more heavily on the dehazed result, while clearer regions should retain more of the original content. To reflect this principle, α_i is formulated as follows:

$$\alpha_i = \begin{cases} 0 & \rho_{INT} < \rho_1 \\ \frac{\bar{\rho}_I - \rho_1}{\rho_2 - \rho_1} & \rho_1 \leq \rho_{INT} \leq \rho_2, \\ 1 & \rho_{INT} > \rho_2 \end{cases} \quad (12)$$

where $\rho_1 = 0.8811$ and $\rho_2 = 0.9344$ represent lower and upper haze-density thresholds for classifying haze levels, and $\bar{\rho}_I$ denotes the average global haze density computed from the entire frame. These threshold values are adopted from our previous work on hazy versus haze-free image classification [32].

This design yields a smooth transition of blending ratios across different haze conditions. Specifically, when $\rho_{INT} > \rho_2$, the patch is heavily hazy, and the dehazed result fully dominates ($\alpha_i = 1$). Intermediate haze levels are handled through linear interpolation to ensure continuity between neighboring patches. The adaptive weighting mechanism enables the system to autonomously balance enhancement strength according to haze densities, without manual control or external inputs.

3.3.2. Image Blending for Autonomous Dehazing

The input image I and its dehazed estimate J are fused using local blending weights α derived in the previous section. This process produces the blended result B according to

$$B = \alpha J + (1 - \alpha)I. \quad (13)$$

The rationale for this blending formulation is threefold.

- **Haze-free regions:** When the input region is haze-free, applying dehazing unnecessarily can amplify noise or introduce color distortion. In such cases, $\alpha = 0$ ensures that the output remains identical to the original image, thereby preserving natural appearance.
- **Mild to moderate haze:** For regions with partial haze, excessive dehazing may lead to over-enhancement or halo artifacts. To mitigate this problem, α varies smoothly between 0 and 1, proportionally to the local haze density, allowing gradual adjustment of dehazing strength.
- **Densely hazy regions:** When the haze density is high, full dehazing is necessary to restore structural visibility and contrast. Here, $\alpha = 1$ fully prioritizes the dehazed image and suppresses the original content.

This autonomous blending strategy allows continuous adaptation across the entire image without manual intervention or scene-specific tuning.

3.4. Adaptive Tone Remapping

This section presents the adaptive tone remapping (ATR) process, corresponding to the adaptive-tone-remapping and self-calibrating-weight modules in Figure 1. Using the fused image B and the estimated global haze density $\bar{\rho}_I$, the system adjusts luminance and chrominance while regulating the enhancement strength to prevent over-correction. Section 3.4.1 details the luminance and chrominance adjustment procedure, and Section 3.4.2 introduces the self-calibrating weight that controls the enhancement strength.

3.4.1. Luminance and Chrominance Adjustment

ATR operates in two sequential phases: (i) luminance enhancement, which compensates for brightness loss caused by haze removal, and (ii) chrominance expansion, which restores the natural color gamut narrowed by luminance correction. The enhancement strength is automatically adjusted using the self-calibrating weight (Γ , introduced in Section 3.4.2), which scales the tone enhancement proportionally to the detected haze level. Therefore, haze-free regions remain unaltered, while regions with heavier haze undergo progressively stronger enhancement.

Denoting the luminance and chrominance components of the image before and after ATR as $\{Y_B, \text{Ch}_B\}$ and $\{Y_{B_f}, \text{Ch}_{B_f}\}$, respectively, ATR can be described as

$$Y_{B_f} = Y_B + \Gamma \cdot g_1(Y_B) \cdot g_2(Y_B), \quad (14)$$

$$\text{Ch}_{B_f} = \text{Ch}_B \left[1 + \frac{Y_{B_f}}{Y_B} \cdot g_3(Y_B) \right], \quad (15)$$

where $g_1(\cdot)$, $g_2(\cdot)$, and $g_3(\cdot)$ denote the nonlinear luminance gain, linear weighting, and chrominance expansion functions, respectively.

The luminance enhancement module combines a nonlinear gain and a linear weight to emphasize structure while preventing over-amplification in bright regions. The nonlinear gain $g_1(\cdot)$ constrains enhanced luminance using the Adaptive Luminance Point (ALP), which is expressed as

$$g_1(Y_B) = \frac{Y_B}{2^{21}} \left[255 \left(1 - \frac{Y_B - \text{ALP}}{255} \right)^\theta \left(\frac{255 - Y_B}{255} \right) \right]^2, \quad (16)$$

$$\text{ALP} = \begin{cases} 0.04 + \frac{0.02}{255} (L_{0.9} - L_{0.1}) & \bar{Y}_B > 128 \\ 0.04 - \frac{0.02}{255} (L_{0.9} - L_{0.1}) & \bar{Y}_B \leq 128 \end{cases}, \quad (17)$$

where the user-defined exponent θ tunes enhancement aggressiveness, \bar{Y}_B represents the average luminance, and L_k denotes the luminance value where the cumulative distribution function $\text{CDF}(L_k) = k$. The linear weighting function $g_2(\cdot)$ modulates enhancement depending on input brightness as

$$g_2(Y_B) = \frac{m}{255} Y_B + b, \quad (18)$$

with slope m and intercept b empirically determined for perceptual balance between dark and bright tones.

While luminance stretching improves brightness, it can inadvertently reduce colorfulness due to the Helmholtz–Kohlrausch effect. To counter this effect, chrominance signals are expanded according to local luminance. The ratio Y_{B_f}/Y_B in Equation (15) provides self-calibration, while $g_3(\cdot)$ follows a piecewise linear mapping:

$$g_3(Y_B) = \begin{cases} 0.7 & Y_B < L_{\text{low}} \\ 0.7 - 0.26 \left(\frac{Y_B - L_{\text{low}}}{L_{\text{high}} - L_{\text{low}}} \right) & L_{\text{low}} \leq Y_B \leq L_{\text{high}}, \\ 0.44 & Y_B > L_{\text{high}} \end{cases} \quad (19)$$

where L_{low} and L_{high} are predefined luminance thresholds. This mapping enhances saturation in dark areas and suppresses excessive color shifts in bright regions, maintaining a visually natural tone distribution.

3.4.2. Self-Calibrating Weight

ATR relies on a self-calibrating weight Γ that modulates enhancement strength according to the overall haze level present in the input image. This weight is formulated as a piecewise function of $\bar{\rho}_I$:

$$\Gamma = \begin{cases} 0 & \bar{\rho}_I \leq \rho_1 \\ \left(\frac{\bar{\rho}_I - \rho_1}{\rho_2 - \rho_1} \right)^n & \rho_1 < \bar{\rho}_I \leq \rho_2, \\ \left(\frac{\Gamma_u - 1}{1 - \rho_2} \right) (\bar{\rho}_I - \rho_2) + 1 & \bar{\rho}_I > \rho_2 \end{cases} \quad (20)$$

where ρ_1 and ρ_2 are user-defined haze-density thresholds, n controls the exponential response, and Γ_u denotes the maximum enhancement factor. As mentioned in Section 3.3.1, $\rho_1 = 0.8811$ and $\rho_2 = 0.9344$ are adopted from our previous work [32].

When the average haze density $\bar{\rho}_I$ is lower than ρ_1 , the image is considered haze-free, and $\Gamma = 0$ disables luminance enhancement, preserving the original appearance. For mild haze levels ($\rho_1 < \bar{\rho}_I \leq \rho_2$), Γ increases smoothly following an exponential profile governed by exponent n (empirically set to 0.1), thereby enabling gradual strengthening of luminance enhancement. When haze density exceeds ρ_2 , Γ transitions to a linear region that caps the enhancement magnitude at $\Gamma_u = 1.2$ to prevent excessive luminance boosts.

3.5. Hardware Accelerator

This section presents the hardware accelerator that implements the proposed dehazing pipeline. All modules are designed using a fully pipelined methodology to sustain high-throughput, pixel-stream processing. Resource usage is minimized through fixed-point arithmetic, where the word length of each signal is carefully selected to ensure that the final output error remains within ± 1 least significant bit (LSB). Figure 6 illustrates the overall hardware design flow, from the algorithm specification to hardware description and verification on the target FPGA.

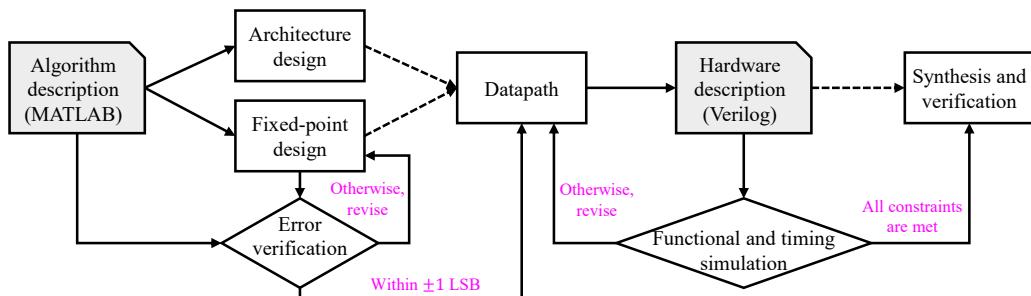


Figure 6. Overview of the hardware design procedure. Dashed arrows indicate that information is forwarded only when the associated condition is satisfied.

Special consideration is given to the interpolation module: As the haze-density map of the current frame becomes available only after two frames, the module is scheduled to operate during the video blanking interval, effectively reducing the latency to a single frame. The high temporal similarity between consecutive frames allows this design choice without affecting output quality.

Due to the complexity of the full system, only simplified datapaths are shown for each module. These diagrams illustrate the intra-module dataflow consistent with RTL (Register Transfer Level)-oriented hardware description and provide a clear view of how each component is realized in hardware. Section 3.5.1 through Section 3.5.4 detail the implementation of all major modules in the accelerator.

3.5.1. Fusion-Based Image Dehazing Module

Figure 7 presents the simplified architecture of the fusion-based dehazing module, consisting of two submodules: (a) detail enhancement and (b) gamma correction, weight calculation, and fusion.

The detail-enhancement submodule (Figure 7a) converts the input RGB stream to the YCbCr domain and enhances only the luminance channel. The enhanced Y component is then recombined with the original Cb and Cr channels to produce the locally enhanced color image. This stage requires two-dimensional filtering, with implementation details provided in Appendix A.

The gamma-correction and fusion submodule (Figure 7b) generates three under-exposed variants of the enhanced image using LUTs corresponding to $\gamma \in \{1.9, 1.95, 2.0\}$. A minimum filter computes the dark-channel response, which is processed through an adder-tree structure to obtain normalized fusion weights \bar{W}_i . Each variant is then multiplied by its corresponding weight, and the weighted results are accumulated through adder trees to produce the final dehazed RGB output $J = \{J^R, J^G, J^B\}$.

An important optimization is applied to reduce hardware cost during weight calculation. As gamma correction is a monotonically increasing operation, it preserves the relative ordering of pixel intensities. As illustrated in Figure 8, the straightforward design—which applies gamma correction first and computes dark-channel responses afterward—requires four minimum filters. By instead applying gamma correction directly to the dark channel

of the input, only one minimum filter is needed. As spatial minimum filters are resource-intensive, this design choice yields a substantial reduction in memory and logic utilization while maintaining the correctness of the weight computation.

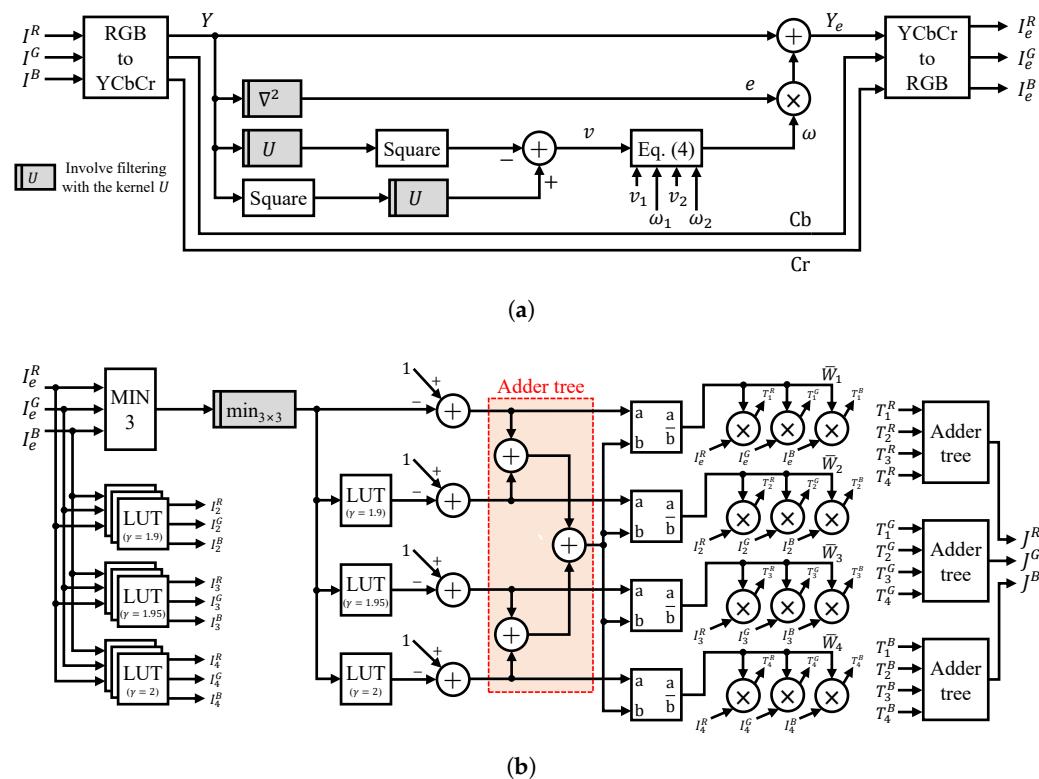


Figure 7. Simplified datapath of the fusion-based dehazing module. As the complete datapath cannot be fully illustrated in a single diagram, it is divided into two parts: (a) detail enhancement and (b) gamma correction, weight calculation, and fusion.

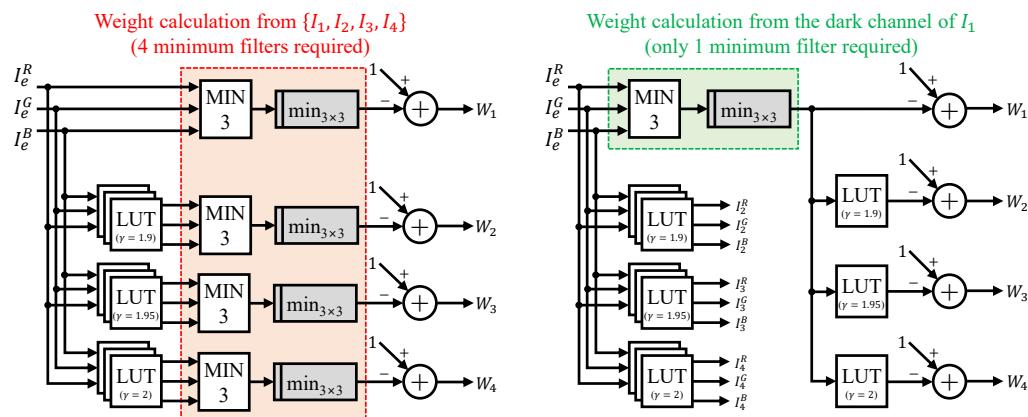


Figure 8. Reduction of minimum-filter resources in weight computation.

3.5.2. Haziness-Degree-Evaluator and Interpolation Modules

The haziness-degree-evaluator and interpolation modules are implemented as a unified, streamlined datapath. As shown in Figure 9, the architecture realizes the operations described in Sections 3.2.1 and 3.2.2.

The first stage computes the local minimum and maximum RGB values to measure color range. These values, together with the local variance v and the inverse regularization factor $\frac{1}{\lambda}$, are used to evaluate the referenceless haze-density model in Equation (11). A square-root unit then produces the pixel-wise haze density ρ_I , which is accumulated and

averaged over the frame to obtain the global haze indicator $\bar{\rho}_I$. This global value acts as a scene-level prior for subsequent tone-remapping and blending stages.

To ensure spatial smoothness and avoid block artifacts, the pixel-wise haze map is processed by a 2×2 low-pass filter followed by a $4 \times$ bilinear interpolator. Both operations are scheduled during the video blanking interval, enabling efficient temporal reuse without reducing throughput in the active video region. The resulting interpolated haze-density map ρ_{INT} provides the continuous spatial guidance required for haze-aware blending in the fusion stage.

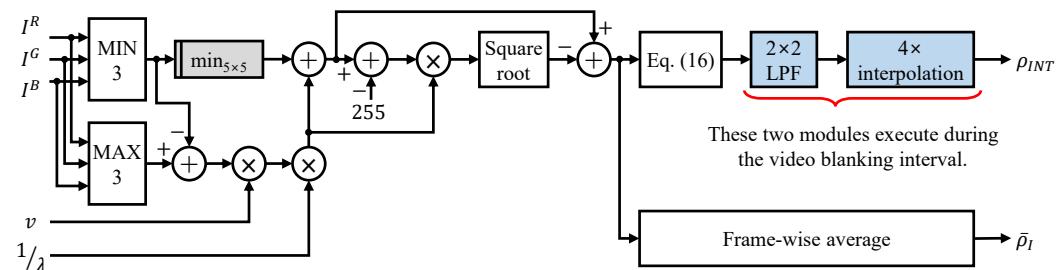


Figure 9. Simplified datapath of the haziness-degree-evaluator and interpolation modules. LPF is the abbreviation of low-pass filter, $\bar{\rho}_I$ denotes the average haze-density value, and ρ_{INT} represents the interpolated haze-density map.

3.5.3. Local-Blending-Weight and Image-Blending Modules

The local-weight calculation and image-blending operations are implemented using a fully pipelined datapath, as shown in Figure 10. In the first stage, the local blending weight α_i is computed according to Equation (12). The interpolated haze-density value ρ_{INT} is compared with thresholds ρ_1 and ρ_2 using a comparator block, and the control logic selects the appropriate branch of the piecewise linear function.

The second stage performs the blending operation defined in Equation (13), where $I = \{I^R, I^G, I^B\}$ and $J = \{J^R, J^G, J^B\}$ denote the original and dehazed RGB channels. Each channel is processed independently through parallel multipliers applying α_i and $1 - \alpha_i$, followed by adders that combine the weighted results to produce the final blended output $B = \{B^R, B^G, B^B\}$.

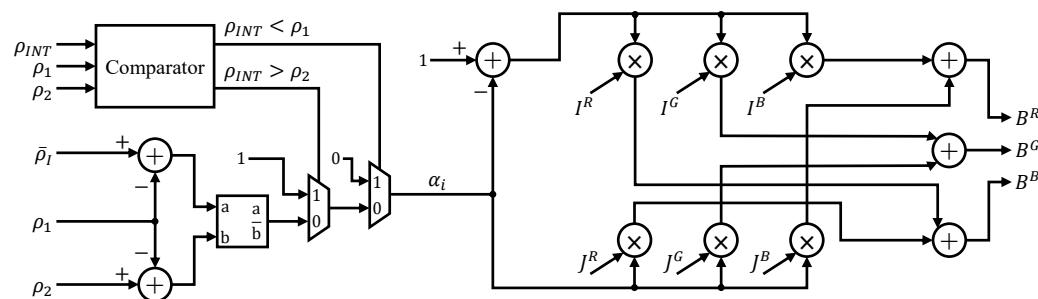


Figure 10. Simplified datapath of local-blending-weight and image-blending modules.

3.5.4. Adaptive-Tone-Remapping and Self-Calibrating-Weight Modules

The hardware implementation of the self-calibrating-weight and ATR modules is shown in Figure 11, where both components are integrated into a unified, fully pipelined datapath.

The pipeline begins by converting the blended RGB input $\{B^R, B^G, B^B\}$ into the YCbCr domain using a 4:2:2 format. The resulting luminance (Y_B) and chrominance (Ch_B) signals are then processed in two parallel branches:

- **Self-Calibrating Weight Generation:** The global average luminance \bar{Y}_B is obtained through an averaging block, while a CDF-based estimator computes the per-

centile luminance levels $L_{0.1}$ and $L_{0.9}$ for ALP (Equation (17)). The average haze density $\bar{\rho}_I$ is processed through Equation (19) to generate the self-calibrating weight Γ , which adaptively modulates the overall tone-remapping strength according to scene haze conditions.

- **Adaptive Tone Remapping Pipeline:** The luminance channel Y_B passes through the nonlinear and linear gain units $g_1(\cdot)$ and $g_2(\cdot)$ (Equations (16) and (18)). Their outputs are scaled by Γ and combined to produce the enhanced luminance Y_{B_f} . In parallel, the chrominance components are adjusted using the chrominance-scaling function $g_3(\cdot)$ and the luminance ratio Y_{B_f}/Y_B (Equation (15)), implementing chrominance expansion.

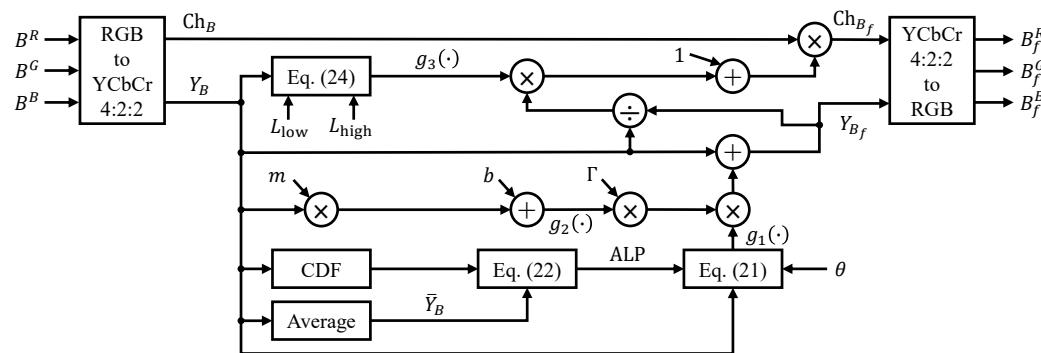


Figure 11. Simplified datapath of the self-calibrating-weight and adaptive-tone-remapping modules. The cumulative distribution function is abbreviated as CDF.

Finally, the enhanced YCbCr signals $\{Y_{B_f}, Ch_{B_f}\}$ are converted back to RGB to yield the adaptively tone-mapped output $\{B_f^R, B_f^G, B_f^B\}$.

4. Evaluation

This section presents a comprehensive evaluation of the proposed dehazing framework and hardware accelerator. Section 4.1 reports the quantitative results on standard benchmarks. Sections 4.2 and 4.3 provide qualitative comparisons on natural and aerial images, demonstrating the method's visual effectiveness across diverse haze conditions. Section 4.4 summarizes the hardware implementation results, highlighting throughput, resource usage, and real-time performance.

4.1. Quantitative Evaluation

To objectively assess the performance of the proposed system, quantitative evaluations were performed on five widely used public datasets: FRIDA2 [34], D-HAZY [35], O-HAZE [36], I-HAZE [37], and Dense-Haze [38]. Table 2 summarizes the characteristics of these datasets, including the numbers of haze-free and hazy image pairs and whether the image scenes are synthetic or real. The evaluation compares the proposed method with five representative algorithms: DCP [1], CAP [2], DehazeNet [18], YOLY [39], and MB-TaylorFormer [20]. For a fair comparison, FCDM [6] was excluded owing to its behavior of resizing input images into square dimensions, whereas other methods process variable-sized images.

Table 2. Summary of the five public datasets used in the quantitative evaluation. The # symbol denotes quantities.

Dataset	Haze-Free (#)	Hazy (#)	Remark
FRIDA2	66	264	Synthetic road scene images
D-HAZY	1472	1472	Synthetic indoor images
O-HAZE	45	45	Real outdoor images
I-HAZE	30	30	Real indoor images
Dense Haze	50	50	Real indoor and outdoor images

Two complementary image quality measures were adopted:

- **The TMQI (Tone-Mapped image Quality Index)** quantifies the structural fidelity and naturalness of tone-mapped or dehazed images with respect to their reference haze-free ground truths. A higher TMQI value indicates better perceptual restoration and tone consistency.
- **The FSIMc (Feature Similarity Index for Color Images)** evaluates the structural and chrominance correspondence between restored images and their references. This metric is particularly sensitive to preservation of edge details and color relationships, making it suitable for assessing dehazed image quality.

Table 3 reports the average TMQI and FSIMc values computed over all datasets. The best and second-best results for each dataset are marked in bold and italic, respectively.

Table 3. Average TMQI and FSIMc values computed on five public datasets. The best and second-best results are boldfaced and italicized, respectively. MB-TF is the abbreviation for MB-TaylorFormer.

Dataset \ Method	DCP	CAP	DehazeNet	YOLY	MB-TF	Proposed
TMQI \uparrow	FRIDA2	0.7291	0.7385	0.7366	0.7176	0.7631
	D-HAZY	0.8631	0.8206	0.7966	0.6817	0.7428
	O-HAZE	0.8403	0.8118	0.8413	0.6566	0.8732
	I-HAZE	0.7319	0.7512	0.7598	0.6936	0.8655
	Dense-Haze	0.6383	0.5955	0.5723	0.5107	0.7237
		Total	0.7357	0.7336	0.7312	0.6520
FSIMc \uparrow	FRIDA2	0.7746	0.7918	0.7963	0.7849	0.7158
	D-HAZY	0.9002	0.8880	0.8874	0.7383	0.7727
	O-HAZE	<i>0.8423</i>	0.7738	0.7865	0.6997	0.8420
	I-HAZE	0.8208	0.8252	0.8482	0.7564	0.8692
	Dense-Haze	0.6419	0.5773	0.5573	0.5763	0.7976
		Total	0.7746	0.7693	0.7725	0.7111
					0.7544	0.7883

For the TMQI, the proposed method achieves competitive or superior performance across most datasets, with a particularly strong improvement on O-HAZE (0.9102), surpassing both prior-based and deep-learning methods. The method also maintains stable performance on challenging D-HAZY and I-HAZE datasets, where the integration of referenceless haze estimation and adaptive tone remapping contributes to preserving structural integrity and visual naturalness. Averaged across all datasets, the proposed system attains an overall TMQI of 0.7496, which is ranked second after the deep transformer-based MB-TaylorFormer (0.7761).

For the FSIMc, the proposed system demonstrates consistent color and structural fidelity, yielding an average FSIMc of 0.7883, the highest among all compared methods. Notably, on FRIDA2, O-HAZE, and I-HAZE, the proposed method exhibits top-ranked

scores, reaffirming its strength in maintaining global contrast and fine-grained color texture. The improved FSIMc over the DCP and CAP confirms that the proposed haze-aware blending and chrominance expansion effectively enhance perceptual color quality without introducing artifacts.

Overall, the quantitative analysis indicates that the proposed system achieves a favorable trade-off between visual fidelity (high TMQI) and structural/color preservation (high FSIMc) while maintaining real-time processing capability (discussed in Section 4.4). The relatively consistent performance across synthetic and real datasets demonstrates that the system generalizes well to diverse imaging conditions.

4.2. Qualitative Evaluation on Natural Images

The visual performance of the proposed dehazing framework was compared against six representative algorithms: the DCP, CAP, DehazeNet, YOLY, MB-TaylorFormer, and FCDM. Figure 12 shows the restoration results across diverse haze conditions, including mild, moderate, and dense haze, haze-free scenes, and challenging failure cases.

The proposed system adapts its enhancement strength using the self-calibrating weight Γ , derived from the average haze density $\bar{\rho}_I$. In this evaluation, thresholds were set to $\rho_1 = 0.8811$ and $\rho_2 = 0.9344$.

In mildly and moderately hazy scenes, the method improves visibility and color contrast while avoiding over-enhancement, halo artifacts, and aggressive tone shifts. Prior-based methods such as the DCP and CAP tend to introduce darkened shadows or overestimated contrast, while deep models may over-brighten distant regions or alter global color balance.

Under dense haze conditions, the proposed approach recovers structural detail and maintains consistent color appearance. The combination of haze-aware blending and adaptive tone remapping enables effective suppression of veiling without color drifting or loss of contrast. Competing approaches often struggle with clipped highlights, oversuppression, and inconsistent chromaticity.

For haze-free inputs, the proposed system correctly disables dehazing operations ($\Gamma = \alpha = 0$), preserving the original brightness and color fidelity. This behavior is in contrast to other approaches, which frequently alter clear scenes due to fixed assumptions about haze presence.

In failure cases involving extreme lighting or reflective surfaces, the method produces stable, visually coherent results without introducing strong artifacts or color distortions. This robustness reflects the advantages of the haze-density-driven, model-free fusion design.

Overall, the qualitative results show that the proposed system achieves a strong balance of haze removal, color naturalness, and stable tone reproduction across a wide range of natural images.

4.3. Qualitative Evaluation on Aerial Images

To further assess generalization across large-scale environments, the method was evaluated using aerial images exhibiting spatially extensive haze. Figure 13 compares the proposed framework with the same set of benchmarks under haze-free, mildly hazy, moderately hazy, and densely hazy conditions, using $\bar{\rho}_I$ and thresholds $\rho_1 = 0.8811$ and $\rho_2 = 0.9344$ for classification.

In haze-free scenes, traditional prior-based and learning-based algorithms frequently modify colors or brightness unnecessarily. The proposed framework preserves the original image by suppressing dehazing when haze density is negligible.

For mild haze, the method restores clarity while maintaining natural tone reproduction. Edge structures such as runway markings and aircraft contours remain consistent without haloing or color oversaturation. Deep learning methods improve detail but occasionally introduce exaggerated saturation.

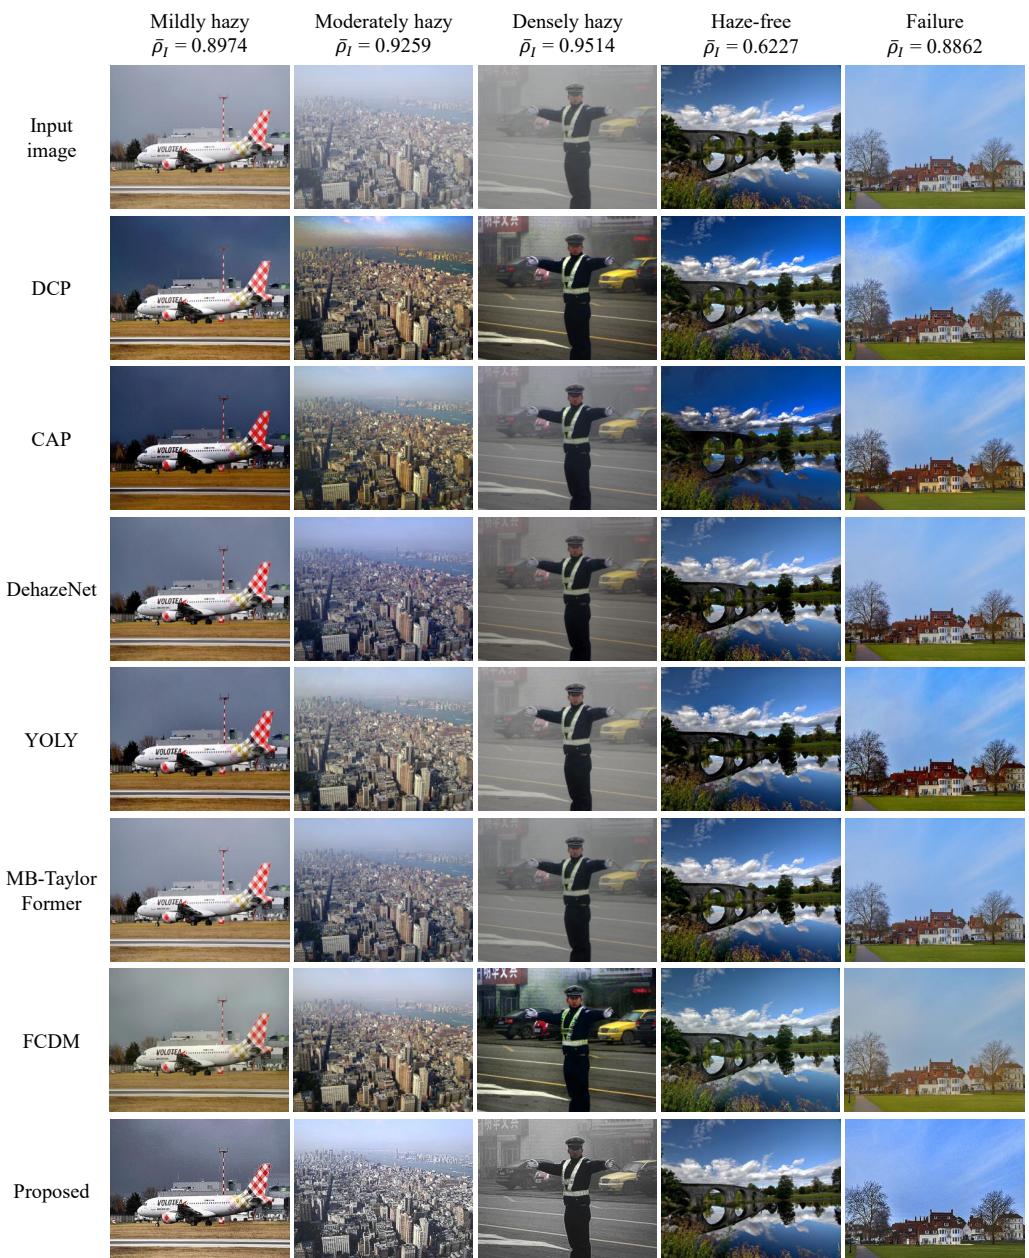


Figure 12. Qualitative comparisons on natural images with varying haze densities. The average haze density $\bar{\rho}_I$ was compared against two thresholds, $\rho_1 = 0.8811$ and $\rho_2 = 0.9344$, to determine the haze condition.

Under moderate haze conditions, the proposed system provides a balanced enhancement of structure and color. Prior-based methods often produce darkened outputs, while some learned models fade global color or produce unnatural tints.

In dense haze, the DCP and FCDM deliver strong dehazing performance but noticeably alter color appearance. In contrast, the proposed framework removes haze while preserving overall brightness and color fidelity.

These results demonstrate that the proposed approach generalizes well to aerial environments, benefiting from both the referenceless haze estimator and the adaptive blending mechanism.

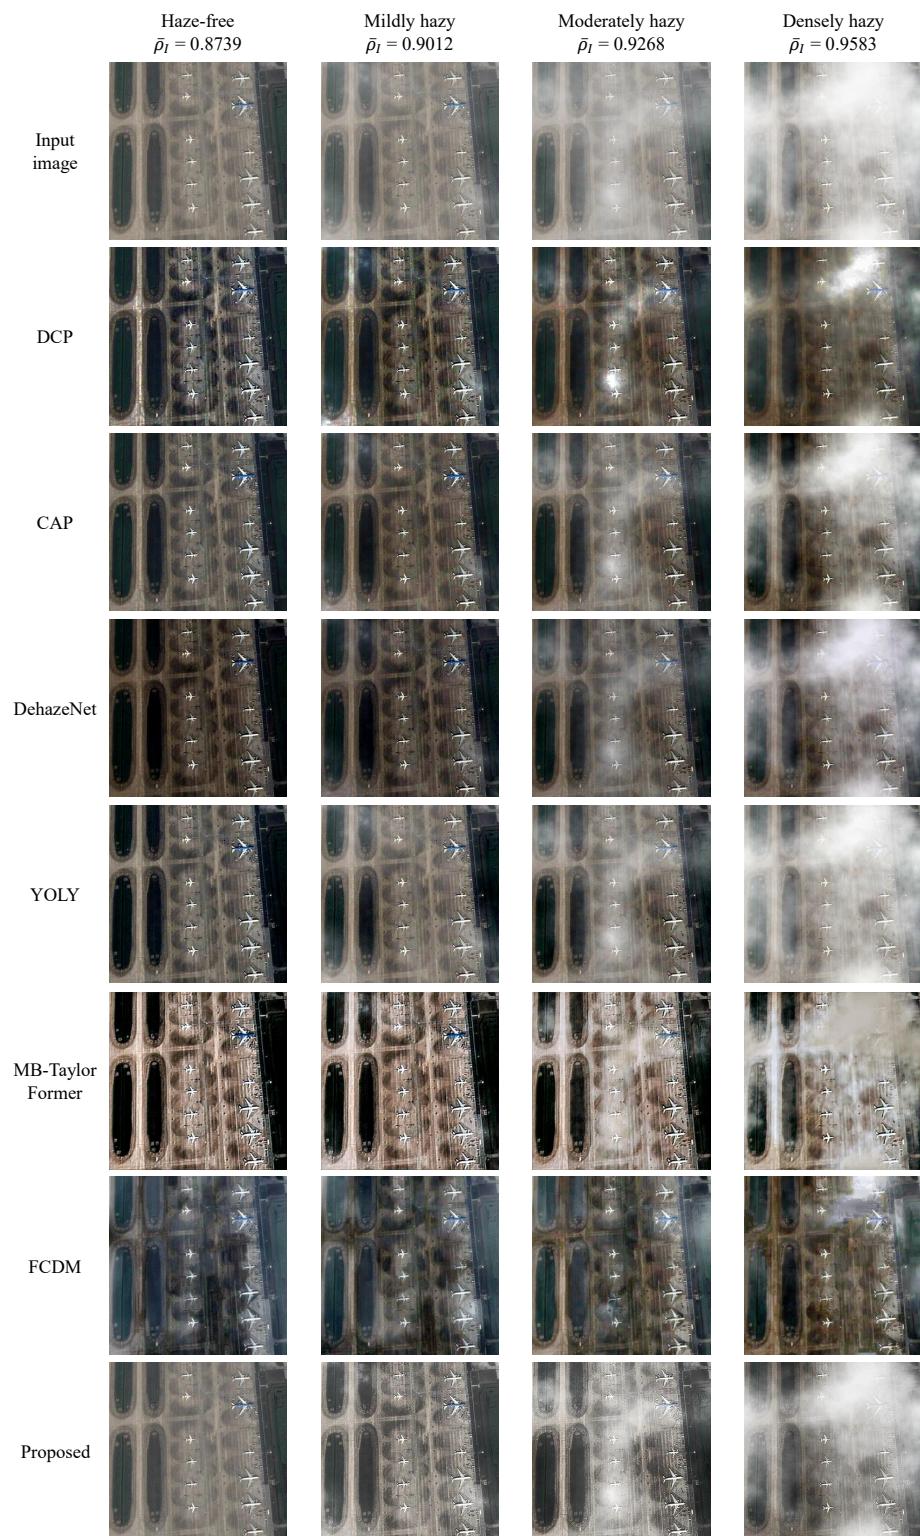


Figure 13. Qualitative comparisons on aerial images across haze levels determined by $\bar{\rho}_I$, $\rho_1 = 0.8811$, and $\rho_2 = 0.9344$.

4.4. Hardware Implementation Results

The proposed autonomous dehazing system was implemented using Verilog HDL following the IEEE Std 1364-2005 specification [40]. The hardware resource utilization on a Xilinx XC7Z-045FFG900-2 MPSoC, designed by AMD (Santa Clara, CA, USA) and manufactured by TSMC (Taichung, Taiwan), device was obtained using Vivado v2023.1. The design occupies 12.33% of slice registers, 26.14% of slice LUTs, and a modest portion of on-chip memory resources (8.26% of 36 Kb RAMs and 2.11% of 18 Kb RAMs). The minimum clock period achieved is 3.68 ns, corresponding to a maximum operating frequency of 271.74 MHz. These utilization rates demonstrate that the system can be seamlessly integrated as a pre- or post-processing module within larger heterogeneous vision pipelines without exceeding typical mid-range FPGA resource budgets.

The maximum processing throughput, expressed in frames per second (fps), can be estimated as

$$\text{FPS} = \frac{f_{\max}}{(H + B_V)(W + B_H)}, \quad (21)$$

where f_{\max} is the maximum frequency and $\{H, W\}$ are the frame height and width; $\{B_V, B_H\}$ represent vertical and horizontal blanking intervals, respectively. The proposed system is designed to operate with the minimum blanking intervals, $B_V = B_H = 1$; thus, substituting the measured f_{\max} into Equation (21) yields the performance for different video resolutions, as summarized in Table 4.

The proposed system achieves 130.86 fps for Full HD (1920×1080) input and 73.63 fps for Quad HD (2560×1440). Even at ultra-high definitions such as UW4K and DCI 4K, real-time performance is maintained with 44.19 fps and 30.69 fps, respectively. These frame rates far exceed the standard 30 fps real-time threshold, confirming the capability of the system to handle high-resolution and low-latency imaging tasks efficiently.

Table 4. Maximum processing speeds in frames per second for different video standards. The # symbol denotes quantities.

Standard	Resolution	Required Clock Cycles (#)	Processing Speed (fps)
Full HD	1920×1080	2,076,601	130.86
Quad HD	2560×1440	3,690,401	73.63
UW4K	3840×1600	6,149,441	44.19
4K UHD TV	3840×2160	8,300,401	32.74
DCI 4K	4096×2160	8,853,617	30.69

The hardware comparison results in Table 5 show clear advantages of the proposed system over three representative real-time FPGA implementations: DCP with Fast Airlight Estimation (DCP-FAE [26]), Fusion-Based Dehazing (FBD [17]), and Saturation-Based Dehazing (SBD [29]). The maximum video resolution reported in the table refers to the highest input resolution each accelerator can sustain at more than 30 fps.

SBD achieves the smallest resource usage but only because its algorithmic design is extremely simple. It estimates airlight from a downsampled version of the input—where the downsampling itself is performed externally—and does not include any meaningful dehazing beyond this operation. Thus, SBD prioritizes speed and minimal resource consumption at the cost of restoration quality.

FBD, on the other hand, illustrates a suboptimal hardware design. It implements image filters using shift registers instead of line memories, leading to excessive logic utilization and a fixed-resolution architecture restricted to 480×270 . Because every resolution change requires re-synthesis, FBD is impractical for real-world deployment.

DCP-FAE is more capable than SBD and FBD, but it remains tied to the DCP, which degrades severely in scenes with large sky regions or bright objects. It also requires more on-chip memory than the proposed system, despite offering lower resolution support.

In contrast, the proposed system delivers the highest throughput with balanced resource usage and is the only design capable of real-time autonomous dehazing at full DCI-4K (4096×2160) resolution. Unlike previous FPGA accelerators—primarily fixed-heuristic, prior-based systems operating at SVGA or FHD resolutions—the proposed architecture uniquely integrates haze-density estimation, adaptive fusion, and self-calibrating tone control into a fully hardware-friendly pipeline. This combination enables true self-adaptation to varying haze levels without manual tuning or physical-model assumptions.

Collectively, these results demonstrate that the proposed design establishes a new performance and capability benchmark among FPGA-based dehazing systems, combining algorithmic autonomy, high-resolution scalability, and efficient hardware utilization.

Table 5. Comparison with other real-time dehazing implementations. The # symbol denotes quantities and NA stands for not available.

Hardware Utilization	DCP-FAE [26]	FBD [17]	SBD [29]	Proposed System
Slice registers (#)	53,400	12,210	547	53,901
Slice LUTs (#)	64,000	79,042	1537	57,146
DSPs (#)	42	69	NA	0
Memory (Mbits)	3.20	NA	NA	1.41
Maximum frequency (MHz)	88.70	108.45	85.20	271.74
Maximum video resolution	SVGA	480×270	FHD	DCI 4K
Autonomous dehazing	Unequipped	Unequipped	Unequipped	Equipped

5. Conclusions

This paper introduced a fast, model-free dehazing framework that departs from traditional transmission estimation and physical-model inversion. The core novelty lies in its haze-density-driven design: a referenceless haze-density estimator provides both local and global guidance for fusion and tone remapping, enabling autonomous adaptation across haze-free, mildly hazy, and densely hazy scenes. Unlike prior fusion-based methods that rely on fixed heuristics, the proposed pipeline incorporates spatially continuous haze-density maps, multi-variant exposure synthesis, and self-calibrated tone control, forming a principled and fully data-independent strategy for robust restoration.

On the hardware side, the paper presented a deeply pipelined, resource-efficient accelerator tailored to the structure of the proposed algorithm. Key architectural innovations include the monotonicity-based optimization that reduces minimum-filter usage, temporally scheduled interpolation to eliminate bottlenecks, and fixed-point datapaths tuned for ± 1 LSB accuracy. Implemented on a Xilinx XC7Z-045FFG900-2 device, the system achieves one-pixel-per-clock processing, a maximum operating frequency of 271.74 MHz, and real-time 30.69 fps throughput at DCI-4K resolution, demonstrating that the algorithm-hardware co-design successfully combines adaptability with high-resolution performance.

Overall, the proposed system delivers a novel and practical dehazing solution by unifying referenceless haze assessment, lightweight model-free fusion, and scalable hardware realization. These characteristics make it well suited for embedded vision, autonomous platforms, and aerial imaging. Future work will explore temporal consistency modeling and extended color-restoration modules to further enhance stability in continuous video operation.

Author Contributions: Conceptualization, B.K. and D.N.; software, J.S.; validation, J.S. and S.A.; data curation, J.S.; writing—original draft preparation, D.N.; writing—review and editing, J.S., D.N., S.A., and B.K.; visualization, J.S. and D.N.; supervision, B.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2023R1A2C1004592).

Data Availability Statement: Data are available in a publicly accessible repository. FRIDA2: <http://perso.lcpc.fr/tarel.jean-philippe/bdd/frida.html> (accessed on 3 May 2025) D-HAZY: https://cs.nyu.edu/~fergus/datasets/nyu_depth_v2.html (accessed on 3 May 2025) and <https://vision.middlebury.edu/stereo/data/scenes2014/> (accessed on 3 May 2025) O-HAZE: <https://data.vision.ee.ethz.ch/cvl/ntire18/i-haze/> (accessed on 3 May 2025) I-HAZE: <https://data.vision.ee.ethz.ch/cvl/ntire18/o-haze/> (accessed on 3 May 2025) Dense-Haze: <https://data.vision.ee.ethz.ch/cvl/ntire19//dense-haze/> (accessed on 3 May 2025).

Acknowledgments: The EDA tool was supported by the IC Design Education Center (IDEC), Korea.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

Two-dimensional filtering operations are realized through a combination of line memories, register chains, and dedicated logic circuits, each customized according to the specific filter kernel. The structural concept is depicted in Figure A1.

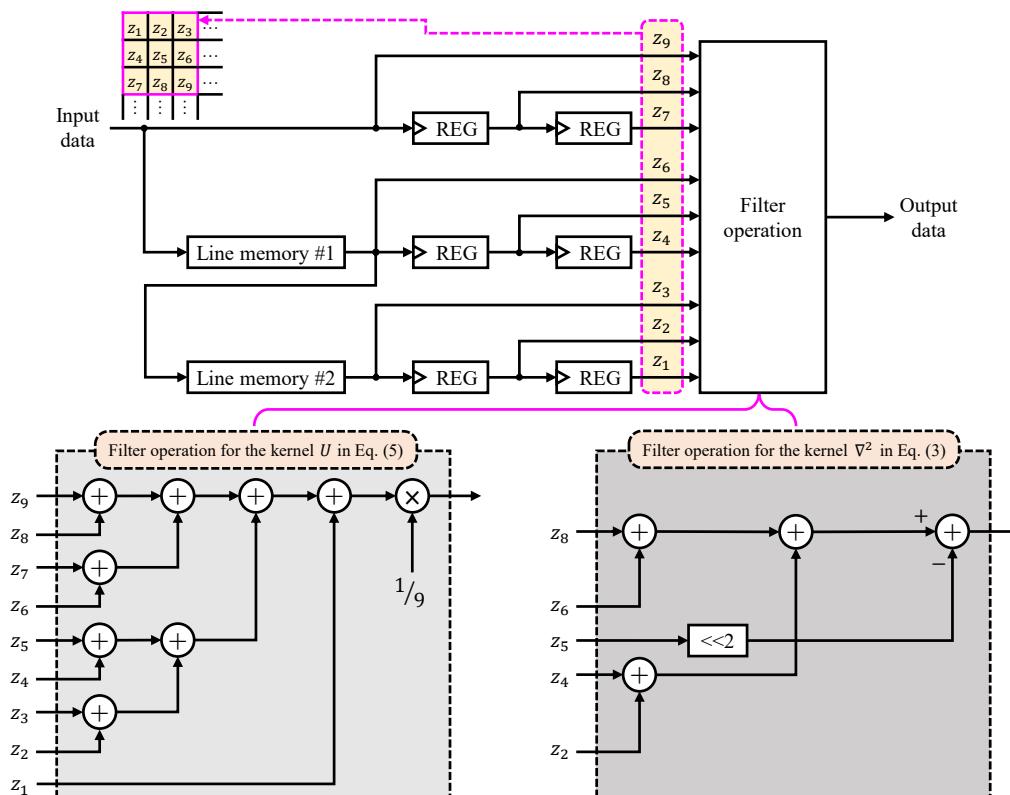


Figure A1. Simplified datapath of 3×3 two-dimensional filters. REG is an abbreviation for register.

In this design, line memories are mapped to on-chip block RAMs on the FPGA, where each memory introduces a one-line vertical delay equal to the image width. Hence, the set of line memories provides vertical pixel buffering, while horizontal delays are achieved through cascaded registers, each introducing a one-pixel shift along the image row.

For example, the 3×3 filter shown in Figure A1 employs two line memories and six registers, allowing simultaneous access to nine pixels $\{z_1, z_2, \dots, z_9\}$ within the current

filter window. These stored pixels are then routed to the filter operation block, which performs kernel-specific arithmetic on the buffered data to produce the corresponding output sample.

The logic circuit responsible for filtering is specialized based on the kernel type. Figure A1 illustrates two representative cases:

- **The moving-average filter** corresponding to the kernel U in Equation (4), and
- **The Laplacian filter** corresponding to ∇^2 in Equation (2).

References

1. He, K.; Sun, J.; Tang, X. Single Image Haze Removal Using Dark Channel Prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **2011**, *33*, 2341–2353. [\[CrossRef\]](#)
2. Zhu, Q.; Mai, J.; Shao, L. A Fast Single Image Haze Removal Algorithm Using Color Attenuation Prior. *IEEE Trans. Image Process.* **2015**, *24*, 3522–3533. [\[CrossRef\]](#) [\[PubMed\]](#)
3. Huang, Y.; Lin, Z.; Xiong, S.; Sun, T. Diffusion Models Based Null-Space Learning for Remote Sensing Image Dehazing. *IEEE Geosci. Remote Sens. Lett.* **2024**, *21*, 8001305. [\[CrossRef\]](#)
4. Bilal, M.; Masud, S.; Hanif, M.S. Efficient Framework for Real-Time Color Cast Correction and Dehazing Using Online Algorithms to Approximate Image Statistics. *IEEE Access* **2024**, *12*, 72813–72827. [\[CrossRef\]](#)
5. Park, S.; Kim, H.; Ro, Y.M. Robust pedestrian detection via constructing versatile pedestrian knowledge bank. *Pattern Recognit.* **2024**, *153*, 110539. [\[CrossRef\]](#)
6. Wang, J.; Wu, S.; Yuan, Z.; Tong, Q.; Xu, K. Frequency compensated diffusion model for real-scene dehazing. *Neural Netw.* **2024**, *175*, 106281. [\[CrossRef\]](#)
7. Chen, Z.; He, Z.; Lu, Z.M. DEA-Net: Single Image Dehazing Based on Detail-Enhanced Convolution and Content-Guided Attention. *IEEE Trans. Image Process.* **2024**, *33*, 1002–1015. [\[CrossRef\]](#)
8. Song, Y.; He, Z.; Qian, H.; Du, X. Vision Transformers for Single Image Dehazing. *IEEE Trans. Image Process.* **2023**, *32*, 1927–1941. [\[CrossRef\]](#)
9. Berman, D.; Treibitz, T.; Avidan, S. Single Image Dehazing Using Haze-Lines. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *42*, 720–734. [\[CrossRef\]](#)
10. Liu, J.; Liu, R.W.; Sun, J.; Zeng, T. Rank-One Prior: Real-Time Scene Recovery. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, *45*, 8845–8860. [\[CrossRef\]](#)
11. Wang, Y.; Hu, J.; Zhang, R.; Wang, L.; Zhang, R.; Liu, X. Heterogeneity constrained color ellipsoid prior image dehazing algorithm. *J. Vis. Commun. Image Represent.* **2024**, *101*, 104177. [\[CrossRef\]](#)
12. Qiu, Z.; Gong, T.; Liang, Z.; Chen, T.; Cong, R.; Bai, H.; Zhao, Y. Perception-Oriented UAV Image Dehazing Based on Super-Pixel Scene Prior. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 5913519. [\[CrossRef\]](#)
13. Liang, S.; Gao, T.; Chen, T.; Cheng, P. A Remote Sensing Image Dehazing Method Based on Heterogeneous Priors. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 5619513. [\[CrossRef\]](#)
14. Wu, R.Q.; Duan, Z.P.; Guo, C.L.; Chai, Z.; Li, C. RIDCP: Revitalizing Real Image Dehazing via High-Quality Codebook Priors. In Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 17–24 June 2023; pp. 22282–22291. [\[CrossRef\]](#)
15. Galdran, A. Image dehazing by artificial multiple-exposure image fusion. *Signal Process.* **2018**, *149*, 135–147. [\[CrossRef\]](#)
16. Ancuti, C.; Ancuti, C.O.; De Vleeschouwer, C.; Bovik, A.C. Day and Night-Time Dehazing by Local Airlight Estimation. *IEEE Trans. Image Process.* **2020**, *29*, 6264–6275. [\[CrossRef\]](#)
17. Lv, T.; Du, G.; Li, Z.; Wang, X.; Teng, P.; Ni, W.; Ouyang, Y. A fast hardware accelerator for nighttime fog removal based on image fusion. *Integration* **2024**, *99*, 102256. [\[CrossRef\]](#)
18. Cai, B.; Xu, X.; Jia, K.; Qing, C.; Tao, D. DehazeNet: An End-to-End System for Single Image Haze Removal. *IEEE Trans. Image Process.* **2016**, *25*, 5187–5198. [\[CrossRef\]](#)
19. Ren, W.; Pan, J.; Zhang, H.; Cao, X.; Yang, M.H. Single Image Dehazing via Multi-scale Convolutional Neural Networks with Holistic Edges. *Int. J. Comput. Vis.* **2020**, *128*, 240–259. [\[CrossRef\]](#)
20. Qiu, Y.; Zhang, K.; Wang, C.; Luo, W.; Li, H.; Jin, Z. MB-TaylorFormer: Multi-branch Efficient Transformer Expanded by Taylor Formula for Image Dehazing. In Proceedings of the 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, 1–6 October 2023; pp. 12756–12767. [\[CrossRef\]](#)
21. Yin, H.; Yang, P. Multistage Progressive Single-Image Dehazing Network With Feature Physics Model. *IEEE Trans. Instrum. Meas.* **2024**, *73*, 5013612. [\[CrossRef\]](#)
22. Ding, H.; Xie, F.; Qiu, L.; Zhang, X.; Shi, Z. Robust Haze and Thin Cloud Removal via Conditional Variational Autoencoders. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 1–16. [\[CrossRef\]](#)

23. Shiau, Y.H.; Yang, H.Y.; Chen, P.Y.; Chuang, Y.Z. Hardware Implementation of a Fast and Efficient Haze Removal Method. *IEEE Trans. Circuits Syst. Video Technol.* **2013**, *23*, 1369–1374. [\[CrossRef\]](#)

24. Zhang, B.; Zhao, J. Hardware Implementation for Real-Time Haze Removal. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2017**, *25*, 1188–1192. [\[CrossRef\]](#)

25. Shiau, Y.H.; Kuo, Y.T.; Chen, P.Y.; Hsu, F.Y. VLSI Design of an Efficient Flicker-Free Video Defogging Method for Real-Time Applications. *IEEE Trans. Circuits Syst. Video Technol.* **2019**, *29*, 238–251. [\[CrossRef\]](#)

26. Park, Y.; Kim, T.H. Fast Execution Schemes for Dark-Channel-Prior-Based Outdoor Video Dehazing. *IEEE Access* **2018**, *6*, 10003–10014. [\[CrossRef\]](#)

27. Wang, L.; Luo, Z.; Gao, L. An Improved Dark Channel Prior Method for Video Defogging and Its FPGA Implementation. *Symmetry* **2025**, *17*, 839. [\[CrossRef\]](#)

28. Qunpeng, G.; Baiquan, Q.; Fengqi, Y.; Liye, C.; Peng, G.; Jiatao, W.; Zonghong, L.; Weixing, W.; Jiaxing, X. Improved adaptive FPGA dark channel prior dehazing algorithm for edge applications in agricultural scenarios. *Smart Agric. Technol.* **2025**, *12*, 101285. [\[CrossRef\]](#)

29. Upadhyay, B.B.; Sarawadekar, K. VLSI Design of Saturation-Based Image Dehazing Algorithm. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2023**, *31*, 959–968. [\[CrossRef\]](#)

30. Kumar, R.; Kaushik, B.K.; Raman, B.; Sharma, G. A Hybrid Dehazing Method and its Hardware Implementation for Image Sensors. *IEEE Sens. J.* **2021**, *21*, 25931–25940. [\[CrossRef\]](#)

31. Ma, K.; Liu, W.; Wang, Z. Perceptual evaluation of single image dehazing algorithms. In Proceedings of the Proceedings of 2015 IEEE International Conference on Image Processing (ICIP), Quebec City, QC, Canada, 27–30 September 2015; pp. 3600–3604. [\[CrossRef\]](#)

32. Ngo, D.; Lee, G.D.; Kang, B. Haziness Degree Evaluator: A Knowledge-Driven Approach for Haze Density Estimation. *Sensors* **2021**, *21*, 3896. [\[CrossRef\]](#)

33. Ngo, D.; Son, J.; Kang, B. VBI-Accelerated FPGA Implementation of Autonomous Image Dehazing: Leveraging the Vertical Blanking Interval for Haze-Aware Local Image Blending. *Remote Sens.* **2025**, *17*, 919. [\[CrossRef\]](#)

34. Tarel, J.P.; Hautiere, N.; Caraffa, L.; Cord, A.; Halmaoui, H.; Gruyer, D. Vision Enhancement in Homogeneous and Heterogeneous Fog. *IEEE Intell. Transp. Syst. Mag.* **2012**, *4*, 6–20. [\[CrossRef\]](#)

35. Ancuti, C.; Ancuti, C.O.; De Vleeschouwer, C. D-HAZY: A dataset to evaluate quantitatively dehazing algorithms. In Proceedings of the Proceedings of 2016 IEEE International Conference on Image Processing (ICIP), Phoenix, AZ, USA, 25–28 September 2016; pp. 2226–2230. [\[CrossRef\]](#)

36. Ancuti, C.O.; Ancuti, C.; Timofte, R.; De Vleeschouwer, C. O-HAZE: A Dehazing Benchmark with Real Hazy and Haze-Free Outdoor Images. In Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Salt Lake City, UT, USA, 18–22 June 2018; pp. 867–8678. [\[CrossRef\]](#)

37. Ancuti, C.; Ancuti, C.O.; Timofte, R.; Van Gool, L.; Zhang, L.; Yang, M.H.; Patel, V.M.; Zhang, H.; Sindagi, V.A.; Zhao, R.; et al. NTIRE 2018 Challenge on Image Dehazing: Methods and Results. In Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Salt Lake City, UT, USA, 18–22 June 2018; pp. 891–901. [\[CrossRef\]](#)

38. Ancuti, C.O.; Ancuti, C.; Timofte, R.; Van Gool, L.; Zhang, L.; Yang, M.H.; Guo, T.; Li, X.; Cherukuri, V.; Monga, V.; et al. NTIRE 2019 Image Dehazing Challenge Report. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Long Beach, CA, USA, 16–17 June 2019; pp. 2241–2253. [\[CrossRef\]](#)

39. Li, B.; Gou, Y.; Gu, S.; Liu, J.Z.; Zhou, J.T.; Peng, X. You Only Look Yourself: Unsupervised and Untrained Single Image Dehazing Neural Network. *Int. J. Comput. Vis.* **2021**, *129*, 1754–1767. [\[CrossRef\]](#)

40. Std 1364-2005; IEEE Standard for Verilog Hardware Description Language. Revision of IEEE Std 1374-2001; IEEE: New York, NY, USA, 2006; pp. 1–590. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.