Sky Segmentation and Dual Atmospheric Light Defogging Algorithm Implemented by FPGA

Hongmin CHEN, Wei SONG, Keqing NING *

School of Artificial Intelligence and Computer Science, North China University of Technology, Beijing 100144, China

18733979225@163.com, songwei@ncut.edu.cn, *ningkq@ncut.edu.cn

Submitted July 12, 2026 / Accepted September 23, 2026 / Online first November 20, 2026

Abstract. *Hazy weather degrades image contrast, obscures details, and causes color distortion, while conventional dark channel prior (DCP) methods often produce inaccurate transmission estimates and boundary artifacts in bright, low-texture regions such as the sky. To address these limitations, this paper proposes an improved DCP-based dehazing method designed for real-time FPGA implementation. The method employs a minimum fusion strategy for multi-scale dark channels to reduce computational complexity. A soft sky-weight map is then constructed using adaptive thresholding followed by Gaussian smoothing. Atmospheric light is estimated separately for the sky and non-sky regions and subsequently combined to generate a spatially varying atmospheric-light map. The transmission map is refined using a fast guided filter, and adaptive Gamma correction is applied to improve the visibility of low-luminance regions. Experiments on the RESIDE SOTS-Outdoor dataset yield average peak signal-to-noise ratio (PSNR), structural similarity (SSIM), entropy (EN), and average gradient (AG) values of 21.3008 dB, 0.8937, 7.2412, and 7.2830, respectively. The proposed algorithm is implemented on an FPGA development board equipped with an Intel (Altera) Cyclone IV E EP4CE10F17C8 device and achieves real-time processing at 60 fps for images with a resolution of 640 × 480. The quantitative and visual results demonstrate that the proposed method provides effective sky-region restoration, smooth boundary transitions, and improved visibility in dark regions while satisfying the real-time processing requirements of a resource-constrained FPGA platform.*

Keywords

Field programmable gate arrays, image restoration, pipeline processing, real-time systems, table lookup

1. Introduction

In real-world scenarios, images acquired under hazy weather conditions are often severely degraded due to the scattering and absorption of light by suspended particles in the atmosphere. This degradation manifests as reduced contrast, blurred details, and color distortion, which not only impairs the visual quality of the images but also significantly undermines the performance of subsequent computer vision algorithms. Therefore, achieving real-time, high-quality image restoration on resource-constrained embedded platforms remains a critical challenge in current research [1].

Existing image dehazing methods mainly include physics-based and deep-learning-based approaches. Deep-learning-based algorithms generally achieve favorable dehazing performance; however, they have high computational and memory requirements and rely heavily on large-scale training data, making them difficult to deploy directly on resource-constrained edge devices. Physics-based methods offer stronger interpretability. Among them, the Dark Channel Prior (DCP) method proposed by He et al. [2] effectively estimates the transmission map and restores clear images by exploiting the statistical dark-channel characteristics of natural haze-free images, and has become one of the representative methods in image dehazing. Subsequently, researchers improved the performance of conventional DCP by introducing improved restoration models and related transmission-estimation strategies [3] and optimizing transmission estimation and filtering procedures. Regarding hardware implementation, existing studies have employed lightweight algorithms, image fusion, multi-scale dark-channel computation, and low-latency streaming architectures to realize real-time image dehazing on FPGAs. These studies have demonstrated the feasibility of real-time FPGA-based dehazing. However, as reported in [4], the assumptions underlying conventional DCP generally do not hold in sky regions. Because sky regions are typically bright and weakly textured, their dark-channel values often do not approach zero, resulting in inaccurate transmission estimation and potentially causing sky darkening, color distortion, and boundary artifacts. Moreover, a single global atmospheric-light value cannot accurately characterize the different imaging properties of sky and non-sky regions.

To address these problems, this paper proposes a real-time FPGA-based image dehazing method based on an improved DCP. The main contributions of this work are as follows:

DOI: 10.13164/re.2026.0509

1. A sky-region segmentation method is proposed. It identifies candidate sky regions through multi-scale dark-channel fusion and normalization, and converts the binary sky mask into a continuous soft weight map through smoothing, thereby reducing abrupt transitions and artifacts at sky boundaries.

2. A dual atmospheric-light estimation strategy is proposed to estimate atmospheric light separately in sky and non-sky regions and construct a spatially varying atmospheric-light map according to the soft sky weight map, thereby improving the stability of transmission estimation and image restoration in different regions.

3. Fast guided filtering and adaptive Gamma correction are incorporated into the image restoration process to refine the transmission map and compensate for low-brightness regions, thereby improving image details, brightness, and overall contrast.

4. A streaming FPGA architecture corresponding to the proposed algorithm is designed. The system is deployed on an Intel Cyclone IV E FPGA platform and achieves real-time processing at 60 fps for a resolution of 640×480.

The remainder of this paper is organized as follows. Section 2 introduces the atmospheric scattering model and the dark channel prior. Section 3 presents the proposed image dehazing method in detail. Section 4 describes the corresponding FPGA architecture. Section 5 presents and analyzes the experimental results. Finally, Section 6 concludes the paper.

2. Theoretical Basis

2.1 Atmospheric Scattering Model

The formation of degraded images in hazy weather is generally described by the atmospheric scattering model. This model posits that the light reaching the camera during the imaging process primarily consists of two components: the direct attenuation term, which represents the scene radiance reaching the imaging plane after atmospheric attenuation, and the airlight term, which is formed by ambient light scattering into the camera. Therefore, the hazy image can be expressed as follows [5]:

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

where A denotes the global atmospheric light, $I(x)$ represents the observed hazy image, $J(x)$ is the haze-free scene radiance image to be recovered, and $t(x)$ signifies the transmission map. The transmission map reflects the degree of light attenuation between the scene point and the imaging device, with its value typically satisfying $0 < t(x) \leq 1$. The atmospheric scattering model is illustrated in Fig. 1.

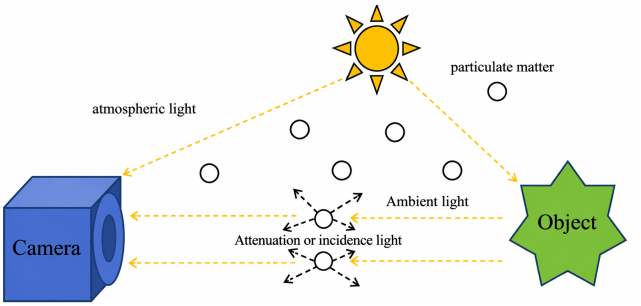


Fig. 1. Schematic diagram of the atmospheric scattering model.

2.2 Dark Channel Prior

The Dark Channel Prior (DCP) serves as one of the benchmark methods in image dehazing. Its core assumption is that, within the local non-sky regions of most haze-free natural images, at least one color channel contains pixels with very low intensity values [6], which are often close to zero. Based on this statistical regularity, the dark channel of a haze-free image J can be defined as [7]:

$$J^{\text{dark}}(x) = \min_{y \in \Omega(x)} \left(\min_{c \in \{\text{R,G,B}\}} J^c(y) \right) \tag{2}$$

where Ω(x) represents the local window centered at pixel x, and $J^c(y)$ denotes the intensity value of pixel y in color channel c. For haze-free natural images, the non-sky regions usually satisfy:

$$J^{\text{dark}}(x) \approx 0. \tag{3}$$

By substituting this prior into the atmospheric scattering model, the transmission map can be estimated. First, the hazy image is normalized by the global atmospheric light as follows [8]:

$$\frac{I^c(x)}{A^c} = t(x)\frac{J^c(x)}{A^c} + 1 - t(x). \tag{4}$$

To estimate the transmission map of the image, He et al. assumed that the global atmospheric light is constant and, based on this assumption, proposed the following formula to calculate the transmission:

$$t(x) = 1 - \omega \min_{y \in \Omega(x)} \left(\min_{c \in \{\text{R,G,B}\}} \frac{I^c(y)}{A^c} \right) \tag{5}$$

where ω is a compensation coefficient introduced to retain a small amount of haze. In high-contrast regions, this coefficient helps preserve sufficient details, while also avoiding excessive dehazing in low-contrast areas, thereby improving the naturalness of the restored image. To avoid the adverse effects caused by errors in global atmospheric-light estimation [9], He et al. estimated the atmospheric light value A by searching for the top 0.1% brightest pixels in the dark channel. Finally, the recovery formula for the dehazed image is expressed as follows [10]:

$$J(x) = \frac{I(x) - A}{\max(t(x), t_0)} + A \tag{6}$$

where t_0 is the lower bound threshold for the transmission map, which is typically set to 0.1 to avoid excessive image

brightening and visual distortion caused by overly low transmission values, thereby ensuring balanced brightness and a natural appearance in the dehazing results.

3. Method

The key computational workflow of the proposed method is summarized in Algorithm 1. The pseudocode presents the sequence of sky-region segmentation, dual atmospheric-light estimation, transmission refinement, image restoration, and adaptive Gamma correction, together with the parameter values used in this study.

3.1 Sky Region Segmentation

The traditional DCP algorithm is generally valid in non-sky regions. However, in sky regions, pixel values are typically high, and the local dark channel values no longer approach zero. Direct application of the DCP can distort transmission-map estimation, causing issues such as dimming and artifacts in sky areas. As illustrated in Fig. 2, this paper first obtains an initial segmentation result for the sky region by utilizing a normalized fused dark channel (Fig. 2(b)). Subsequently, the sky mask is transformed into a continuous weight map via Gaussian smoothing (Fig. 2(c)), thereby achieving a transition from hard segmentation to soft segmentation (Fig. 2(d)) and effectively avoiding abrupt transition artifacts at the sky boundaries. Let the input hazy image be I, and its initial dark channel is defined as:

$$D_0(x) = \min_{c\in\{\mathrm{R,G,B}\}} I^c(x). \tag{7}$$

To balance texture details and algorithm stability, this paper adopts multi-scale windows to perform minimum filtering on the dark channel [11], obtaining $D_{\mathrm{small}}(x)$ and $D_{\mathrm{large}}(x)$. However, the primary objectives of this study are sky-region segmentation and dual atmospheric-light estimation. Therefore, instead of the continuous-weight fusion scheme based on local statistics adopted by Wang et al. [12], which serves as the reference approach for this fusion comparison, this work employs a simpler minimum fusion strategy to construct the fused dark channel, as illustrated in Fig. 4(a). We define it as:

$$D_{\mathrm{fused}}(x) = \min\left(D_{\mathrm{small}}(x), D_{\mathrm{large}}(x)\right). \tag{8}$$

Rather than directly performing sky discrimination on the original dark channel, this paper first calculates the global initial atmospheric light scalar of the entire image:

$$A_0 = \max_x \left(\max_{c\in\{\mathrm{R,G,B}\}} I^c(x) \right). \tag{9}$$

To avoid division by zero, a lower bound is imposed on the initial atmospheric-light scalar by defining $\widehat{A}_0 = \max(A_0, 10^{-6})$. Each color channel is then normalized as $I_n^c(x) = I^c(x)/\widehat{A}_0$, where $c \in \{\mathrm{R,G,B}\}$. The channel-wise minimum of the normalized image is subsequently processed

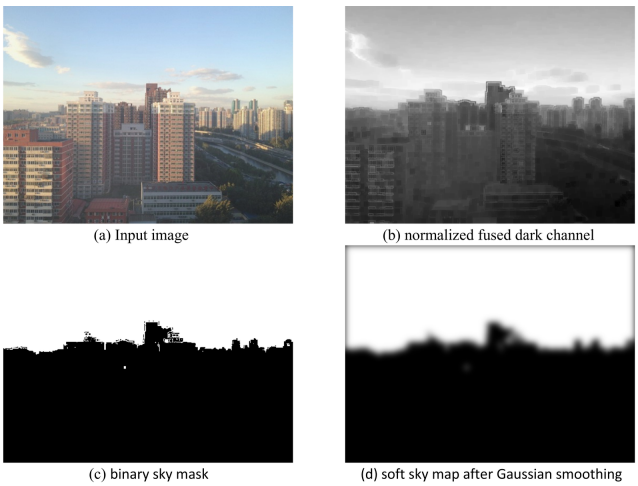
Algorithm 1. Proposed method.

Input: Hazy RGB image $I \in [0,1]^{H\times W\times 3}$.

Output: Restored haze-free image J.

Parameter: $w_s = 3$, $w_l = 5$, $\omega = 0.85$, $t_0 = 0.10$, $\rho = 0.72$, $\sigma_s = 6.0$, $p = 0.001$, $\lambda_{\mathrm{obj}} = 0.4$, $\lambda_{\mathrm{sky}} = 0.25$, $\eta = 0.55$, $r = 6$, $\varepsilon = 10^{-3}$, $s = 4$, $T_\gamma = 0.38$, and $\gamma = 0.80$.

Compute $D_0(x) = \min_{c\in\{\mathrm{R,G,B}\}} I^c(x)$.
Apply $w_s \times w_s$ and $w_l \times w_l$ minimum filters to D_0 to obtain D_{small} and D_{large}.
Compute $D_{\mathrm{fused}}(x) = \min(D_{\mathrm{small}}(x), D_{\mathrm{large}}(x))$.
Estimate $A_0 = \max_x(\max_{c\in\{\mathrm{R,G,B}\}} I^c(x))$ and set $A_0^{\mathrm{vec}} = [A_0, A_0, A_0]$.
Normalize I using A_0 and compute $D_{n,\mathrm{fused}}$ using the same multi-scale procedure.
Determine $\alpha = \mathrm{Otsu}(D_{n,\mathrm{fused}})$.
Set $M_{\mathrm{sky}}(x) = 1$ if $D_{n,\mathrm{fused}}(x) > \alpha$; otherwise, set $M_{\mathrm{sky}}(x) = 0$.
Set $M_{\mathrm{sky}}(x) = 0$ outside the upper $\rho = 72\%$ of the image.
Compute $S(x) = G(M_{\mathrm{sky}}(x), \sigma_s)$.
Compute $B(x) = \max_{c\in\{\mathrm{R,G,B}\}} I^c(x)$.
Select the top $p = 0.1\%$ candidate pixels from D_{fused} in the non-sky region and from B in the sky region.
Average the RGB values of the brightest 30% of the selected candidates to obtain $A_{\mathrm{obj}}^{\mathrm{raw}}$ and $A_{\mathrm{sky}}^{\mathrm{raw}}$.
Replace a missing regional candidate with A_0^{vec}.
Compute $A_{\mathrm{obj}} = (1-\lambda_{\mathrm{obj}})A_0^{\mathrm{vec}} + \lambda_{\mathrm{obj}} A_{\mathrm{obj}}^{\mathrm{raw}}$.
Compute $A_{\mathrm{sky}} = (1-\lambda_{\mathrm{sky}})A_0^{\mathrm{vec}} + \lambda_{\mathrm{sky}} A_{\mathrm{sky}}^{\mathrm{raw}}$.
Construct $A_{\mathrm{map}}(x) = (1-\eta S(x))A_{\mathrm{obj}} + \eta S(x) A_{\mathrm{sky}}$.
Compute $t_{\mathrm{sky}} = \max(t_0, 1-\omega\alpha)$.
Compute $t_{\mathrm{non}}(x) = \max(1-\omega D_{n,\mathrm{fused}}(x), t_0)$.
Construct $t_{\mathrm{init}}(x) = S(x)t_{\mathrm{sky}} + (1-S(x))t_{\mathrm{non}}(x)$.
Refine the tx map as $t(x) = \mathrm{FastGuidedFilter}(I_{\mathrm{gray}}, t_{\mathrm{init}}, r, \varepsilon, s)$.
Restore $J_0(x) = [I(x) - A_{\mathrm{map}}(x)]/t(x) + A_{\mathrm{map}}(x)$.
Calculate the average luminance $\bar{Y}$ of J_0.
if $\bar{Y} < T_\gamma$ **then**
    $J_1(x) = J_0(x)^\gamma$.
else
    $J_1(x) = J_0(x)$.
end if
Obtain $J(x) = \mathrm{clip}(J_1(x), 0, 1)$.
return J



(a) Input image (b) normalized fused dark channel (c) binary sky mask (d) soft sky map after Gaussian smoothing

Fig. 2. The process of sky segmentation.

using the same 3×3 and 5×5 minimum filters and pixel-wise minimum fusion described above, yielding $D_{n,\text{fused}}(x)$. A threshold α is then employed to construct the sky mask. Specifically, α is adaptively determined for each input image using Otsu's method, which selects the threshold by maximizing the between-class variance of the normalized fused dark-channel histogram. We define the sky mask as follows:

$$M_{\text{sky}}(x) = \begin{cases} 1, & D_{n,\text{fused}}(x) > \alpha, \\ 0, & D_{n,\text{fused}}(x) \le \alpha. \end{cases} \tag{10}$$

The candidate sky region is further restricted to the upper $\rho = 0.72$ of the image. This empirically selected positional prior retains most sky candidates while reducing false detections caused by bright buildings and ground regions.

Furthermore, to prevent the hard segmentation from causing abrupt artifacts at the sky boundaries, this paper does not directly apply the binary sky mask to subsequent calculations. Instead, Gaussian smoothing is applied to adjust its spatial distribution, thereby constructing a soft sky weight map:

$$S(x) = \text{G}(M_{\text{sky}}(x), \sigma_{\text{s}}) \tag{11}$$

where $\text{G}(\cdot)$ represents the Gaussian smoothing operator, and σ_{s} denotes the smoothing scale and is set to 6.0. This soft sky weight map is applied to the subsequent atmospheric light mapping and image restoration stages, ensuring a smoother and more natural transition between regions [13].

3.2 Dual Atmospheric Light Estimation Strategy

When the brightness distribution and imaging characteristics differ significantly between the sky and non-sky regions, the single global atmospheric light model used in the traditional DCP fails to accurately describe the real-world scene. This limitation easily leads to over-restoration in sky regions or under-restoration in terrestrial areas. To address this issue, rather than employing a single atmospheric-light estimate, this paper separately estimates the atmospheric light for the sky and non-sky regions.

First, the global initial atmospheric light scalar A_0 obtained in the previous section is expanded into a three-channel vector. In the non-sky region, this paper uses the fused dark channel $D_{\text{fused}}(x)$ as the basis to select candidate pixels (as shown by the blue dots in Fig. 3(b)) and uses their RGB values to obtain the atmospheric light candidate value $A_{\text{obj}}^{\text{raw}}$ for the non-sky region; in the sky region, this paper uses the luminance map

$$B(x) = \max_{c \in \{\text{R,G,B}\}} I^c(x) \tag{12}$$

as a reference and selects candidate pixels with higher luminance within the sky region (as indicated by the red dots in Fig. 3(c)) to estimate the candidate atmospheric light value for the sky region $A_{\text{sky}}^{\text{raw}}$ [14]. However, since localized candidate estimation is susceptible to noise, extreme highlights, and segmentation errors, this paper does not directly adopt $A_{\text{obj}}^{\text{raw}}$ and $A_{\text{sky}}^{\text{raw}}$. Instead, each is shrunk toward the global atmospheric light A_0^{vec} to obtain more stable regional atmospheric light estimates:

$$\begin{aligned} A_{\text{obj}} &= (1 - \lambda_{\text{obj}}) A_0^{\text{vec}} + \lambda_{\text{obj}} A_{\text{obj}}^{\text{raw}}, \\ A_{\text{sky}} &= (1 - \lambda_{\text{sky}}) A_0^{\text{vec}} + \lambda_{\text{sky}} A_{\text{sky}}^{\text{raw}} \end{aligned} \tag{13}$$

where λ_{obj} and λ_{sky} control the degrees of dependence of the atmospheric light in the non-sky and sky regions on their respective local candidate values and are set to 0.4 and 0.25, respectively. After obtaining A_{obj} and A_{sky}, we combine them with the soft sky weight map $S(x)$ obtained in the previous section and use the following equation to construct a spatially varying atmospheric light map, where η controls the contribution of the sky-region atmospheric light to the spatially varying atmospheric light map and is set to 0.55:

$$A_{\text{map}}(x) = (1 - \eta S(x)) A_{\text{obj}} + \eta S(x) A_{\text{sky}}. \tag{14}$$

As illustrated in Fig. 3, the blue candidate points in (b) are mainly distributed across the building and terrestrial regions, corresponding to the atmospheric light estimation of the non-sky region. Conversely, the red candidate points in (c) are concentrated in the highly illuminated sky region, corresponding to the atmospheric light estimation of the sky region. Figure 3(d) displays the constructed spatially varying atmospheric light map $A_{\text{map}}(x)$, and this map exhibits a continuous and smooth transition between the sky and non-sky regions, rather than a simple binary switching. This observation demonstrates that the proposed dual atmospheric light strategy can effectively characterize the atmospheric light discrepancies across different regions, thereby providing a reliable atmospheric-light estimate for subsequent image restoration. However, to avoid the interference of regional atmospheric light modeling on the stability of the main transmission-estimation pipeline, the dual atmospheric light strategy in this paper is used only in the final image-recovery stage and is not directly used for transmission estimation.

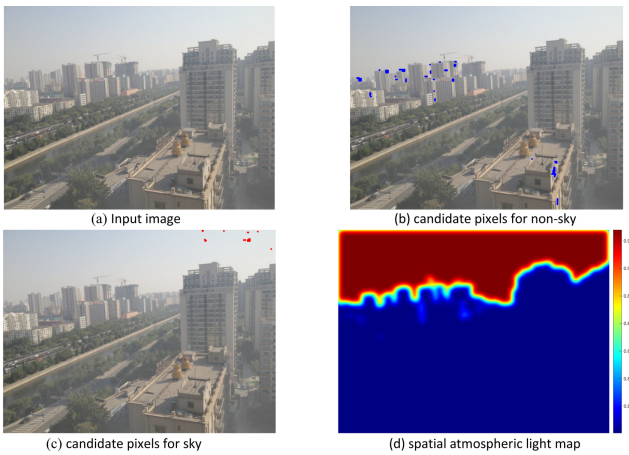


Fig. 3. Schematic diagram of the dual atmospheric light estimation process.

3.3 Transmittance Estimation and Image Restoration

After obtaining the normalized fused dark channel $D_{n,\text{fused}}(x)$ and the soft sky weight map $S(x)$, this paper further constructs the transmission map and completes the image restoration. Similar to the adoption of different atmospheric lights for the sky and non-sky regions, this paper also employs distinct transmission expressions for these two types of regions [15].

Since the dark channel prior is generally invalid in sky regions, directly applying the conventional DCP transmission estimation method tends to cause over-restoration or luminance distortion in the sky area. Therefore, for the sky region, we employ a constrained sky transmission formulation:

$$t_{\text{sky}} = \max\left(t_0, 1 - \omega\alpha\right) \tag{15}$$

where t_0 is the minimum transmission threshold, ω represents the haze preservation coefficient, and α denotes the sky segmentation threshold. Within the sky region, this expression provides a relatively stable transmission estimation, thereby mitigating the sensitivity of the sky area to abnormal dark channel responses.

In the non-sky region, this paper retains the basic framework of the dark channel prior and uses the normalized fused dark channel to estimate the transmission, while further imposing a lower-bound constraint using the following equation:

$$t_{\text{non}}(x) = \max\left(1 - \omega D_{n,\text{fused}}(x), t_0\right) \tag{16}$$

where $1 - \omega D_{n,\text{fused}}(x)$ represents the transmission map estimated using the normalized fused dark channel. To prevent numerical instability issues caused by excessively small transmission values, the maximum value between this transmission estimate and t_0 is taken to enforce a lower bound constraint.

After obtaining the transmission maps for the sky and non-sky regions, we combine them with the soft sky weight map $S(x)$ and define the following equation to construct the initial transmission map for the entire image:

$$t_{\text{init}}(x) = S(x)t_{\text{sky}} + (1 - S(x))t_{\text{non}}(x). \tag{17}$$

According to the above equation, the sky region primarily adopts a relatively stable transmission representation, while the non-sky region mainly relies on the estimation results of the traditional dark channel prior. Since $S(x)$ is a continuously varying soft weight map, the transmission map can achieve a smooth transition at the boundary between the sky and non-sky regions, thereby reducing abrupt boundary transitions caused by hard segmentation.

To further improve the smoothness of the transmission map and preserve edge structures, this paper introduces the fast guided filter to refine the initial transmission map. Taking the grayscale version of the input image $I_{\text{gray}}(x)$ as the guidance image [16], the final transmission map can be obtained as follows:

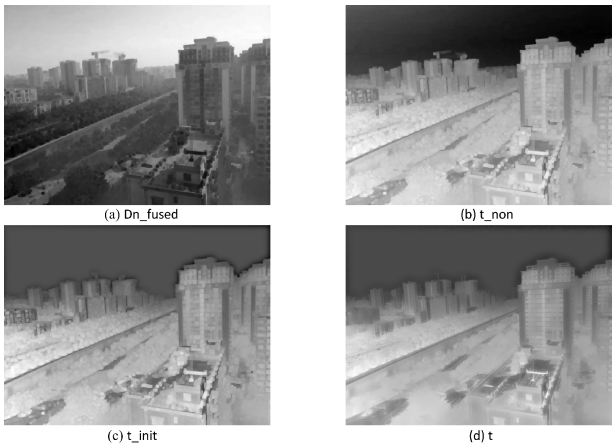


Fig. 4. Intermediate results of transmission estimation: (a) normalized fused dark channel $D_{n,\text{fused}}(x)$; (b) non-sky transmission map $t_{\text{non}}(x)$; (c) initial fused transmission map $t_{\text{init}}(x)$; (d) refined transmission map $t(x)$.

$$t(x) = \mathcal{F}\left(I_{\text{gray}}(x), t_{\text{init}}(x)\right) \tag{18}$$

where $\mathcal{F}(\cdot)$ represents the fast guided filtering operator. The filtering radius r, regularization parameter ε, and subsampling factor s are set to 6, 10^{-3}, and 4, respectively. This step can effectively preserve edge details and suppress block artifacts, providing a more stable input for the subsequent image restoration stage.

As illustrated in Fig. 4, the normalized fused dark channel (Fig. 4(a)) can effectively reflect the spatial distribution characteristics of haze density in the scene. Consequently, the constructed non-sky transmission map $t_{\text{non}}(x)$ (Fig. 4(b)) exhibits a distinct hierarchy in terrestrial regions. By further incorporating $t_{\text{non}}(x)$ with the soft sky weight map, a continuously varying initial transmission map $t_{\text{init}}(x)$ (Fig. 4(c)) is obtained. After refinement via the fast guided filter, the final transmission map (Fig. 4(d)) effectively mitigates blocking artifacts while preserving edge structures.

Finally, after obtaining the refined transmittance $t(x)$ and the spatially varying atmospheric light map $A_{\text{map}}(x)$ constructed in the previous section, the initial fog-free image can be restored based on the atmospheric scattering model:

$$J_0(x) = \frac{I(x) - A_{\text{map}}(x)}{t(x)} + A_{\text{map}}(x). \tag{19}$$

3.4 Adaptive Gamma Correction

After the aforementioned transmission estimation and dual atmospheric light recovery, the degradation caused by haze is effectively mitigated. However, for certain heavily hazed scenes or scenes with low overall illumination, the recovered results are still prone to local dimness and insufficient visual contrast. Under such circumstances, directly applying a fixed brightness compensation tends to cause over-enhancement in regions with already normal brightness, thereby inducing color distortion or highlight clipping.

Therefore, this paper introduces an adaptive gamma correction strategy, which determines whether to perform brightness enhancement based on the overall luminance of the recovered image.

As described in Sec. 3.3, the obtained haze-free image is denoted as $J_0(x)$. Here, the average luminance of its corresponding grayscale image is first calculated as follows:

$$\bar{Y} = \frac{1}{N} \sum_{x \in \Omega} Y(x) \tag{20}$$

where $Y(x)$ denotes the corresponding grayscale value of $J_0(x)$ and N represents the total number of pixels in the image. When the average luminance $\bar{Y}$ falls below a predefined threshold T_γ, the current restored result is considered to be overall dim; at this point, gamma correction is applied:

$$J_1(x) = J_0(x)^\gamma \tag{21}$$

where γ is the gamma coefficient. In this paper, a value of $\gamma < 1$ is selected to enhance the brightness and visibility of low-luminance regions. When $\bar{Y} \geq T_\gamma$, the overall brightness of the image is already within a reasonable range. In this case, gamma correction is no longer performed, which can be expressed as:

$$J_1(x) = J_0(x). \tag{22}$$

Therefore, the adaptive gamma correction in this paper can be uniformly expressed as:

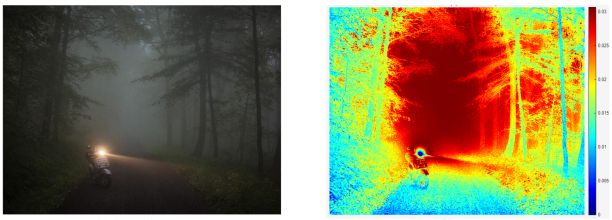
$$J_1(x) = \begin{cases} J_0(x)^\gamma, & \bar{Y} < T_\gamma, \\ J_0(x), & \bar{Y} \geq T_\gamma. \end{cases} \tag{23}$$

This strategy effectively prevents the over-enhancement issues induced by fixed enhancement methods. When the image is globally dim, gamma correction increases the luminance of dark regions, making image details clearer. When the image brightness is already sufficient, the original restored result remains unchanged, thereby preventing over-brightness distortion in the sky or highlight regions.

In the practical implementation, this paper sets $T_\gamma = 0.38$ and $\gamma = 0.80$. This parameter setting is relatively mild, primarily aimed at compensating for mild underexposure in the restored results without significantly altering the original color distribution of the image. After completing the gamma correction, the pixel values are further clipped to the valid range $[0, 1]$ to yield the final restored result:

$$J(x) = \text{clip}\left(J_1(x), 0, 1\right). \tag{24}$$

The adaptive gamma correction is mainly used to improve the visual quality of low-luminance restored images [17]; compared with fixed brightness compensation, this strategy adaptively determines whether enhancement is required based on the overall image brightness. As illustrated in Fig. 5, panel (a) shows the input hazy image, while panel (b) shows the spatial distribution of brightness enhancement as a heat map. For instance, the headlight area itself has a relatively high brightness, so it is marked in blue, indicating that no brightness enhancement is needed. This selective processing enhances the visibility of dark areas while maintaining the natural appearance of bright areas [18].



(a) Input image (b) difference map

Fig. 5. Adaptive gamma correction: (a) input image; (b) brightness-enhancement map.

4. Hardware Implementation

The overall hardware block diagram of the real-time dehazing system proposed in this paper is illustrated in Fig. 6. It primarily consists of the sky segmentation module, the dual atmospheric-light estimation module, the dark-channel fusion module, the subsampled fast guided-filter module, and the adaptive gamma-correction module.

4.1 Dark Channel Computing and Multi-Scale Fusion

The traditional DCP algorithm typically employs a fixed-size window to calculate the local minimum for estimating the haze-density distribution in an image. For a fixed window scale, a smaller window can better preserve image details but is more susceptible to noise; a larger window provides stronger regional smoothing but may cause edge blurring and detail loss. To address these problems, Wang et al. utilized an adaptive multi-scale weighted fusion strategy [12]. In contrast to this strategy, this paper adopts a lightweight multi-scale fusion method and obtains the fused dark channel through the pixel-wise minimum operation shown in (8).

First, a channel-wise minimum operation is applied to the input RGB image to obtain the initial dark channel [19] defined in (7). Because direct implementation of large two-dimensional window buffers is costly on an FPGA, local minimum filtering is implemented using a line-buffer-based sliding-window architecture. Shift registers and multi-level FIFO buffers construct the 3×3 and 5×5 neighborhood windows, whose data are updated continuously with the input pixel stream. This design avoids full-frame buffering and reduces on-chip storage consumption [20].

For the two dark-channel scales, 3×3 and 5×5 minimum windows are used for local erosion, yielding the small-scale dark channel $D_{\text{small}}(x)$ and the large-scale dark channel $D_{\text{large}}(x)$. The small-scale window preserves local details, while the large-scale window improves regional stability and noise immunity. Afterward, as formulated in (8), a lightweight fusion is performed on $D_{\text{small}}(x)$ and $D_{\text{large}}(x)$.

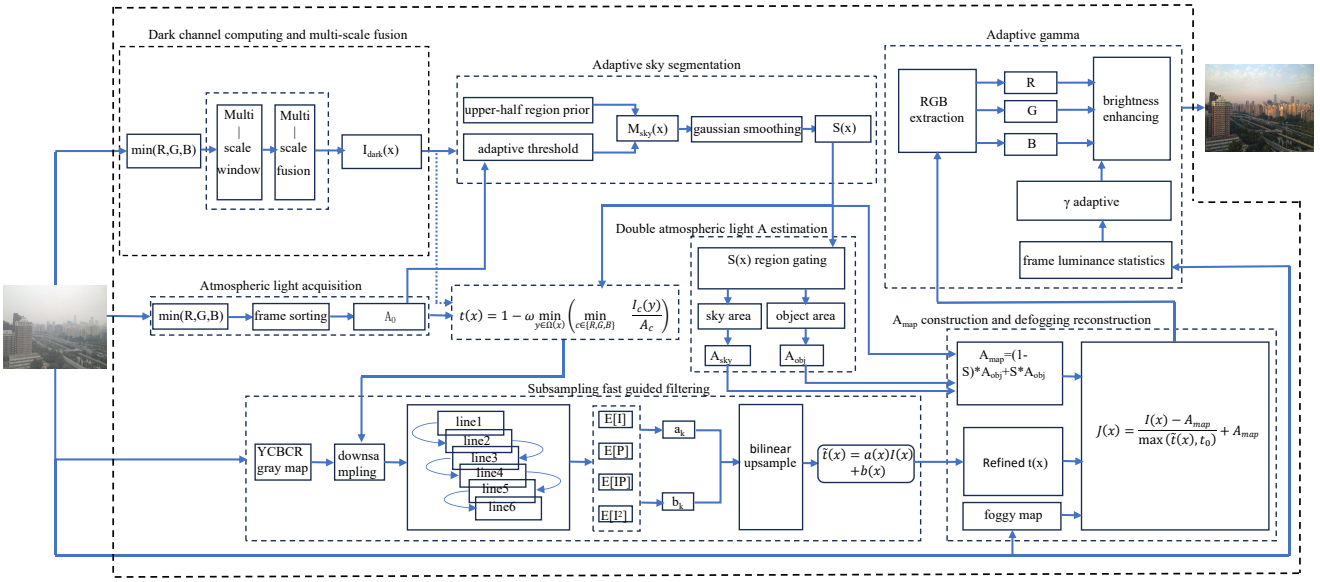


Fig. 6. Overall block diagram of the real-time dehazing hardware architecture.

Compared with weighted fusion methods, this approach eliminates the need for extra multipliers and weight update modules, thereby reducing hardware overhead and making it better suited to FPGA pipeline implementation. Concurrently, by retaining the darker response regions from both scales of dark channels, it effectively enhances the representation capability of haze information in sky boundary areas and distant building regions, providing a more stable input for subsequent sky segmentation and transmission estimation.

During the hardware implementation process, the dark channel computation module adopts a streaming pipeline architecture. The window operations and minimum comparison processes at each stage are executed in parallel, enabling the system to continuously output the fused dark-channel result at the pixel clock rate and thereby improving real-time processing performance [21].

4.2 Sky Area Segmentation Module

Because sky regions usually exhibit high brightness, low texture, and strong local uniformity, their dark-channel values are often relatively high. Therefore, when the traditional dark channel is directly adopted to estimate the transmission map, the transmission in sky regions may be underestimated. This further causes distortions, such as over-darkening, color distortion, and boundary artifacts in the sky regions. To improve the dehazing stability in sky areas, this paper incorporates a sky region segmentation module. This module achieves rapid detection of the sky region through a rule-based discrimination method, thereby providing regional prior information for subsequent dual atmospheric light estimation and transmission-map estimation [22].

In the hardware implementation, sky-candidate detection is restricted to the upper 72% of the image according to the same spatial prior used in the proposed algorithm. The hardware sky-segmentation module then applies the decision rule in (10) to the normalized fused dark-channel map to obtain the initial binary sky mask $\text{sky_mask}(x)$.

Thereafter, as formulated in (11), this paper further smoothes the initial sky mask to construct the soft segmentation result $S(x)$. Although Gaussian filters are conventionally used in theoretical algorithms, implementing a large-scale Gaussian kernel convolution on an FPGA platform consumes a prohibitive amount of multiplier (DSP) resources. To circumvent this bottleneck, this paper employs two cascaded 3×3 box filters in hardware to approximate the Gaussian smoothing effect. The mean calculation of a box filter only involves additions and divisions. Since the window size is 9, the hardware can replace the complex divider with an approximate shift operation and a very small number of adders. This cascaded mean filtering design not only substantially reduces the utilization of DSP resources but also generates a smoothly varying soft sky weight map $S(x)$ that improves visual continuity [23], thereby providing input for downstream atmospheric-light mapping. After the line buffers are filled, the two filter stages operate as a cascaded pipeline and produce one soft-weight value per clock cycle. Although the resulting smoothing kernel differs from an exact large-scale Gaussian kernel, it retains the required low-pass smoothing and continuous transition at sky boundaries while avoiding multiplier-intensive convolution.

4.3 Dual Atmospheric Light Estimation Module

Based on the results of the sky region segmentation, this paper introduces a dual atmospheric light estimation strategy to separately estimate the atmospheric light for the sky and non-sky regions, which are subsequently combined via weighted fusion to construct a spatially varying atmospheric light map.

Conventional algorithms require sorting the Top-N pixels in the full-frame dark-channel map to estimate the global atmospheric light. On an FPGA, this requires buffering a full frame of high-resolution image data and thus consumes substantial BRAM resources. To circumvent this bottleneck, this paper develops a parallel statistical method based on an inter-frame delayed update mechanism. In the system, two mutually independent sets of Top-N sorting register chains are instantiated: one is utilized to screen candidate pixels in the non-sky region with the highest fused dark channel values, while the other is dedicated to screening candidate pixels in the sky region with the highest intensities from the original luminance map. After the active video data of the current frame have been transmitted and the system enters the vertical blanking interval (VBLANK), the hardware finite state machine (FSM) uses the accumulators to compute the average intensities of these two sets of candidate pixels, thereby obtaining A_{obj} and A_{sky} for the current frame. Subsequently, DSP multiply-accumulate (MAC) units synthesize the spatially varying atmospheric-light map $A_{\text{map}}(x)$ according to (14).

4.4 Subsampled Fast Guided-Filter Module

The raw transmission map obtained through local minimum filtering is prone to significant block artifacts. Therefore, the system introduces a fast guided filtering mechanism and implements a hardware-level minimalist architecture via fourfold downsampling and time-division multiplexing (TDM) techniques to reduce the block artifacts caused by the dark channel prior. The downsampling operation is implemented by simple row-and-column data decimation, reducing the amount of data processed by the filter to one-sixteenth of the original size.

During the core computation phase of the guided filter, statistical features such as the mean, variance, and covariance of the guidance image and transmission map must be computed using large-window box-filtering operations on four input channels. To enhance resource utilization, this module develops a hardware architecture based on time-division multiplexing, instantiating a single set of shared multi-line data buffer arrays. Under the control of a finite state machine (FSM) scheduler [24], the module alternately processes the four channels, namely $E[I]$, $E[P]$, $E[IP]$, and $E[I^2]$, thereby reducing the utilization of DSP and LE resources by nearly 75%. Subsequently, lookup tables (LUTs) are employed instead of dividers to compute the linear coefficient a and bias b. Finally, the system leverages a bilinear upsampling unit based on a multiply-accumulate (MAC) array to reconstruct a refined transmission map $t(x)$ aligned with the original image size through two-dimensional weighted interpolation of four adjacent low-resolution coefficients.

In the FPGA implementation, the RGB input channels and the principal pixel-level intermediate results, including the dark channel, sky mask, soft sky weight, transmission map, and atmospheric-light map, are represented as 8-bit unsigned data. Normalized fixed-point coefficients are encoded with a scale factor of 256; for example, 0.20, 0.30, and 0.55 are approximately represented by 51, 76, and 141, respectively. In the guided-filter module, the linear coefficient a and bias b are represented as 16-bit signed Q4.12 and Q8.8 values, respectively. Extended-width accumulators and multiplication results are retained during intermediate calculations, after which right shifts and bit truncation restore the required fixed-point scale. The implementation primarily uses truncation rather than additional rounding. The refined transmission and restored RGB outputs are saturated to the valid range of 0–255, and the transmission is further constrained by a lower bound of $t_{\min} = 0.20$. In the current hardware configuration, the dual atmospheric-light estimates are also constrained between $\max(A_0 - 31, 0)$ and A_0. These extended-width, truncation, and output-saturation operations limit overflow effects and keep the final outputs within the valid 8-bit range.

4.5 Adaptive Gamma Module

This paper introduces an adaptive gamma-correction module to address low overall brightness and insufficient visual contrast in the restored image. Based on the luminance statistics of each frame, the module determines whether to apply brightness enhancement, thereby improving the visibility of dark regions while avoiding overexposure in highlight regions.

To compute luminance statistics without introducing additional frame latency, this design adopts a pipeline accumulation scheme. While the image reconstruction module outputs the pixel stream, the luminance-statistics unit operates in parallel. A high-bit-width accumulator is instantiated inside the unit to accumulate the grayscale value of each pixel in real time while the frame-valid signal is asserted, thereby obtaining the total luminance sum of the current frame. When the falling edge of the frame-valid signal indicates the end of the current frame, the hardware finite state machine (FSM) immediately invokes a timing-optimized divider to calculate the average luminance of the current frame. According to the algorithm design, the threshold for the average luminance is set to 0.38. In the hardware logic control, when the average luminance computed from the previous frame falls below this threshold, the FSM enters a compensation-enabled state; otherwise, no compensation is performed.

Directly implementing nonlinear exponential operations in an FPGA incurs extremely high resource consumption; therefore, this module adopts a lookup-table (LUT)-based optimization scheme. Prior to hardware synthesis, a 256-level grayscale mapping file with $\gamma = 0.80$ is pregenerated, which is loaded by the RTL initialization logic from a hexadecimal file. To support parallel non-blocking processing of the three RGB channels, the system utilizes synthesis directives to force the mapping of the LUTs for these three color channels into three independent on-chip M9K Block RAMs. The mapping network is designed as

a two-stage pipeline: the first clock cycle captures the high bits of the input pixel to serve as the address for reading the M9K memory, and the second clock cycle determines the final output through a bypass multiplexer (MUX). If the statistical analysis of the preceding frame identifies the image as dim, the system outputs the gamma-corrected value from the M9K; otherwise, if the brightness is deemed sufficient, the raw registered input data are bypassed to the output.

5. Experimental Results And Analysis

5.1 Experimental Platform and Testing Environment

The hardware verification platform employs an FPGA development board equipped with an Intel (Altera) Cyclone IV E EP4CE10F17C8 device. RTL design, synthesis, place-and-route, and hardware resource utilization analysis are performed using the 64-bit Quartus II 13.0 software. Algorithm simulation is conducted in MATLAB R2025a on a computer running the 64-bit Windows 11 operating system and equipped with an AMD Ryzen 7 5800U processor operating at 1.90 GHz.

The hardware system uses a 50 MHz board clock as the PLL input. The PLL generates a 100 MHz memory clock and a 25 MHz image-processing clock, and the dehazing pipeline operates in the 25 MHz clock domain. For same-clock paths in this image-processing domain, the post-place-and-route timing report gives a maximum operating frequency of 32.98 MHz, which is higher than the selected 25 MHz operating frequency. The VGA timing uses 800 horizontal clock periods and 525 vertical lines; therefore, the output frame rate is $25\,\text{MHz}/(800 \times 525) = 59.52\,\text{fps}$, which is reported as approximately 60 fps. The RTL implements a streaming pipeline in which active image data are processed at the pixel-clock rate. The register-level alignment latency from the input valid signal to the output valid signal is approximately 25 clock cycles, or 1.0 µs. This value denotes the pixel-path alignment latency and excludes the frame-level parameter-update interval. Frame-level statistics and atmospheric-light-related parameters are updated between frames and applied to the subsequent frame, without interrupting the active pixel stream.

5.2 Objective Evaluation Index Analysis

This paper utilizes the SOTS-Outdoor test set from the RESIDE dataset as the objective evaluation benchmark. This test set comprises 500 pairs of hazy and corresponding ground-truth haze-free images, which are used to evaluate the dehazing performance of the algorithm. We select metrics such as PSNR (Peak Signal-to-Noise Ratio), SSIM (Structural Similarity), EN (Entropy), and AG (Average Gradient) for comprehensive analysis. Specifically, PSNR is employed to evaluate the reconstruction fidelity of the dehazed image relative to the ground truth, where a higher value indicates less distortion. SSIM assesses image similarity across three dimensions: luminance, contrast, and structure, with larger values representing superior similarity. EN is utilized to measure the information content within an image, where a higher value implies richer information in the dehazed results. Lastly, AG reflects the sharpness of the image, and a larger value denotes stronger edge and detail information.

Metric	[2]	[25]	[12]	Ours
PSNR (dB)	17.5764	20.6539	19.1373	21.3008
SSIM	0.8544	0.8733	0.8822	0.8937
EN	6.7381	7.1071	7.1864	7.2412
AG	5.9894	7.3833	6.4595	7.2830
TIME (s)	9.2670	0.1490	0.0641	0.1089

Tab. 1. Quantitative comparison of different dehazing methods.

Resource type	Used	Available	Utilization (%)
Logic elements	8116	10320	79%
I/O pins	78	180	43%
Memory bits	263304	423936	62%
DSP blocks	45	46	98%
PLLs	1	2	50%

Tab. 2. FPGA resource utilization.

Table 1 illustrates the performance comparison of various image dehazing algorithms. The proposed algorithm achieves the highest performance in terms of PSNR, SSIM, and EN. Compared with [2], [25], and [12], the PSNR is improved by 3.7244 dB, 0.6469 dB, and 2.1635 dB, respectively. The SSIM is improved by 0.0393, 0.0204, and 0.0115, respectively, while the EN is improved by 7.47%, 1.89%, and 0.7626%, respectively. Although the AG value is slightly lower than that of Huang [25], it shows improvements of 21.6% and 12.75% over [2] and [12], respectively. In terms of processing time, the proposed algorithm is slightly slower than the method of Wang et al. [12]; however, considering the overall balance among all metrics, the difference remains within an acceptable range.

Table 2 presents the FPGA resource utilization summary. Specifically, the design uses 8,116 logic elements (79%), 78 I/O pins (43%), 263,304 on-chip memory bits (62%), 45 DSP blocks (98%), and one phase-locked loop (PLL) (50%). Benefiting from the lookup-table (LUT)-based optimization and the fully pipelined design, the system successfully implements the proposed algorithm within limited hardware resources while achieving real-time processing. When dealing with the 640×480 resolution video stream output by the OV5640 camera, the system sustains a processing rate of 60 fps.

5.3 Subjective Evaluation Analysis

As illustrated in Fig. 7, we select four different outdoor hazy images to subjectively compare the dehazing performance of different methods. Since the primary innovations of this paper include the specialized processing of sky regions, all four selected images are outdoor scenes containing sky areas. From left to right, the columns sequentially display

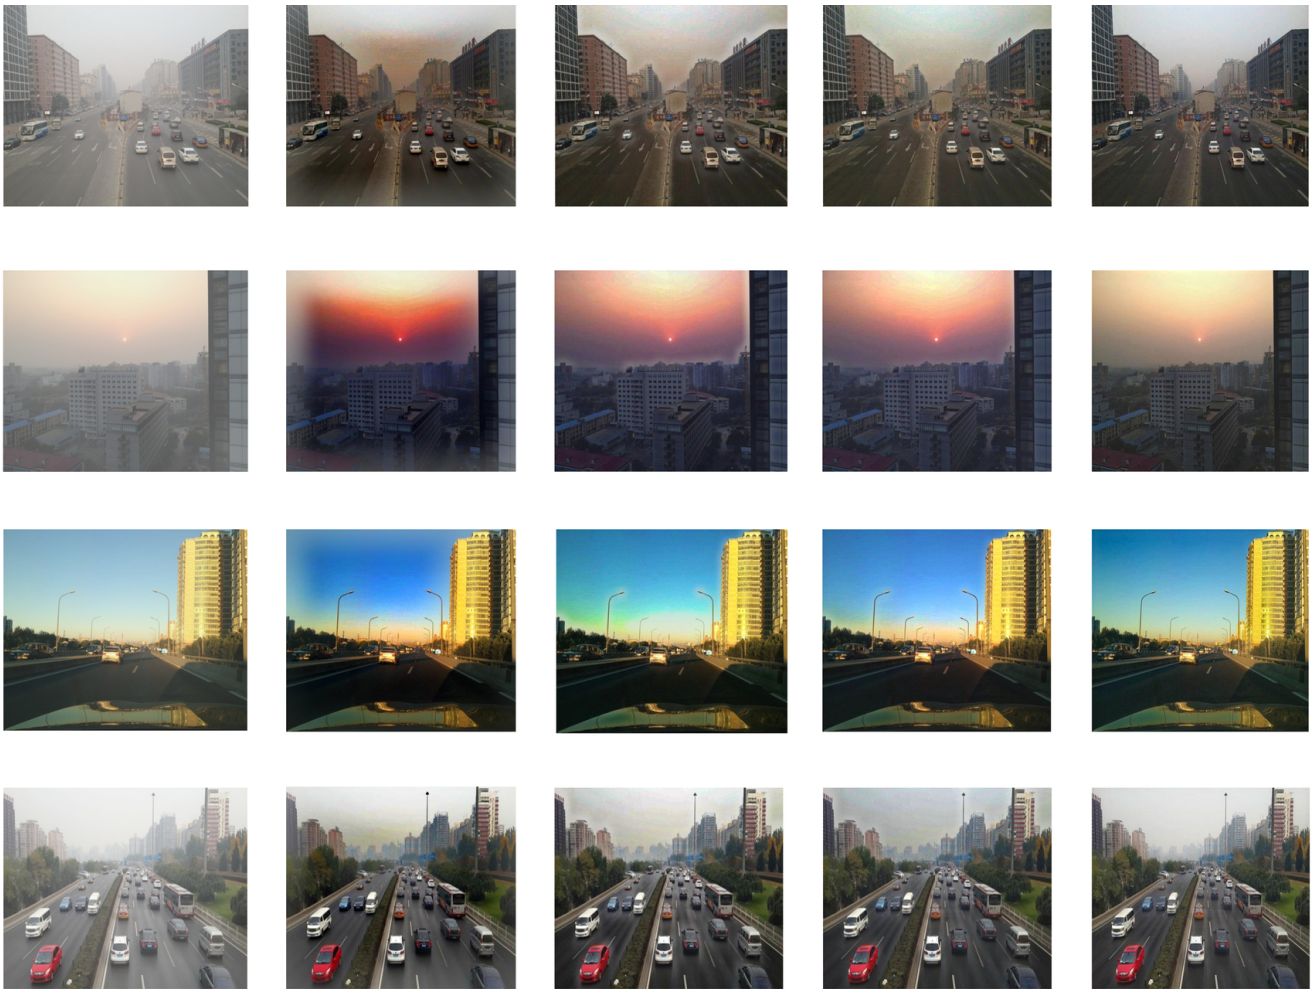


Fig. 7. Visual comparison on four outdoor scenes: input, References [2], [25], [12], and ours from left to right.

the original hazy images and the processing results of [2], Huang's method, Wang's method, and our proposed method. As shown in Fig. 7, the results of [2] and Huang's method exhibit visible halo artifacts around some object boundaries, whereas these artifacts are less pronounced in the results of Wang's method and the proposed method. Within the sky regions, Reference [2] exhibits the most prominent overexposure artifacts. In particular, obvious color distortion is visible in the second image. Compared with Wang et al. [12], the dual-atmospheric-light strategy employed in this work alleviates overexposure and color distortion in sky regions. Compared with Huang et al. [25], which uses binary hard segmentation as its sky-region representation, this work constructs a continuous soft sky weight map that enables a smooth transition between sky and non-sky regions, thereby reducing boundary halo artifacts and improving color consistency. Furthermore, adaptive gamma correction is introduced to better recover the overall brightness and details in low-luminance regions. Across these four examples, our proposed method achieves the best overall visual balance among the compared methods. In the sky regions, our framework suppresses overexposure and color distortion, resulting in more natural color reproduction. Thanks to the integrated soft sky segmentation strategy, visible halo artifacts are reduced at the junction between the sky and non-sky regions. In addition, no obvious halo artifacts are observed at the edges of objects in non-sky areas.

Beyond the visual comparison, the proposed work provides an FPGA-oriented realization of the sky-aware dehazing algorithm. The sky-region processing, dual-atmospheric-light estimation, fast guided filtering, and adaptive gamma correction are integrated into a streaming pixel-processing architecture. Atmospheric-light parameters are updated at frame boundaries and reused in the following frame, which avoids full-frame buffering while maintaining continuous active-pixel processing. Compared with conventional FPGA implementations of DCP-based dehazing, the proposed design provides a complete hardware pipeline for sky-aware restoration rather than accelerating only the basic DCP computation. Therefore, the contribution of this work includes both the restoration strategy and its real-time implementation on a resource-constrained FPGA platform.

5.4 Sky-Region Segmentation Validation

To validate the effectiveness of the sky-region segmentation module, 50 representative outdoor hazy images were selected from the test set. The sky regions were manually

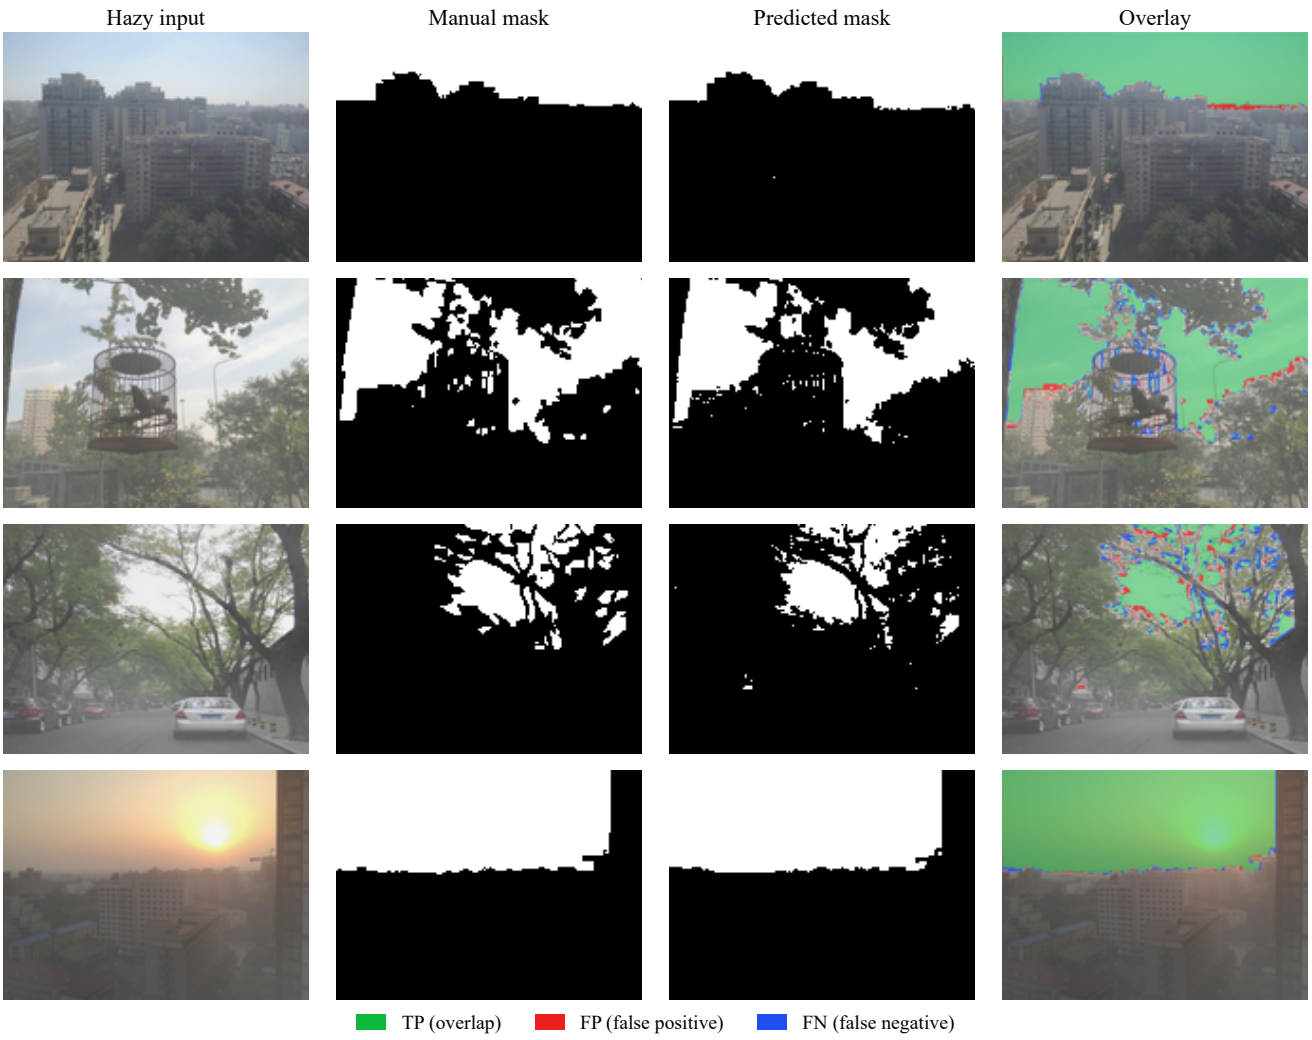


Fig. 8. Representative results of sky-region segmentation validation: hazy input images, manually annotated sky masks, predicted binary sky masks, and error overlays from left to right. Green, red, and blue denote true-positive, false-positive, and false-negative pixels, respectively.

annotated according to the corresponding haze-free reference images, and the manual annotations were used as reference masks. The reference masks were compared pixel by pixel with the binary sky masks generated by the MATLAB implementation. Precision, recall, F1-score, and intersection over union (IoU) were used as the evaluation metrics.

Table 3 shows that the average precision, recall, F1-score, and IoU are 0.9253, 0.9707, 0.9453, and 0.8996, respectively. The high recall indicates that the proposed method covers most of the manually annotated sky regions, while the IoU close to 0.90 demonstrates good spatial agreement between the predicted and reference masks. The pixel accuracy is 0.9644. Some deviations remain in images with small sky regions or complex sky boundaries, which is also reflected by the lower IoU compared with the recall. Representative segmentation results are shown in Fig. 8.

As shown in Fig. 8, four images are selected to illustrate the segmentation results. From the first to the fourth columns, the figure presents the original hazy images, manually annotated sky masks, sky masks predicted by the proposed method, and overlays of the manually annotated and predicted sky regions, respectively. The first row shows a typical hazy scene; the second and third rows show scenes with complex sky boundaries formed by foliage and other objects; and the fourth row shows a dusk scene. Because manual delineation of intricate sky boundaries may not be pixel-perfect, the manual masks in the second column and the corresponding overlays in the fourth column may contain minor annotation errors. Therefore, the original images in the first column and the predicted masks in the third column provide additional visual evidence for assessing the segmentation performance.

Metric	Value
Precision	0.9253
Recall	0.9707
F1-score	0.9453
IoU	0.8996

Tab. 3. Pixel-level evaluation of the sky-region segmentation results on 50 manually annotated images.

6. Conclusion

This paper proposes and implements an FPGA-based real-time image dehazing system through algorithmic optimization and hardware architecture design. At the algorithmic level, to address the limitations of the traditional dark channel prior algorithm, this study introduces a spatial prior and an adaptive threshold to construct a soft sky weight map. Concurrently, a dual atmospheric light estimation strategy is proposed to model the atmospheric light in different regions, and an adaptive gamma-correction mechanism is integrated during the image restoration phase. This methodology reduces artifacts in high-luminance regions while improving shadow details and overall contrast. At the hardware implementation level, this paper designs a streaming, fully pipelined FPGA architecture for real-time video processing. By using the on-chip M9K block RAMs to implement lookup tables (LUTs), direct hardware implementation of nonlinear exponential operations is avoided, reducing the associated arithmetic-resource overhead. The real-time dehazing system is successfully deployed on a Cyclone IV E development board, sustaining a processing rate of 60 fps at a resolution of 640×480.

The proposed method is primarily applicable to daytime outdoor hazy images, particularly in road surveillance, vehicle vision, and long-distance urban imaging scenarios that contain a clearly identifiable sky region and require real-time processing on resource-constrained platforms. However, the detection of sky candidate regions employs a positional prior that restricts the search to the upper 72% of the image. When the sky is located in the lower part of the image, heavily occluded by buildings or trees, or absent from the image, the accuracy of sky segmentation and dual-atmospheric-light estimation may decrease. In addition, scenes captured entirely at night may not satisfy the current sky-discrimination criteria or the assumptions of the dark channel prior, potentially resulting in incorrect segmentation or color bias. Accordingly, future work will focus on extending the proposed framework to nighttime image dehazing.

References

[1] NARASIMHAN, S. G., NAYAR, S. K. Contrast restoration of weather degraded images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2003, vol. 25, no. 6, p. 713–724. DOI: 10.1109/TPAMI.2003.1201821

[2] HE, K., SUN, J., TANG, X. Single image haze removal using dark channel prior. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Miami (FL, USA), 2009, p. 1956–1963. DOI: 10.1109/CVPR.2009.5206515

[3] DAI, Z., JALAB, H. A., LIU, S. Haze removal algorithm using improved restoration model based on dark channel prior. In *4th International Visual Informatics Conference (IVIC)*. Bangi (Malaysia), 2019, p. 157–169. DOI: 10.1007/978-3-030-34032-2_15

[4] XIAO, J., ZHU, L., ZHANG, Y., et al. Scene-aware image dehazing based on sky-segmented dark channel prior. *IET Image Processing*, 2017, vol. 11, no. 12, p. 1163–1171. DOI: 10.1049/iet-ipr.2017.0058

[5] NARASIMHAN, S. G., NAYAR, S. K. Vision and the atmosphere. *International Journal of Computer Vision*, 2002, vol. 48, no. 3, p. 233–254. DOI: 10.1023/A:1016328200723

[6] YANG, G., EVANS, A. N. Improved single image dehazing methods for resource-constrained platforms. *Journal of Real-Time Image Processing*, 2021, vol. 18, no. 6, p. 2511–2525. DOI: 10.1007/s11554-021-01143-6

[7] LIU, Y., LI, H., WANG, M. Single image dehazing via large sky region segmentation and multiscale opening dark channel model. *IEEE Access*, 2017, vol. 5, p. 8890–8903. DOI: 10.1109/ACCESS.2017.2710305

[8] CHEN, E., GONG, Y., TIE, Y. (eds.) *Advances in Multimedia Information Processing – PCM 2016: 17th Pacific-Rim Conference on Multimedia, Xi'an, China, September 15–16, 2016, Proceedings, Part II*. Cham (Switzerland): Springer International Publishing, 2016. . DOI: 10.1007/978-3-319-48896-7

[9] RU, Y., TANAKA, G. Proposal of dehazing method and quantitative index for evaluation of haze removal quality. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 2017, vol. E100.A, no. 4, p. 1045–1054. DOI: 10.1587/transfun.E100.A.1045

[10] XU, H., GUO, J., LIU, Q., et al. Fast image dehazing using improved dark channel prior. In *IEEE International Conference on Information Science and Technology (ICIST)*. Wuhan (China), 2012, p. 663–667. DOI: 10.1109/ICIST.2012.6221729

[11] YANG, G., YANG, H., YU, S., et al. A multi-scale dehazing network with dark channel priors. *Sensors*, 2023, vol. 23, no. 13, p. 1–13. DOI: 10.3390/s23135980

[12] NING, K., WANG, T., GUO, W., et al. FPGA-accelerated real-time dehazing using multi-scale dark channel (in Chinese). *Electronics Optics & Control*, 2026, p. 1–9. [Online] Cited 2026-07-11. Available at: https://link.cnki.net/urlid/41.1227.TN.20260325.1058.002

[13] WU, X. A real-time framework for HD video defogging using modified dark channel prior. *Journal of Real-Time Image Processing*, 2024, vol. 21, no. 2, p. 1–55. DOI: 10.1007/s11554-024-01432-w

[14] LIU, S., CHEN, P., LAN, J., et al. Underwater image restoration via multiscale optical attenuation compensation and adaptive dark channel dehazing. *Computers and Electrical Engineering*, 2025, vol. 123, p. 1–18. DOI: 10.1016/j.compeleceng.2025.110228

[15] LEE, Y.-L., HWANG, Y.-T. A light-weight image dehazing scheme and its hardware accelerator design in FPGA. In *IEEE International Conference on Consumer Electronics - Taiwan (ICCE-Taiwan)*. Kaohsiung (Taiwan), 2025, p. 347–348. DOI: 10.1109/ICCE-Taiwan66881.2025.11207769

[16] DU, G. A real-time effective fusion-based image defogging architecture on FPGA. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 2021, vol. 17, no. 3, p. 1–21. DOI: 10.1145/3446241

[17] WU, X., LI, Y., HAN, X., et al. Design of HD real-time video defogging system based on FPGA. In *9th International Conference on Integrated Circuits and Microsystems (ICICM)*. Wuhan (China), 2024, p. 831–836. DOI: 10.1109/ICICM63644.2024.10814178

[18] CHARUGUNDLA, N. S., INGUVA, A., REDDY, C. V., et al. FPGA and ASIC realization of a saturation-driven image dehazing technique. In *IEEE DELCON - International Conference on Recent Smart Technologies in Engineering for Sustainable Development*. New Delhi (India), 2025, p. 1–8. DOI: 10.1109/DELCON68055.2025.11400136

[19] ZHOU, Y., XU, H., JIANG, G., et al. UIE-SFIFormer: Underwater image enhancement based on physical-guided spatial–frequency interaction transformer. *IEEE Journal of Oceanic Engineering*, 2025, vol. 50, no. 2, p. 727–742. DOI: 10.1109/JOE.2024.3458109

[20] XU, L. C., HAN, J., XIE, Y. M. Low-latency defogging algorithm for camera monitor systems implemented on ZYNQ. In *Photonics & Electromagnetics Research Symposium - Fall (PIERS-Fall)*. Chiba (Japan), 2025, p. 1–7. DOI: 10.23919/PIERS-Fall62445.2025.11394346

[21] UPADHYAY, B., YADAV, S. K., SARAWADEKAR, K. P. VLSI architecture of saturation based image dehazing algorithm and its FPGA implementation. In *IEEE 65th International Midwest Symposium on Circuits and Systems (MWSCAS)*. Fukuoka (Japan), 2022, p. 1–4. DOI: 10.1109/MWSCAS54063.2022.9859535

[22] LI, J., GUO, Y., LI, G. FPGA implementation of dark channel prior and defogging for video images. Chapter in *Electronic Engineering and Informatics*. Amsterdam (The Netherlands): IOS Press, 2024, p. 705–712. DOI: 10.3233/ATDE240136

[23] YOONG, N. K. J. *CLAHE HLS Implementation on Zynq SoC FPGA*. Final Year Project, Nanyang Technological University, Singapore, 2024. [Online] Cited 2025-03-28. Available at: https://dr.ntu.edu.sg/handle/10356/176412

[24] RAHMAWATI, L., RUSTAD, S., MARJUNI, A., et al. Transmission map refinement using Laplacian transform on single image dehazing based on dark channel prior approach. *Cybernetics and Information Technologies*, 2024, vol. 24, no. 4, p. 126–142. DOI: 10.2478/cait-2024-0039

[25] HUANG, H., HU, K., YANG, L., et al. FPGA design of image dehazing system for intelligent dashcam (in Chinese). *Journal of Shanghai Jiao Tong University*, 2024, vol. 58, no. 4, p. 565–578. DOI: 10.16183/j.cnki.jsjtu.2023.026

About the Authors . . .

Hongmin CHEN was born in 2003. Hongmin Chen is currently pursuing the M.S. degree at the School of Artificial Intelligence and Computer Science, North China University of Technology, Beijing, China. The research interest is FPGA-based image dehazing.

Wei SONG was born in 1980. Wei Song holds a Ph.D. degree and is currently a Professor at the School of Artificial Intelligence and Computer Science, North China University of Technology, Beijing, China. The research interests include data mining and recommender systems.

Keqing NING (corresponding author) was born in 1983. Keqing Ning holds a Ph.D. degree and is currently a Senior Laboratory Engineer at the School of Artificial Intelligence and Computer Science, North China University of Technology, Beijing, China. The research interests include embedded systems and image processing.