

Article

Adaptive Image Enhancement Method for Object Recognition Based on Statistical Photometric Characteristics

Chae-yeong Kim and Soon-kak Kwon * 

Department of Computer Software Engineering, Dong-eui University, Busan 47340, Republic of Korea; rainbow545454@naver.com

* Correspondence: skkwon@deu.ac.kr; Tel.: +82-51-890-1727

Abstract

We propose an adaptive image enhancement method based on photometric statistics to improve object detection under adverse illumination conditions. Conventional image enhancement methods primarily target perceptual quality and may alter recognition-relevant features, potentially degrading detection performance. In contrast, the proposed method adaptively determines the enhancement intensity by combining a predefined domain-specific preset with a photometric risk score calculated from the photometric statistics of the input image to estimate the risk of enhancement-induced photometric risks. Based on these estimates, the enhancement intensity is adaptively controlled, and unnecessary transformation is conditionally bypassed. Experiments on the Berkeley DeepDrive 100K dataset using YOLOv11 demonstrate that the proposed method improves F1-score and mean Average Precision by 0.101 and 0.195, respectively, compared with unprocessed images under adverse conditions, including low-light environments. These results demonstrate that detector-oriented adaptive enhancement can improve robustness while reducing performance degradation caused by unnecessary or excessive image transformation.

Keywords: image enhancement; object recognition; photometric statistic; task-oriented enhancement

1. Introduction

Object recognition systems should operate reliably under adverse and dynamic lighting conditions ranging from insufficient illumination to extreme overexposure for safety-critical applications including autonomous driving [1,2]. However, these object recognition systems exhibit performance drops on these degraded images. Although applying image enhancement methods can improve perceptual quality, they often alter the feature distribution during the recognition process, causing it to deviate from the training data and consequently diminishing recognition accuracy [3,4].

Various methods to enhance image quality have been widely adopted, including global pixel-wise transformations such as gamma correction and histogram equalization, as well as contrast-limited adaptive histogram equalization (CLAHE). Adaptive image enhancement methods adjust enhancement strength according to local image characteristics such as brightness distribution or exposure level. Recently, deep learning models for image enhancement have achieved impressive visual quality improvements. However, most existing image enhancement focuses on improving perceptual image quality rather than recognition performance. Several studies have shown that image transformations



Academic Editor: Atsushi Mase

Received: 19 June 2026

Revised: 29 July 2026

Accepted: 30 July 2026

Published: 1 August 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 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

aimed at improving visual quality do not necessarily improve object recognition performance and may even lead to an unexpectedly decline in accuracy [5–7]. This degradation occurs because image transformations can alter the distribution of visual features, causing them to deviate from the feature distribution learned during detector training. This presents the fundamental challenge of controlling image enhancement to improve visibility without disrupting the feature distributions required for object detection. This study assesses the risk of enhancement-induced feature distribution shifts using photometric statistics and investigates whether controlling transformation intensity accordingly can consistently improve object detection performance under adverse imaging conditions. Based on this assessment, the proposed method aims to adaptively control the image transformation intensity to achieve an improvement in object detection performance under adverse imaging conditions.

We propose an adaptive image enhancement method based on photometric statistics. In this study, a photometric risk refers to the potential for enhancement-induced changes in image photometric properties such as brightness and contrast to adversely affect the detector responses. The proposed method adaptively determines the enhancement intensity by estimating this photometric risk from the input image statistics. In addition, we define a photometric domain to represent luminance distribution characteristics such as overexposure, underexposure, and non-uniform illumination. This photometric domain of an image is determined by statistical measures of luminance and its left–right intensity distribution. Simultaneously, a photometric risk score quantifying the likelihood of photometric risk is computed based on the measured statistics. The image is enhanced using gamma correction, CLAHE, and horizontal luminance balancing. The image enhancement intensity is adaptively determined based on the photometric domain and the photometric risk score. In addition, the photometric risk score is used to adjust the model’s confidence threshold to compensate for the reduction in confidence caused by image transformations across various object-related vision tasks. The proposed method is evaluated using the Berkeley DeepDrive 100K that includes adverse-condition images, YOLO object detector, and standard detection metrics, including precision, recall, F1-score, and mean average precision. The experiment results demonstrate that the proposed method significantly improves object detection performance under adverse imaging conditions. This work makes the following contributions:

- We propose an adaptive image enhancement method optimized for recognition tasks, where the enhancement intensity is regulated based on the photometric risk derived from photometric statistics.
- We present a domain-driven enhancement intensity control approach that effectively prevents over-enhancement by establishing domain-specific upper bounds. Specifically, we determine the target image domain based on photometric statistics, thereby applying corresponding presets to constrain the enhancement intensity.
- We introduce a skip mechanism that bypasses the enhancement process for already optimized images for recognition tasks. This prevents performance degradation caused by unnecessary image distortion while significantly improving computational efficiency.

The remainder of this paper is organized as follows. Section 2 reviews the characteristics of existing image enhancement methods. Section 3 presents the proposed method. Section 4 describes the experimental setup and compares the object detection performance of the proposed method with that of existing enhancement methods. Finally, Section 5 concludes the paper.

2. Related Works

Traditional image enhancement methods improve image visibility by transforming pixel intensities. Representative techniques include gamma correction [8], histogram equalization [9], and contrast-limited adaptive histogram equalization (CLAHE) [10,11]. These methods enhance contrast within global or local regions at a low computational cost. However, these methods typically apply unconditional global transformations without considering spatial illumination variability. In practical environments, such uniform adjustments may introduce artifacts such as noise amplification, structural distortion, or exaggerated contrast boundaries. From the perspective of object recognition, these transformations can impair performance by creating a mismatch between the actual feature distribution and the one learned during training [5]. Adaptive and content-aware enhancement methods dynamically adjust transformation parameters based on image characteristics. Examples include adaptive gamma correction [12] and adaptive CLAHE [13], which optimize enhancement strength according to brightness distribution or histogram statistics. However, many conventional adaptive methods lack an explicit switching or gating mechanism to determine the necessity of preprocessing. Therefore, they may apply redundant transformations to already stable frames, potentially introducing unnecessary statistical perturbations.

According to Retinex theory, the intensity of an image pixel is determined by the product of the illumination component, which describes the spatial distribution of light, and the reflectance component, which represents the inherent physical properties of an object. Retinex-based image enhancement decomposes an image into these two components and suppresses non-uniform illumination that hinders visibility to restore the original details and color. However, applying a small spatial scale often leads to color distortion, while a large scale results in image blurring or loss of detail. Furthermore, halo artifacts, characterized by a light-smearing effect, tend to occur at boundaries with abrupt brightness changes, such as the edge between a dark building and a bright sky. To address these limitations, the Multi-scale Retinex methods were proposed to combine the results at various scales. Petro et al. [14] addressed color restoration and post-processing limitations in Retinex-based enhancement by introducing RGB- and luminance-based strategies for balancing color correction and local contrast enhancement. The Multi Scale Retinex with Color Restoration (MSRCR) algorithm [15] applies Multiscale Retinex independently to each RGB channel and incorporates a color restoration step to compensate for saturation loss while enhancing local contrast under non-uniform illumination. Nevertheless, the Retinex-based methods focus primarily on perceptual restoration and do not explicitly consider detector-oriented input stability.

Recently, it has become popular to incorporate deep learning into image enhancement research. Deep learning-based methods are generally divided into supervised learning and unsupervised learning approaches. The methods based on supervised learning are trained using paired datasets consisting of degraded input images and their corresponding clean ground-truth images. The network is optimized to minimize the discrepancy between the enhanced output and the ground-truth image. Wang et al. [16] proposed DeepUPE, an end-to-end underexposed image enhancement network that learns an image-to-illumination mapping rather than regressing an enhanced image. The method estimates a multi-channel illumination map using local and global image features and employs illumination-aware reconstruction, smoothness, and color losses to recover natural exposure, contrast, and color. Xu et al. [17] proposed an SNR-aware low-light image enhancement framework that combines a transformer-based long-range branch with a convolution-based short-range branch. The method uses an estimated SNR map to adaptively emphasize non-local information in severely degraded low-SNR regions and local information in relatively high-SNR

regions, while its SNR-guided attention suppresses contributions from noise-dominated tokens. Cui et al. [18] proposed the Illumination Adaptive Transformer, which combines local pixel-wise adjustment with global ISP-related parameter estimation to address low-light enhancement and exposure correction within a unified framework. Cai et al. [19] proposed a one-stage Retinex-based Transformer that jointly estimates illumination and restores enhancement-induced corruptions. This approach generally leads to stable training behavior and superior enhancement performance. However, paired datasets containing perfectly aligned ground-truth images are notoriously difficult to acquire in practice.

To overcome the limitation of the supervised learning networks, unsupervised approaches that can be trained without reference images have been proposed. Representative methods include generative model-based approaches such as EnlightenGAN [20], CycleGAN [21], and diffusion model-based frameworks, which learn to transform degraded images into clean images. In GAN-based approaches, the generator produces visually enhanced images while the discriminator learns to distinguish generated images from real images, thereby encouraging perceptually realistic results. Nevertheless, these generative approaches may produce hallucinated textures or non-existent objects during the enhancement process. These hallucination artifacts can degrade downstream recognition performance by disrupting the structural and geometric consistency critical for object recognition [6]. In addition, the image enhancement methods based on the ground-truth images or the reference images often suffer from limited generalization to out-of-distribution domains due to the domain gap between training data and real-world applications. To address these issues, zero-reference methods have been explored that perform enhancement using only the input image itself. Instead of relying on paired labels, these methods utilize loss functions based on constraints such as exposure control, color consistency, and spatial coherence. This guides the network to enhance image quality while preserving the structural and statistical characteristics of the input, thereby improving generalization across diverse domains. Guo et al. [22] proposed Zero-Reference Deep Curve Estimation (Zero-DCE), which formulates low-light enhancement as pixel-wise, image-specific curve estimation. The network is trained without paired or unpaired reference images using non-reference loss functions that guide exposure, color, and spatial consistency. Zero-DCE++ [23] is a lightweight extension of Zero-DCE that estimates reusable pixel-wise enhancement curves from downsampled inputs while retaining the zero-reference learning framework. Zheng and Gupta [24] proposed a semantic-guided zero-shot framework that combines pixel-wise enhancement factor estimation with recurrent low-light image enhancement. However, these approaches remain limited in that they focus primarily on visual enhancement and do not account for performance gains in downstream object recognition tasks.

3. Proposed Method

We propose the adaptive image enhancement method to improve object recognition performance. The proposed method analyzes the input image by utilizing photometric statistics in order to determine its domain and adjusts the intensity of image enhancement according to domain-specific characteristics. This is motivated by the observation that improvements in perceptual image quality do not necessarily translate into improved detector robustness. The input image is analyzed to measure various metrics of illumination and saturation distributions. Based on these metrics, the image domain is determined to facilitate the classification of image quality degradation caused by varying imaging conditions. In addition, a photometric risk score is computed to quantify photometric risk induced by image transformation. Subsequently, it is determined whether to skip the image transformation step to prevent excessive enhancement from degrading object recognition

performance on images with already sufficient quality. During object recognition, the corresponding image transformations can induce a decline in confidence scores. To compensate for this, the model's confidence threshold is adjusted based on the photometric risk score. Figure 1 illustrates the flow of the proposed method.




Figure 1. Flow of the proposed method.

3.1. Photometric Statistics Measurement

The proposed method utilizes a set of illumination statistics to analyze photometric conditions and estimate the photometric risk induced by image transformation. For this analysis, each input image is converted into the LAB color space, and only the Lightness (L) channel is used. The range of the pixel values is [0, 255]. The median (m) and 95th percentile (p_{95}) of luminance values are computed to analyze the global photometric condition. These statistics characterize the effective luminance distribution through its central and upper portions of the luminance distribution. These statistics provide measures of illumination deficiency across the luminance distribution. For analyzing structural instability, the coefficient of luminance variation (CV) is computed as follows:

$$CV = \sigma / (\mu + \varepsilon), \quad (1)$$

where σ and μ are the standard deviation and mean of luminance, respectively, and ε is a stabilization constant to prevent division by zero. In this paper, ε is set to 1×10^{-6} . The CV serves as an indicator of potential degradation in detector feature stability under contrast collapse or noise-dominated conditions. To assess left–right illumination imbalance (dLR), the image is divided into two equally sized regions along its vertical centerline. This partition is intended to capture global lateral illumination asymmetry while reducing sensitivity to local objects, textures, and structural variations. Then, the illumination imbalance is calculated from the difference between the mean luminance values of the left and right halves as follows:

$$dLR = |\mu_L - \mu_R| / (\mu_L + \mu_R), \quad (2)$$

where μ_L and μ_R are the mean luminance of the left and right regions, respectively. The illumination imbalance quantitatively represents the likelihood of spatial asymmetry in feature representation arising from luminance imbalance between the left and right regions.

3.2. Photometric Domain Determination and Photometric Risk Estimation

To categorize images based on the stability of feature representations in object detectors, we define the photometric domains as summarized in Table 1. The photometric domains are categorized based on types of adverse photometric variations, such as high exposure, low exposure, and non-uniform illumination. These variations degrade object recognition performance by introducing a photometric risk between the learned feature representations and those extracted from input images [25]. The photometric domains are determined in a rule-based manner based on the metrics for the spatial photometric characteristics because these domains are defined based on statistical conditions rather than semantic

categories. The rationale underlying the photometric domain determination criteria in Table 1 is described in Appendix A. If an image satisfies the criteria of multiple domains, the domain with the highest priority is selected.

Table 1. Photometric domain definitions and criteria.

Priority	Domain Name	Criteria	Description
1	Night-dark	$m < 28$ and $p95 < 80$	extremely low exposure, blurring of structural details
2	Indoor	$dLR > 0.40$ and $m < 130$ and $p95 < 200$ and $CV > 0.18$	lateral luminance imbalance caused by local lighting
3	Night	$m < 70$ and $p95 < 180$	low global illumination, contrast degradation
4	Day	$CV < 0.30$ and $p95 > 200$	high exposure and clear visibility
5	Dusk	Otherwise	intermediate illumination

Along with the domain determination, a photometric risk score is calculated to predict a photometric risk by image transformation. Partial risk scores for the photometric statistics are calculated as follows:

$$\begin{aligned}
 s_m &= \text{clamp}\left(\frac{70-m}{70}\right) \\
 s_{p95} &= \text{clamp}\left(\frac{180-p95}{180}\right) \\
 s_{CV} &= \text{clamp}\left(\frac{CV-0.3}{0.4}\right) \\
 s_{dLR} &= \text{clamp}\left(\frac{dLR}{0.5}\right),
 \end{aligned} \tag{3}$$

where $\text{clamp}(\cdot)$ is a clamping function that constrains the values to the range $[0, 1]$. Consequently, each score is bounded between 0 and 1. The terms s_m and s_{p95} quantify the photometric risks associated with insufficient median and upper-percentile luminance, respectively. Accordingly, the reference values of 70 and 180 were adopted from the criteria used to determine the ‘Night’ domain. The term s_{CV} represents the photometric risk associated with excessive spatial luminance variation. Its reference value was set to 0.3 based on the criterion of the ‘Day’ domain, indicating that luminance variation begins to contribute to the risk score when it exceeds the range associated with relatively uniform daytime illumination. The denominator of 0.4 defines the normalization range over which the risk score increases from 0 to 1. The overall risk score s is computed as a weighted sum of the photometric risk scores of individual metrics as follows:

$$s = \sum_{k \in \{m, p95, CV, dLR\}} s_k w_k, \tag{4}$$

where w_k is the weight for each metric. The w_m , w_{p95} , w_{CV} , and w_{dLR} are set to 0.35, 0.25, 0.25, and 0.15, respectively, based on the assumed relative contribution of each illumination factor to the overall risk score. The overall risk score s is also bounded between 0 and 1. The photometric risk score determines the intensity of the image transformation.

3.3. Image Enhancement Process Based on Domain Preset

The input image is processed using a diverse set of enhancement techniques including gamma correction, CLAHE, and horizontal luminance balancing. Gamma correction mitigates contrast compression, ensuring that details in dark areas are effectively enhanced. CLAHE recovers structural details in dark environments by boosting local contrast. Hori-

horizontal luminance balancing mitigates left–right illumination imbalance to prevent biased feature extraction caused by asymmetrical lighting conditions. The strength of each image enhancement technique is modulated by its corresponding parameter, θ_γ , θ_{clip} , and θ_{gain} . To prevent color distortion, these enhancement processes are applied solely to the L channel. Before the transformation, the pixel values are normalized to $[0, 1]$.

The gamma correction is applied as follows:

$$O_\gamma(x, y) = I(x, y)^{1/(\theta_\gamma + \epsilon)}, \quad (5)$$

where $I(x, y)$ is the luminance value of (x, y) pixels in the input image. CLAHE is applied with a grid size of 8×8 and an adaptive clip limit determined by θ_{clip} . Horizontal luminance balancing is performed using column-wise gain balancing as follows:

$$\begin{aligned} O_{gain}(x, y) &= \text{clamp}(g(x)I(x, y)) \\ g(x) &= (\mu / (\mu_{col}(x) + \epsilon))^{\theta_{gain}}, \end{aligned} \quad (6)$$

where $g(x)$ is a gain function and μ_{col} is the luminance mean of the x -th column.

Although image transformation can visually enhance an image, it can also amplify noise or cause oversaturation. Such artifacts may distort feature representations and consequently degrade object recognition performance. To counteract this, domain-specific upper bounds are established for each technique to preserve feature-level robustness. In the ‘Night-dark’ domain, the luminance range is severely limited. Therefore, the gamma correction and CLAHE are not applied due to the risk of noise amplification during contrast enhancement. In the ‘Night’ domain, aggressive CLAHE is employed to mitigate the adverse effects of local illumination sources. Images in the ‘Indoor’ domain are characterized by severe local contrast distortions caused by locally dominant light sources. Accordingly, horizontal luminance balancing is intensively applied to this domain. The ‘Dusk’ domain is characterized by mixed illumination and abrupt brightness variations. To address this, all the image enhancement techniques are applied aggressively. For images in the ‘Day’ domain, strong gamma correction is applied to correct oversaturated illumination.

We design domain-specific presets that define the maximum enhancement intensity based on the characteristics of each domain as shown in Table 2. These presets mitigate photometric risks caused by excessive enhancement and ensure that the feature representation after image enhancement remains within the learned distribution of the detector. In Table 2, Θ_γ , Θ_{clip} , and Θ_{gain} are the preset parameters defining the operational limits for θ_γ , θ_{clip} , and θ_{gain} , respectively, to cap the maximum enhancement intensity. The derivation of these presets is elaborated in Appendix B.

Table 2. Domain-specific upper bounds of image enhancement parameters.

Domain	Θ_γ	Θ_{clip}	Θ_{gain}
Night-dark	1.00	0.0	0.06
Night	1.10	1.8	0.04
Dusk	1.20	2.0	0.04
Indoor	1.00	1.4	0.04
Day	1.30	1.4	0.08

The θ_γ is determined based on the photometric risk score for global luminance s_m as follows:

$$\theta_\gamma = \min(0.05s_m + 1.08, \Theta_\gamma(\text{domain})). \quad (7)$$

In (7), the intensity of gamma correction θ_γ is adaptively adjusted from a base level according to the illumination risk score s . The intensity of gamma correction is limited by the domain-specific constant Θ_γ to prevent excessive brightness enhancement. The parameters θ_{clip} and θ_{gain} are calculated as follows:

$$\begin{aligned}\theta_{clip} &= As_{CV}\Theta_{clip}(domain) \\ \theta_{gain} &= As_{dLR}\Theta_{gain}(domain),\end{aligned}\quad (8)$$

where A is an attenuation factor which is adaptively determined based on s_m , s_{CV} , and s_{dLR} as follows:

$$A = 1 - (\alpha_m s_m + \alpha_{CV} s_{CV} + \alpha_{dLR} s_{dLR}), \quad (9)$$

where α_m , α_{CV} , and α_{dLR} are weights for s_m , s_{CV} , and s_{dLR} , respectively. In this study, α_m , α_{CV} , and α_{dLR} are set to 0.35, 0.25, and 0.40, respectively. In (8), the domain-specific constants Θ_{clip} and Θ_{gain} determine the maximum intensities of the CLAHE and horizontal luminance balancing, respectively, while the corresponding partial risk scores control the proportion of each intensity applied. These intensity parameters are further modulated by the attenuation factor A . As the photometric risk increases, A decreases and suppresses the intensities of the corresponding image transformations.

3.4. Risk Score-Guided Confidence Threshold Adjustment

Even though the proposed image transformations are designed to preserve object recognition performance, they can still introduce photometric risks that reduce confidence scores. This reduction acts as a primary factor in increasing false negatives. To mitigate this effect, an adaptive confidence thresholding strategy is introduced to recalibrate the decision criteria based on the illumination and structural characteristics of each image. The adjusted confidence threshold is derived by applying a luminance-aware factor f_l and a risk-aware factor f_r . The luminance-aware factor is adaptively lowered according to the median pixel intensity of the input image. By default, f_l is set to 1.00. However, for images with extremely low illumination ($m < 50$), f_l is adjusted to 0.85. For moderately dark images ($m < 70$), it is set to 0.90. This adjustment lowers the confidence threshold in poorly lit environments to minimize potentially missed objects. The risk-aware factor f_r is calculated to reflect the structural instability and noise levels present in the input image as follows:

$$f_r = \begin{cases} \hat{f}_r = 0.94 - 0.10s & \\ 0.88\hat{f}_r, & \text{if } s > 0.6 \\ \hat{f}_r, & \text{otherwise} \end{cases} \quad (10)$$

In (10), the base factor $\hat{f}_r$ provides a linear response to the quantified environmental risk, ensuring that the confidence threshold is gradually relaxed as structural instability increases. The conditional formulation of f_r introduces an additional scaling penalty, which is triggered only in high-risk environments. This dual-stage relaxation encourages the detector to focus on maximizing recall in images with high photometric risks induced by the image enhancement. The adjusted confidence threshold, which is applied after detector inference to determine whether each candidate prediction, is determined by selecting the smaller factor between f_l and f_r as follows:

$$T_{conf}^* = T_{conf} \times \min(f_l, f_r), \quad (11)$$

where T_{conf} and T_{conf}^* are the base and adjusted confidence thresholds, respectively.

3.5. Selective Skip Mechanism

A skip mechanism is introduced to maintain object recognition performance by avoiding unnecessary feature modification. The enhancement process is selectively executed based on the following criteria:

$$G_{skip} = \begin{cases} 1, & s < T_{skip}(domain) \\ 0, & \text{otherwise} \end{cases}, \quad (12)$$

where G_{skip} is a binary indicator that determines whether the transformation process is bypassed. The domain-specific threshold $T_{skip}(domain)$ controls the sensitivity of the skip decision. A larger threshold increases the range of risk score s for which skipping is allowed, whereas a smaller threshold makes transformation activation more likely. This skip mechanism prevents unnecessary image transformations when the input image exhibits a low photometric risk. Table 3 presents the domain-specific threshold values, which were validated through the sensitivity analysis detailed in Appendix B. Images in the ‘Night’ and ‘Night-dark’ domains generally produce higher risk scores because of their lower median and upper-percentile luminance values. Therefore, relatively higher domain-specific thresholds are assigned to these domains to account for their elevated risk-score distributions. This setting allows images with comparatively low risk within each low-light domain to bypass unnecessary transformation, thereby reducing the likelihood of background-noise amplification and other transformation-induced artifacts, while ensuring that images with sufficiently high photometric risk are still enhanced.

Table 3. Domain-specific thresholds for skip determination.

Domain Name	Night-Dark	Night	Dusk	Indoor	Day
T_{skip}	0.61	0.58	0.52	0.52	0.51

4. Experiment Results and Discussion

We conducted experiments to evaluate the object detection performance achieved through the proposed image enhancement. We adopted YOLOv11n [26] as the baseline object detector, conducting experiments on the Berkeley DeepDrive 100K (BDD100K) object detection subset [4]. BDD100K is a large-scale, open-source autonomous driving dataset curated by the UC Berkeley AI Research lab, which extensively encompasses diverse weather conditions and time-of-day variations. The images in the BDD100K dataset have a resolution of 1280×720 pixels. YOLO11n, the nano-scale variant of the YOLO11 family, was adopted as the baseline object detector because it provides a favorable balance between detection accuracy and computational efficiency. As a single-stage detector, YOLO11n performs feature extraction, bounding-box regression, and object classification within a unified network. Before being fed into the network, the input images were rescaled to the configured YOLO input size of 640×640 pixels, and the corresponding bounding-box coordinates were transformed accordingly. For this input resolution, the model contains approximately 2.6 million parameters and requires approximately 6.5 billion floating-point operations (FLOPs). All experiments were conducted on a workstation equipped with an AMD Ryzen 7 3700X 8-Core Processor (Advanced Micro Devices, Inc., Santa Clara, CA, USA), an NVIDIA GeForce RTX 3080 Ti (NVIDIA Corporation, Santa Clara, CA, USA) with 12 GB of memory, and 32 GB of system memory, running Ubuntu 18.04 LTS.

For detector training and evaluation, we used a publicly available third-party derivative of BDD100K [27] which comprises 70,000 training images, 10,000 validation images, and 20,000 test images and provides object detection annotations in YOLO format. The

training configuration is provided in Table 4. After fine-tuning, the detector weights were fixed and applied identically to the unenhanced images and all enhancement methods throughout the comparative evaluation. The test split was treated as a development set and used to determine the domain-detection criteria and conduct sensitivity analyses for selecting the optimal parameter values, as described in Appendices A and B, respectively. For performance evaluation, 7000 images were randomly sampled from the 10,000-image validation split, and the same fixed subset was used consistently throughout all experiments. Each experimental configuration was evaluated three times using this identical subset, and the mean performance across the three runs was reported.

Table 4. Configuration used to training YOLO.

Parameter	Setting
Epoch	200
Batch size	64
Initial learning rate	0.01
Optimizer setting	Auto
optimizer	SGD with Nesterov momentum
Learning-rate schedule	Linear decay
Momentum	0.9
Weight decay	0.0005
Warm-up epochs	3
Warm-up momentum	0.8

The proposed method was compared against standard gamma correction, CLAHE, adaptive gamma correction, adaptive CLAHE, Retinex-based MSRCR [15], and deep learning-based Zero-DCE++ [23]. In this comparative setup, a gamma value in the standard gamma correction and a clip limit for CLAHE were fixed exclusively at 1.2 and 3.0, respectively. In addition, adaptive gamma correction and adaptive CLAHE were introduced for comparative evaluation to overcome the limitations of standard approaches whose fixed parameters result in poor adaptability under diverse illumination conditions. These methods dynamically adjust the gamma value and clip limit, respectively, based on the median value of the L channel in the LAB color space. The gamma value θ_γ for adaptive gamma correction and the clip limit θ_{clip} for the adaptive CLAHE are defined as follows:

$$\theta_\gamma = \begin{cases} 1.40, & m < 60 \\ 1.20 & 60 \leq m < 100, \\ 1.05, & \text{otherwise} \end{cases} \quad (13)$$

$$\theta_{clip} = \begin{cases} 3.0, & m < 60 \\ 2.0 & 60 \leq m < 100. \\ 1.2, & \text{otherwise} \end{cases} \quad (14)$$

These adaptive approaches aggressively enhance the luminance and contrast to maximize the visibility of hidden details in low-light scenes, while conversely suppressing exposure in bright images to prevent over-saturation. Similar to the proposed method, all standard and adaptive variants adjust only the values of the L channel to prevent color distortion. For Zero-DCE++, we utilized the pre-trained weights provided by the authors in their official repository [28].

For evaluation, Precision (P), Recall (R), F1-score, mean average precision at an IoU threshold of 0.5 (mAP@0.5), mean average precision averaged over IoU thresholds from 0.5 to 0.95 (mAP@0.5:0.95), and mean of intersection over union (mIoU) were employed

as the performance metrics for object detection. The base confidence threshold applied in the calculation of Precision, Recall, and the F1-score was set to 0.5. A non-maximum suppression (NMS) IoU threshold of 0.50 was consistently applied to all evaluated methods. Additionally, Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) were introduced to evaluate the proposed method's capability to preserve the quality and structural similarity of the original images during the image enhancement process. All measurements were obtained using a Python-based evaluation pipeline. The image transformation algorithms were implemented using the OpenCV library. The versions of Python and OpenCV were 3.11 and 4.13.0, respectively.

For each method, every input image was first processed using the corresponding transformation algorithm, and the resulting image set was subsequently evaluated using the same pretrained YOLO object detection model. A fixed low confidence floor of 0.05 was applied uniformly to all methods for mAP evaluation. Predictions with confidence scores below 0.05 were discarded before AP calculation. The adaptive confidence threshold defined in Equation (11) was not applied to mAP calculation. The same dataset splits, model weights, inference settings, and evaluation procedure were used for all methods to ensure a fair comparison.

Figure 2 illustrates representative samples from each domain classified based on the domain criteria detailed in Table 1. Images in the 'Night-dark' domain are characterized by extremely low illumination. Images in the 'Night' domain exhibit relatively higher illumination than those in the 'Night-dark' domain and typically contain localized light sources. The 'Dusk' domain represents ordinary outdoor environments under dim natural lighting conditions. Images in the 'Indoor' domain show spatially imbalanced illumination, often caused by localized light sources, shadows, or other lighting effects. Finally, images in the 'Day' domain are characterized by high global illumination. Table 5 details the number of images classified into each domain. It also specifies the subset of images to which the proposed method was applied excluding those skipped by (12). Out of a total of 7000 images, the proposed method was applied to 4143 images, while the remaining images bypassed the proposed image enhancement.




Figure 2. Image samples from the domains determined by the proposed method.

Table 5. Domain determination results on experiment dataset.

Domain	Night-Dark	Night	Dusk	Indoor	Day	Total
All images	990	1565	290	1775	2380	7000
Applied images	759	1219	122	961	1082	4143
Ratio (%)	77%	78%	42%	54%	45%	59%

Tables 6 and 7 compare the object detection accuracies between conventional image enhancement methods and the proposed approach for all the experiment images and for the images where the proposed method was actually applied, respectively. Specifically, the proposed method was selectively applied based on the condition defined in (5). In Table 6, the experiment results indicate that standard Gamma Correction and CLAHE do not yield a tangible improvement in actual detection performance. Regarding gamma correction, the brightness-related augmentation included in the YOLO training pipeline may already expose the model to similar illumination variations, thereby making gamma-based preprocessing largely redundant. For CLAHE, the enhanced local contrast may reduce ambiguous background responses and thereby improve precision. However, it may also alter or suppress subtle object-related features. This may increase the number of missed detections and consequently reduce recall and the F1-score. For Adaptive Gamma Correction and Adaptive CLAHE, performance even deteriorated. The conventional methods like MSRCR and Zero-DCE++ exhibit a severe degradation in detection accuracy. Despite their significant contributions to perceptual enhancement, the excessive image transformation leads to a photometric risk, which consequently undermines the capabilities of the object detector. Conversely, the proposed method preserves the original structural integrity while achieving an improvement of approximately 0.0071 in the F1-score and approximately 0.0123 in the mAP@0.5. Although precision slightly decreases, this is effectively compensated for by a substantial increase in recall. This implies that the reduction in false negatives under adverse conditions, such as low-light and foggy environments, contributes to ensuring overall detection stability and reliability. However, the improvement in mAP@0.5:0.95 remains marginal. This may be because input-level image transformation primarily improves the detectability of objects, whereas performance at higher IoU thresholds depends on fine-grained localization capability. We conducted a paired bootstrap analysis using detection metrics obtained from 7000 paired images, comprising the original images without enhancement and the corresponding images enhanced using the proposed method. The paired bootstrap analysis yielded a 95% confidence interval of [0.00231, 0.00336] for the mean performance difference. Because the interval did not include zero, the result provides statistical support for the observed performance improvement and indicates that it is unlikely to be explained solely by sampling variability. Table 7 compares the object detection accuracy specifically for the subset of experiment images where the image enhancement process was not skipped. These images specifically represent environments adverse to object detection, characterized by severe conditions such as low-light or over-saturation. The proposed method yields improvements in the F1-score and mAP@0.5 of approximately 0.0101 and 0.195, respectively. These results imply that the proposed image enhancement method can improve object detection performance, even under adverse conditions for object detection. The proposed method exhibits an average per-image latency of 36.8 ms, introducing an additional latency of 22.7 ms compared with detection performed without the image transformation process. Although the proposed method achieves lower latency than the Retinex-based MSRCR and the deep learning-based Zero-DCE++, its process still incurs a non-negligible processing delay. The observed latency may be attributed to the value-sorting operation required to compute the median-based illumination metric used

for risk-score estimation. One potential approach to alleviating this computational burden is to employ streaming quantile estimation algorithms, such as T-Digest and Greenwald–Khanna. Although these algorithms produce an approximation of the median, they may reduce computational costs while maintaining reasonable illumination estimation accuracy and overall performance.

Table 6. Comparison of object detection performance for all experiment images.

Enhancement Method	<i>P</i>	<i>R</i>	F1-Score	mIoU	mAP@0.5	mAP@0.5:0.95	PSNR	SSIM	Latency (ms)
None	0.7048	0.3114	0.4320	0.7554	0.2997	0.1956	-	-	14.1
Gamma correction	0.7011	0.3129	0.4327	0.7554	0.3006	0.1953	27.44	0.9349	14.9
CLAHE	0.7044	0.3119	0.4324	0.7524	0.3006	0.1963	19.18	0.7335	26.5
Adaptive gamma correction	0.6928	0.3121	0.4303	0.7533	0.2990	0.1940	25.35	0.8394	15.6
Adaptive CLAHE	0.7040	0.3130	0.4333	0.7530	0.3015	0.1969	20.09	0.7544	28.5
MSRCR [24]	0.6853	0.2831	0.4007	0.7374	0.2708	0.1766	8.63	0.4311	1928.6
Zero-DCE++ [16]	0.6519	0.2786	0.3904	0.7262	0.2637	0.1698	12.07	0.5200	40.3
Proposed method	0.6536	0.3306	0.4391	0.7507	0.3120	0.1984	42.43	0.9892	36.8

Table 7. Comparison of object detection performance for experiment images where the proposed method was applied.

Enhancement Method	<i>P</i>	<i>R</i>	F1-Score	mIoU	mAP@0.5	mAP@0.5:0.95	PSNR	SSIM
None	0.7118	0.3122	0.4341	0.7773	0.2978	0.1868	-	-
Gamma correction	0.7000	0.3154	0.4349	0.7758	0.2998	0.1871	27.44	0.9349
CLAHE	0.7047	0.3157	0.4360	0.7741	0.3012	0.1888	19.18	0.7335
Adaptive gamma correction	0.6906	0.3139	0.4316	0.7733	0.2974	0.1851	25.35	0.8394
Adaptive CLAHE	0.7047	0.3167	0.4370	0.7742	0.3022	0.1893	20.09	0.7544
MSRCR [24]	0.6718	0.2797	0.3950	0.7545	0.2637	0.1649	8.63	0.4311
Zero-DCE++ [16]	0.6321	0.2745	0.3828	0.7405	0.2548	0.1568	12.07	0.5200
Proposed method	0.6328	0.3422	0.4442	0.7765	0.3173	0.1917	42.43	0.9892

Figure 3 illustrates the F1-score and mAP@0.5 gains achieved by the proposed method for the images grouped within each risk-score interval. The risk-score range from 0 to 0.7 was divided into intervals with a step size of 0.1. The detection performance of the original and enhanced images was then compared for the images within each interval. The performance gain is calculated as the difference between the metric value obtained after applying the proposed method and that obtained from the original images. To ensure a fair comparison, both the original and enhanced images were evaluated using the same object detector under identical detection settings. For images with risk scores below 0.2, the proposed method yields only marginal improvements in object detection performance. This indicates that images in these intervals generally retain sufficient visual quality for the detector to identify objects reliably without substantial enhancement. A noticeable increase in the F1-score and mAP@0.5 gains are observed for images with risk scores ranging from 0.2 to 0.5. This trend suggests that the proposed method becomes increasingly beneficial as image degradation begins to adversely affect object detection. For risk scores of 0.5 and above, both performance gains show slight decreases. Images in these intervals have relatively high photometric risks caused by image transformation. Therefore, the proposed method adaptively reduces the enhancement intensity to preserve detection-relevant features and prevent excessive image modification. Overall, Figure 3 demonstrates that the proposed method adjusts its enhancement effect according to the estimated risk score, providing greater detection benefits for degraded images while limiting potentially detrimental transformations for high-risk images.




Figure 3. Performance gains achieved for images within risk score intervals: (a) for F1-score; (b) for mAP@0.5.

Table 8 compares the object detection performance between the original and transformed images across different domains. The symbols ΔP , ΔR , $\Delta F1$, and ΔmAP denote the gains in Precision, Recall, F1-score, and mAP@0.5, respectively, relative to the baseline performance obtained by directly applying the detector to the original images. The reduction in Precision was particularly pronounced in the ‘Night-dark’ and ‘Night’ domains. Although higher confidence thresholds were applied, the enhancement process amplified background noise and low-contrast structures, causing some of them to be incorrectly detected as objects and consequently increasing false positives. Nevertheless, Recall increased consistently across all domains. The resulting gains in Recall were sufficient to offset the corresponding losses in Precision, leading to overall improvements in both F1-score and mAP@0.5 across every domain. In particular, the ‘Indoor’ and ‘Day’ domains exhibited more pronounced improvements than the other domains.

Table 8. Detection performance gain across different domains.

Domain	ΔP	ΔR	$\Delta F1$	ΔmAP
Night-dark	−0.1019	0.0291	0.0014	0.0097
Night	−0.0897	0.0266	0.0026	0.0121
Dusk	−0.0620	0.0267	0.0061	0.0196
Indoor	−0.0658	0.0338	0.0180	0.0281
Day	−0.0672	0.0311	0.0156	0.0258

To quantitatively analyze the individual contributions of the core control elements within the proposed method, we conducted a stepwise ablation study. In Table 9, ‘No-skip’, ‘No-attenuation’, and ‘No-confidence threshold adjustment’ denote the omission of the skip mechanism, of the attenuation factor A from (9), and of the confidence threshold adjustment. The configuration of ‘Only confidence threshold adjustment’ indicates that only the confidence threshold adjustment was applied. The image enhancements utilizing intensity derived only from risk scores degraded object detection accuracy due to excessive image transformation. This implies that the attenuation factor successfully suppresses excessive image transformation based on the photometric risk scores. The unconditional application of the proposed method without the skip mechanism introduced redundant transformations to already stable imagery, which consequently lowered accuracy. Conversely, the proposed method leverages both the attenuation factor and the skipping mechanism to prevent excessive image enhancement, effectively preserving structural features and yielding an improvement in overall object detection performance. When image enhancement was applied without confidence adjustment, Precision and Recall remained close to the baseline values obtained by directly applying the detector to the original, unprocessed images.

However, the corresponding F1-score did not substantially improve because the confidence score degradation induced by the image transformation was not compensated. Applying only confidence threshold adjustment without image enhancement substantially improved recall, but it markedly reduced precision and consequently lowered the F1-score. Changing the confidence threshold without image transformation affects only the acceptance criterion for existing predictions. By contrast, the confidence threshold adjustment in the proposed method adjusts the confidence threshold in accordance with the transformed images, thereby improving detection performance.

Table 9. Ablation study of the proposed method.

Configuration	<i>P</i>	<i>R</i>	<i>F1-Score</i>	<i>mAP@0.5</i>
No-attenuation	0.6166	0.3300	0.4299	0.3018
No-skip	0.6343	0.3286	0.4329	0.3104
No-confidence threshold adjustment	0.7085	0.3116	0.4330	0.3124
Only confidence threshold adjustment	0.5915	0.3246	0.4191	0.2951
Proposed method	0.6536	0.3306	0.4391	0.3120

Table 10 presents evaluations of the proposed method using the YOLO26n detector and the COCO dataset [29]. When evaluated with YOLO26n [30], the proposed method continued to improve both the F1-score and mAP compared with the baseline results obtained by directly applying the detector to the original, unprocessed images. The improvement observed with YOLO26n raises the possibility that the proposed method could retain some effectiveness across related detector variants. On the evaluated COCO dataset, the proposed method produced a slight improvement in mAP@0.5, whereas the F1-score decreased. This mixed result suggests that the effectiveness of the proposed method depends on the characteristics of the evaluated data. In particular, adverse illumination conditions targeted by the proposed method appear to be relatively underrepresented in COCO, limiting the number of images that benefit from enhancement. Overall, these additional experiments demonstrate limited generalization under the evaluated settings and indicate that further validation using more diverse detector architectures and datasets is required.

Table 10. Detection performance gain across different detectors and datasets.

Detector Model	Dataset	ΔP	ΔR	$\Delta F1$	ΔmAP
YOLOv11n [26]	BDD100k	−0.0512	0.0192	0.0071	0.0123
YOLOv11n [26]	COCO	−0.0325	0.0183	−0.0230	0.0075
YOLO26n [30]	BDD100k	−0.1494	0.0496	0.0168	0.0316

Figure 4 presents images from each domain, including the original images and their enhanced versions obtained using gamma correction, CLAHE, adaptive gamma correction, adaptive CLAHE, MSRCR, Zero-DCE++, and the proposed method, along with the corresponding YOLO object detection results. The standard gamma correction, CLAHE, gamma correction, and adaptive CLAHE did not lead to a noticeable improvement in object detection compared with the original images. While ZERO-DCE++ and MSRCR enhance the perceived visibility of objects, some objects remain undetected. This suggests that visually pleasing enhancement does not necessarily translate to improved machine learning performance. In contrast, the proposed method successfully enhances object detection performance through targeted, minimal image enhancements that preserve essential features.

Most conventional image enhancement methods, including Zero-DCE++ and MSRCR, are designed to improve perceptual image quality by increasing visibility under diverse illumination conditions. As these methods are primarily optimized for visual enhancement, their transformations may modify contrast and intensity distributions in ways that are not necessarily aligned with the feature representations of a deep learning detector. The proposed method pursues a different objective: rather than maximizing visual enhancement, it selectively adjusts photometric conditions while preserving detector-relevant features. This task-oriented design accounts for the observation that visually enhanced images do not always yield higher object detection accuracy.




Figure 4. Comparison of image enhancement results. The green boxes indicate the bounding boxes of detected objects.

Figure 5 visualizes the YOLO activation maps obtained after applying various image enhancement methods, including the proposed method. Compared with the results obtained from the unenhanced image, several conventional enhancement methods were observed to produce weaker or more diffuse activations in the object regions, while increasing the detector's responses to background noise. In contrast, the proposed method limits excessive enhancement intensity, resulting in more concentrated and distinct activation responses in the actual object regions.




Figure 5. Activation map visualizations for different image enhancement methods. Blue indicates pixels with little to no activation, whereas colors closer to yellow indicate pixels with higher activation levels.

5. Conclusions

We proposed an adaptive image enhancement method based on photometric statistics to improve object recognition performance. The photometric domain was determined based on spatial photometric characteristics, such as the luminance median value, the coefficient of luminance variation, and illumination imbalance, to represent its photometric characteristics. In conjunction with this, the photometric risk score was calculated to quantify the likelihood of its photometric risk caused by image enhancement. The proposed method combined three techniques—Gamma correction, CLAHE, and horizontal luminance balancing—to achieve image enhancement. The intensities of these techniques were adaptively set according to the domain-specific preset and the photometric risk score. Subsequently, the confidence threshold was adjusted to mitigate the reduction in confidence scores caused by the image transformations. If the image was considered sufficiently optimal for object recognition, this enhancement process was skipped to prevent degrading object recognition performance from excessive enhancement. The experiment results on the BDD100K dataset demonstrated that the proposed method improved F1-score by 0.0101 and the mAP@0.5 by 0.0195 on images under adverse conditions. This detection

accuracy enhancement was particularly pronounced in the ‘Indoor’ and ‘Day’ domains that represent localized and global overexposure conditions, respectively. Specifically, this meaningful improvement in Recall demonstrated the capability of the proposed method to mitigate false negative issues in object recognition under adverse imaging conditions. The proposed method suppresses excessive image enhancement and consequently improves object detection performance under adverse imaging conditions by adaptively regulating transformation intensity based on the photometric statistics. However, the proposed method still has limitations. Although it effectively improves object detection performance for images captured under adverse conditions, its performance gains are relatively limited for images captured under normal conditions. Future work will therefore investigate a deep learning-based image analysis approach that adaptively determines transformation intensity according to the characteristics of each input image, with the aim of achieving more consistent performance improvements across diverse imaging conditions. In addition, some of the parameters used to determine the image transformation intensity were empirically selected. Future work will investigate deep learning-based methods for automatically estimating these parameters according to the photometric characteristics of the input images. The proposed method has the potential to significantly enhance visual perception capabilities across a wide range of real-world applications, including low-light autonomous driving, mobile robotics, and intelligent surveillance frameworks.

Author Contributions: Conceptualization, C.-y.K. and S.-k.K.; software, C.-y.K.; writing—original draft preparation, C.-y.K. and S.-k.K.; supervision, S.-k.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partly supported by the Innovative Human Resource Development for Local Intellectualization program through the Institute of Information & Communications Technology Planning & Evaluation (IITP), grant funded by the Korean government (MSIT) (IITP-2026-RS-2020-II201791, 100%) and by the BB21plus funded by Busan Metropolitan City and Busan Techno Park.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

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

Appendix A

Appendix A describes the rationale for the photometric domain determination criteria presented in Table 1. The criteria were established using 20,000 images from the test split of the publicly available third-party derivative of BDD100K [30]. Based on this dataset, the criteria were determined through a combination of statistical analysis of the photometric feature distributions and empirical examination of the characteristic illumination patterns observed in the images.

The criteria for the ‘Night-dark’ and ‘Night’ domains were primarily determined from the distributions of the median luminance m and the 95th-percentile luminance p_{95} . The empirical distributions were visualized using histograms and smoothed using kernel density estimation (KDE), from which local minima were identified as candidate boundaries between distinct illumination regimes, as shown in Figure A1. For the entire image set, local minima were observed at approximately 62.7 and 160 for m and for p_{95} , respectively, indicating a transition between low-light and relatively well-exposed images. The images satisfying these initial low-light conditions were then analyzed separately to further distinguish severely underexposed scenes. In this second-stage analysis, local

minima were identified at approximately 20.4 and 90.3 for m and $p95$, respectively. Based on these density valleys and the visual examination of images near the candidate boundaries, the thresholds were adjusted to practical integer values, resulting in $m < 70$ and $p95 < 180$ for ‘Night’ domain and $m < 28$ and $p95 < 80$ for ‘Night-dark’ domain.




Figure A1. Histogram and kernel density estimate for median luminance and 95th-percentile luminance metrics: (a) distributions for the entire image set; (b) distributions for images identified as low-light in the first-stage analysis. The red dashed lines indicate the local minima of the corresponding kernel density estimation curves.

The criteria of ‘Indoor’ and ‘Day’ domains were empirically defined according to the characteristic illumination patterns repeatedly observed in the training data. The Indoor domain was designed to capture spatially nonuniform illumination caused by directional or localized light sources. Thus, a large left–right luminance difference ($dLR > 0.40$) and high luminance variation ($CV > 0.18$) were jointly considered, while the upper bounds on m and $p95$ prevented globally bright scenes from being assigned to this domain. The Day domain was characterized by the presence of sufficiently bright regions ($p95 > 200$) together with relatively uniform illumination ($CV < 0.30$). This combination prevented scenes containing only localized bright areas from being classified as Day. These empirical boundaries were selected by visually examining representative images around the candidate thresholds and assessing whether the resulting domain assignments consistently reflected the intended illumination characteristics. The representative examples in Figure 2 shows that the images assigned to each domain generally exhibit the intended photometric characteristics.

Appendix B

Appendix B presents the sensitivity analysis conducted to determine the optimal domain-specific preset values for controlling the image transformation intensity, threshold for the skip mechanism, and confidence adjustment parameters. The identical YOLOv11 detector from the experiments was employed for this evaluation. To accommodate scenarios

where false negatives result in more critical losses than false positives, we adopted a recall-oriented weighted score defined as follows:

$$score = 0.4P + 0.6R. \quad (A1)$$

Tables A1–A3 present the detection performance across varying enhancement intensity parameters. Object detection performance was evaluated on the BDD100K test set by varying the parameters θ_γ , θ_{clip} , and θ_{gain} for gamma correction, CLAHE, and horizontal luminance balancing, respectively. For gamma correction, the detection accuracies of the ‘Night-dark’ and ‘Indoor’ domains were not improved by increase in θ_γ . Specifically, in the extremely low-light ‘Night-dark’ domain, gamma correction tends to either obscure object features or amplify background noise. Similarly, gamma correction proves ineffective in the ‘Indoor’ domain, which is characterized by spatially uneven illumination. For the remaining domains, the optimal parameters are assigned to the domain-specific preset. For CLAHE, performance improvements were observed across all domains except for the ‘Night-dark’ domain. This enhanced detection efficacy is attributed to CLAHE’s capability to ameliorate locally over-saturated or low-light regions. Conversely, in the ‘Night-dark’ domain, the technique amplified background noise, thereby degrading detection performance. Notably, the ‘Indoor’ and ‘Day’ domains did not utilize the parameters yielding the best performance for the final presets. In these domains, the performance discrepancy between the best and sub-best parameters was marginal. A higher intensity of image transformation poses a greater risk of disrupting feature distributions. Therefore, we selected the sub-best parameters for the final presets, featuring lower values than the best parameters. Regarding horizontal luminance balancing, illuminating balancing across all domains effectively suppresses local over-saturation caused by left-right illumination asymmetry, thereby improving detection performance.

The sensitivity analysis for the skip mechanism was also conducted. In this sensitivity analysis, the domain-specific skip threshold T_{skip} in (12) was varied from 0.10 to 1.00 in increments of 0.01, and the corresponding object detection performance was evaluated. Table A4 presents the results within the range of 0.50–0.63, which contains the optimal threshold for each domain. Images belonging to the ‘Night-dark’ and ‘Night’ domains generally exhibit low global illumination and therefore tend to produce higher risk scores than those in the other domains. Consequently, relatively high thresholds are required in these domains to prevent unnecessary transformation of images with comparatively low illumination risk. When the threshold is set too low, most images fail to satisfy the skip condition and are therefore subjected to illumination enhancement. This excessive transformation may amplify background noise in some images and consequently degrade object detection performance. Conversely, an excessively high threshold may cause even images requiring illumination enhancement to be skipped, thereby preventing sufficient visibility improvement. The best detection performance was achieved at thresholds of 0.61 and 0.58 for the ‘Night-dark’ and ‘Night’ domains, respectively. In contrast, images in the ‘Dusk’, ‘Indoor’, and ‘Day’ domains generally exhibit higher illumination levels and therefore produce lower average risk scores. As a result, many images in these domains satisfy the skip condition even when relatively low thresholds are applied, thereby avoiding unnecessary image transformation. If the threshold is set too low, however, transformation may still be applied to some adequately illuminated images, potentially causing artifacts such as brightness oversaturation and contrast distortion. Nevertheless, the resulting degradation in detection performance was less pronounced than the degradation caused by background noise amplification in the low-light domains. The optimal thresholds for the ‘Dusk’, ‘Indoor’, and ‘Day’ domains were 0.52, 0.52, and 0.51, respectively.

Table A1. Detection performance across different gamma correction intensities. Underlined values represent the performance corresponding to the selected preset parameters.

θ_γ	Night-Dark	Night	Dusk	Indoor	Day
1.00	<u>0.3619</u>	0.5014	0.5495	<u>0.4414</u>	0.4093
1.05	0.3586	0.5007	0.5501	0.4388	0.4093
1.10	0.3546	<u>0.5025</u>	0.5500	0.4390	0.4093
1.15	0.3498	0.5014	0.5510	0.4371	0.4093
1.20	0.3465	0.5015	<u>0.5513</u>	0.4377	0.4093
1.25	0.3419	0.5012	0.5500	0.4336	0.4165
1.30	0.3411	0.4999	0.5498	0.4324	<u>0.4212</u>
1.35	0.3369	0.4995	0.5507	0.4323	0.4070

Table A2. Detection performance across different clip limits of contrast-limited adaptive histogram equalization. Underlined values represent the performance corresponding to the selected preset parameters.

θ_{clip}	Night-Dark	Night	Dusk	Indoor	Day
0.0	<u>0.3717</u>	0.4992	0.5485	0.4359	0.3960
0.2	<u>0.3565</u>	0.4998	0.5482	0.4345	0.3960
0.4	0.3567	0.5009	0.5488	0.4347	0.3960
0.6	0.3602	0.4999	0.5495	0.4371	0.3960
0.8	0.3579	0.5001	0.5498	0.4353	0.3951
1.0	0.3596	0.5026	0.5498	0.4360	0.3951
1.2	0.3575	0.5028	0.5510	0.4373	0.3951
1.4	0.3540	0.5028	0.5501	<u>0.4379</u>	<u>0.4070</u>
1.6	0.3533	0.5026	0.5498	0.4375	0.4070
1.8	0.3533	<u>0.5031</u>	0.5510	0.4374	0.4070
2.0	0.3501	<u>0.5017</u>	<u>0.5514</u>	0.4380	0.3808
2.2	0.3500	0.5028	0.5511	0.4375	0.3903
2.4	0.3509	0.5028	0.5506	0.4380	0.3903
2.6	0.3469	0.5014	0.5510	0.4357	0.4093
2.8	0.3473	0.5006	0.5501	0.4377	0.3903

Table A3. Detection performance across different column gains. Underlined values represent the performance corresponding to the selected preset parameters.

θ_{gain}	Night-Dark	Night	Dusk	Indoor	Day
0.00	0.3503	0.5013	0.5512	0.4373	0.4165
0.02	0.3503	0.5010	0.5512	0.4378	0.4165
0.04	0.3498	<u>0.5022</u>	<u>0.5513</u>	<u>0.4391</u>	0.4093
0.06	<u>0.3511</u>	0.5021	0.5509	0.4360	0.3998
0.08	0.3494	0.5019	0.5508	0.4380	<u>0.4308</u>
0.10	0.3498	0.5014	0.5510	0.4371	0.4093
0.12	0.3498	0.5016	0.5506	0.4367	0.4093
0.14	0.3502	0.5016	0.5506	0.4366	0.4093
0.16	0.3510	0.5017	0.5507	0.4358	0.3998
0.18	0.3508	0.5010	0.5503	0.4349	0.4093
0.20	0.3493	0.5016	0.5504	0.4352	0.4093

Table A4. Detection performance across different skip thresholds. Underlined values represent the performance corresponding to the selected skip threshold.

T_{skip}	Night-Dark	Night	Dusk	Indoor	Day
0.50	0.34384	0.38388	0.51245	0.52677	0.53230
0.51	0.34413	0.38427	0.51252	0.52704	<u>0.53232</u>

Table A4. *Cont.*

T_{skip}	Night-Dark	Night	Dusk	Indoor	Day
0.52	0.34623	0.38453	<u>0.51253</u>	<u>0.52704</u>	0.53232
0.53	0.34684	0.38478	0.51253	0.52704	0.53229
0.54	0.34654	0.38517	0.51253	0.52677	0.53229
0.55	0.34701	0.38517	0.51200	0.52677	0.53229
0.56	0.34745	0.38525	0.51176	0.52677	0.53229
0.57	0.34761	0.38547	0.51176	0.52677	0.53229
0.58	0.34808	<u>0.38566</u>	0.51176	0.52677	0.53229
0.59	0.34807	0.38566	0.51176	0.52677	0.53229
0.60	0.34875	0.38560	0.51176	0.52677	0.53226
0.61	<u>0.34905</u>	0.38560	0.51176	0.52677	0.53226
0.62	0.34899	0.38560	0.51176	0.52677	0.53226
0.63	0.34894	0.38550	0.51176	0.52677	0.53226

Table A5 presents sensitivity analyses of the luminance-aware factor (f_l) for low-illumination images ($m < 50$) and moderately dark images ($50 \leq m < 70$), respectively. Table A6 presents the results within the range of 0.81–1.00, which contains the optimal parameters. This sensitivity analysis identified the optimal values of f_l as 0.85 and 0.90 for low-illumination and moderately dark images, respectively. Table A6 shows the sensitivity analyses of the risk-aware factor (f_r). To facilitate the sensitivity analysis, the numerical constants in (10) is parameterized as follows.

$$f_r = \begin{cases} \hat{f}_r = f_{r0} - \lambda_r s \\ \eta_r \hat{f}_r, & \text{if } s > 0.6 \\ \hat{f}_r, & \text{otherwise} \end{cases} \quad (\text{A2})$$

Specifically, f_{r0} and η_r were evaluated over the range of 0.85–1.00, whereas λ_r was varied from 0 to 0.15. The sensitivity analysis confirmed that the constants originally specified in (10) correspond to the optimal parameter values.

Table A5. Detection performance across different luminance-aware factor. Underlined values represent the performance corresponding to the selected luminance-aware factor.

f_l	$m < 50$	$50 \leq m < 70$
0.80	0.45331	0.45202
0.81	0.45332	0.45202
0.82	0.45337	0.45231
0.83	0.45342	0.45231
0.84	0.45342	0.45261
0.85	<u>0.45344</u>	0.45287
0.86	0.45339	0.45318
0.87	0.45342	0.45327
0.88	0.45337	0.45334
0.89	0.45337	0.45339
0.90	0.45337	<u>0.45344</u>
0.91	0.45337	0.45332
0.92	0.45337	0.45332
0.93	0.45337	0.45307
0.94	0.45337	0.45307
0.95	0.45337	0.45307
0.96	0.45337	0.45307
0.97	0.45337	0.45307
0.98	0.45337	0.45307
0.99	0.45337	0.45307
1.00	0.45337	0.45307

Table A6. Detection performance across different constant values for risk-aware factor. Underlined values represent the performance corresponding to the selected constant value.

f_{r0}	Score	λ_r	Score	η_r	Score
0.85	0.45143	0.00	0.45421	0.85	0.45337
0.86	0.45164	0.01	0.45413	0.86	0.45337
0.87	0.45187	0.02	0.45390	0.87	0.45337
0.88	0.45209	0.03	0.45380	0.88	<u>0.45344</u>
0.89	0.45236	0.04	0.45361	0.89	0.45339
0.90	0.45254	0.05	0.45361	0.90	0.45328
0.91	0.45274	0.06	0.45330	0.91	0.45318
0.92	0.45298	0.07	0.45331	0.92	0.45318
0.93	0.45333	0.08	0.45332	0.93	0.45318
0.94	<u>0.45344</u>	0.09	0.45332	0.94	0.45311
0.95	0.45333	0.10	<u>0.45334</u>	0.95	0.45311
0.96	0.45332	0.11	0.45334	0.96	0.45311
0.97	0.45331	0.12	0.45332	0.97	0.45311
0.98	0.45331	0.13	0.45332	0.98	0.45311
0.99	0.45331	0.14	0.45332	0.99	0.45311
1.00	0.45331	0.15	0.45332	1.00	0.45311

References

- Huang, Y.; Chen, Y. Survey of state-of-art autonomous driving technologies with deep learning. In Proceedings of the IEEE International Conference on Software Quality, Reliability and Security Companion, Macau, China, 11–14 December 2020.
- Su, H.; Gao, H.; Wang, X.; Fang, X.; Liu, Q.; Huang, G.; Li, X.; Cao, Q. Object detection in adverse weather for autonomous vehicles based on sensor fusion and incremental learning. *IEEE Trans. Instrum. Meas.* **2024**, *73*, 2535510. [\[CrossRef\]](#)
- Sun, Y.; Wang, X.; Liu, Z.; Miller, J.; Efros, A.A.; Hardt, M. Test-time training with self-supervision for generalization under distribution shifts. In Proceedings of the International Conference on Machine Learning, Virtual Event, 13–18 July 2020.
- Yu, F.; Chen, H.; Wang, X.; Xian, W.; Chen, Y.; Liu, F.; Madhavan, V.; Darrell, T. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Virtual Event, 14–19 June 2020.
- Dodge, S.; Karam, L. Understanding how image quality affects deep neural networks. In Proceedings of the IEEE International Conference on Quality of Multimedia Experience, Lisbon, Portugal, 6–8 June 2016.
- Li, M.; Zhao, H.; Guo, X. LIME-Eval: Rethinking low-light image enhancement evaluation via object detection. *arXiv* **2024**, arXiv:2410.08810.
- Zheng, S.; Ma, Y.; Pan, J.; Lu, C.; Gupta, G. Low-light image and video enhancement: A comprehensive survey and beyond. *arXiv* **2022**, arXiv:2212.10772.
- Gonzalez, R.C.; Woods, R.E. *Digital Image Processing*, 3rd ed.; Pearson Education: Upper Saddle River, NJ, USA, 2008.
- Pizer, S.M.; Amburn, E.P.; Austin, J.D.; Cromartie, R.; Geselowitz, A.; Greer, T.; ter Haar Romeny, B.; Zimmerman, J.B.; Zuiderveld, K. Adaptive histogram equalization and its variations. *Comput. Vis. Graph. Image Process.* **1987**, *39*, 355–368. [\[CrossRef\]](#)
- Musa, P.; Al Rafi, F.; Lamsani, M. A review: Contrast limited adaptive histogram equalization (CLAHE) methods to help the application of face recognition. In Proceedings of the International Conference on Informatics and Computing, Palembang, Indonesia, 17–18 October 2018.
- Reza, A.M. Realization of the contrast limited adaptive histogram equalization for real-time image enhancement. *J. VLSI Signal Process. Syst.* **2004**, *38*, 35–44. [\[CrossRef\]](#)
- Rahman, S.; Rahman, M.M.; Abdullah-Al-Wadud, M.; Al-Quaderi, G.D.; Shoyaib, M. An adaptive gamma correction for image enhancement. *EURASIP J. Image Video Process.* **2016**, *2016*, 35. [\[CrossRef\]](#)
- Chang, Y.; Jung, C.; Ke, P.; Song, H.; Hwang, J. Automatic contrast-limited adaptive histogram equalization with dual gamma correction. *IEEE Access* **2018**, *6*, 11782–11792. [\[CrossRef\]](#)
- Petro, A.B.; Sbert, C.; Morel, J.M. Multiscale retinex. *Image Process. Line* **2014**, *4*, 71–88. [\[CrossRef\]](#)
- Wang, J.; Lu, K.; Xue, J.; He, N.; Shao, L. Single Image Dehazing Based on the Physical Model and MSRCR Algorithm. *IEEE Trans. Circuits Syst. Video Technol.* **2018**, *28*, 2190–2199. [\[CrossRef\]](#)
- Wang, R.; Zhang, Q.; Fu, C.W.; Shen, J.; Zheng, W.S.; Jia, J. Underexposed photo enhancement using deep illumination estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019.

17. Xu, X.; Wang, R.; Fu, C.W.; Jia, J. SNR-aware low-light image enhancement. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022.
18. Cui, Z.; Li, K.; Gu, L.; Su, S.; Gao, P.; Jiang, Z.; Qiao, Y.; Harada, T. You only need 90k parameters to adapt light: A light weight transformer for image enhancement and exposure correction. In Proceedings of the British Machine Vision Conference, London, UK, 21–24 November 2022.
19. Cai, Y.; Bian, H.; Lin, J.; Wang, H.; Timofte, R.; Zhang, Y. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 1–6 October 2023.
20. Jiang, Y.; Gong, X.; Liu, D.; Cheng, Y.; Fang, C.; Shen, X.; Yang, J.; Zhou, P.; Wang, Z. Enlightengan: Deep light enhancement without paired supervision. *IEEE Trans. Image Process.* **2021**, *30*, 2340–2349. [[CrossRef](#)] [[PubMed](#)]
21. Zhu, J.; Park, T.; Isola, P.; Efros, A.A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017.
22. Guo, C.; Li, C.; Guo, J.; Loy, C.C.; Hou, J.; Kwong, S.; Cong, R. Zero-reference deep curve estimation for low-light image enhancement. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020.
23. Li, C.; Guo, C.; Loy, C.C. Learning to enhance low-light image via zero-reference deep curve estimation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *44*, 4225–4238. [[CrossRef](#)] [[PubMed](#)]
24. Zheng, S.; Gupta, G. Semantic-guided zero-shot learning for low-light image/video enhancement. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2022.
25. Rodríguez-Rodríguez, J.A.; López-Rubio, E.; Ángel-Ruiz, J.A.; Molina-Cabello, M.A. The impact of noise and brightness on object detection methods. *Sensors* **2024**, *24*, 821. [[CrossRef](#)] [[PubMed](#)]
26. Ultralytics YOLO11. Available online: <https://github.com/ultralytics/ultralytics> (accessed on 24 July 2026).
27. BDD100k YOLO-Format Dataset. Available online: <https://www.kaggle.com/datasets/a7madmostafa/bdd100k-yolo> (accessed on 29 June 2026).
28. Zero-DCE++ GitHub Repository. Available online: https://github.com/Li-Chongyi/Zero-DCE_extension (accessed on 17 June 2026).
29. Lin, T.Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Zitnick, C.L.; Dollár, P. Microsoft COCO: Common objects in context. In Proceedings of the European Conference on Computer Vision, Zürich, Switzerland, 6–12 September 2014.
30. Jocher, G.; Qiu, J.; Liu, M.; Lyu, S.; Akyon, F.C.; Kalfaoglu, M.E. Ultralytics YOLO26: Unified real-time end-to-end vision models. *arXiv* **2026**, arXiv:2606.03748.

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.