FEATURE EXTRACTION TO DETECT DEGRADATION OF ROAD SIGNS USED IN AUTOMATED DRIVING SCENARIOS

Cristiana-Mirela LUPU
Communications Department, Technical University of Cluj-Napoca, Cluj-Napoca, Romania
lupu.cristianamirela@gmail.com

Abstract: ***This paper proposes a lightweight method for detecting visual degradation in road signs, based on a Threshold-Based Color Quantization algorithm. HSV color space was chosen for its strong alignment with human visual perception.. The method extracts dominant color clusters from the region of interest and matches them between an undegraded reference sign and a potentially degraded one using centroid distance and relative pixel-proportion similarity. The method exhibits several advantages, including high processing speed, independence from a predefined number of clusters, algorithmic simplicity, and effective feature representation. A normalized degradation score is obtained by combining the number of unmatched clusters, the Euclidean distance between centroids, and the proportional difference between corresponding clusters. The approach was evaluated on a synthetic dataset of approximately 377 road signs with eight types of simulated degradation, achieving an RMSE of 0.12 when predicting degradation severity, which corresponds to an approximate prediction accuracy of 88%. The study discusses these limitations and outlines potential improvements, including the integration of multimodal features and learning-based approaches to enhance performance in real life environments.***

Key words: *Road sign degradation, color segmentation, color clustering, self-driving car, smart city, computer vision*

I. INTRODUCTION

Self-driving cars are currently among the most discussed topics in modern technology, representing a rapidly evolving field driven by continuous research and development. According to [1], autonomous driving relies on four major technological pillars: real-time and embedded systems, machine learning, edge computing, and cloud computing. The primary objective of autonomous vehicles is to reduce traffic accidents caused by human error, thereby increasing overall road safety.

This work is also connected to the broader concept of Smart Cities. As industrial and technological advancements accelerate, expectations regarding quality of life, efficient resource usage, and the automation of labor-intensive tasks increasingly shape the way urban environments are designed. According to [2], a smart city operates as a dynamic and coherent metabolic system interacting with its built and social environment. Key components include smart buildings, mobility, energy, healthcare, and governance.

Although extensive research exists on traffic sign detection and recognition, only a limited number of studies address the degradation of traffic signs as a visual and functional problem. In [3], it is proposed a method for measuring damage using SSIM, providing a perceptual metric of structural degradation. However, SSIM primarily captures changes in luminance, contrast, and local structure, making it highly sensitive to illumination variations, exposure, and viewpoint changes. Moreover, SSIM does not explicitly model chromatic information, limiting its ability to quantify color-based degradation such as fading, discoloration, or pigment loss: phenomena that are central to traffic sign deterioration. Additional studies examine robustness to degradation from a recognition perspective. For example, the work in [7] demonstrates that signs affected by occlusion, vandalism, or discoloration can still be recognized through cognitive and perceptual modeling approaches. While this confirms the practical relevance of visually degraded signs in real-world environments, such methods focus on maintaining recognition performance rather than estimating or quantifying the degradation itself.

Other research directions rely on machine-learning-based solutions. Paper [4] classify degraded signs using a flexible mixture model and transfer learning, while paper [5] employ ensemble YOLOv5 models with test-time augmentation to detect visually corrupted signs. These approaches use supervised learning and deep architectures, differing fundamentally from the classical computer vision techniques targeted in this work. An earlier classical contribution is provided by [6], who introduced a small real-world database of degraded signs and proposed a detection baseline using RGB color segmentation, shape detection, and heuristics, achieving an F-score of 0.91. While valuable as evidence that degraded signs require specialized treatment, their system focuses on detection rather than quantifying the degree of visual deterioration.

Despite these contributions, existing approaches predominantly address detection, classification, or structural similarity assessment, whereas few methods focus on estimating color-based visual degradation through lightweight, classical, and computationally efficient techniques. This gap motivates the approach proposed in this study.

The core objective of this work is to develop a method capable of estimating the degradation of a road sign over time using classical computer vision techniques. A combination of edge detection, corner detection, and color-based analysis is employed to extract features relevant to visual deterioration. The resulting prototype produces a normalized degradation score, which could be integrated into a video perception module (for example, in an autonomous vehicle) to enable real-time map updates via cloud-connected systems.

Manuscript received October 28, 2025; revised December 16, 2025

The remainder of this paper is organized as follows: Section II reviews related work and theoretical background, Section III describes the proposed method, Section IV presents the experimental results and Section V concludes the paper and outlines directions for future research.

II. RELATED WORK and THEORETICAL BACKGROUND

In the context of SLAM (Simultaneous Localization and Mapping), landmarks are visual features used as reference points for localization and map construction [3]. Traffic signs serve as stable landmarks in many urban environments, but their visual appearance can change over time due to environmental exposure, vandalism, dirt, or partial occlusion. Such changes may alter the landmark's signature, leading to incorrect associations in SLAM or recognition modules.

For this reason, a **quantitative degradation score** is valuable. Values close to zero indicate that the current sign closely matches the reference (no aging), moderate increases correspond to progressive chromatic deterioration, and large values suggest that the observed sign differs significantly from the reference and may represent a different landmark.

Road signs can suffer from several types of degradation, including, as it can be seen:

- Graffiti or stickers applied to their surface
- Accumulation of dirt, mud, or dust
- Obstructions such as vegetation or other objects
- Scratched or faded paint.

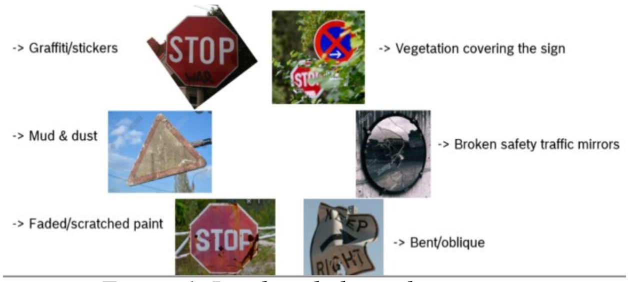

Figure 1. Landmark degradation scenarios.

Since these forms of degradation primarily affect the chromatic characteristics of the sign, color information is a key indicator of deterioration. Geometric distortions (e.g., bent or tilted signs) may also occur but fall outside the primary scope of this work.

Selecting a suitable color space is essential for evaluating chromatic changes. Prior comparative studies demonstrate that HSV offers advantages over RGB and CIELAB for color-based segmentation because it more effectively decouples chromaticity from brightness [10–12]. RGB mixes luminance and chrominance, causing visually similar shades to appear artificially distant under Euclidean distance, while CIELAB, although perceptually uniform, is more sensitive to noise in low-saturation regions and can amplify small illumination fluctuations. By contrast, HSV allows the chromatic components (Hue and Saturation) to be analyzed independently of brightness, which is particularly important when assessing degradation such as fading or discoloration. Focusing on these chromatic channels reduces sensitivity to illumination variations and yields more stable and meaningful color quantization, making HSV a more suitable and less restrictive choice for detecting color-based deterioration in traffic signs.

Existing research confirms the relevance of degraded signs but does not provide a lightweight, color-based degradation classical method for quantification. In paper [3] it is measured the damage using SSIM, a structural similarity metric sensitive to illumination changes and poorly suited for capturing chromatic fading or discoloration. Studies [6] and [7] are based on classical approaches, but do not quantify degradation.

While papers [4], [5] and [8] propose machine and transfer learning approaches, not classical ones.

These studies demonstrate the importance of handling visual degradation, yet they pursue different goals: SSIM-based approaches focus on structural similarity, ML-based approaches target classification or detection, and classical baselines focus on robust sign detection. None of these works explicitly measure color-based degradation using a simple clustering mechanism, which forms the central objective of this paper.

Because color degradation is the primary cue, the segmentation step in this work is limited to isolating the sign region and quantizing the colors within it. General segmentation categories such as thresholding, region-based, edge-based, or clustering methods are well documented; however, only clustering-based color quantization is directly relevant to our method.

Classical clustering algorithms such as K-means were initially evaluated due to their simplicity and widespread use in color quantization. However, K-means requires specifying the number of clusters in advance [8], which is unsuitable for traffic signs whose number of perceptually distinct colors varies with fading, dirt coverage, or external occlusions. If is set too low, distinct colors merge and meaningful degradations are masked; if set too high, noise or illumination variations create artificial clusters, exaggerating degradation. These limitations make fixed-clustering unstable and highly sensitive to parameter tuning. Nonetheless, K-means served as a natural baseline and confirmed that chromatic clustering is an effective direction, motivating the shift toward a more flexible alternative.

To overcome the need for a predefined, this work employs a **threshold-based color quantization method (TBCQ)**, conceptually related to online or sequential clustering techniques. In this approach, new clusters are created only when a pixel is sufficiently distant from all existing centroids; thus, the final number of clusters emerges naturally from the true chromatic variability of the sign. Each cluster is represented by its centroid and pixel proportion, forming a compact descriptor of the sign's color composition. By comparing the descriptor of a reference sign with that of a current observation, a normalized degradation score is computed that reflects both chromatic displacement and changes in color proportions.

This method offers multiple key advantages and will be explained in more detail in the next section.

III. PROPOSED METHOD

The proposed method estimates the visual degradation level of a traffic sign by comparing a reference image (undamaged sign) with a current observation (potentially degraded). The core of the method is a **Threshold-Based Color Quantization (TBCQ)** algorithm operating in the HSV color space. Unlike traditional clustering methods such as K-Means—where the number of clusters must be fixed in advance—TBCQ dynamically creates clusters based solely on color variability in the image. This property is essential for degradation assessment, since a faded or partially occluded sign may contain a different number of dominant color regions compared to its undamaged counterpart.

The method requires only the sign region itself. The dataset provided bounding-box detections and semantic class labels, allowing the approximate geometric shape of each sign (octagon, triangle, circle, etc.) to be inferred. A shape-constrained crop was then applied to remove background pixels. No explicit edge- or corner-detection algorithms were used.

The pipeline consists of four stages:

1. Preprocessing step: Sign extraction from bounding box and semantic shape.
2. Color quantization using TBCQ in HSV space.
3. Cluster matching between reference and observed sign.
4. Computation of a normalized degradation score.

3.1. Preproccessing step: sign extraction

The dataset was collected entirely in-house. Public datasets were unsuitable due to mismatches in annotation format, inconsistent degradation levels, or the absence of bounding boxes compatible with the acquisition pipeline. Nonetheless, the proposed algorithm is compatible with any dataset that provides cropped sign regions.

The original input frame consisted exclusively of images captured in real driving conditions, with bounding boxes and class labels provided by the acquisition system. From the class label, the corresponding **semantic shape** (e.g., circle for speed limits, triangle for warnings) was inferred.

The sign is then cropped from the bounding box to eliminate background pixels, based on its shape. This is an approximation rather than a precise contour extraction, but is sufficient for isolating the sign interior, which is the only region used for color degradation analysis.

The cropped region was then converted from **RGB to HSV**, where color tone (Hue) is decoupled from brightness (Value), providing a more perceptually meaningful representation for degradation-based quantization.

3.2. Threshold-Based Color Quantization (TBCQ)

While K-Means can be used for vector quantization, the proposed threshold-based method differs in how cluster initialization and adaptation are handled, particularly regarding the dynamic determination of the number of clusters.

For traffic signs, this is problematic:

- the number of dominant colors changes when the sign fades,
- dirt or graffiti introduce new colors,
- illumination and aging modify the cluster distribution.

Fixing k would artificially force mismatches and hide or exaggerate degradation. TBCQ avoids this by creating clusters dynamically based on a predefined color distance threshold.

Thus, the number of clusters reflects the actual chromatic variability, making it directly informative for degradation assessment. The threshold value controls the granularity of the color quantization. A small threshold produces many clusters (over-segmentation), while a large threshold merges distinct colors. To determine a stable value, extensive experiments were performed on over 100 traffic signs, across multiple degradation scenarios and camera conditions. The threshold was varied and its impact on reconstruction accuracy and cluster stability was evaluated. The chosen value (=0.25) offered the best compromise between accuracy and computational efficiency across nearly all cases. This value was therefore fixed for all experiments.

The TBCQ algorithm operates as follows, pseudocode was translated from [9]:

Given the training sequence, we set *a priori* an upper bound (a **threshold**) ε for the radius of each "ball" that constitutes a class C_i.

- The first vector x_1 creates the first class C_1;
- Upon receiving each vector x_j, a decision is made:
  - if $d(x_j, y_i) > \varepsilon$ for all existing classes C_i, then x_j creates a new class;
  - otherwise, x_j is assigned to the class whose centroid is the closest, and the position of the new centroid is recalculated.

The algorithm is applied in the HSV space, following five main steps:

Step 1: The algorithm takes as input an HSV image representing the cropped traffic sign region.

Step 2: A matrix, denoted as "centroid", is initialized with the first color cluster centroid, corresponding to the first valid pixel in the input image. The pixel's Hue, Saturation, and Value components are stored, and the associated pixel count is initialized to zero.

Step 3: The algorithm iteratively processes each pixel in the image by performing the following operations:

- **Distance Calculation:** Compute the distance between the current pixel and each existing cluster center. The formula used for this computation is provided below, along with additional supporting equations.
- Minimum Distance Selection: Identify the minimum distance among all computed distances.
- Threshold Comparison: Compare the minimum distance to a predefined threshold:
- If the distance exceeds the threshold, a new cluster is created with the current pixel as its center.
- If the distance is within the threshold, the pixel is assigned to the closest cluster, the cluster's pixel count is incremented, and the centroid is updated as the centroid (gravity center) of all assigned pixels.

Step 4: After all pixels have been clustered, the image is reconstructed by assigning to each pixel the HSV values of its corresponding cluster center.

Step 5: The proportion of each cluster is computed as the percentage of pixels relative to the total number of pixels in the cropped image.

The reconstructed HSV image is subsequently converted back to the RGB color space. Finally, the results are visualized using multiple output representations for analysis and validation. The output consists of dominant color clusters, each characterized by a centroid (representing the dominant color) and the proportion of pixels belonging to that cluster.

To achieve perceptually meaningful distances, the Hue channel (circular variable) was embedded into a 2-D representation. This transforms the HSV pixel (H, S, V) into a 4-dimensional vector:

$$new\ HSV\_pixel=(cos(H),sin(H),S,V) \qquad (1)$$

Euclidean distance is then computed in this space to preserve the circularity of Hue and ensures more stable clustering under gradual fading:

$$distance\ ((H_1,S_1,V_1),(H_2,S_2,V_2)) = sqrt((cos(H_1)-cos(H_2))^2 + ((sin(H_1)-sin(H_2))^2 + (S_1-S_2)^2+(V_1-V_2)^2\ ) \quad (2)$$

3.3. Cluster Matching Between Reference and Degraded

TBCQ algorithm is applied to both the original and degraded images, the resulting feature sets (defined as clusters of dominant colors) are forwarded to a matching module in order to determine if there is a degradation.

This method attempts to pair clusters of similar color characteristics across the two input frames, based on proximity in the HSV color space and pixel distribution.

Subsequently, a comparison is performed between the matched cluster pairs, as well as unmatched clusters from each image. The differences in color composition, cluster size, and presence or absence of corresponding features contribute to the computation of a degradation score, which uantifies the level of visual deterioration.

The input to the next stage of the algorithm consists of a list of matched cluster pairs, each pair representing similar color clusters identified in the original and degraded images. The output generated includes the computed distances between the centers of these matched clusters.

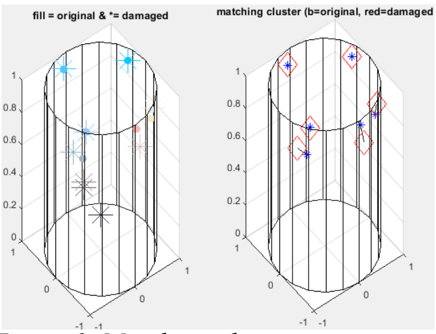

Figure 2. Matching clusters output step.

To compute the final degradation score, four key metrics were defined, encompassing both paired and unpaired clusters:

- Number of non-paired clusters,
- Distance between the centers of matched clusters,
- Difference in color proportion between paired clusters,
- Total percentage of non-paired cluster pixels.

To derive a scalar degradation score in the range [0, 1], we evaluated four candidate formulations based on the normalized quantities extracted from the cluster comparison module. Let M denote the set of matched cluster pairs, U - the set of unmatched clusters, D - the mean centroid distance across matched pairs, P - the mean absolute proportion difference for matched clusters, and U - the normalized proportion of pixels belonging to unmatched clusters. All quantities were normalized to [0, 1] prior to combination. The four candidate scoring formulas are defined as follows.

Variant 1: Non-Paired Cluster Count
This formulation uses only the normalized count of non-paired clusters:

$$S_1 = K = |n_r - n_d| \,/\, max(n_r, n_d) \quad (3)$$

This metric penalizes discrepancies in the number of color clusters; however, it frequently overestimates damage when small spurious clusters arise due to noise or minor graffiti marks.

Variant 2: Distance-based pair metric
This formulation combines centroid displacement and proportion deviation for matched clusters:

$$S_2 = (D_{norm} + P) \,/\, 2 \quad (4)$$

where $D_{\text{norm}} = D/d_{\max}$ and $d_{\max}$ is a conservative upper bound for the feature-space distance. Although this captures structural changes within matched colors, it fails when few or no matched clusters exist and remains insensitive to newly introduced colors.

Variant 3: Mean of first three parameters
This formulation integrates discrepancies in cluster count, centroid position, and proportional mass:

$$S_3 = (K+D_{norm}+P)/3 \quad (5)$$

While more comprehensive, this metric often underestimates damage in cases of minor but visually noticeable surface alterations, because centroid shifts and proportion differences remain small even when new colors appear.

Variant 4: Unmatched-clusters proportion
The fourth formulation computes the fraction of pixel mass corresponding to unmatched clusters:

$$S_4 = U = \frac{\sum_{i\in\mathcal{U}^R} n_i^R + \sum_{j\in\mathcal{U}^D} n_j^D}{N} \quad (6)$$

where N is the total number of pixels in the sign region. This quantity directly represents the visible area that does not match the reference sign's chromatic structure.

3.4. Selection of the Final Formula (Based on Human Perception Study)

To determine which formulation aligns most closely with human perception, subjective degradation scores were collected through a survey involving licensed drivers. Participants evaluated the visual deterioration of each degraded sign on a normalized scale, providing a perceptual reference for comparison. Among the four candidates, the unmatched-pixel formulation S_4 exhibited the highest consistency with human judgments. Specifically, it achieved the lowest average deviation from the subjective scores across a range of degradation types, including fading, graffiti, dirt accumulation, and partial occlusions. The measure implicitly reflects the proportion of the sign's surface that appears altered—an attribute strongly correlated with human perception of damage. As a result, the final degradation score adopted in this work is: S = U.

This formulation offers strong perceptual alignment, robustness to matching variability, and computational simplicity, rendering it suitable for real-time embedded applications in automotive environments.

The final result represents a normalized damage level ranging from 0 to 1, where 0 indicates no degradation and 1 corresponds to complete deterioration, implying no visual similarity between the two representations of the traffic sign. When the damage level reaches 0.5 or higher (typically meaning that more than half of the traffic sign is covered or deteriorated), it indicates significant degradation, and the precise calculation of the value becomes less reliable.

IV. EXPERIMENTS AND RESULTS

4.1. Experimental Setup

The evaluation dataset consists of traffic sign images extracted from video sequences recorded with an in-vehicle stereo camera operated under controlled geometry: the camera remained fixed relative to the vehicle, ensuring approximately constant viewing angle during acquisition. All experiments, synthetic and real, were conducted under the assumption that two images compared for degradation correspond to **the same traffic sign**, captured under **similar illumination and similar viewpoint**, reflecting the operational conditions of bus-mounted camera systems used for infrastructure monitoring.

These preconditions are fundamental to the proposed feature. When they are not satisfied, illumination shifts distort the HSV descriptors and directly affect the cluster-matching pipeline, as discussed later.

Two complementary evaluation scenarios were considered:

1. **Synthetic degradation**, where controlled alterations were added in MATLAB to simulate realistic deterioration mechanisms.
2. **Real degradation**, where images of the same physical sign taken one or two years apart were compared.

4.2. RMSE and Accuracy

The algorithm produces a **numerical degradation score**, and accuracy is quantified by:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_i \left(d_i^{est} - d_i^{exp} \right)^2} \qquad (7)$$

where d_i^{est} is the estimated value and d_i^{exp} is the human expected reference.

An RMSE of 0 indicates perfect agreement. An RMSE of 0.12 corresponds to approximately 88% agreement on a 0--1 scale, since the average absolute deviation is below 0.12.

Thus, RMSE directly expresses how close the algorithm's scores are to human perception.

4.3. Synthetic Degradation

Eight types of synthetic deterioration were applied to a dataset of 377 traffic signs, producing over 3000 (original, degraded) image pairs. For each case, human observers provided an **expected degradation score in [0,1]**, where:

- **0 = no visible damage**,
- **1 = extremely heavy degradation or severe occlusion (≈ "unusable sign")**.

These human judgements were used as ground-truth targets for evaluating the proposed numerical damage metric.

4.4. Synthetic Results Summary

Most degradation types produced low RMSE values (<0.07), indicating strong agreement.

Only severe occlusions (>50%) produced high errors, as the algorithm is not designed for cases where more than half the sign surface becomes invisible, a scenario where even human observers often assign values close to 1.

Cases exceeding 50% occlusion produce discrepancies because they break an implicit assumption of the method: at least part of the original chromatic structure must remain visible for a meaningful comparison.

When excluding these two extreme categories, the global RMSE drops from **0.12 to 0.08**, corresponding to an accuracy above **92%**.

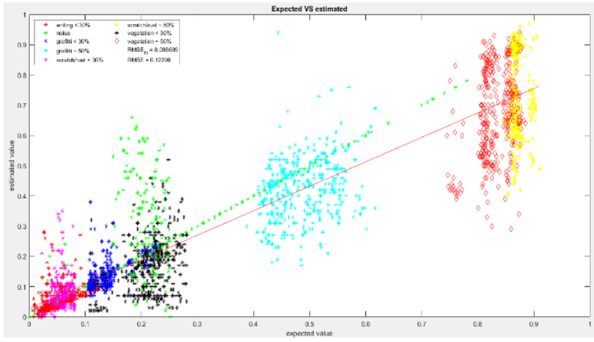

Figure 3. Figure representing Statistics of Expected vs. Estimated Damage Value on synthetic degradation.

Synthetic degradation	RMSE value	Comments
Writing < 30%	0.003	Good agreement
Noise	0.037	Good agreement
Grafitti < 30%	0.002	Good agreement
Grafitti ≈ 50%	0.064	Good agreement
Rust < 30%	0.011	Good agreement
Rust > 50%	0.12	Large deviation
Vegetation < 30%	0.045	Good agreement
Vegetation>50%	0.155	Large deviation

Table 1: RMSE values for each category from Fig.4

4.5. Real Degradation

Real traffic sign acquisition proved challenging: only 12 valid (original, aged) pairs were obtained, and none exhibited strong physical degradation. Furthermore, these images violated the required preconditions:

- The signs were photographed in different seasons,
- Under different lighting conditions (direct sunlight vs. overcast sky),
- With non-identical camera orientation due to vehicle stopping position.

Under these deviations, the algorithm often failed to match clusters between the two images, producing degradation values close to 1 (not because the sign was damaged, but because the color distributions were fundamentally different after illumination changes).

For real images with no visible degradation, the human expected value is naturally close to 0.

This result is not an indicator that the feature is fundamentally flaweed, rather it confirms that the method is not illumination invariant and requires controlled acquisition conditions.

This is consistent with its intended deployment:

- Bus-mounted cameras reviewing the same street at the same time of day,
- Same viewpoint, same weather and lighting context,
- Frames captured from consistent angles during routine infrastructure monitoring.

Under such operational constraints, illumination variation is minimal, and the method is expected to behave similarly to the synthetic case.

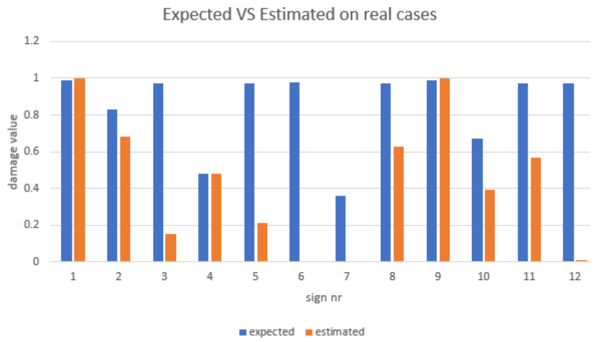

Figure 4. Figure representing Statistics of Expected vs. Estimated Damage Value on real degradation.

V. CONCLUSIONS AND FUTURE WORK

The proposed method relies on a classical, threshold-based vector quantization algorithm adapted to the HSV color space to estimate traffic sign degradation. This choice was motivated by its compatibility with resource-constrained embedded environments and its ability to operate without a predefined number of clusters. Its key advantages include:

• High Processing Speed: Enables real-time execution on hardware with limited computational capacity.
• No Need for a Predefined Number of Clusters: Dynamically adjusts to the number of distinct colours present in each input image.
• Simplicity and Flexibility: Straightforward to implement and adapt to specific application constraints.
• Effective Feature Extraction: Identifies dominant color clusters that form a compact descriptor of the sign's appearance.

Experimental results show that the method performs reliably under the predefined operating conditions and captures perceptually meaningful degradation levels. However, accuracy decreases in scenarios involving complex real-world alterations, as the method relies exclusively on chromatic cues that cannot fully describe all types of visual deterioration.

Future work will therefore continue to emphasize classical, computationally lightweight techniques. Several realistic extensions are envisioned.

First, additional non-deep-learning features, such as local texture descriptors or simple gradient- and edge-based statistics, may help capture structural degradation phenomena that are not purely chromatic, while preserving explainability and low computational cost.

Second, cluster post-processing strategies (e.g., centroid-based merging or noise suppression) could reduce over-segmentation and improve robustness across varying input resolutions.

Third, evaluating the system on a broader real-world dataset, ideally containing time-spaced captures of the same signs, would provide a more accurate assessment of long-term degradation patterns.

In addition, machine learning approaches may be explored as a complementary direction. Rather than replacing the proposed classical framework, future work may investigate learning-based models that operate on top of the extracted color–texture descriptors. Such models could learn to predict degradation levels from a compact feature vector rather than raw images, allowing ML to enhance robustness while preserving computational efficiency. This incremental integration would enable the use of data-driven without abandoning the interpretability and resource efficiency of the current system. In summary, this work demonstrates that classical vector quantization, combined with perceptually oriented color representation, provides a viable and efficient foundation for estimating traffic sign degradation. The proposed method highlights the continued relevance of lightweight, interpretable algorithms in an era increasingly dominated by deep learning, offering a complementary solution suitable for embedded perception systems where computational budget and transparency remain essential.

REFERENCES

[1] Z. Dong, W. Shi, G. Tong, and K. Yang, "Collaborative autonomous driving: Vision and challenges," 2020 Int. Conference on Connected and Autonomous Driving (MetroCAD), pp. 17–26.

[2] S. Koutra, V. Becue, and C. S. Ioakimidis, "A multiscalar approach for 'smart city' planning," in 2018 IEEE International Smart Cities Conference (ISC2), 2018, pp.1–7.

[3] Ersü et al., "Automatic Visual Traffic Sign Damage Detection and Measurement of Damaged Area", 2024 18th International Conference on Control, Automation, Robotics and Vision

[4] A. Mannan, K. Javed, A. Ur Rehman, H. A. Babri and S. K. Noon, "Classification of Degraded Traffic Signs Using Flexible Mixture Model and Transfer Learning," in *IEEE Access*, vol. 7, pp. 148800-148813, 2019, doi: 10.1109/ACCESS.2019.2947069

[5] Garta, I.Y.; Tai, S.-K.; Chen, R.-C., "Improved Detection of Multi-Class Bad Traffic Signs Using Ensemble and Test Time Augmentation Based on Yolov5 Models". *Appl. Sci.* 2024, *14*, 8200. https://doi.org/10.3390/app14188200

[6] G. Floros, K. Kyritsis and G. Potamianos, "Database and baseline system for detecting degraded traffic signs in urban environments," *2014 5th European Workshop on Visual Information Processing (EUVIP)*, Paris, France, 2014, pp. 1-5, doi: 10.1109 /EUVIP.2014.7018395

[7] Mannan, A. and Javed, K. and ur Rehman, Atta and Babri, H. A. and Noon, S. K.,Cognition based Recognition of Partially Occluded Traffic Signs, journal Scientia Iranica, volume 29,number 4,pages 1881-1897, doi 10.24200/sci.2022.55200.4111

[8] "Clustering into an unknown number of clusters," Clustering Into an Unknown Number of Clusters | Baeldung on Computer Science, accessed: 2025-03-03.

[9] H. B. Rene Boite, Traitement de la parole, Presse Polythechnique et Universitaire Romandes, 2000, pp. 146.

[10] S. R. P. B. Amrita Mohanty, Zameer Muzaffar Mir,"Color image segmentation using automated k-means clustering with rgb and hsv color spaces,"International Journal of Computer Applications" (0975– 8887), vol. 79, no. 2, pp. 31–43, 2013.

[11] Chengjun Liu, Guojun Wang, "A Comparative Study of Color Spaces for Color Image Segmentation", Journal Pattern Recognition, 2003

[12] T. I. Md. Rakib Hassan, Romana Rahman Ema, "Analysis of color images using cluster based segmentation techniques," Global Journal of Computer Science and Technology, vol. 1, no. 2, 2017.