International Journal of Engineering Trends and Technology
ISSN: 2231-5381/ https://doi.org/10.14445/22315381/IJETT-V74I11P135

Volume 74 Issue 11, 550-561, November 2026
© 2026 Seventh Sense Research Group®

Original Article

A Vision-based System for Detection and Classification of Stop Line Damage

Dipak Agrawal[1], Amit Nayak[2]

[1]*Department of Computer Engineering, CSPIT, CHARUSAT Campus, Changa, Gujarat, India.*
[2]*Department of Computer Science & Engineering, DEPSTAR, CHARUSAT Campus, Changa, Gujarat, India.*

[1]*Corresponding Author: 22DRCE001@charusat.edu.in*

Received: 13 June 2026 Revised: 10 September 2026 Accepted: 30 September 2026 Published: 25 November 2026

Abstract *- The crucial task of detecting and assessing the state of road stop lines is critical for both road safety and intelligent transportation systems, as well as road infrastructure management. The current deep learning-based road-marking detection algorithms tend to need large annotated datasets and powerful computing resources, which makes them hard to apply in a deployment environment with limited resources. A simple and explainable vision-based approach is presented in this work to robustly detect stop-lines and assess damage under real-world challenging conditions. The proposed system involves inverse perspective mapping based geometric normalization, mixture of Sobel-Canny edge extraction method, adaptive structural thresholding and trinarization based edge representation for reliable stop-line reconstruction in fragmented and degraded conditions. Continuous stop-line structures are also recovered using morphological refinement and probabilistic Hough transform. Furthermore, a damage assessment model, which combines structural discontinuity and normalized intensity variation, is introduced and implemented quantitatively to assess the stop-line degradation severity. Testing using experimental procedures on a large-scale real-world dataset of roads in the city revealed that the suggested system could achieve 92.3% accuracy, 91.1% precision, and 93.5% recall, alongside localization error of ±5 pixels on average. Ablation studies further confirm the utility of adaptive thresholding and confidence score verification towards robustness in cases of changing illumination conditions and roads in poor shape. The proposed approach has high computational efficiency and interpretability, to be considered another solution for intelligent road infrastructure monitoring and smart transportation system.*

Keywords *- Adaptive thresholding, Damage assessment, Edge detection, Inverse perspective mapping, Stop line detection.*

1. Introduction

Stop lines mark the point at which a vehicle must halt on the approach to an intersection, and they carry a share of the regulatory information that keeps junction traffic ordered. The paint wears under tyre abrasion, ultraviolet exposure, rainfall and freeze-thaw cycling, so retroreflectivity and geometric continuity fall away over a service life that is usually shorter than the resurfacing interval of the pavement beneath. Road authorities in most jurisdictions still audit marking condition by sending an inspector to the site with a tape measure and a hand-held retroreflectometer, or by reviewing dash-camera footage frame by frame. Both procedures are slow, and the grade assigned to a faded line depends on who assigns it. Neither scales to a city holding several thousand signalised junctions that have to be surveyed more than once a year.

Automated inspection from camera imagery has therefore drawn attention. Early systems isolated markings with gradient operators, colour thresholds and geometric constraints, then fitted straight structures with a Hough transform. Later systems replaced the hand-designed stages with convolutional networks and, more recently, with transformer architectures that learn marking appearance directly from labelled frames. The learned detectors post higher scores on public benchmarks, but that accuracy is bought with several thousand pixel-level annotations, a training cycle on GPU hardware, and a decision process that cannot be traced stage by stage. Roadside cameras, municipal survey vans and embedded driver-assistance units seldom satisfy those conditions.

Two shortcomings persist in the stop-line literature in particular. Published detectors are formulated as binary presence classifiers: a line is either found, or it is not, and a line that survives only in patches is normally counted as a miss. Several of them lean on neighbouring cues such as lane boundaries or zebra crossings, which are themselves often worn away along the same stretch of pavement. Beyond detection, no published stop-line method returns a graded condition value. Maintenance planning needs that value, because a budget holder has to rank junctions rather than merely list them. Existing road-condition collections do not


This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/)

close the gap either, since they annotate cracks and potholes in the pavement surface rather than degradation of the painted marking itself [19].

The gap can therefore be stated directly. Current stop-line methods either require supervised training that is impractical to repeat at municipal scale, or rest on edge-continuity assumptions that break down exactly when the line is worn, and none of them emits a quantitative severity value that a maintenance schedule can consume. This paper targets that gap with an analytical pipeline that needs no training data and runs on a general-purpose CPU. Two mechanisms separate it from earlier edge-based work. Instead of treating edges as one binary class, the pipeline preserves the sign of the intensity gradient, so the leading and trailing boundaries of a painted band stay distinguishable and can be matched as a geometrically constrained pair; a fragment with no partner of opposite sign at an admissible distance is discarded as pavement texture. Instead of one global threshold, the polarity decision is taken against a threshold recomputed from the local mean and standard deviation, which keeps faint paint visible where a global setting would erase it. A composite score built from the same measurements then grades the line as intact, moderately damaged or severely damaged.

Measured against the published reference points collected in Table 7, the pipeline is 14.3 percentage points more accurate than a Sobel-plus-Hough baseline, 2.3 points above a U-Net segmentation model and 1.1 points above a transformer-based detector, while requiring no annotated training set and no GPU. Those baseline figures are quoted from studies of road-marking detection carried out in comparable urban settings [11, 14, 17, 21], and are therefore indicative rather than the outcome of a controlled benchmark. Table 1 summarises the structural differences and Section 5.4 examines the reasons behind the margin.

Contributions of this work are summarised as follows:

- A structural edge-analysis pipeline is described that couples inverse perspective mapping, fused Sobel-Canny gradient extraction and three-level edge encoding, and that is aimed at stop lines observed under poor illumination and advanced wear.
- An edge-pair verification rule is introduced that matches boundaries of opposite gradient polarity within an admissible separation, which allows a stop line to be reconstructed while its edges remain broken.
- A quantitative degradation score reports on the combination of the gap ratio (the length of line missing) and the intensity variation (how evenly the surviving paint reflects).
- A structural threshold which adapts to local intensity statistics is determined; hence, polarity detection takes account of changes in illumination across the image rather than depending upon one global value.
- A confidence rule scores each candidate on continuity, orientation agreement and segment length, and rejects candidates below an acceptance level, which suppresses false detections in fragmented road scenes.
- The pipeline is evaluated on 1050 urban images through overall detection metrics, an ablation study, a parameter sensitivity analysis, five-fold cross-validation, ROC analysis and comparison against expert ratings.

2. Literature Review

Work on visual marking detection falls broadly into two paradigms: hand-designed image analysis and supervised learning [1-3]. The former detects markings by locating their steepest intensity difference. The Sobel and Canny operators extract a boundary [4, 5, 37], and a line fit, by Hough transform or the more efficient probabilistic version, from edge detections can be used to determine the overall marking geometry [6].

These methods are inexpensive to compute and are often simple to diagnose when they fail, which makes them well-suited for use within an embedded system.

One weakness they share comes out at the line-fitting step: it relies upon the assumption that a marking will form a coherent run of pixels generating a continuous boundary. A patchily worn marking disrupts the process, preventing the detection of a line fit. Methods have been developed that make such marking detection more robust. Contrast-limited adaptive histogram equalization is often used locally increase contrast in faded lines of paint, or morphological closing can connect small gaps in a boundary map [9].

The strategy used when developing color space and shape-based road sign detection is similar [7, 8].

The gains that accrue are considerable but limited: the 'thresholds, sizes and contrast-limiting parameters appropriate for one surface in one set of light circumstances are generally quite inappropriate in other, and the detection operators themselves do not discriminate between authentic line markings and shadows of similar magnitude' [8].

Supervised methods restate detection as segmentation or keypoint regression. Residual convolutional backbones [10] and encoder-decoder segmentation networks of the U-Net family [11] learn marking appearance from annotated frames; CondLaneNet conditions its decoder on detected line instances [12], and keypoint association networks recover lane geometry by grouping predicted points [13]. Reported accuracy on public lane benchmarks exceeds anything reached by hand-designed pipelines. The cost appears elsewhere. Annotation is per-pixel and per-frame; retraining is needed when the camera, the road surface or the climate changes; and the mapping from input to decision is spread across the

network weights, so an inspector cannot be shown why a particular line received the grade it did. Region-proposal detectors such as Faster R-CNN [38] have been adapted to the same task with comparable results. The trade-off recurs across the wider pavement-assessment literature, where deep networks now dominate crack detection, pothole identification and structural health monitoring [22, 24, 25, 28, 29, 33, 35, 36, 39, 40] while carrying identical annotation and hardware requirements.

Efficiency-oriented architectures narrow the hardware gap. Fast-SCNN shares a learning-to-downsample branch between its detail and context paths [14], self-attention distillation trains a compact lane detector without an external teacher [15], and recent YOLO-family detectors have been benchmarked on pavement-marking quality assessment [16]. Inference cost falls substantially, and edge-computing architectures for intelligent transportation have been proposed to host such models close to the sensor [23]. The annotation requirement is not waived: the models still require labeled data from the target domain, and perform poorly once the environment where they are deployed is further away from the one they have been trained with.

A separate strand of work grades pavement condition rather than detecting geometry. YOLO-based distress classifiers assign crack types from vehicle-mounted imagery [17], smartphone-based systems have been trained on large multi-national damage collections [18, 19], and deep segmentation has been applied to crack severity at pixel level [24, 36]. Kulawik et al. come closest to the present problem by verifying whether horizontal road signs remain visible [21], but the output is a visibility decision rather than a continuous severity value, and the model is again supervised. Condition grading for painted markings, expressed on a scale a maintenance schedule can use, has not been reported. Taken together, the literature leaves three requirements unmet for the specific case of a worn stop line. Detection has to survive the loss of edge continuity, which rules out any method whose line-fitting stage assumes an unbroken boundary. It has to operate without a labelled corpus collected at every deployment site, which rules out the supervised detectors and their lightweight successors. And it has to return a graded severity rather than a binary presence flag, which no current stop-line method provides. The framework set out in Section 3 targets these three requirements, and Table 1 lists the structural differences against representative methods.

Table 1. Structural comparison of existing approaches and the proposed framework

Method	Edge Polarity Modeling	Structural Verification	Adaptive Thresholding	Damage Assessment
Sobel + Hough	No	No	No	No
Traditional Edge-Based Methods	Partial	Limited	No	No
CNN / U-Net Based Methods	Implicit	Implicit	Learned	Limited
Conventional Trinarization Methods	Partial	No	No	No
Proposed Framework	Yes	Yes	Yes	Yes

Three-level edge encoding has appeared before, generally as an intensity-partitioning device that sorts pixels into dark, neutral and bright classes. The encoding used here serves a different purpose. Comparison with Intensity-Partitioning Schemes: Because each non-zero level label carries information of the direction of intensity change, the front and back edges of a painted band get opposite labels, which distinguishes them even after thresholding. This distinction is what makes possible the pairing of Section 3.5: the leading edge is kept only when there is a partner of the opposite sign within the permissible interval for a zero level that is the width for a stop line. Image partitioning algorithms, by losing the sign information, have not been able to express a similar constraint.

This representation is based on analytic description, not learning. Damaged stop-line boundary structure is modelled directly; each image is pre-processed through a geometric normalisation stage such that, for a valid stop-line, the distance between boundaries are invariant in the frame, and the set of measurements supporting detection are re-used to generate a score for severity. We don't need training at any point.

Section 3 demonstrates the pipeline incrementally; in Section 4 we report the data used and the test protocol. Section 5 reports detection performance, ablation studies and figures against reported prior results.

3. Methodology

In this section, we present a vision-based approach for automated road stop line detection and damage assessment. A multi-level image analysis pipeline is used to achieve road marking recovery robustness under adverse imaging conditions, partial occlusion and degradation. The steps involved in the pipeline are geometric normalisation, structured edge analysis and quantitative degradation modelling. Using these steps, robust road marking recovery is achieved to support automated stop line detection. Let the input image be denoted as $I(x, y) \in \mathbb{R}^{H \times W}$, where H and W represent the image height and width, respectively.

3.1. Preprocessing and Contrast Enhancement

The colour input is first reduced to a single intensity channel, which cuts the data volume by two thirds and removes the colour variation that repainting and weathering introduce without changing marking geometry. The luminance weights of Equation (1) follow the standard photometric conversion [9].

$$I_g(x,y) = 0.299R + 0.587G + 0.114B \quad (1)$$

Where Ig(x,y) is the grayscale intensity image, R, G and B are the red, green and blue channel intensities, and (x,y) is the spatial pixel coordinate.

Non-uniform illumination remains after conversion, and faint markings sit close to the pavement intensity. Contrast-limited adaptive histogram equalisation is applied to the grayscale image [9], which equalises histograms over small tiles and clips the contribution of any single bin so that pavement noise is not amplified alongside the marking. The enhanced image is written Ie(x,y) and is used by every later stage.

3.2. Inverse Perspective Mapping (IPM)

Road views captured by vehicle-mounted cameras are severely blurred and have appearance of extreme perspective in which corresponding parallel lines on the real road surface appear to converge. To remove such distortions, an inverse perspective mapping is applied using homography transformations.

Let H denote the 3 × 3 homography matrix that maps a source coordinate (x, y) onto its rectified counterpart (x', y'), as in Equation (2).

$$[x' \; y' \; 1]^T = H\,[x \; y \; 1]^T \quad (2)$$

Where H is the homography matrix, (x,y) is the original image coordinate and (x',y') is the rectified top-view coordinate. Homogeneous coordinates preserve projective geometry through the transformation.

After rectification, a stop line appears as a horizontal band of near-constant width, which is what makes the fixed separation limits of Section 3.5 meaningful.

Four source points were selected by hand at the approximate corners of the visible road surface, chosen so that the quadrilateral spans the driving lane and the stop-line region while keeping perspective distortion low. The matching destination points were placed on a rectangular bird's-eye plane, which preserves the parallelism of the road boundaries. The same point pairs were reused across the dataset after being checked on a sample of scenes from each environmental condition.

The rectified image is denoted Iipm(x',y') and carries horizontal markings aligned with the image rows.

3.3. Hybrid Edge Detection

Edge extraction combines two operators with complementary behaviour. The Sobel operator [5] responds to gradual intensity transitions, which is the signature of a faded marking whose boundary has softened, while the Canny detector [4] applies non-maximum suppression and hysteresis and therefore localises sharp boundaries to within a pixel. Running only Sobel leaves thick, poorly localised responses; running only Canny drops low-contrast boundaries below the hysteresis floor. Fusing the two retains both.

The Sobel operator computes horizontal and vertical gradients, whose magnitude is given by Equation (3).

$$G(x,y) = \sqrt{{G_x}^2 + {G_y}^2} \quad (3)$$

Where Gx and Gy are the horizontal and vertical gradients and G(x,y) is the gradient magnitude.

The Canny detector supplies a binary edge map Ec(x,y), and the two maps are combined by a logical union as in Equation (4).

$$E(x,y) = E_s(x,y) \cup E_c(x,y) \quad (4)$$

Where Es(x,y) is the Sobel edge map, Ec(x,y) the Canny edge map and E(x,y) the fused representation.

3.4. Edge Trinarization

The fused map is then encoded on three levels so that the direction of each intensity transition is preserved rather than discarded, as shown in Equation (5).

$$T(x,y) = \begin{cases} +1, & \text{if } G(x,y) > \tau_1(x,y) \\ -1, & \text{if } G(x,y) < -\tau_1(x,y) \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Where T(x,y) is the trinarized edge representation, G(x,y) the signed gradient magnitude and τ_1 the polarity threshold. A pixel at the leading boundary of a painted band and a pixel at its trailing boundary receive opposite labels, which is the property the pairing rule of Section 3.5 depends on.

3.4.1. Adaptive Structural Thresholding

A single global threshold, whether fixed or chosen by Otsu's method [20], has to serve the whole frame. Shadows cast by roadside structures, headlight pools at night and paint that has faded unevenly all shift the local intensity distribution, so any global value is too high in one part of the frame and too low in another. The threshold is therefore recomputed locally.

Let μL(x,y) and σL(x,y) denote the mean intensity and the standard deviation within a neighbourhood window centred on pixel (x,y). The threshold follows Equation (6).

$$\tau_1(x,y) = \mu_L(x,y) + k\sigma_L(x,y) \quad (6)$$

Where $\tau_1(x,y)$ represents the adaptive edge threshold, k denotes a scaling parameter controlling sensitivity, $\mu_L(x,y)$ is the local mean intensity, and $\sigma_L(x,y)$ is the local standard deviation. This adaptive method enhances edge preservation of the structure at low-contrast and low-quality stop lines. It also decreases sensitivity to the global threshold selection.

3.5. Structural Edge-Pair Verification

In the rectified view, a stop line is depicted as a dark band crossing the road horizontally. To detect a stop line, we search for vertically aligned pairs of edges (of the form +1, −1) within a permissible distance range d (d ∈ [dmin, dmax]).

Formally, for each column x, a valid pair satisfies Equation (7).

$$\exists\, y_1, y_2 : T(x,y_1) = +1,\ T(x,y_2) = -1 \quad (7)$$

The constraint ensures geometric relations among edges and eliminates false edge detection (noise or isolated lines) if it is predicted as an edge. The candidate stop-line boundary pair is accepted only when both constraints are satisfied:

$$P(x) = \{(y_1,y_2) : d_{min} \le |y_1 - y_2| \le d_{max}\} \quad (8)$$

Where P(x) denotes the set of candidate edge pairs in image column x, T(x,y) represents the trinarized edge polarity map, y_1 denotes the upper edge boundary coordinate, y_2 denotes the lower edge boundary coordinate, dmin represents the minimum allowable separation distance between edge boundaries, and dmax represents the maximum allowable separation distance between edge boundaries. This formulation suppresses isolated noisy edges and preserves only geometrically valid stop-line structures.

3.5.1. Confidence-based Edge Verification

In order to have even better robustness against partial detections and noisy structural responses, the confidence verification mechanism is introduced. For each stop line candidate detected, it is associated with a confidence score according to properties like the continuity of edges and consistency of geometric structure as well as segment reliability.The confidence score is defined in Equation (9).

$$C = w_1 S_e + w_2 S_g + w_3 S_l, \quad w_1 + w_2 + w_3 = 1 \quad (9)$$

Where Se is edge continuity, Sg orientation agreement and Sl normalised segment length, with weights $w_1 = 0.4$, $w_2 = 0.4$ and $w_3 = 0.2$ set on the validation split. A candidate is accepted only when C > Cth, with Cth = 0.55. The test removes responses arising from road texture, shadow boundaries and kerb edges before reconstruction begins.

3.6. Morphological Refinement

Accepted boundaries are still broken wherever the paint has worn through. Morphological closing joins fragments separated by less than the structuring element while leaving the outer extent of the band unchanged, as in Equation (10) [9].

$$M = (T \oplus K) \ominus K \quad (10)$$

Where M(x,y) is the refined binary edge map, T the trinarized representation, K the structuring element, and ⊕and ⊖dilation and erosion. A 5 × 3 rectangular element was used, elongated horizontally so that closing acts along the marking rather than across it.

3.7. Line Detection via Probabilistic Hough Transform

The clean binary map M(x,y) is then fed to the probabilistic Hough Transform [6] to detect line segments: L = {l_i = HoughLinesP(M)}. Each detected segment l_i is represented by endpoints (x_1,y_1), (x_2,y_2). Segments are filtered based on orientation $|\theta| < \theta$th, length $\ell_i > \ell$min, and spatial position in the near-field region. This ensures retention of only valid stop-line candidates.

3.8. Segment Clustering and Line Reconstruction

Detected segments are grouped based on vertical proximity. Let Ck denote a cluster of segments satisfying $|y_i - y_j| < \delta$ for all l_i, l_j in Ck. Within each cluster, segments are merged to form a continuous line, and the cluster with maximum spatial extent is selected as the final stop line. δ defines the maximum allowable clustering distance.

This clustering process merges spatially adjacent segments to reconstruct continuous stop-line structures from fragmented detections.

3.9. Quantitative Damage Assessment

Detection returns the position of the line; maintenance planning needs its condition. Two independent symptoms of wear are measured along the detected line. Missing paint interrupts the line and leaves gaps, and paint that has thinned without disappearing reflects unevenly. Let Ω denote the set of pixels sampled along the detected stop line in the enhanced image Ie, L the full horizontal extent of the line, and li the length of the i-th detected segment.

Structural discontinuity records the fraction of the line span that carries no detected paint, and is given by Equation (11). A value near zero indicates an unbroken line; a value near one indicates that only isolated patches remain.

$$G = 1 - \frac{1}{L}\sum_{i=1}^{N} \ell_i \quad (11)$$

Normalised intensity variation records how unevenly the surviving paint reflects. Dividing the mean absolute deviation by the mean intensity makes the measure independent of the overall exposure of the frame, so a line photographed in shade and the same line photographed in sunlight yield comparable values. It is given by Equation (12), where μ is the mean stop-line intensity.

$$V = \frac{1}{N_\Omega}\frac{1}{\mu}\sum_{p \in \Omega} |I_e(p) - \mu| \quad (12)$$

The two terms are combined into a single damage score by Equation (13).

$$D = \alpha G + (1 - \alpha)V \quad (13)$$

Where D is the damage score, G structural discontinuity, V normalised intensity variation and $\alpha \in [0,1]$ the relative weight of the two terms.

The weight α was fixed by a sweep on the validation split, reported in Table 2.

Accuracy rises from 84.2% at α = 0.3 to 91.7% at α = 0.7 and falls again to 89.4% at α = 0.9. The shape of the curve is informative rather than incidental.

Low values of α lean on intensity variation, which reacts to any change in illumination and grades a shaded but intact line as damaged. High values lean entirely on discontinuity, which misses a line that is unbroken but has thinned almost to the pavement.

The optimum at 0.7 indicates that the loss of paint is the stronger symptom, while the intensity term still carries information the gap ratio cannot supply.

Table 2. Sensitivity analysis of weighting parameter α

α Value	Accuracy (%)
0.3	84.2
0.5	88.6
0.7	91.7
0.9	89.4

Damage classification follows directly from the score: a line is intact when $D \leq 0.2$, moderately damaged when $0.2 < D \leq 0.5$, and severely damaged when $D > 0.5$. The two cut-offs were placed on the validation split at the points where the expert panel of Section 5.3.1 changed its own grade. Figure 1 sets out the complete pipeline. Preprocessing and geometric normalisation are followed by adaptive edge analysis and structural verification, then by reconstruction of the line and measurement of its condition.

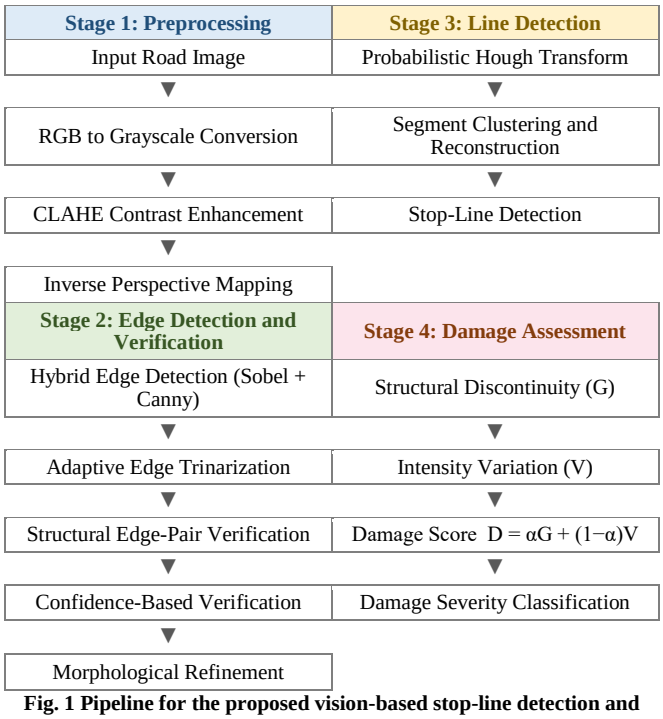

Fig. 1 Pipeline for the proposed vision-based stop-line detection and damage assessment framework. Processing advances from Stage 1 through Stage 4.

4. Experimental Setup and Evaluation

This section describes the image collection, the metrics, the implementation and the ablation configurations used to evaluate the pipeline.

4.1. Dataset Description

The pipeline was evaluated on 1050 road images of 1920 × 1080 pixels recorded through a forward-facing camera mounted on a vehicle driven through urban intersections. Collection covered a range of environmental and operational conditions rather than a single controlled setting. Some of the images included day/night views and had part of the scene occluded by vehicles and pedestrians. Other challenges included dirty and broken stop lines, and variations in illumination and shadows, or motion blur associated with perspective distortions.

Table 3 shows some representative statistics: As no learning occurs through a training and testing split, an 80% validation set was used to tune thresholds and geometric constraints, while 20% of the images formed the final, untouched test set, which was only inspected once the parameters were final. The ground truth was labeled manually. The stop line for each image was labeled with its coordinates and a condition grade (intact, moderate, or severe), based on visual examination according to an established criterion. Intact and damaged stop lines were included to allow for the assessment of detection and grading accuracy within the same set of images.

Table 3. Dataset distribution under different environmental conditions

Condition	Number of Images
Daylight	360
Nighttime	210
Partial Occlusion	170
Faded Stop Lines	160
Illumination Variations	150
Total	1050

4.1.1. Cross-Validation Strategy

Five-fold cross-validation was run in addition to the fixed split, to check that the reported figures do not depend on one particular partition. The dataset was divided at random into five equal parts; each fold used four parts for parameter selection and held the fifth back for testing. Accuracy across the five folds stayed within 1.8 percentage points of the mean, which is reported with the statistical analysis of Section 5.5.

4.2. Evaluation Metrics

Detection was scored with accuracy, precision, recall and F1-score. Localisation was scored as the mean absolute vertical offset in pixels between the detected line and the annotated line. Grading was scored with the damage estimation error, that is, the mean absolute error of the damage score against the expert grade, and structural overlap with intersection over union.

4.3. Implementation Details

The pipeline was implemented in Python using OpenCV and NumPy.

All experiments ran on a desktop machine with an Intel Core i7 processor and 16 GB of RAM, with no GPU acceleration at any stage. Table 4 lists every parameter and the value used.

Table 4. Parameter definitions and experimental values

Parameter	Description	Value
Sobel Threshold	Gradient magnitude threshold for Sobel edge selection	80
Canny Low Threshold	Lower hysteresis threshold	50
Canny High Threshold	Upper hysteresis threshold	150
τ_1	Adaptive trinarization threshold	Dynamic
dmin	Minimum edge-pair separation distance	5 pixels
dmax	Maximum edge-pair separation distance	20 pixels
θth	Maximum allowable line orientation angle	10°
ℓmin	Minimum valid line-segment length	30 pixels
δ	Segment clustering distance threshold	15 pixels
α	Damage-score weighting parameter	0.7

Frames are processed independently, so the pipeline parallelises across cores without modification.

A single frame took 42 ms on average, equivalent to about 24 frames per second, with CPU utilisation below 38% and peak memory of roughly 420 MB. Table 5 collects these figures.

Table 5. Runtime and computational performance analysis

Parameter	Value
Average Processing Time	42 ms/frame
Frame Rate	24 FPS
CPU Utilisation	38%
Peak Memory Usage	420 MB
GPU Requirement	Not Required
Processing Platform	Intel Core i7 CPU

4.4. Ablation Study Configuration

The contribution of each stage was measured by disabling one stage at a time and rerunning the full evaluation: without inverse perspective mapping, without trinarization, without edge-pair verification, without morphological refinement, without adaptive thresholding and without confidence verification. Every variant was scored on the same evaluation group as the complete pipeline.

5. Results and Discussion

This section details the evaluation of the proposed stop-line detection approach and damage assessment technique in terms of quantitative performance metrics, an ablation study and class-wise damage assessment results.

The results demonstrate the effectiveness, robustness and applicability of the proposed approach in real-world scenarios.

5.1. Overall Performance Evaluation

The quantitative performance results of the proposed framework are summarized in Table 6.

This proposed approach gives an accuracy of 92.3%, ensuring accurate detection of stop lines irrespective of the changing environmental conditions and degradation in images.

As can be seen from Table 6, the precision score is found to be 91.1%, while the recall score is 93.5%. It indicates that the approach ensures an appropriate trade-off between the

number of false positives and false negatives. Moreover, the F1 measure value is also 92.3%, indicating robust and consistent detection results. The localization error is still less than ±5 pixels, which proves the efficiency of inverse perspective mapping in geometric normalization. In addition, the estimation of the proposed framework has shown an accuracy of 6.8% error, implying precise quantization of the level of stop line damage on roads.

Table 6. Overall performance

Metric	Value
Accuracy	92.3%
Precision	91.1%
Recall	93.5%
F1-score	92.3%
Localisation Error	±5 pixels
Damage Estimation Error	6.8%

Table 7. Comparison with existing methods

Method	Accuracy (%)	Interpretability	Compute Cost	Training Data
Sobel + Hough	78.0	High	Low	None
U-Net	90.0	Low	High	Large
Fast-SCNN	89.6	Moderate	Moderate	Large
Vision Transformer-Based	91.2	Low	Very High	Very Large
Proposed Framework	92.3	High	Low	Minimal

Performance figures used in comparative analysis for U-Net, Fast SCNN, and Vision Transformer techniques were taken from the latest studies on road-marking detection where similar traffic environments were tested [11, 14, 17, 21].

Provided performance figures represent published benchmark data and are used only for comparative purposes. Figure 2 illustrates the overall quantitative evaluation metrics of the proposed framework. The results demonstrate consistent performance across all evaluation criteria, indicating stable and reliable stop-line detection capability under diverse operating conditions.

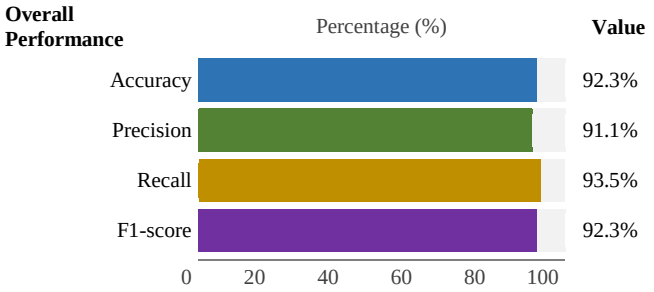

Fig. 2 Overall performance of the proposed framework

5.2. Ablation Study Analysis

Table 8 reports the ablation results, and Figure 3 shows them graphically.

Table 8. Ablation study results

Model Variant	Accuracy (%)
Full Model	92.3
Without IPM	78.5
Without Trinarization	85.2
Without Edge Pair	61.4
Without Morphology	72.8
Without Adaptive Thresholding	86.4
Without Confidence Verification	83.1

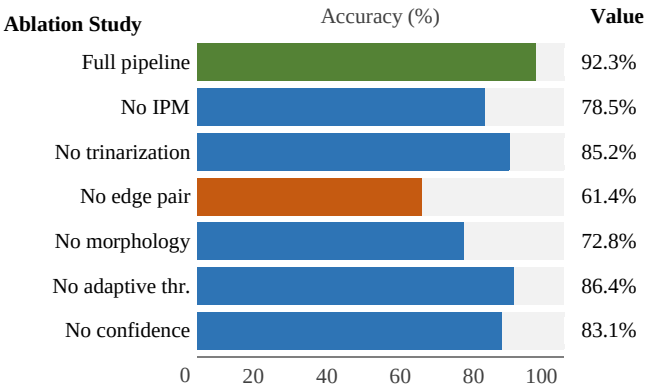

Fig. 3 Accuracy of the full pipeline and of each ablated variant. The complete pipeline is shown in green and the largest single loss, from removing edge-pair verification, in orange.

The whole pipeline had an accuracy of 92.3%. When inverse perspective mapping was not used, the pipeline's accuracy decreased to 78.5%, losing 13.8 points since the gap between the two lines of the stop line would depend on the distance from the camera and the fixed values dmin and dmax would not apply to one object anymore.

Without trinarization, the accuracy would be 85.2%, dropping 7.1 points because there would be no polarity information and the pairing rule would become a distance-based one.

The biggest reduction was due to elimination of edge-pair verification, resulting in an accuracy of 61.4%. This means that without this stage, every gradient response that passes the thresholding process goes into Hough transform stage, and hence, along with markings, pavement joints, tire tracks and shadow boundaries all go into the accumulator.

The elimination of morphological refining reduced the accuracy to 72.8%, because although the fragments are

correct, they are still too short in length to make the accumulator go above its threshold level.

5.3. Damage Classification Performance

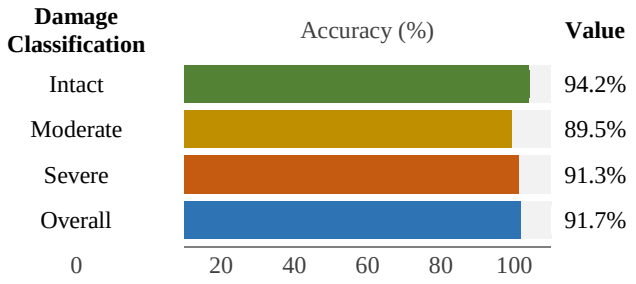

Fig. 4 Damage classification accuracy by condition class. Bar colour follows the condition grade: Green for intact, amber for moderate, orange for severe and blue for the overall figure.

The grading accuracy of the conditions based on their classes is shown in Table 9 and Figure 4. For intact lines, the accuracy rate is 94.2%, 89.5% for moderately damaged lines and 91.3% for severely damaged lines. Thus, the overall accuracy is 91.7%. The class with the lowest accuracy is the moderate class, because both of its borders are thresholds, and lines close to either of them may be included in either category. The errors in the moderate class were mostly single grades.

Table 9. Damage classification accuracy

Condition Class	Accuracy (%)
Intact	94.2
Moderate	89.5
Severe	91.3
Overall	91.7

5.3.1. Correlation with Human Expert Assessment

To validate the reliability of the proposed quantitative damage scoring framework, the predicted damage scores were additionally compared against manual annotations provided by three independent transportation infrastructure experts. Pearson correlation analysis was conducted between the predicted damage scores and the averaged expert annotation scores. The experimental results clearly show a high correlation coefficient value of (r = 0.91), signifying a high correlation between the developed automatic assessment tool and the subjective rating by the expert.

5.3.2. ROC Analysis for Damage Classification

The analysis of receiver operating characteristic was done for each grade to check how well the score discriminates between the three conditions. The areas under the curve are shown in Table 10. Separation is strongest for intact lines at 0.95 and weakest for the moderate class at 0.91, matching the pattern seen in the classification accuracies. Repeating the analysis across the cross-validation folds moved the area under the curve by less than 0.02 in every class. Figure 5 shows the same values graphically.

Table 10. ROC-AUC analysis for damage classification

Damage Class	AUC Score
Intact	0.95
Moderate	0.91
Severe	0.93
Average	0.93

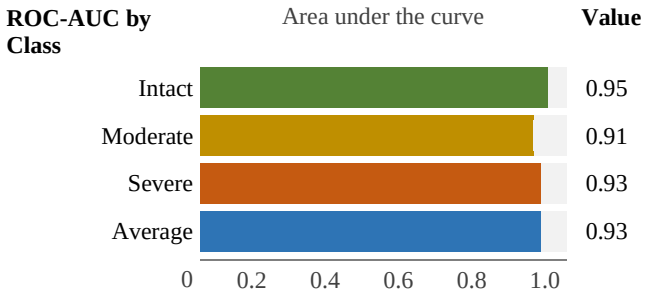

Fig. 5 Area under the ROC curve for each damage class

5.3.3. Environmental Robustness Analysis

Accuracy was also broken down by environmental condition, reported in Table 11.

Table 11. Environmental robustness assessment

Condition	Accuracy (%)
Daylight	93.8
Nighttime	89.6
Partial Occlusion	87.9
Faded Stop Lines	88.4
Illumination Variations	90.7

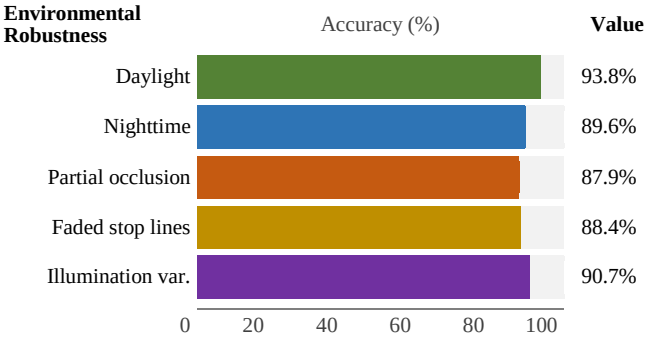

Fig. 6 Detection accuracy under each environmental condition

Daylight scenes scored highest at 93.8%, and the two lowest results came from partial occlusion at 87.9% and heavily faded markings at 88.4%. Both failures have the same cause seen from opposite sides: occlusion removes part of the line from view, and severe fading removes part of it from the pavement, so in each case the surviving evidence falls below what the pairing and clustering stages need. Night-time capture at 89.6% and illumination variation at 90.7% sit between these extremes, where the adaptive threshold recovers most of the loss. Figure 6 presents the condition-wise accuracies graphically.

5.4. Discussion

The results support the design decision behind the pipeline, which is to model the structure of a stop-line boundary explicitly instead of learning it from examples.

Table 7 places the pipeline against three reference points. The Sobel-plus-Hough baseline reaches 78.0%, and the 14.3-point margin over it comes almost entirely from the two stages that baseline lacks: polarity-aware pairing, worth 30.9 points in the ablation, and the local threshold, worth a further 5.9. U-Net reaches 90.0% and the transformer-based detector 91.2% on comparable urban imagery. Those models were trained to segment markings under conditions represented in their training sets, and the figures they report come from benchmarks in which markings are largely intact. The dataset used here is weighted the other way, holding 160 heavily faded and 170 partially occluded frames, and it is on exactly those images that a learned appearance model has least support.

The pipeline described in this paper does not model appearance at all. It tests a geometric relation that a worn stop line still satisfies wherever any paint survives, which is why the margin holds on degraded scenes. No claim is made that it would outperform a well-trained network on clean, densely annotated imagery.

Across the 1050 images, the two stages that carry the most weight are the adaptive threshold and the confidence test. The first keeps low-contrast markings above the polarity cut-off in frames where illumination is uneven, and the second discards the linear structures that a lower threshold inevitably admits. Their contributions are complementary, which is visible in the ablation results: each costs between 5 and 10 points on its own, and the pipeline retains most of its accuracy when either is removed but not both.

Interpretability follows from the same design [30, 34]. Every intermediate result is an image or a number that can be examined, so a rejected candidate can be traced to the specific test that rejected it. A supervised detector offers no equivalent, which matters when the output feeds a maintenance budget and has to be defended.

The computational profile suits roadside and in-vehicle deployment [27]. No GPU is required, memory stays around 420 MB, and no labelled corpus has to be collected before the system runs at a new site. This aligns with road traffic safety management practice [32] and with emerging expectations for traceable and ethically governed automation in transport [26, 31]. Limits remain. Trinarization still depends on the scaling parameter k, and performance falls when occlusion or fading removes most of the line from the frame. Video sequences would supply the temporal redundancy that a single frame lacks, and this is the direction taken in continuing work.

5.5. Statistical Validation

Repeated runs across different dataset partitions were used to quantify the variability of the reported figures. Table 12 gives the mean, standard deviation and 95% confidence interval for each metric.

Table 12. Statistical performance analysis

Metric	Mean	Standard Deviation	95% Confidence Interval
Accuracy	92.3	1.8	±1.2
Precision	91.1	1.5	±1.0
Recall	93.5	1.6	±1.1
F1-score	92.3	1.4	±0.9

Paired t-tests and Wilcoxon signed-rank tests were then run against the Sobel-plus-Hough baseline on the same images. Both returned $p < 0.05$ for accuracy, precision and recall, so the margin over the baseline is not attributable to partition noise.

6. Conclusion

This paper has presented an analytical pipeline for detecting stop lines and grading their condition on degraded urban roads. The pipeline chains inverse perspective mapping, fused Sobel and Canny edge extraction, adaptive structural thresholding, three-level edge encoding, edge-pair verification, morphological closing and a probabilistic Hough transform, and requires no training data at any stage. Its distinguishing feature is the retention of gradient polarity through thresholding, which allows the two boundaries of a painted band to be matched as a geometrically constrained pair. Earlier edge-based methods discard the sign and can only test for continuity, which is precisely the property a worn marking no longer has. On 1050 urban images, the pipeline reached 92.3% accuracy, 91.1% precision and 93.5% recall, with localisation within ±5 pixels and a damage estimation error of 6.8%.

The ablation study showed edge-pair verification to be the single most important stage, worth 30.9 percentage points of accuracy on its own. Compared with the deep learning methods surveyed in Section 2, the pipeline reaches comparable accuracy at a fraction of the computational and data cost, which places it within reach of embedded traffic monitoring hardware. Two limitations remain. Heavy occlusion and near-total fading both reduce the evidence below what the pairing and clustering stages require, and the scaling parameter k has to be set once per camera installation. Continuing work is directed at estimating k automatically from scene statistics, at accumulating evidence across video frames rather than treating each frame independently, and at extending the same polarity-based analysis to arrows, crosswalks and lane markings. The pipeline offers an accurate, traceable and computationally light route to automated stop-line inspection, and is applicable to intelligent

transportation systems, municipal asset management and road safety auditing.

Acknowledgements

None.

Author's Contribution

DA: Drafting the manuscript; AN: Supervision, reviewing and editing. All authors read and approved the final version of the manuscript.

Funding

None.

Conflicts of Interest

Dipak Agrawal declares that he has no conflict of interest. Amit Nayak declares that he has no conflict of interest.

Ethical Approval and Consent to Participate

Not applicable.

Consent for Publication

Not applicable.

Data Availsability

Data sharing will be made available on request.

References

[1] Tom B.J. Coenen, and Amir Golroo, "A Review on Automated Pavement Distress Detection Methods," *Cogent Engineering*, vol. 4, no. 1, pp. 1-23, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[2] Yang Xing et al., "Advances in Vision-based Lane Detection: Algorithms, Integration, Assessment, and Perspectives on ACP-based Parallel Vision," *IEEE/CAA Journal of Automatica Sinica*, vol. 5, no. 3, pp. 645-661, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[3] Noor Jannah Zakaria et al., "Lane Detection in Autonomous Vehicles: A Systematic Review," *IEEE Access*, vol. 11, pp. 3729-3765, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[4] John Canny, "A Computational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679-698, 1986. [CrossRef] [Google Scholar] [Publisher Link]

[5] Milan Sonka, Vaclav Hlavac, and Roger Boyle, *Image Processing, Analysis, and Machine Vision with MindTap*, 4th ed., Cengage, 2017. [Google Scholar] [Publisher Link]

[6] Mehrez Marzougui et al., "A Lane Tracking Method based on Progressive Probabilistic Hough Transform," *IEEE Access*, vol. 8, pp. 84893-84905, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[7] Any Gupta, and Ayesha Choudhary, "A Framework for Camera-based Real-Time Lane and Road Surface Marking Detection and Recognition," *IEEE Transactions on Intelligent Vehicles*, vol. 3, no. 4, pp. 476-485, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[8] Toan Minh Hoang et al., "Road Lane Detection by Discriminating Dashed and Solid Road Lanes using a Visible Light Camera Sensor," *Sensors*, vol. 16, no. 8, pp. 1-23, 2016. [CrossRef] [Google Scholar] [Publisher Link]

[9] Rafael C. Gonzalez, and Richard E. Woods, *Digital Image Processing*, 4th ed., Pearson Education, 2019. [Google Scholar] [Publisher Link]

[10] Kaiming He et al., "Deep Residual Learning for Image Recognition," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, pp. 770-778, 2016. [CrossRef] [Google Scholar] [Publisher Link]

[11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Munich, Germany, vol. 9351, pp. 234-241, 2015. [CrossRef] [Google Scholar] [Publisher Link]

[12] Lizhe Liu et al., "CondLaneNet: A Top-to-Down Lane Detection Framework based on Conditional Convolution," *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Montreal, QC, Canada, pp. 3773-3782, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[13] Jinsheng Wang et al., "A Keypoint-based Global Association Network for Lane Detection," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA, pp. 1392-1401, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[14] Rudra P.K. Poudel, Stephan Liwicki, and Roberto Cipolla, "Fast-SCNN: Fast Semantic Segmentation Network," *arXiv*, pp. 1-9, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[15] Yuenan Hou et al., "Learning Lightweight Lane Detection CNNs by Self Attention Distillation," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Seoul, South Korea, pp. 1013-1021, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[16] Gian Antariksa et al., "Comparative Analysis of Advanced AI-based Object Detection Models for Pavement Marking Quality Assessment During Daytime," *2025 IEEE Conference on Artificial Intelligence (CAI)*, Santa Clara, CA, USA, pp. 1023-1028, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[17] Yuchuan Du et al., "Pavement Distress Detection and Classification based on YOLO Network," *International Journal of Pavement Engineering*, vol. 22, no. 13, pp. 1659-1672, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[18] Hiroya Maeda et al., "Road Damage Detection and Classification using Deep Neural Networks with Smartphone Images," *Computer-Aided Civil and Infrastructure Engineering*, vol. 33, no. 12, pp. 1127-1141, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[19] Deeksha Arya et al., "RDD2022: A Multi-National Image Dataset for Automatic Road Damage Detection," *Geoscience Data Journal*, vol. 11, no. 4, pp. 846-862, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[20] Nobuyuki Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62-66, 1979. [CrossRef] [Google Scholar] [Publisher Link]

[21] Joanna Kulawik, Mariusz Kubanek, and Sebastian Garus, "The Verification of the Correct Visibility of Horizontal Road Signs using Deep Learning and Computer Vision," *Applied Sciences*, vol. 13, no. 20, pp. 1-13, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[22] Young-Jin Cha, Wooram Choi, and Oral Büyüköztürk, "Deep Learning-based Crack Damage Detection using Convolutional Neural Networks," *Computer-Aided Civil and Infrastructure Engineering*, vol. 32, no. 5, pp. 361-378, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[23] Suresh Chavhan et al., "Edge Computing AI-IoT Integrated Energy-Efficient Intelligent Transportation System for Smart Cities," *ACM Transactions on Internet Technology*, vol. 22, no. 4, pp. 1-18, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[24] Wooram Choi, and Young-Jin Cha, "SDDNet: Real-Time Crack Segmentation," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 9, pp. 8016-8025, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[25] Markus Eisenbach et al., "How to Get Pavement Distress Detection Ready for Deep Learning? A Systematic Approach," *2017 International Joint Conference on Neural Networks (IJCNN)*, Anchorage, AK, USA, pp. 2039-2047, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[26] European Parliament and Council of the European Union, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA Relevance), Official Journal of the European Union, Brussels, 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

[27] "ETSI EN 302 665: Intelligent Transport Systems (ITS); Communications Architecture, European Telecommunications Standards Institute, Sophia Antipolis, France, 2010. [Google Scholar] [Publisher Link]

[28] Kasthurirangan Gopalakrishnan, "Deep Learning in Data-Driven Pavement Image Analysis and Automated Distress Detection: A Review," *Data*, vol. 3, no. 3, pp. 1-19, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[29] Jinchao Guan et al., "Automated Pixel-Level Pavement Distress Detection based on Stereo Vision and Deep Learning," *Automation in Construction*, vol. 129, pp. 1-18, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[30] David Gunning, and David W. Aha, "DARPA's Explainable Artificial Intelligence (XAI) Program," *AI Magazine*, vol. 40, no. 2, pp. 44-58, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[31] Joanna Isabelle Olszewska, "IEEE Standard Model Process for Addressing Ethical Concerns During System Design: IEEE Standard 7000-2021," *IEEE*, 2021. [Google Scholar] [Publisher Link]

[32] ISO 39001:2012, Road Traffic Safety (RTS) Management Systems - Requirements with Guidance for Use, International Organization for Standardization, *ISO*, 2012. [Google Scholar] [Publisher Link]

[33] Dongho Kang, and Young-Jin Cha, "Autonomous UAVs for Structural Health Monitoring using Deep Learning and an Ultrasonic Beacon System with Geo-Tagging," *Computer-Aided Civil and Infrastructure Engineering*, vol. 33, no. 10, pp. 885-902, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[34] Anton Kuznietsov et al., "Explainable AI for Safe and Trustworthy Autonomous Driving: A Systematic Review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 12, pp. 19342-19364, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[35] Vishal Mandal, Lan Uong, and Yaw Adu-Gyamfi, "Automated Road Crack Detection using Deep Convolutional Neural Networks," *2018 IEEE International Conference on Big Data (Big Data)*, pp. 5212-5215, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[36] Qipei Mei, Mustafa Gül, and Riasat Azim, "Densely Connected Deep Neural Network Considering Connectivity of Pixels for Automatic Crack Detection," *Automation in Construction*, vol. 110, pp. 1-13, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[37] Mark Nixon, and Alberto S. Aguado, *Feature Extraction and Image Processing for Computer Vision*, 3rd ed., Academic Press, 2012. [Google Scholar]

[38] Shaoqing Ren et al., "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137-1149, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[39] Nima Sholevar, Amir Golroo, and Sahand Roghani Esfahani, "Machine Learning Techniques for Pavement Condition Evaluation," *Automation in Construction*, vol. 136, pp. 1-17, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[40] Qin Zou et al., "CrackTree: Automatic Crack Detection from Pavement Images," *Pattern Recognition Letters*, vol. 33, no. 3, pp. 227-238, 2012. [CrossRef] [Google Scholar] [Publisher Link]