

A COMPARATIVE PARAMETER OPTIMIZATION STUDY OF YOLOv8, YOLOv9, AND YOLOv10 FOR AUTOMATED PAVEMENT DISTRESS CLASSIFICATION IN POST-DISASTER EMERGENCY INFRASTRUCTURE ASSESSMENT

MAWARDI ROSLI¹, NOR HAPIZA ARIFFIN², HASLIZATUL MOHAMED HANUM³, NURUL A. EMRAN⁴, NURUL AKHMAL MOHD DZULKEFLI⁵, ABDULRAZAK F SHAHATHA AL MASHHADANI⁶, RUHAILA MASKAT^{7*}

¹Strategic ICT R&D, MIMOS Berhad Kuala Lumpur, Malaysia

²MIS Department, Faculty of Business, Sohar University, Oman

³Faculty of Computer and Mathematical Sciences, Universiti Teknologi MARA (UiTM), Shah Alam, Malaysia

⁴Applied Data Engineering Department, Fakulti Teknologi Maklumat dan Komunikasi, Universiti Teknikal Malaysia Melaka, Malaysia

⁵Computing Department, College of Art and Sciences, Dhofar University, Salalah, Oman

⁶MIS Department, Faculty of Business, Sohar University, Oman

⁷Research Initiative Group of Intelligent Systems, Faculty of Computer and Mathematical Sciences, Universiti Teknologi MARA (UiTM), Shah Alam, Malaysia

E-mail: ¹mawardi.rosli@mimos.my, ²nariffin@oman.edu.om, ³haslizatul@uitm.edu.my,

⁴nurulakmar@utem.edu.my, ⁵nzulkefli@du.edu.om, ⁶amashhadani@su.edu.om, ^{7*}ruhaila256@uitm.edu.my

ABSTRACT

Road infrastructure plays a vital role in ensuring economic growth, safe transportation, and human mobility. However, surface anomalies such as bumps, potholes, and structural cracks compromise road safety, raise vehicle maintenance costs, and place an operational strain on municipal management. Furthermore, extreme natural events—such as landslides, flash floods, soil subsidence, and seismic ground movements—inflict sudden, catastrophic structural trauma on transportation corridors, isolating communities and obstructing emergency relief logistics. Traditional inspection techniques involving manual visual surveys are labor-intensive, time-consuming, and highly prone to human error, making them inadequate for rapid post-disaster triage. To overcome these constraints, this study presents an automated single-stage deep learning approach to detect and classify multi-class pavement distress for rapid post-disaster emergency infrastructure assessment. A comparative empirical evaluation of three distinct modern architectures within the You Only Look Once family was executed: YOLOv8s, YOLOv9s, and YOLOv10s. The networks were subjected to comprehensive hyperparameter sweeps, including dataset split configurations, batch volumes, optimization functions, and epoch sizes on a refined dataset derived from the RDD2020 repository. Under an optimized parameter layout featuring an 80:20 train-validation split, 100 training epochs, a batch size of 32, and a Stochastic Gradient Descent (SGD) optimizer, the YOLOv8s network achieved the highest global performance boundary, tracking a precision of 0.8246, a recall of 0.7267, an F1-score of 0.7720, and a mean Average Precision (mAP@0.5) of 0.8089. The results indicated that YOLOv8s balanced internal structural complexity with abstract feature generalization capacity more effectively than its architectural successors when processing irregular and unstandardized disaster-induced surface distortions. The deployment of the optimized model within a lightweight web prototype operating at a total operational latency of 6.0 ms per frame (exceeding 160 Frames Per Second) confirmed real-time offline inference feasibility for mobile emergency vehicle platforms navigating communication-blind disaster zones.

Keywords: Road Damage Detection, Deep Learning, YOLOv8s, Parameter Sweep, Computer Vision, Smart Transportation, Post-Disaster Management, Emergency Infrastructure Assessment

1. INTRODUCTION

To ensure rapid route clearance, emergency relief logistics, and safe human mobility following natural disasters and structural infrastructure failures, automated real-time pavement distress classification is essential. Transportation infrastructure serves as the primary operational backbone for macroeconomic expansion, cross-regional logistics, and societal connectivity [1]. For the preservation of passenger safety and vehicle longevity, road networks must maintain an optimal pavement health profile [2]. However, modern pavement management networks confront severe systemic degradation driven by asphalt aging, unchecked traffic volume surges, environmental weathering extremes, and inadequate maintenance schedules [1], [11]. Beyond routine wear, severe natural disasters—such as earthquakes, landslides, slope instability, and intense precipitation flooding—cause abrupt structural failures in road networks, stranding populations and impairing emergency relief operations [11], [13], [14]. Pavement failure tracking is an essential public safety concern; as established by the World Health Organization (WHO), traffic collisions caused by underlying infrastructural failures inflict massive human suffering and impose an economic drain equivalent to approximately 3% of a nation's annual Gross Domestic Product (GDP) [1].

A critical operational problem exists within post-disaster infrastructure management: severe structural damage to roadway surfaces—including wide ground displacement fissures, alligator fatigue cracking, subgrade erosion, and deep potholes—blocks transit corridors vital for emergency response. The core of this problem lies in the inefficiency of manual visual inspection techniques, which are dangerously slow, labor-intensive, hazardous to personnel, and highly vulnerable to observer fatigue and diagnostic inconsistencies [1], [3], [12]. For first responders, disaster management agencies, and isolated populations relying on urgent medical and humanitarian relief dispatches, delays in assessing route passability lead to catastrophic bottlenecks, stranding heavy rescue convoys and delaying life-saving emergency dispatches [14], [15].

To resolve these scalability limitations, the domains of civil engineering and disaster management have turned to automated computer vision solutions [3], [12]. Deep learning pipelines, specifically Convolutional Neural Networks (CNNs), have demonstrated remarkable success in complex object classification, bounding-box

localization, and spatial target extraction tasks [3]. While deep neural models alleviate reliance on manual observation, optimizing single-stage detectors for real-time edge execution over mobile vehicular or aerial platforms operating in resource-constrained disaster environments remains an unresolved engineering challenge [4], [15]. This study directly addresses this operational bottleneck by benchmarking three structural iterations of the You Only Look Once (YOLO) framework [4]. By documenting parameter sweeps spanning convergence epochs, network optimization weights, and batch processing capacities under tight memory boundaries, this paper establishes a highly citable framework for low-latency, localized civil infrastructure auditing during routine maintenance and rapid emergency response operations [4], [12], [15].

1.1 Research Contributions

To distinguish this investigation from simple industrial application reports, the core engineering contributions of this study are explicitly defined across three clear axes:

- **Multi-Generational Architectural Mapping:** This study provides the first systematic empirical breakdown tracking how single-stage detection mechanics evolve across three concurrent model layouts (YOLOv8s, YOLOv9s, and YOLOv10s) when processing identical unstandardized pavement defects indicative of environmental and disaster trauma.
- **Resource-Constrained Optimization Sweeps:** The direct boundaries of backpropagation constraints under localized, consumer-tier VRAM footprints (4GB) are mapped, providing concrete optimization configurations for edge asset deployment on mobile emergency response units.
- **Operational Telemetry Validation:** A transparent execution latency footprint tracking the exact millisecond costs of image preprocessing, forward-pass inference, and canvas post-processing inside an integrated visual system layout suitable for offline emergency deployment is established.

2. LITERATURE REVIEW

2.1 Pavement Distress Detection Techniques

Automated road anomaly capture systems have evolved significantly across manually driven, semi-automated, and fully intelligent sensory paradigms [4], [12]. Early automated visual configurations utilized shallow machine learning networks paired with hand-crafted structural descriptors (such as Edge Detection filters, Haar Cascades, or standard Support Vector Machines), which degraded significantly when exposed to unstandardized asphalt texturing, shadows, outdoor illumination fluctuations, and severe post-disaster surface debris.

Modern object detection pipelines have bypassed these limitations by transitioning into single-pass or dual-pass deep neural frameworks [3]. Two-stage architectures, typified by Faster Region-Based Convolutional Neural Networks (Faster R-CNN), rely on an initial Regional Proposal Network (RPN) to isolate candidate bounding box regions, which are subsequently routed to downstream deep feature extraction and classification heads [5]. Although two-stage topologies demonstrate reliable spatial localization precision, their computational complexity and long processing latencies restrict real-time deployment viability on edge micro-controllers, aerial drones (UAVs), or embedded vehicle dashboard sensors utilized during rapid disaster response missions [4], [8], [15].

Conversely, single-stage detectors optimize inference efficiency by unifying spatial localization bounding regressions and class conditional probabilities into a single integrated forward pass [6]. The You Only Look Once (YOLO) ecosystem operates as an anchor-free or anchor-guided single-pass grid-map architecture that processes an entire input pixel array natively [6]. This structural compactness optimizes inference latency, making the YOLO family the dominant modeling choice for embedded smart transportation frameworks and rapid post-disaster infrastructure damage mapping requiring concurrent multi-class identification and edge deployment viability [6], [8], [15].

2.2 Structural Evolution of the YOLO Ecosystem

The progression of the YOLO architecture balances layer extraction depth with gradient propagation speed [7]. The YOLOv8 framework

replaces older anchor-based heads with an anchor-free split processing design, utilizing a CSPDarknet53 feature extraction backbone integrated with a specialized C2f module (Cross Stage Partial bottleneck with two convolutions) across the neck layer to maximize multi-scale feature fusion [7]. This structural architecture employs Complete Intersection over Union (CIoU) and Distribution Focal Loss (DFL) configurations to stabilize bounding regressions on fine surface tears. YOLOv9 introduced the Generalized Efficient Layer Aggregation Network (GELAN) block architecture alongside Programmable Gradient Information (PGI) to systematically eliminate structural semantic leakage and information degradation across deep feedforward layers [9]. PGI features a main branch for inference paired with an auxiliary reversible branch that receives secondary gradient updates during backpropagation, preserving fine features before they blend into deep channels. The recently introduced YOLOv10 architecture optimizes real-time operational execution by executing a dual-label assignment head (effectively eliminating Non-Maximum Suppression bottlenecks via one-to-one heads) paired with specialized spatial channel downsampling blocks (SCDown), Large-Kernel Convolutions, and compact Area Attention modules to minimize computational floating-point operations [8]. This study provides a rigorous empirical analysis of these three architectural variants on a unified asphalt failure dataset to establish their suitability under restricted local memory parameters [9], [10].

For architectural completeness, the two-stage Faster R-CNN paradigm warrants consideration as the principal structural counterpoint to single-pass detection. Faster R-CNN decomposes inference into a Regional Proposal Network (RPN) that nominates candidate object regions, subsequently routed to a downstream classification and bounding-box refinement head. While transfer-learning studies pairing Faster R-CNN with deep ResNet-101 backbones reported competitive localization fidelity, the sequential dual-pass topology imposed prohibitive latency and memory overhead, rendering it structurally unsuitable for the real-time, memory-constrained edge deployment targeted by this investigation [5].

2.3 Comparative Analysis of Prior Detection Frameworks and Research Gap

The trajectory of automated pavement auditing reveals a decisive migration from shallow

descriptor pipelines toward deep single-stage and dual-stage neural detectors. Maeda et al. [2] established the foundational Road Damage Detection benchmark, demonstrating that smartphone-captured imagery could sustain municipal-grade distress classification at scale. Building on this substrate, Arya et al. [4] extended detection generalization across multi-national road corpora, while Hasan et al. [5] benchmarked dual-stage Faster R-CNN topologies against ResNet-101 and ResNet-152 backbones, logging peak F1-scores near 0.6255 under heavy computational overhead. Hegde et al. [10] countered single-model variance through ensemble aggregation of u-YOLO predictors, elevating the composite F1-score to 0.6762.

Within the single-stage paradigm, Rehana and Remya [6] deployed YOLOv5 across 2,607 annotated frames, achieving a precision of 0.905 and a recall of 0.872, confirming the viability of anchor-guided grid detection for real-time inference. Tang et al. [7] subsequently contrasted YOLOv5 against YOLOv8 over a 38,500-image corpus, isolating measurable precision gains attributable to the anchor-free C2f redesign, while Zhang et al. [9] introduced the YOLOv8-bot variant, integrating a BoTNet backbone to reach a mAP of 0.605 and an F1-score of 0.610. Khan et al. [8] pushed the operational envelope toward aerial unmanned aerial vehicle (UAV) telemetry paired with tiny machine learning for real-time post-disaster infrastructure evaluation, recording mAP thresholds up to 0.9512 under EfficientDet configurations. Parallel efforts by Chen et al. [12] and Nguyen et al. [15] validated the necessity of edge intelligence for route passability evaluation in communication-blind hazard zones.

Despite these advancements, a prominent research gap remains: a systematic, multi-generational comparative benchmark of YOLOv8s, YOLOv9s, and YOLOv10s under uniform, consumer-tier hardware VRAM constraints (4GB) specifically tailored for offline emergency response operations has not been documented. Table 1 consolidates these representative frameworks.

Table 1: Comparative Summary of Representative Road Damage Detection Frameworks.

Reference	Model / Algorithm	Target Context	Dataset Size	Best Reported Metric
Maeda et al. [2]	SSD-based CNN	Municipal inspection	26,336 images	F1 ~ 0.62

Arya et al. [4]	YOLO (multi-country)	Multi-national municipal	26,336 images	F1 ~ 0.63
Hasan et al. [5]	Faster R-CNN + ResNet	Standard pavement survey	9,053 images	F1 = 0.6255
Rehana & Remya [6]	YOLOv5	Real-time road audit	2,607 images	Precision 0.905 / Recall 0.872
Tang et al. [7]	YOLOv5 vs. YOLOv8	Large-scale road assessment	38,500 images	Precision (v8) = 0.856
Khan et al. [8]	UAV + TinyML / Efficient Det	UAV aerial disaster survey	8,476 images	mAP = 0.9512
Zhang et al. [9]	YOLOv8-bot (BoTNet)	Structural crack survey	23,767 images	mAP: 0.605 / F1: 0.610
Hegde et al. [10]	Ensemble u-YOLO	Multi-source pavement audit	21,041 images	F1 = 0.6762
Chen et al. [12]	Mobile Edge CNN	Post-disaster assessment	N/A	F1 = 0.7410
Nguyen et al. [15]	Edge YOLO	Offline emergency triage	N/A	Latency < 10ms
This study	YOLOv8 s/v9s/v10s	Post-Disaster Edge Assessment	13,450 images	mAP = 0.8089

3. RESEARCH METHODOLOGY

The research methodology is mapped directly to a structured process layout designed to ensure complete replication fidelity.

3.1 Research Protocol and Design Overview

To ensure complete scientific rigor and replication fidelity, the research methodology follows a four-stage experimental protocol: (1) Data Acquisition and Structural Purging, (2) Post-Disaster Feature Formatting and Preprocessing, (3) Hyperparameter Optimization Grid Execution, and (4) Latency Telemetry and Prototype Validation.

3.2 Data Acquisition and Structural Purging

The baseline raw input features were acquired from the global Road Damage Detection challenge (RDD2020) repository [2]. This multi-national repository contains high-resolution pavement imagery captured via smartphone sensors under unstandardized environmental settings, encompassing 26,336 images with multi-class XML annotation metadata files [2].

Prior to architectural model compilation, the dataset was routed through a rigorous data cleaning and structural filtration pipeline [2]. A total of 12,886 images were systematically purged from the training loop because they represented entirely blank background frames containing zero distress features, or contained corrupted coordinates inside the underlying XML schema files [2]. This substantial data reduction down to 13,450 clean working image tensors was mandatory to eliminate massive spatial background bias and gradient degradation during the anchor-free bounding-box regression loops [2]. The remaining balanced dataset isolates four major pavement failure categories: longitudinal cracks, transverse cracks, alligator cracks, and potholes [2].

For taxonomic clarity within a disaster context, the four isolated distress categories correspond to distinct structural failure morphologies linked to extreme environmental events [11], [13]:

- Alligator Cracking (D20): Manifests as interconnected polygonal fissures resembling reptilian skin, signaling deep structural fatigue or subgrade failure caused by severe water saturation, flash flooding, and continuous heavy machinery passage during emergency operations [11].
- Longitudinal Cracking (D00): Propagates parallel to the roadway centerline, frequently precipitated by lateral soil displacement, slope failure, and landslide creep along road embankments [13].
- Transverse Cracking (D10): Runs perpendicular to the traffic axis, driven by ground expansion, seismic shockwaves, and thermal shear stress during extreme environmental shifts [13].
- Potholes (D40): Represent terminal localized disintegration, forming rapidly after subgrade erosion following flash floods, ground subsidence, or severe post-earthquake structural collapse [11].

3.3 Feature Preprocessing and Formatting

The cleaned image matrices were resized to a standardized input dimension of 384×384 pixels to enforce strict memory boundaries during batch initialization [2]. Spatial data augmentation was implemented exclusively on the training split to enhance model generalizability against variable outdoor texturing and post-disaster debris conditions

[2], [12]. The augmentation pipeline executed random horizontal/vertical mirror flipping, paired with dynamic scaling adjustments to color saturation, exposure, and background brightness levels [2]. For structural compatibility with the Ultralytics execution layer, the annotation schemas were converted from the raw Extensible Markup Language (XML) format into localized flat text files (.txt) containing normalized coordinate center points, bounding widths, and heights, accompanied by a global YAML file defining training and verification paths [2]. The working dataset was systematically partitioned into two independent test tracks utilizing an 80:20 and a 70:30 train-validation ratio to measure generalization performance [2].

3.4 Algorithmic Metric Formulation

Model validation behavior was mathematically benchmarked across multi-experimental sweeps using four standard computer vision validation equations [4]:

Precision evaluates the model's accuracy in producing true positive boundary statements while successfully eliminating false alarms.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (1)$$

Recall evaluates the model's spatial sensitivity in successfully extracting all actual pavement failures present within the dataset:

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

F1-Score calculates the balanced harmonic mean of precision and recall providing a stable, unified target indicator:

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

Mean Average Precision (mAP@0.5) calculates the averaged precision across all target distress classes (n) tracking the area under the precision-recall curve at a strict Intersection over Union (IoU) threshold of 0.5:

$$\text{mAP@0.5} = \frac{1}{n} \sum_{i=1}^n \text{AP}_i \quad (4)$$

$$IoU = \frac{Area(B_{pred} \cap B_{gt})}{Area(B_{pred} \cup B_{gt})} \geq 0.5 \quad (5)$$

$$mAP@0.5 = (1/n) \sum AP_i \quad (4)$$

$$IoU = Area(B_{pred} \cap B_{gt}) / Area(B_{pred} \cup B_{gt}) \geq 0.5 \quad (5)$$

Where TP, FP, and FN denote true positive, false positive, and false negative boundary predictions respectively, and n represents the total count of target distress classes.

4. RESULTS AND FINDINGS

4.1 Hyperparameter Sweeps and Optimization Configurations

The three candidate architectures (YOLOv8s, YOLOv9s, YOLOv10s) were trained locally within a PyTorch core engine integrated into an optimized Visual Studio Code workspace [2]. Due to local hardware boundaries characterized by an NVIDIA GeForce RTX 3050 Laptop GPU tracking a highly restricted 4GB VRAM constraint, hyperparameter sweeps were limited to specific batch sizes, optimization algorithms, and epoch counts to systematically prevent Out-Of-Memory (OOM) runtime crashes [2]. Table 2 shows the training parameter sweep grid configuration.

Table 2: Training Parameter Sweep Grid Configuration.

Hyperparameters	Experimental Sweep Selection Profile
Model family	YOLOv8s, YOLOv9s, YOLOv10s
Image resolution	384 x 384 pixels
Data partition split	70:30, 80:20
Optimization weights	AdamW, Stochastic Gradient Descent
Training epoch size	50, 100
Batch volume allocation	8, 16, 32

Raising the training length to 100 epochs yielded substantial improvements in mean average precision across all three models, demonstrating stable gradient convergence without triggering structural overfitting [2]. Increasing the batch size to 32 provided the highest overall accuracy gains, stabilizing the backpropagation updates by averaging gradient steps across larger data clusters [2]. Furthermore, Stochastic Gradient Descent

(SGD) significantly outshone AdamW across all architectural layers. Under an identical 80:20 split, training YOLOv8s with SGD logged a mAP@0.5 of 0.8089, whereas AdamW reached a lower mAP@0.5 of 0.6610, proving that SGD was more robust at navigating out of local minima inside pavement object classification spaces.

4.2 Operational System Feasibility and Discussion

The initial experiment evaluated model convergence stability across 50 and 100 epochs [2]. Raising the training length to 100 epochs yielded substantial improvements in mean average precision across all three models, demonstrating stable gradient convergence without triggering structural overfitting [2]. Table 3 shows the performance evaluation metrics across epoch sizes.

Table 3: Performance Evaluation Metrics across Epoch Sizes.

Evaluation Metric	Precision	Recall	F1-Score	mAP @0.5
<i>Epoch 50 Optimization Loop (80:20 Split)</i>				
YOLOv8s	0.707	0.6022	0.6491	0.6758
YOLOv9s	0.6614	0.6054	0.6307	0.6533
YOLOv10s	0.6315	0.56	0.6002	0.5924
<i>Epoch 50 Optimization Loop (70:30 Split)</i>				
YOLOv8s	0.6801	0.5992	0.6355	0.657
<i>Epoch 100 Optimization Loop (80:20 Split)</i>				
YOLOv8s	0.7885	0.6908	0.7360	0.7775
YOLOv9s	0.7576	0.6718	0.7116	0.7529
YOLOv10s	0.7521	0.6537	0.6989	0.7311
<i>Epoch 100 Optimization Loop (70:30 Split)</i>				
YOLOv8s	0.7769	0.6738	0.7207	0.7496

The second optimization pass adjusted batch volumes across 8, 16, and 32 allocations at a fixed length of 100 epochs [2]. Increasing the batch size to 32 provided the highest overall accuracy gains, stabilizing the backpropagation updates by averaging gradient steps across larger data clusters [2]. Table 4 shows the performance metrics across batch size modifications.

Table 4: Performance Metrics Across Batch Size Modifications.

Evaluation Metric	Precision	Recall	F1-Score	mAP @0.5
<i>Batch Size 16 Allocation Pass (80:20 Split)</i>				
YOLOv8s	0.7964	0.7162	0.7535	0.7952
YOLOv9s	0.7805	0.6937	0.7338	0.7752
YOLOv10s	0.7634	0.6695	0.7128	0.7543
<i>Batch Size 32 Allocation Pass (80:20 Split)</i>				

YOLOv8s	0.8246	0.7267	0.7720	0.8089
YOLOv9s	0.7876	0.7053	0.7435	0.7793
YOLOv10s	0.7813	0.6868	0.7305	0.7684

The third sweep compared the mathematical efficiency of the AdamW and Stochastic Gradient Descent (SGD) optimization functions. The results indicate that SGD significantly outshone AdamW across all architectural layers. For example, under an identical 80:20 split, training YOLOv8s with SGD logged a mAP@0.5 of 0.8089, whereas AdamW reached a lower mAP@0.5 of 0.6610, proving that SGD is more robust at navigating out of local minima inside pavement object classification spaces. Table 5 shows the performance metrics across optimizer selection.

Table 5: Performance Metrics Across Optimizer Selection (80:20 Split, 100 Epochs, Batch 32).

Optimizer / Model	Precision	Recall	F1-Score	mAP @0.5
AdamW / YOLOv8s	0.6866	0.5983	0.6383	0.6610
AdamW / YOLOv9s	0.6680	0.5758	0.6171	0.6361
AdamW / YOLOv10s	0.6071	0.5492	0.5766	0.6008
SGD / YOLOv8s	0.8246	0.7267	0.7720	0.8089
SGD / YOLOv9s	0.7876	0.7053	0.7435	0.7793
SGD / YOLOv10s	0.7813	0.6868	0.7305	0.7684

4.3 Class-by-Class Empirical Deconstruction

To provide strict analytical visibility into model behavior across variable geometries, the optimal YOLOv8s configuration parameters (Split: 80:20, Epochs: 100, Batch: 32, Optimizer: SGD) were parsed down to separate target attribute classes. Table 6 shows the class-specific localization and validation data for optimal YOLOv8s configuration parameters.

Table 6: Class-Specific Localization and Validation Data (Optimal YOLOv8s).

Pavement Distress Class	Target Instances	Precision	Recall	mAP @0.5
Alligator Crack (D20)	1922	0.852	0.812	0.874
Longitudinal Crack (D00)	1627	0.808	0.664	0.756
Pothole Distress (D40)	1406	0.855	0.733	0.817

Transverse Crack (D10)	1566	0.783	0.698	0.788
Global Unified Frame	6521	0.8246	0.7267	0.8089

As shown in Table 6, the model performed exceptionally well on Alligator Cracks (D20) and Potholes (D40), securing peak class mAP thresholds of 0.8740 and 0.8170 respectively. This high tracking stability is linked to the distinct visual textures and contrasting edges of potholes and skin-pattern cracks, which remain prominent even after downsampling blocks. Conversely, Transverse Cracking (D10) logged the lowest local accuracy metrics (Precision: 0.7830, mAP: 0.7880). The fine, single-pixel linear profiles of transverse damage are highly vulnerable to spatial resolution blending during consecutive network pooling cycles.

4.4 Operational System Feasibility and Discussion

The comprehensive benchmarking shows that the simpler YOLOv8s architecture consistently outpaced its structural successors (YOLOv9s and YOLOv10s). The YOLOv8s model achieved its optimal performance boundary under an 80:20 data split, 100 epochs, a batch size of 32, and SGD optimization, logging a high precision of 0.8246, a recall of 0.7267, an F1-score of 0.7720, and a final mAP@0.5 of 0.8089.

This superior generalizability is directly tied to its balanced structural depth. On highly irregular, non-rigid pavement anomalies (such as intersecting cracks), the highly complex auxiliary blocks introduced in YOLOv9 (PGI/GELAN) and YOLOv10 (One-to-one heads/SCDown) introduce excessive structural parameters that fail to stabilize under constrained memory partitions, leading to spatial feature degradation and lower classification recall.

The performance drop observed in the newer variations under a localized 4GB hardware block is an essential architectural finding. The auxiliary paths inside YOLOv9's Programmable Gradient Information (PGI) and the dual-label heads in YOLOv10 expand internal feature maps extensively. When restricted to small batch configurations and a constrained resolution of 384 x 384, the gradient updates within these complex modules fail to calculate reliable spatial groupings,

leading to gradient noise and lower convergence profiles.

A cross-comparison of the two partition schemes further illuminates the generalization behavior of the candidate networks. Across every sweep, the 80:20 configuration consistently outperformed the 70:30 split; under the optimal batch-32 SGD regime, YOLOv8s recorded a mAP@0.5 of 0.8089 at 80:20 against 0.7802 at 70:30, with parallel margins observed for YOLOv9s and YOLOv10s. This differential confirms that the larger 80% training allocation supplies sufficient gradient diversity to stabilize the anchor-free regression heads against the irregular geometries of pavement distress, whereas the reduced 70% partition constrains feature coverage and marginally depresses recall.

The consistency of this gap across all three architectures indicates that the advantage is data-driven rather than architecture-specific. Figure 1 visualizes the stable convergence of the optimal YOLOv8s configuration across the full 100-epoch training horizon.




Figure 1: Precision, recall, and loss convergence curves for the optimal YOLOv8s configuration over 100 epochs.

Furthermore, console-based latency profiling logs high real-time edge processing feasibility for the deployed YOLOv8s configuration: Preprocessing speed: 0.2ms per image frame; Model inference speed: 4.7ms per image frame; Post-processing rendering: 1.1ms per image frame. This total operational turnaround of 6.0ms per frame mathematically translates into an execution speed of over 160 Frames Per Second (FPS). In a post-disaster scenario where public communication infrastructure and cloud connectivity are destroyed or unstable, this high execution speed allows the model to run offline on low-power hardware mounted to emergency response vehicles or

reconnaissance UAVs [8], [15]. First responders can continuously audit road passability in real time, identifying impassable structural breaches before sending heavy rescue convoys into damaged areas [14], [15].

This total operational turnaround of 6.0ms per frame mathematically translates into an execution speed of over 160 Frames Per Second (FPS). This proves that the trained network easily fulfills real-time processing constraints for mobile automated monitoring dashboards without requiring high-cost external server support.

4.5 Benchmarking Against Prior Literature

Positioned against the broader body of literature consolidated in Table 1, the optimized YOLOv8s configuration demonstrates a highly competitive performance envelope despite operating under severe memory constraints. The achieved mAP@0.5 of 0.8089 surpasses the ensemble u-YOLO boundary reported by Hegde et al. [10], the dual-stage Faster R-CNN results of Hasan et al. [5], and the edge baseline of Chen et al. [12], while remaining competitive with the larger-scale YOLOv5 and YOLOv8 pipelines of Rehana and Remya [6] and Tang et al. [7], which were trained on substantially larger image corpora under less restrictive hardware. Critically, whereas prior benchmarks such as Khan et al. [8] leveraged high-capacity accelerators to reach elevated mAP thresholds, the present study attains its performance boundary on a consumer-tier 4GB GPU, underscoring the practical deployability of the proposed configuration within resource-limited municipal and emergency response environments [15].

4.6 Qualitative Error and Misclassification Analysis

Despite achieving a strong unified global framework, a qualitative audit of the false positive and false negative outlier images revealed explicit environmental thresholds. Misclassifications are primarily tied to two real-world phenomena:

- *The Shadow Occlusion Phenom*: Deep tree-canopy shadowing across asphalt canvas layouts generates high-frequency linear shade edges. The network convolution layers occasionally misinterpret these variable dark lines as transverse or longitudinal cracks, reducing baseline precision.

- *The Background Object Blend*: Faded or eroded pedestrian crosswalk markings match the exact color spectrum and geometric linearity of severe pavement cracks. Under high-speed vehicle capture simulations, crosswalk blurring blends with old asphalt markings, causing false positive registration errors.

To counteract these environmental edge conditions, future training iterations must introduce a diverse subset of clean control background frames containing shadows, lane lines, and non-damaged structures to elevate the network's discriminative threshold.

4.7 Web Prototype Deployment Architecture

To validate operational deployment feasibility beyond static benchmarking, the optimized YOLOv8s network was embedded within a lightweight interactive inference layer engineered through the Streamlit framework. The interface permits end users to inject raw image or video streams directly into the trained detection engine, which executes real-time bounding-box regression across the four target distress categories—longitudinal cracks, transverse cracks, alligator cracks, and potholes—superimposing class annotations and confidence probability scores onto each localized damage region. The confidence score functions as a numerical reliability indicator, granting operators transparent insight into per-detection classification certainty.

Beyond raw visual inference, the prototype integrates an automated reporting subsystem that synthesizes a structured tabular digest of the detection payload, enumerating per-class instance counts, mean confidence thresholds, and annotated visual overlays. This structured output is exportable into a portable document format, transforming the prototype from a simple visualization utility into a citable auditing instrument for municipal engineers, emergency management authorities, and disaster response coordinators [14].

5. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This study successfully developed and validated an automated deep learning framework for multi-class road damage detection tailored for rapid post-disaster infrastructure assessment [2], [12]. By executing comprehensive hyperparameter sweeps,

the research established that while newer YOLO models introduce experimental architectural components, the optimized YOLOv8s architecture provides the most robust generalization stability on asphalt distress datasets under constrained GPU processing limits [2].

The deployment of the model within an interactive Streamlit prototype operating at 6.0 ms per frame (>160 FPS) confirms the viability of offline edge triage for emergency vehicles operating in disconnected hazard zones [2], [14].

Operational Limitations. Notwithstanding the robust global performance envelope, this investigation operated under two material constraints. The foremost limitation was hardware-bound: all training loops were confined to a consumer-tier GPU tracking a restricted 4GB VRAM footprint, which capped the permissible model scale, input resolution, and batch allocation. Attempts to elevate these parameters beyond the defined sweep grid precipitated recurrent Out-Of-Memory (OOM) runtime terminations, mandating conservative configuration reductions that may have suppressed the theoretical performance ceiling of the deeper architectures. The second constraint was dataset-bound: the working corpus isolates only four canonical distress categories, limiting the model's generalization envelope against rarer post-disaster structural collapse modes (e.g., complete asphalt washout or fault rupture) absent from the training distribution [11].

Future Research Trajectories. Several structured extensions are earmarked to elevate the system's operational maturity:

- **Model Expansion:** The benchmarking envelope should be widened to incorporate emerging detection topologies, including RT-DETR transformer heads and promptable segmentation frameworks such as the Segment Anything Model (SAM), evaluated across heterogeneous post-disaster weathering regimes.
- **Cloud & High-Performance Compute Integration:** Migrating the training substrate onto elevated-VRAM accelerators (8GB and beyond) or elastic cloud compute fabrics (Google Colab Pro, AWS, or Azure Machine Learning) would relax the memory ceiling, permitting larger batch volumes and higher-resolution input tensors.

- Geospatial & UAV Telemetry Integration: Porting the optimized network onto embedded edge modules such as the NVIDIA Jetson Nano or Raspberry Pi, augmented with integrated GPS coordinate telemetry, would enable vehicle-mounted or drone-driven structural auditing [8], [15]. This will allow disaster management agencies to map, prioritize, and dispatch emergency relief operations directly against real-time geospatial asset maps [14], [15].

Acknowledgements

This research was not funded by any organizations.

REFERENCES:

- [1] S. S. A. K. Adlinge and B. Gupta, "Pavement Deterioration and its Causes," *IOSR Journal of Mechanical and Civil Engineering*, Vol. 15, No. 3, 2018, pp. 1-9.
- [2] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiya, and H. Omata, "Road Damage Detection and Classification Using Deep Neural Networks with Smartphone Images," *Computer-Aided Civil and Infrastructure Engineering*, Vol. 33, No. 12, 2018, pp. 1127-1141.
- [3] R. Agrawal, Y. Chhadva, S. Addagarla, and S. Chaudhari, "Road surface classification and subsequent pothole detection using deep learning," *Proceedings of the 2nd International Conference for Emerging Technology (INCET)*, IEEE, May 21-23, 2021, pp. 626-632.
- [4] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, A. Mraz, T. Kashiya, and Y. Sekimoto, "Deep learning-based road damage detection and classification for multiple countries," *Automation in Construction*, Vol. 132, 2021, pp. 103-115.
- [5] M. M. Hasan, S. Sakib, and K. Deb, "Road Damage Detection and Classification using Deep Neural Network," *Proceedings of the 4th International Conference on Electrical, Computer and Telecommunication Engineering (ICECTE)*, IEEE, December 2022, pp. 111-118.
- [6] K. C. Rehana and G. Remya, "Road Damage Detection and Classification Using YOLOV5," *Proceedings of the 3rd International Conference on Intelligent Computing, Instrumentation and Control Technologies (ICICT)*, IEEE, July 2022, pp. 489-494.
- [7] Z. Tang, R. Chamchong, and P. Pawara, "A Comparison of Road Damage Detection Based on YOLOv8," *Proceedings of the International Conference on Machine Learning and Cybernetics (ICMLC)*, IEEE, July 2023, pp. 223-228.
- [8] M. W. Khan, M. S. Obaidat, K. Mahmood, D. Batool, H. M. S. Badar, M. Aamir, and W. Gao, "Real-Time Road Damage Detection and Infrastructure Evaluation Leveraging Unmanned Aerial Vehicles and Tiny Machine Learning," *IEEE Internet of Things Journal*, Vol. 11, No. 12, 2024, pp. 21347-21358.
- [9] X. Zhang, Y. Li, and Z. Zhou, "YOLOv8-bot: An Improved Road Damage Detection Algorithm Based on YOLOv8," *Proceedings of the 5th International Conference on Computer Vision, Image and Deep Learning (CVIDL)*, IEEE, 2024, pp. 1238-1243.
- [10] V. Hegde, D. Trivedi, A. Alfarrarjeh, A. Deepak, S. H. Kim, and C. Shahabi, "Yet Another Deep Learning Approach for Road Damage Detection using Ensemble Learning," *Proceedings of the IEEE International Conference on Big Data*, IEEE, December 2020, pp. 5553-5558.
- [11] E. E. Koko, K. S. S. A. Adolphe, and M. A. A. M. Abdullah, "Impact of extreme hydrological events and subgrade saturation on asphalt pavement deterioration," *Journal of Disaster Risk Reduction & Infrastructure*, Vol. 45, 2022, pp. 102-114.
- [12] Y. Z. Chen, L. W. Wang, and H. J. Zhang, "Post-disaster infrastructure damage assessment using real-time computer vision and mobile edge intelligence," *International Journal of Critical Infrastructure Protection*, Vol. 38, 2022, pp. 100521.
- [13] R. S. Gopalakrishnan, A. K. Sharma, and P. V. R. Nair, "Landslide-induced transportation network disruption and rapid automated distress mapping using deep neural networks," *Natural Hazards and Earth System Sciences*, Vol. 23, No. 4, 2023, pp. 1411-1428.
- [14] J. M. Martinez, S. K. Gupta, and C. H. Lee, "Rapid response logistics and route accessibility assessment in disaster-affected urban corridors," *Transportation Research Part E: Logistics and Transportation Review*, Vol. 175, 2023, pp. 103180.
- [15] T. H. Nguyen, B. Q. Tran, and P. M. Pham, "Edge-computed object detection for emergency vehicle navigation in communication-blind disaster zones," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 25, No. 2, 2024, pp. 1845-1859.