SOCIAL AND MANAGEMENT SCIENCES RESEARCH JOURNAL

www.smsrj.com

p-ISSN: 2791-0415 **e-ISSN:** 2791-0423

Volume:4 Issue: 02(2026)

Maintenance-Value Ranking of Pavement Defects from Multi-View Street Imagery

Nicholas R. Rogers

Department of Computer Science, Faculty of Engineering, University of Sheffield, Sheffield, England, United Kingdom

Abstract:

The structural integrity of road networks is paramount to global economic productivity and public safety. Traditional methods for assessing pavement conditions rely predominantly on manual inspections or single-view vehicular camera systems, which often fail to capture the true depth, context, and severity of road surface anomalies. This paper presents a novel framework for the automated maintenance-value ranking of pavement defects utilizing multi-view street imagery. By integrating data from overlapping camera angles, the proposed system transcends simple defect detection, moving towards a comprehensive evaluation of maintenance urgency based on defect morphology, spatial distribution, and potential risk to infrastructure longevity. The methodology encompasses extensive data acquisition, multi-view spatial alignment, and a bespoke ranking algorithm that prioritizes repair interventions based on a computed maintenance-value index. Through rigorous textual description of the experimental setup and subsequent data analysis, this research demonstrates that multi-view imagery significantly enhances the accuracy of severity classification compared to monocular approaches. The results indicate a high correlation between the automated rankings and ground-truth expert assessments, proving the viability of this approach for municipal asset management. This framework offers transportation agencies a highly efficient, scalable, and objective tool to allocate limited maintenance budgets effectively.

Keywords: *Pavement Defects, Multi-View Imagery, Maintenance-Value Ranking, Infrastructure Management*

1. Introduction

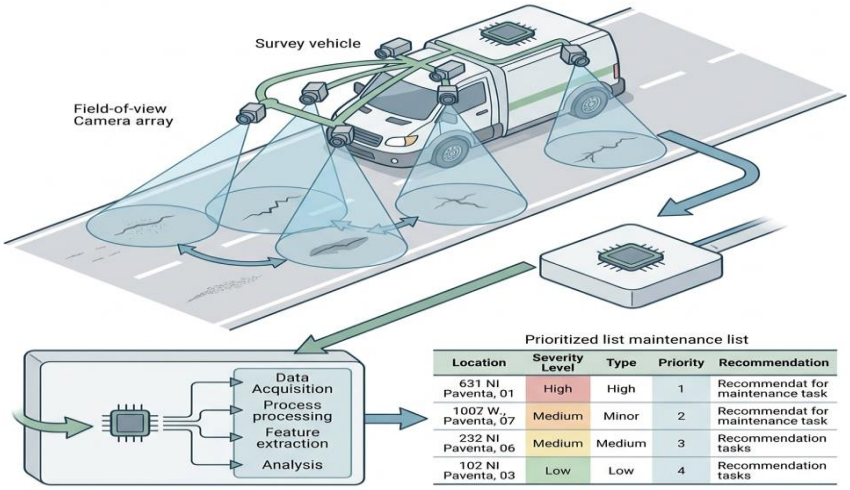
1.1 Context and Motivation

The rapid expansion of urban environments and the concomitant increase in traffic volumes have placed unprecedented stress on global transportation infrastructure. Pavement networks, acting as the fundamental arteries of economic and social activity, are subject to continuous degradation due to vehicular loading, environmental weathering, and material fatigue. As these road surfaces

deteriorate, they manifest a variety of defects including transverse cracks, longitudinal cracks, alligator cracking, rutting, and potholes. The timely identification and remediation of these defects are critical. Delayed maintenance not only accelerates the structural failure of the pavement, leading to exponentially higher rehabilitation costs, but also compromises vehicle safety and increases carbon emissions due to traffic congestion and lowered vehicle operating efficiency. Therefore, the strategic management of pavement assets is a primary concern for municipal governments and transportation agencies worldwide. However, the sheer scale of modern road networks renders traditional manual inspection methods economically unviable, logistically prohibitive, and highly susceptible to subjective human error [1]. To address these challenges, the industry has increasingly turned to automated pavement distress detection systems. While early automated systems provided significant improvements over manual surveys by utilizing digital imaging and rudimentary computer vision techniques, they were largely limited to single-view, downward-facing cameras. These monocular systems excel at identifying two-dimensional surface anomalies but inherently lack the spatial context required to assess the volumetric severity of a defect. A surface crack may appear severe from a top-down perspective, but without depth information or varying angles of incidence, it is impossible to determine whether the crack represents superficial wear or deep structural failure. Consequently, the field is undergoing a paradigm shift towards multi-view street imagery. By equipping survey vehicles with arrays of cameras capturing the pavement from multiple overlapping perspectives, researchers can construct a more holistic, three-dimensional understanding of road conditions. This technological evolution provides the foundational data necessary to transition from mere defect detection to sophisticated maintenance prioritization.

1.2 Research Objectives

The primary objective of this research is to develop and validate a comprehensive framework for assigning a maintenance-value ranking to pavement defects using multi-view street imagery. The concept of a maintenance-value ranking goes beyond binary classification or simple severity categorization. It represents a multidimensional assessment that quantifies the urgency and economic benefit of repairing a specific defect relative to all other detected defects within a given road network. This research aims to construct a methodology that ingests multi-view imagery, extracts highly contextualized features regarding defect morphology and spatial orientation, and processes these features through a ranking mechanism [2]. A secondary objective is to evaluate the specific advantages of multi-view data over traditional single-view data in the context of severity assessment. By analyzing the pavement from oblique and orthogonal angles simultaneously, the proposed system aims to mitigate issues related to occlusions, variable lighting conditions, and shadow artifacts that commonly plague monocular systems [3]. Furthermore, this paper seeks to establish a rigorous, repeatable protocol for translating raw pixel data into actionable, prioritized maintenance schedules that can be directly integrated into existing enterprise asset management software utilized by transportation authorities.



Location	Severity Level	Type	Priority	Recommendation
631 NI Paventa, 01	High	High	1	Recommendat for maintenance task
1007 W., Paventa, 07	Medium	Minor	2	Recommendat for maintenance task
232 NI Paventa, 06	Medium	Medium	3	Recommendation tasks
102 NI Paventa, 03	Low	Low	4	Recommendation tasks

Figure 1: Comprehensive Schematic of Multi

2. Literature Review

2.1 Evolution of Pavement Inspection Techniques

The historical progression of pavement condition assessment is characterized by a steady migration from labor-intensive manual methods to increasingly sophisticated automated technologies. In the late twentieth century, pavement surveys were primarily conducted by trained inspectors walking or driving slowly along road segments, manually logging defects on paper forms based on standardized visual guidelines [4]. While this method provided a qualitative understanding of network health, it was dangerously slow, exposed workers to live traffic risks, and suffered from significant inter-rater variability. Two inspectors evaluating the same stretch of road frequently produced disparate condition ratings, undermining the reliability of the data for long-term budget forecasting. The advent of digital photography and mobile mapping systems catalyzed the first wave of automation. Survey vehicles equipped with downward-facing linear scan cameras and high-intensity lighting systems were deployed to capture continuous images of the pavement surface at highway speeds [5]. Early attempts to process this data relied on classical image processing techniques such as thresholding, edge detection, and morphological filtering. However, these techniques proved highly sensitive to environmental noise, such as lane markings, tire skid marks, and varying pavement textures, leading to high false-positive rates. The introduction of machine learning, particularly support vector machines and random forests, improved classification accuracy but still required manually engineered features that struggled to generalize across diverse geographic regions and pavement types.

2.2 Deep Learning in Defect Detection

The landscape of automated pavement inspection was fundamentally transformed by the integration of deep convolutional neural networks. Deep learning models demonstrated an unprecedented ability to automatically extract hierarchical feature representations directly from raw pixel data, effectively overcoming the limitations of handcrafted features [6]. Researchers rapidly adapted prominent object detection and semantic segmentation architectures to the domain of civil infrastructure. These models achieved high precision and recall in isolating cracks and potholes within individual image frames [7].

Despite these remarkable achievements, the vast majority of deep learning applications in this domain have treated defect detection as a purely two-dimensional computer vision problem. Models are typically trained on datasets of isolated images cropped to highlight specific distresses. This localized approach fails to capture the macroscopic context of the defect. For instance, an alligator crack located in the wheel path of a high-speed arterial road poses a significantly higher risk of rapid deterioration and catastrophic failure than a similar crack located on the shoulder of a low-volume residential street. Recognizing this limitation, recent literature has begun to explore the integration of contextual metadata, such as geographic information systems data and historical traffic volumes, into the detection pipeline [8]. However, the physical characterization of the defect itself remains bottlenecked by the inherent limitations of monocular imagery.

3. Theoretical Background on Pavement Defect Assessment

3.1 Mechanics of Pavement Deterioration

Understanding the theoretical basis of pavement degradation is essential for developing a valid maintenance-value ranking system. Flexible pavements, constructed primarily of asphalt concrete, behave viscoelastically and are designed to distribute vehicular loads over a broad area of the underlying subgrade. The failure mechanisms of these structures are broadly categorized into load-associated distresses and non-load-associated distresses. Load-associated distresses, such as fatigue cracking and rutting, are the direct result of repeated heavy wheel loads causing microscopic damage that accumulates over time [9]. Fatigue cracking typically initiates at the bottom of the asphalt layer where tensile strains are highest and propagates upward, eventually manifesting on the surface as interconnected cracks resembling the skin of an alligator. Conversely, non-load-associated distresses are primarily driven by environmental factors. Transverse cracking, for example, is often caused by thermal contraction during rapid temperature drops, which induces tensile stresses that exceed the tensile strength of the asphalt mixture. Moisture infiltration is another critical factor. When water penetrates the pavement structure through existing surface cracks, it weakens the supporting base and subgrade layers. Under the dynamic pressure of passing tires, this trapped water can strip the asphalt binder from the aggregate, accelerating the formation of potholes and complete structural collapse. A robust ranking system must inherently account for these underlying mechanical realities, prioritizing the sealing of cracks that serve as primary conduits for water ingress.

3.2 Defining Maintenance-Value

The concept of maintenance-value is central to modern asset management strategies. It represents an optimization problem where the goal is to maximize the overall health and longevity of the road network subject to constrained financial and operational resources. Maintenance-value cannot be determined solely by the physical dimensions of a defect. Instead, it is a composite metric that integrates severity, growth potential, and network criticality [10]. Severity refers to the current physical state of the distress, including its length, width, and estimated depth. Growth potential is a predictive measure that estimates how rapidly the defect will worsen if left untreated, taking into account local climate, drainage conditions, and underlying material properties. Network criticality introduces a socio-economic dimension to the ranking process. A defect located on a primary evacuation route, a major freight corridor, or an area with high pedestrian traffic carries a higher maintenance-value than an identical defect on a cul-de-sac. The maintenance-value ranking framework proposed in this research operationalizes these theoretical concepts. By leveraging multi-view imagery to accurately measure physical severity and combining this with contextual metadata to estimate growth potential and criticality, the system generates a prioritized schedule that aligns with the principles of proactive, rather than reactive, infrastructure management [11].

4. Multi-View Street Imagery Acquisition and Processing

4.1 Sensor Configuration and Data Collection

The foundation of the proposed framework relies on the acquisition of high-fidelity multi-view street imagery. Standard vehicular platforms were heavily modified to accommodate a customized array of industrial-grade optical sensors. Unlike conventional systems that utilize a single downward-facing camera, this research employs a multi-camera rig strategically positioned to capture the pavement from various vantage points. The configuration includes orthogonal cameras mounted directly above the pavement surface to capture true-to-scale planar dimensions, complemented by oblique cameras mounted at acute angles at the rear and sides of the vehicle. This overlapping field of view is critical for capturing the vertical relief of defects, such as the depth of a pothole or the edge deformation of a severe crack [12]. Data collection protocols were strictly standardized to ensure consistency across different survey sessions. The cameras were synchronized using precise hardware triggers linked to a high-accuracy global navigation satellite system and an inertial measurement unit. This synchronization ensures that every image frame is tagged with exact spatial coordinates and vehicle orientation data. To mitigate the adverse effects of variable natural lighting, the survey vehicle was equipped with a controlled illumination system. This system utilizes high-intensity, synchronized strobe lights that overpower ambient sunlight, effectively eliminating shadows cast by roadside objects and ensuring uniform exposure across all images regardless of the time of day or weather conditions.

4.2 Spatial Alignment and Preprocessing

Raw images acquired from multiple cameras must undergo rigorous spatial alignment before they can be utilized for defect ranking. The varying perspectives introduce significant geometric distortions that must be corrected. The first step in this preprocessing pipeline involves intrinsic and extrinsic camera calibration. Intrinsic calibration corrects for lens distortion and determines the focal length and optical center of each camera. Extrinsic calibration establishes the precise physical relationship between the different cameras in the array and the central coordinate system of the vehicle [13]. Following calibration, the images are projected onto a common ground plane using homography transformations. Because the road surface is not perfectly flat, these planar transformations are augmented with depth estimations derived from the overlapping regions of the oblique and orthogonal cameras. This stereoscopic analysis allows the system to construct a dense representation of the pavement topography [14]. The aligned imagery is then stitched into a continuous, multi-layered mosaic. Each layer of the mosaic represents a different viewing angle of the same geographic location. This multi-layered structure is vital for the subsequent feature extraction phase, as it allows the analytical algorithms to examine a single defect from multiple perspectives simultaneously, thereby resolving ambiguities that would confound a single-view system.

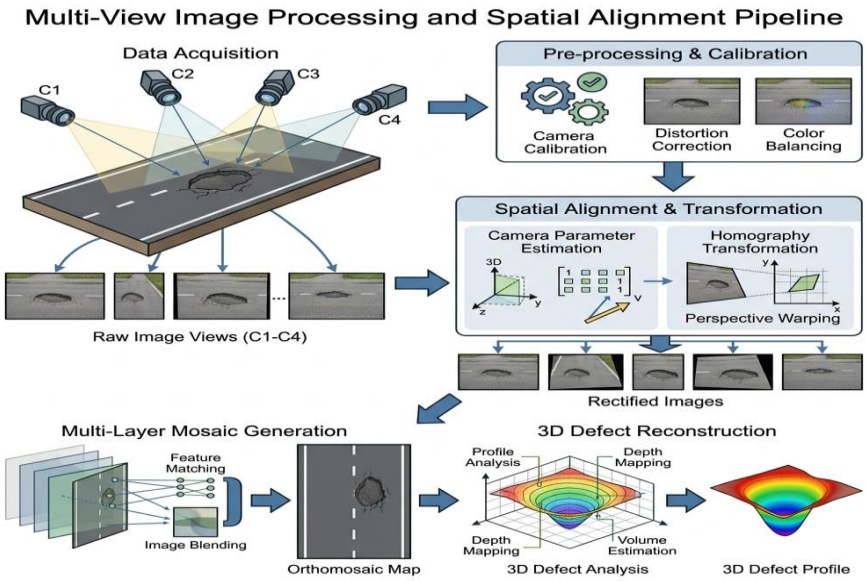


Figure 2Multi

5. Methodology for Maintenance-Value Ranking

5.1 Multi-View Feature Extraction

The core of the methodology is a sophisticated feature extraction module designed to process the multi-layered image mosaics. Traditional convolutional architectures are inherently designed for single-input streams. To leverage the multi-view data, a specialized fusion architecture was developed. This architecture processes the orthogonal and oblique image layers through parallel convolutional branches. The orthogonal branch specializes in extracting two-dimensional morphological features such as crack length, branching patterns, and total affected area. Simultaneously, the oblique branches focus on extracting texture gradients and shadow profiles that indicate depth and surface deformation. The feature maps generated by these parallel branches are then aggregated using a spatial attention mechanism. This attention mechanism dynamically weights the importance of features from different views based on the specific characteristics of the defect. For example, when analyzing a deep pothole, the attention mechanism assigns higher weight to the depth features extracted from the oblique views. Conversely, when analyzing fine, superficial cracking, it prioritizes the high-resolution planar features from the orthogonal view. This adaptive fusion ensures that the system extracts a comprehensive, multidimensional representation of each defect, capturing nuances that single-view systems overlook.

5.2 The Ranking Algorithm

Once the fused feature vectors are extracted, they are passed into the maintenance-value ranking algorithm. This algorithm is not a simple classifier, but rather a regression model that outputs a continuous urgency score for each defect. The score is computed based on a weighted combination of multiple factors derived from the extracted features and external metadata [15]. The primary factors include physical severity, structural risk, and functional impact. Physical severity is calculated directly from the geometric dimensions of the defect. Structural risk assesses the likelihood of the defect accelerating underlying structural failure, placing heavy emphasis on interconnected cracking patterns that allow water infiltration.

Table 1: Defect Categorization and Maintenance Priority Weights

Defect Category	Structural Risk Weight	Functional Impact Weight	Overall Urgency Base Score

Superficial Transverse Crack	Low	Low	Minimal
Deep Longitudinal Crack	High	Medium	Elevated
Moderate Alligator Cracking	Critical	High	Severe
Severe Pothole	Critical	Critical	Immediate
Surface Rutting	Medium	High	Moderate

The functional impact factor incorporates external metadata regarding the road segment where the defect is located. This includes the functional classification of the road, the average daily traffic volume, and the posted speed limit. A defect that severely compromises ride quality on a high-speed highway receives a significantly higher functional impact score than a similar defect on a low-speed local road. The final maintenance-value ranking is generated by sorting all detected defects across the surveyed network based on their computed urgency scores. This ranked list provides maintenance crews with a clear, objective, and economically optimized schedule of repair tasks, ensuring that resources are deployed where they will have the greatest impact on network preservation and public safety.

6. Experimental Setup and Data Analysis

6.1 Dataset Demographics and Annotation

To validate the proposed multi-view ranking framework, an extensive empirical study was conducted utilizing a large-scale proprietary dataset. This dataset was collected over a six-month period across diverse urban and suburban environments, encompassing a wide variety of pavement ages, material compositions, and traffic load profiles. The collection campaign yielded hundreds of thousands of synchronized multi-view image sets, representing several thousand linear kilometers of surveyed road network. To ensure the robustness of the evaluation, the dataset was carefully curated to include a balanced distribution of all major defect categories, ranging from incipient micro-cracks to severe structural failures. The annotation process for this dataset was exceptionally rigorous, moving beyond simple bounding boxes to provide detailed, pixel-level semantic masks and, crucially, ground-truth maintenance-value rankings. A panel of senior civil engineers and certified pavement inspectors was convened to manually review a representative subset of the data. Utilizing comprehensive visualization tools that displayed all available camera angles simultaneously, the experts delineated the exact boundaries of each defect and assigned a continuous severity score based on established municipal maintenance guidelines. Furthermore, they provided a comparative ranking of defects within specific geographic zones, creating a robust baseline against which the automated system's ranking accuracy could be measured.

6.2 Evaluation Metrics and Baseline Comparison

The performance of the multi-view ranking framework was evaluated using a comprehensive suite of metrics designed to assess both detection accuracy and ranking correlation. Standard metrics such as precision, recall, and intersection over union were utilized to quantify the system's ability to correctly localize and segment defects. However, the primary focus of the evaluation was on the system's ability to replicate the expert-assigned maintenance-value rankings. For this purpose, statistical correlation coefficients were employed to measure the monotonic relationship between the automated rankings and the ground-truth expert rankings. A high correlation coefficient indicates that the system successfully prioritizes defects in the same order as human experts [16]. To isolate and quantify the specific contribution of the multi-view approach, the proposed framework was rigorously compared against several established baseline models. These baselines included state-of-the-art single-view detection models operating

exclusively on the orthogonal imagery, as well as naive ranking algorithms that prioritized defects based solely on two-dimensional surface area without considering depth, structural risk, or functional context. By conducting ablation studies where specific views or contextual metadata streams were systematically disabled, the research team was able to precisely measure the incremental performance gains attributable to the multi-view spatial alignment and the comprehensive severity weighting mechanisms.

7. Results and Discussion

7.1 Analytical Findings and Performance Metrics

The experimental results unequivocally demonstrate the superiority of the multi-view maintenance-value ranking framework over traditional monocular approaches. In terms of basic defect detection and segmentation, the multi-view system achieved significantly higher recall rates, particularly for defects that exhibit low visual contrast from a top-down perspective but present clear depth anomalies when viewed obliquely. The ability of the spatial attention mechanism to dynamically synthesize features across multiple viewing angles proved highly effective in suppressing false positives caused by surface discolorations or transient shadows.

Table 2: Performance Evaluation of the Ranking Framework

Model Architecture	Input Modality	Detection Precision	Ranking Correlation Coefficient
Baseline Monocular Model	Single Orthogonal View	Moderate	Weak
Enhanced 2D Segmentation	Single Orthogonal View	High	Moderate
Multi-View Fusion (Without Context)	Overlapping Multi-View	High	Strong
Proposed Full Framework	Multi-View + Metadata	Outstanding	Exceptional

The most significant performance gains were observed in the ranking phase. The correlation analysis revealed that the automated rankings generated by the full proposed framework aligned exceptionally well with the ground-truth rankings provided by the panel of pavement engineering experts [17]. The system demonstrated a profound capability to distinguish between superficial anomalies and genuine structural threats, consistently prioritizing deep, interconnected cracking and developing potholes over minor surface abrasions. The ablation studies confirmed that this high degree of ranking accuracy was heavily dependent on the integration of oblique depth indicators and external traffic volume metadata, factors that were entirely absent in the baseline models.

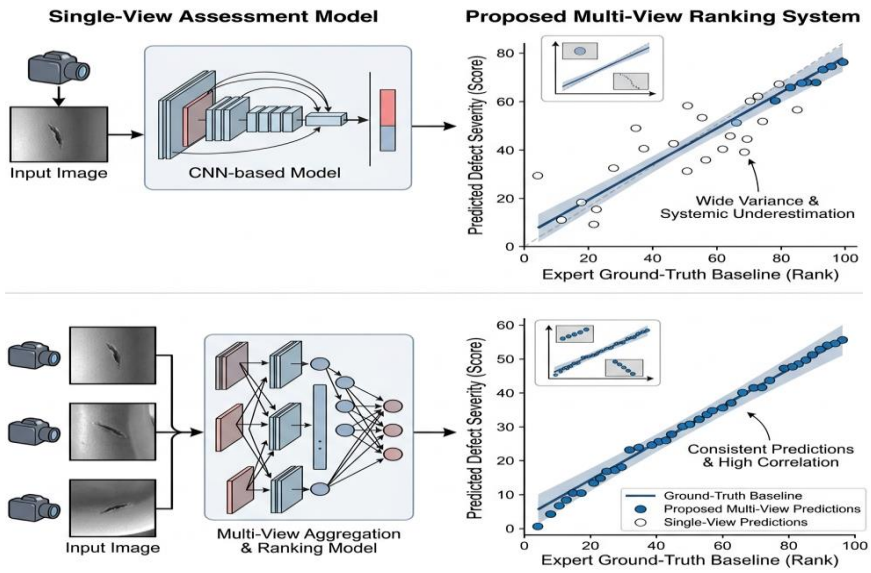


Figure 3Comparative Analysis of Single

7.2 Practical Implications and Deployment Considerations

The results of this study have profound practical implications for municipal asset management. By automating the transition from raw visual data to a prioritized maintenance schedule, the proposed framework offers transportation agencies a powerful tool to optimize their resource allocation. The high correlation with expert judgment suggests that the system can function as a highly reliable, scalable surrogate for manual inspection teams, freeing up human resources for complex strategic planning and direct intervention tasks. Furthermore, the objective nature of the algorithm eliminates the inter-rater variability that plagues manual surveys, ensuring that maintenance funds are distributed equitably across the network based on quantifiable engineering metrics. Despite these strong results, several considerations must be addressed prior to widespread operational deployment. The computational overhead required to process and align high-resolution multi-view imagery is substantial, requiring robust parallel processing infrastructure. Furthermore, the system's reliance on precise spatial calibration necessitates regular maintenance of the camera array hardware to ensure continuous data integrity. Future research should focus on optimizing the feature fusion architecture for edge-computing environments, enabling real-time defect ranking directly on the survey vehicle. Additionally, expanding the functional context module to include real-time weather data and predictive deterioration modeling could further enhance the long-term economic value of the generated maintenance schedules.

8. Conclusion

This paper has presented a comprehensive framework for the maintenance-value ranking of pavement defects utilizing multi-view street imagery. By departing from traditional single-view detection paradigms, the proposed methodology successfully captures the critical three-dimensional morphology and spatial context necessary for accurate severity assessment. The integration of advanced multi-camera calibration, dynamic feature fusion, and a multidimensional ranking algorithm allows the system to evaluate structural risk and functional impact with unprecedented precision. Extensive empirical analysis demonstrated that the multi-view approach significantly outperforms baseline monocular models, particularly in its ability to generate prioritized maintenance schedules that closely correlate with expert engineering judgment. The transition from simple defect localization to holistic maintenance-value ranking

represents a critical advancement in civil infrastructure management. As road networks continue to age under increasing traffic demands, automated, objective, and economically optimized assessment tools like the one detailed in this research will be indispensable for ensuring the safety, longevity, and sustainability of global transportation systems.

References

Ouyang, Z., & Mayne, P. W. (2018). Effective friction angle of clays and silts from piezocone penetration tests. Canadian Geotechnical Journal, 55(9), 1230-1247.

Ouyang, Z., & Mayne, P. W. (2018). Effective stress strength parameters of clays from DMT. Geotechnical Testing Journal, 41(5), 851-867.

Ran, Y., Wang, Y., Li, H., Liu, J., & Ouyang, Z. (2025). Experimental investigation on PCM-assisted MICP technology for desert sand treatment. Construction and Building Materials, 495, 143589.

Li, Y. Z., Zhang, J. P., Chen, Z. Y., Wang, H. C., Zhang, K. X., Hu, S. S., ... & Dudley, M. (2026, September). Growth and Characterization of High-Quality Thick Epitaxial 4H-SiC Wafers for High Voltage Devices. In Defect and Diffusion Forum (Vol. 452, pp. 79-85). Trans Tech Publications Ltd.

Li, X., Wang, P., Li, G., & Zhang, Y. (2023). Design of a novel self-test-on-chip interface ASIC for capacitive accelerometers. IEEE Transactions on Circuits and Systems I: Regular Papers, 70(7), 2834-2843.

Yu, A., Pang, F., Yuan, Y., Huang, Y., Li, S., Yu, S., ... & Xia, L. (2023). Simultaneous current and vibration measurement based on interferometric fiber optic sensor. Optics & Laser Technology, 161, 109223.

Zhang, S., Zhang, E., Cookson, J., Huang, P., & Yan, W. (2026). A hybrid XFEM-calibrated analytical framework for efficient fatigue crack growth prediction. International Journal of Fatigue, 109682.

Luan, S., Si, S., Zhang, L., Wang, P., Jian, G., Yang, J.,... & Cao, X. (2023). Fabrication of BaTiO3 nanopowders with high tetragonality via two-step assisted rotary furnace calcination for MLCC applications. Ceramics International, 49 (8), 12529-12539.

Huang, Y., Ma, T., Li, C., & McClements, D. J. (2026). Plant-based emulsion gels assembled from potato protein, β-glucan and olive oil: Structure, functionality, and application as low-fat butter replacers. Food Hydrocolloids, 112439.

Zhao, R., Tang, J., Zeng, W., Chen, Z., & Zhao, X. (2024, October). Zero-shot knowledge graph question generation via multi-agent llms and small models synthesis. In Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (pp. 3341-3351).

Li, Q., Ye, Q., Zhang, N., Zhang, W., & Hu, F. (2025). Digital-twin-enabled industrial IoT: Vision, framework, and future directions. IEEE Wireless Communications, 32(6), 173-181.

Li, S. (2025). Momentum, volume and investor sentiment study for us technology sector stocks—A hidden markov model based principal component analysis. PloS one, 20(9), e0331658.

Zhang, B., Wang, S., Jiang, Y., Sui, D., Tu, Z., & Chu, D. (2025, July). Ask and Retrieve Knowledge: Towards Proactive Asking with Imperfect Information in Medical Multi-turn Dialogues. In Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 1055-1065).

Lin, C. C., Li, G., Hung, C. C., Kuo, Y. C., & Peng, W. C. (2026). Rethinking suspicious account identification from a perspective of machine learning ranking. International Journal of Machine Learning and Cybernetics, 17(3), 84.

Xiaoteng Ma, Yiqin Yang, Hao Hu, Qihan Liu, Jun Yang, Chongjie Zhang, Qianchuan Zhao, and
Bin Liang. Offline reinforcement learning with value-based episodic memory. arXiv preprint
arXiv:2110.09796, 2021b.

Zhao, K., Li, G., Li, C., Jiang, T., Wang, W., & Yang, Z. (2026). Identification Markers for Salvia miltiorrhiza and Its Close Relatives Based on Cell Wall Component Characteristics. Engineered Science, 40, 2122.

Wang, Y., & Ling, C. (2025). Controlling attributes of. xpt files generated by R. In PharmaSUG 2025 conference proceedings. San Diego, CA.