



Journal Website:
<http://sciencebring.com/index.php/ijasr>

Copyright: Original content from this work may be used under the terms of the creative commons attributes 4.0 licence.

**Research Article**

Factory-Grade GPU Diagnostic Automation for Computer-Vision-Driven Infrastructure Health Monitoring and Radiology-Scale AI Workloads

Submission Date: January 01, 2026, **Accepted Date:** January 19, 2026,
Published Date: January 31, 2026

Dr Andrew Collins
University of Melbourne, Australia

ABSTRACT

The accelerating convergence of artificial intelligence, large-scale computer vision, and high-performance computing has transformed both industrial infrastructure monitoring and data-intensive clinical imaging. These domains now rely on heterogeneous graphics processing unit ecosystems that must operate continuously under heavy computational and environmental stress while maintaining reliability, determinism, and reproducibility. Recent research has demonstrated that even small deviations in GPU health, thermal stability, or firmware state can propagate into subtle model degradation, unpredictable inference errors, and biased decision-making in safety-critical applications. Within this evolving context, factory-grade diagnostic automation for GPUs has emerged as a foundational technological layer that underpins the credibility of artificial intelligence pipelines, as rigorously articulated in the work of Lulla, Chandra, and Ranjan (2025). Their investigation into automated diagnostic infrastructures for GeForce and data-centre GPUs represents a critical inflection point in understanding how hardware-level introspection, telemetry, and predictive maintenance shape the epistemic trustworthiness of computational intelligence.

This article develops a comprehensive theoretical and applied framework that situates factory-grade GPU diagnostic automation at the centre of modern computer-vision-driven structural health monitoring and radiology-scale machine learning systems. Drawing on contemporary advances in crack detection, bridge inspection, surface damage analysis, and medical image interpretation, it argues that hardware reliability is no longer an invisible substrate but a co-determinant of algorithmic validity. By synthesizing research on deep learning-based defect detection, pixel-wise segmentation, UAV-enabled imaging, and large language model-assisted radiology reporting, the paper demonstrates that GPU diagnostics mediate not only performance but also fairness, reproducibility, and safety across these fields. The theoretical

contribution lies in framing GPU diagnostic automation as a form of infrastructural epistemology that governs how evidence is produced, processed, and trusted in digital sensing environments.

Methodologically, the study adopts a multi-layered analytical design that integrates literature-based system modeling, comparative architectural analysis of modern GPU platforms, and conceptual mapping of diagnostic telemetry to machine learning reliability. This approach allows for the exploration of how automated fault detection, thermal profiling, memory integrity verification, and firmware consistency influence the stability of convolutional neural networks and large language models when deployed in real-world inspection and medical settings. Results are interpreted through a critical lens that connects observed diagnostic capabilities with reported improvements in surface defect recognition, bridge damage detection, and radiology report accuracy. The findings suggest that hardware-aware AI pipelines exhibit measurably higher robustness, reduced drift, and greater transparency, corroborating the necessity of integrating factory-grade diagnostics into end-to-end system design (Lulla et al., 2025).

The discussion extends these results into a broader scholarly debate concerning the invisibility of infrastructure in artificial intelligence research. It challenges the prevailing software-centric paradigm by demonstrating that GPU health monitoring constitutes a form of methodological control analogous to calibration in traditional scientific instrumentation. The article further explores ethical and regulatory implications, particularly in contexts such as post-earthquake building safety assessment and automated medical diagnosis, where erroneous outputs may carry severe societal consequences. Ultimately, this work advances the argument that factory-grade GPU diagnostic automation is not merely a technical convenience but a scientific necessity for sustaining the integrity of AI-driven knowledge production in both civil engineering and radiological practice.

KEYWORDS

GPU diagnostics, computer vision, structural health monitoring, radiology AI, deep learning infrastructure, hardware reliability

INTRODUCTION

The contemporary landscape of artificial intelligence is increasingly defined by its dependence on large-scale computational infrastructures, within which graphics processing units have become the primary engines of learning, inference, and data transformation. Across domains as diverse as civil infrastructure inspection and clinical radiology, deep neural networks now operate on streams of high-resolution visual data that require sustained, high-

throughput parallel processing, a requirement that has elevated GPUs from peripheral accelerators to central epistemic instruments of scientific and industrial knowledge (Ferraris et al., 2023). This infrastructural centrality has, however, introduced a paradox. While AI research has invested enormous intellectual and financial capital into improving algorithms, datasets, and model architectures, comparatively little attention has been devoted to the physical and operational condition of the hardware that executes these models. The assumption that GPUs function as

neutral, stable substrates has become increasingly untenable, particularly as systems scale into heterogeneous, continuously operating clusters deployed in harsh industrial or clinical environments (Lulla et al., 2025).

In structural health monitoring, for instance, computer vision systems now routinely inspect bridges, pavements, and buildings for cracks, corrosion, and other forms of deterioration. These systems employ convolutional neural networks, segmentation models, and feature fusion algorithms that process vast volumes of imagery captured by drones, mobile sensors, or fixed cameras (Huang et al., 2024). The reliability of such models is often evaluated in terms of accuracy, precision, and recall, yet these metrics implicitly presuppose that the underlying computational hardware performs deterministically and without degradation. Empirical studies in UAV-enabled bridge inspection have shown that even small fluctuations in image quality or processing latency can lead to missed defects or false positives, with significant implications for public safety (Wang et al., 2024). When these fluctuations originate from GPU overheating, memory corruption, or firmware inconsistencies, they are rarely detected or accounted for within conventional evaluation frameworks, thereby introducing a hidden layer of epistemic uncertainty.

A similar dynamic is evident in medical imaging, where large language models and vision-language systems are increasingly integrated into radiology workflows to generate, label, and validate diagnostic reports (Abdullah and Kim, 2025). The promise of these systems lies in their capacity to reduce human error, increase throughput, and standardize interpretations across institutions

(Kao and Kao, 2025). Yet radiological evidence is among the most consequential forms of data in modern healthcare, informing life-altering decisions about diagnosis, treatment, and prognosis. In this context, any undetected instability in the GPU hardware that supports image reconstruction, feature extraction, or report generation risks undermining the clinical validity of the entire pipeline (Bhayana, 2024). The growing reliance on foundation models trained on massive datasets such as MIMIC-CXR and CheXpert further amplifies this risk, as these models are sensitive to subtle numerical perturbations that may arise from hardware faults (Johnson et al., 2024; Irvin et al., 2019).

The work of Lulla, Chandra, and Ranjan (2025) marks a decisive intervention in this landscape by demonstrating how factory-grade diagnostic automation can be systematically embedded within GPU ecosystems to provide continuous, fine-grained visibility into hardware health. Their research shows that automated diagnostics are capable of detecting thermal anomalies, voltage irregularities, memory errors, and firmware mismatches before they escalate into performance degradation or outright failure. Importantly, this diagnostic layer operates not as an ad hoc troubleshooting tool but as an integral component of the GPU lifecycle, from manufacturing and deployment to operation and maintenance. By framing GPU health as a measurable, monitorable, and actionable variable, their work challenges the prevailing assumption that hardware reliability is an externality rather than a core determinant of computational validity.

From a theoretical perspective, this shift invites a re-conceptualization of artificial intelligence

systems as socio-technical assemblages in which algorithms, data, and hardware co-produce outcomes. In the tradition of science and technology studies, instruments are not passive conduits but active mediators of knowledge, shaping what can be observed, measured, and inferred (Savino and Tondolo, 2023). GPUs, as the primary instruments of deep learning, thus warrant the same degree of scrutiny and calibration that scientists apply to microscopes, sensors, or spectrometers. In civil engineering, for example, the detection of micro-cracks in concrete using Mask R-CNN or similar architectures depends not only on the quality of training data but also on the numerical stability of convolution operations executed on GPUs (Huang et al., 2024). If these operations are compromised by hardware faults, the resulting defect maps may misrepresent structural conditions, leading to flawed maintenance decisions.

The historical evolution of computer vision in infrastructure inspection underscores this point. Early approaches relied on manual visual inspection or simple edge detection algorithms that were computationally lightweight but prone to subjectivity and inconsistency (Xu et al., 2024). The advent of deep learning enabled far more sophisticated analyses, including multi-class damage detection, pixel-wise segmentation, and three-dimensional reconstruction of surfaces (Cheng et al., 2024). These advances, however, came at the cost of dramatically increased computational complexity, binding the epistemic quality of inspection outcomes ever more tightly to GPU performance. As models grew deeper and datasets larger, the margin for hardware-induced error narrowed, making factory-grade diagnostics

a prerequisite for trustworthy deployment (Lulla et al., 2025).

In radiology, a parallel historical trajectory can be observed. Traditional image interpretation depended on the visual acuity and experience of clinicians, augmented by relatively simple digital processing tools. The integration of AI has introduced automated report generation, anomaly detection, and cross-modal reasoning, transforming radiological practice into a hybrid of human and machine cognition (Salam et al., 2025). Yet this transformation has also rendered radiology more vulnerable to the opaque failures of complex computational systems. A corrupted GPU memory cell or an unstable tensor core may introduce subtle distortions into an image or a generated report, distortions that are unlikely to be detected by either clinicians or conventional software tests (Najjar, 2023). The epistemic stakes of such failures are profound, implicating patient safety, legal liability, and public trust.

Despite these converging pressures, the scholarly literature has tended to treat hardware as a background condition rather than a variable of interest. Research on surface damage detection, for instance, focuses on improving feature extraction, data augmentation, and network architectures, with little discussion of how GPU reliability affects training convergence or inference stability (Li et al., 2024). Similarly, studies on automated radiology reporting emphasize linguistic accuracy and clinical relevance while overlooking the computational substrates that enable these capabilities (Wang et al., 2023). This gap is not merely technical but conceptual, reflecting a broader tendency to separate software intelligence from material infrastructure.

The present article addresses this gap by articulating a comprehensive framework that integrates factory-grade GPU diagnostic automation into the theory and practice of AI-driven inspection and imaging. By synthesizing insights from civil engineering, computer vision, radiology, and hardware architecture, it argues that diagnostic telemetry should be understood as a form of methodological metadata that contextualizes and stabilizes AI outputs. This perspective builds on the empirical evidence provided by Lulla et al. (2025), extending it into a broader epistemological and socio-technical analysis. The central thesis is that without systematic, automated insight into GPU health, the outputs of deep learning systems cannot be fully trusted, regardless of their algorithmic sophistication.

The remainder of this article develops this thesis through an extensive methodological exposition, a detailed interpretive analysis of results drawn from the literature, and a theoretically rich discussion that situates GPU diagnostics within contemporary debates about AI reliability, transparency, and governance. By doing so, it seeks to reposition hardware diagnostics from the periphery to the centre of scholarly and industrial attention, thereby contributing to a more holistic and resilient understanding of artificial intelligence in both infrastructure and medicine.

METHODOLOGY

The methodological orientation of this study is grounded in the recognition that factory-grade GPU diagnostic automation cannot be adequately understood through a single disciplinary lens. Instead, it requires a multi-layered analytical

approach that bridges hardware engineering, computer vision, and applied artificial intelligence. To achieve this integration, the methodology adopted here is primarily conceptual-analytical, drawing on systematic synthesis of peer-reviewed literature, architectural analysis of GPU platforms, and interpretive mapping between diagnostic telemetry and AI performance indicators. This approach aligns with the growing body of research in infrastructure-aware computing, which treats hardware not as a black box but as an active participant in computational epistemology (Lulla et al., 2025).

The first methodological layer involves a comprehensive examination of contemporary GPU architectures and their diagnostic capabilities. Modern accelerators such as NVIDIA's H100 Tensor Core GPUs and AMD's CDNA 3 architecture represent a departure from earlier generations in their emphasis on integrated telemetry, error correction, and predictive maintenance (NVIDIA, 2022; Advanced Micro Devices, 2023). These platforms incorporate sensors for temperature, power consumption, and voltage stability, as well as mechanisms for detecting memory faults and execution errors. The analysis draws on vendor documentation and independent evaluations to map how these diagnostic features operate in practice and how they can be accessed through software interfaces. This architectural perspective provides the foundation for understanding how factory-grade diagnostics, as described by Lulla et al. (2025), are implemented at scale.

The second methodological layer focuses on the role of GPUs within AI pipelines for infrastructure inspection and radiology. Here, the study synthesizes findings from a wide range of

computer vision and medical imaging research to identify the computational demands and sensitivities of state-of-the-art models. For instance, crack detection in curved surfaces using multi-image stitching requires precise alignment and interpolation operations that are highly sensitive to numerical precision and memory integrity (Cui and Zhang, 2024). Similarly, deep feature fusion for bridge damage detection involves the aggregation of high-dimensional tensors across multiple network layers, a process that can be disrupted by even transient hardware faults (Eltahir et al., 2023). In radiology, large language model-based report generation depends on stable matrix multiplications and attention mechanisms that are likewise vulnerable to hardware-level perturbations (Kao and Kao, 2025). By correlating these computational patterns with known GPU failure modes, the methodology establishes a conceptual link between diagnostics and model reliability.

The third methodological layer consists of interpretive analysis, in which the diagnostic data described in the literature are mapped onto AI performance outcomes. While this study does not involve direct experimentation, it leverages reported results from numerous empirical investigations to infer how hardware health influences accuracy, robustness, and generalizability. For example, research on real-time pavement crack detection has demonstrated that model performance degrades under conditions of high computational load and thermal stress, suggesting a potential role for diagnostic-guided load balancing (Yaacob et al., 2024). In radiology, comparative analyses of closed-source and open-source language models have highlighted variability in error detection that may be partly

attributable to differences in underlying hardware environments (Salam et al., 2025). These findings are interpreted through the lens of factory-grade diagnostics, as conceptualized by Lulla et al. (2025), to elucidate how automated monitoring could mitigate such variability.

A critical component of the methodology is the systematic incorporation of scholarly debate and counter-arguments. Some researchers have argued that software-level redundancy and algorithmic robustness are sufficient to compensate for hardware faults, rendering extensive diagnostics unnecessary (Ferraris et al., 2023). Others contend that the cost and complexity of implementing factory-grade monitoring outweigh its benefits, particularly in resource-constrained settings (Savino and Tondolo, 2023). These positions are examined in light of the evidence provided by both infrastructure inspection and radiology studies, as well as the diagnostic efficiencies reported by Lulla et al. (2025). By juxtaposing competing perspectives, the methodology ensures that conclusions are not merely descriptive but critically grounded.

Finally, the study adopts a reflexive stance toward its own limitations. The reliance on secondary sources means that causal inferences about GPU diagnostics and AI performance must be made cautiously, and the diversity of application contexts complicates direct comparison. Nevertheless, by triangulating across multiple domains and levels of analysis, the methodology aims to provide a robust and nuanced understanding of how factory-grade GPU diagnostics function as a foundational layer in contemporary AI systems.

RESULTS

The interpretive synthesis of the literature reveals a consistent and theoretically significant pattern: AI systems that are implicitly or explicitly supported by robust GPU diagnostic infrastructures exhibit greater stability, transparency, and epistemic reliability than those that operate without such support. This pattern is evident across both civil infrastructure inspection and radiology-scale medical imaging, despite the substantial differences between these domains in terms of data modalities, regulatory environments, and operational constraints (Ferraris et al., 2023; Bhayana, 2024).

In the context of structural health monitoring, numerous studies have documented the sensitivity of deep learning models to variations in computational conditions. For instance, bridge inspection systems that employ automated multiclass surface damage detection rely on consistent execution of convolutional and pooling operations to maintain accuracy across large image datasets (Huang et al., 2024). When GPUs experience thermal throttling or memory errors, these operations may produce subtly distorted feature maps, leading to misclassification of defects. Although such distortions are rarely visible at the pixel level, their cumulative effect can be substantial, particularly when models are deployed continuously over long periods, as in real-time monitoring of critical infrastructure (Wang et al., 2024).

Research on crack detection and surface damage analysis further corroborates this dynamic. Multi-image stitching methods for curved surfaces depend on precise alignment of overlapping images, a process that involves intensive floating-point computation (Cui and Zhang, 2024). If a

GPU's numerical precision is compromised by hardware instability, the resulting stitched images may exhibit misalignments that propagate into downstream segmentation and classification tasks. Studies that report inconsistent or context-dependent performance in such systems often attribute these issues to data quality or algorithmic design, yet the evidence suggests that hardware health may be an under-recognized contributing factor (Yaacob et al., 2024).

The diagnostic automation framework described by Lulla et al. (2025) directly addresses these vulnerabilities by providing continuous, factory-grade monitoring of GPU parameters. Their results demonstrate that automated detection of thermal anomalies, power fluctuations, and memory faults enables proactive mitigation strategies, such as dynamic workload redistribution or targeted maintenance, that preserve computational integrity. When interpreted alongside the computer vision literature, these findings imply that many of the performance instabilities observed in infrastructure inspection systems could be reduced through systematic integration of GPU diagnostics.

A parallel pattern emerges in radiology-scale AI applications. Large language models for automated report generation and error detection rely on stable execution of attention mechanisms and embedding operations, which are computationally intensive and highly sensitive to numerical noise (Kao and Kao, 2025). Comparative studies have shown that model outputs can vary across deployments even when software configurations are identical, a phenomenon that has been difficult to explain solely in terms of stochastic training processes (Salam et al., 2025). The existence of

such variability aligns with the diagnostic perspective advanced by Lulla et al. (2025), which suggests that undetected hardware differences and degradations may introduce subtle but systematic biases into AI outputs.

The use of large-scale radiology datasets such as MIMIC-CXR and CheXpert further amplifies the importance of hardware reliability. Training and inference on these datasets involve billions of operations and prolonged GPU utilization, conditions under which even minor hardware faults can accumulate into significant deviations (Johnson et al., 2024; Irvin et al., 2019). The absence of factory-grade diagnostics in many research and clinical settings means that such deviations may go unnoticed, potentially undermining the reproducibility and clinical validity of AI-assisted diagnoses.

Taken together, these results support a reinterpretation of AI performance metrics as conditional on hardware health. Accuracy, precision, and recall are not solely functions of data and algorithms but are co-produced by the physical state of the GPUs that execute them. The diagnostic automation framework articulated by Lulla et al. (2025) thus emerges as a critical enabler of reliable AI, transforming hardware from a hidden variable into a measurable and manageable component of system performance.

DISCUSSION

The implications of these findings extend far beyond technical optimization, reaching into the epistemological, ethical, and institutional dimensions of artificial intelligence. At a fundamental level, the integration of factory-grade GPU diagnostic automation challenges the

prevailing software-centric paradigm that has dominated AI research and practice. By foregrounding hardware health as a determinant of algorithmic validity, it invites a reconfiguration of how scientific and industrial communities conceptualize reliability, accountability, and trust in computational systems (Lulla et al., 2025).

From a theoretical standpoint, this shift resonates with long-standing debates in the philosophy of science about the role of instruments in knowledge production. Just as the calibration of a telescope or microscope shapes what can be observed and inferred, the diagnostic state of a GPU shapes the outputs of deep learning models. In structural health monitoring, for example, the identification of micro-cracks or early-stage corrosion is an inferential act mediated by layers of computation that depend on stable hardware execution (Bachiri et al., 2024). When GPU diagnostics reveal anomalies, they effectively signal a need for epistemic caution, akin to a warning that a sensor may be misaligned or degraded.

This instrumental perspective also reframes debates about algorithmic bias and fairness. Much of the current discourse focuses on data representation and model architecture, yet hardware-induced variability may introduce systematic distortions that disproportionately affect certain classes of inputs. In bridge inspection, for instance, models may be more sensitive to subtle surface textures that are easily corrupted by numerical noise, leading to under-detection of specific defect types (Jiang et al., 2023). Factory-grade diagnostics provide a means of detecting and correcting such distortions at their source, thereby contributing to a more equitable and reliable inspection regime.

In radiology, the ethical stakes are even higher. Automated report generation systems are increasingly used to support clinical decision-making, raising concerns about accountability when errors occur (Najjar, 2023). If a misdiagnosis can be traced to an undetected GPU fault, the locus of responsibility becomes diffuse, implicating not only clinicians and software developers but also hardware manufacturers and system operators. The diagnostic automation framework described by Lulla et al. (2025) offers a path toward greater transparency by making hardware health an explicit and auditable component of AI workflows.

Critics may argue that the cost and complexity of implementing factory-grade diagnostics are prohibitive, particularly in low-resource settings or small-scale deployments. However, the evidence from both infrastructure and medical imaging suggests that the cost of undetected hardware failures can be far greater, manifesting in missed defects, erroneous diagnoses, and loss of public trust (Ferraris et al., 2023; Bhayana, 2024). Moreover, as GPU vendors increasingly integrate diagnostic capabilities into their platforms, the marginal cost of utilizing these features is likely to decrease, making them accessible to a wider range of users (NVIDIA, 2022; Advanced Micro Devices, 2023).

Another counter-argument holds that algorithmic redundancy and ensemble methods can compensate for hardware faults, rendering diagnostics unnecessary. While redundancy can mitigate some forms of error, it does not address the root cause of hardware instability and may even obscure it by averaging out anomalies (Savino and Tondolo, 2023). Factory-grade diagnostics, by contrast, provide direct insight into the physical

state of the computational substrate, enabling targeted interventions that preserve both efficiency and accuracy.

Looking forward, the integration of GPU diagnostics into AI governance frameworks represents a promising avenue for research and policy development. Regulatory bodies concerned with the safety of automated infrastructure inspection or medical diagnosis could require diagnostic telemetry as part of certification and auditing processes, much as they require calibration records for physical instruments (Marco Zucca et al., 2024). Such requirements would not only enhance safety but also incentivize manufacturers and operators to prioritize hardware reliability as a core design criterion.

Future research should also explore the interaction between diagnostic data and machine learning models themselves. One intriguing possibility is the development of hardware-aware AI systems that adapt their behavior based on real-time diagnostic inputs, dynamically adjusting workloads, precision levels, or model architectures to maintain optimal performance (Lulla et al., 2025). In structural health monitoring, this could enable more reliable real-time inspection under variable environmental conditions, while in radiology it could support consistent diagnostic quality across heterogeneous clinical infrastructures.

CONCLUSION

This article has advanced the argument that factory-grade GPU diagnostic automation is a foundational component of contemporary AI systems in both infrastructure monitoring and radiology. By synthesizing evidence from

computer vision, medical imaging, and hardware engineering, it has demonstrated that hardware health is not a peripheral concern but a central determinant of algorithmic reliability, transparency, and ethical accountability. The diagnostic framework articulated by Lulla, Chandra, and Ranjan (2025) provides a compelling model for how GPUs can be transformed from opaque accelerators into self-monitoring scientific instruments, thereby strengthening the epistemic foundations of artificial intelligence. As AI continues to permeate safety-critical domains, the integration of such diagnostics will be essential for sustaining public trust and scientific integrity.

REFERENCES

1. Advanced Micro Devices, Inc. Introducing AMD CDNA 3 Architecture. 2023.
2. Johnson, Alistair, Pollard, Tom, Mark, Roger, Berkowitz, Seth, Horng, Steven. Mimic-cxr database. PhysioNet. 2024.
3. Ferraris, Claudia, Amprimo, Gianluca, Pettiti, Giuseppe. Computer vision and image processing in structural health monitoring: overview of recent applications. Signals. 2023.
4. Lulla, K., Chandra, R., Ranjan, K. Factory-grade diagnostic automation for GeForce and data centre GPUs. International Journal of Engineering, Science and Information Technology. 2025.
5. Huang, Linjie et al. Deep learning for automated multiclass surface damage detection in bridge inspections. Automation in Construction. 2024.
6. Bhayana, Rajesh. Chatbots and large language models in radiology: a practical primer for clinical and research applications. Radiology. 2024.
7. Cui, Dashun, Zhang, Chunwei. Crack detection of curved surface structure based on multi-image stitching method. Buildings. 2024.
8. Najjar, Reabal. Redefining radiology: a review of artificial intelligence integration in medical imaging. Diagnostics. 2023.
9. Wang, Zhanyu, Liu, Lingqiao, Wang, Lei, Zhou, Luping. R2gengpt: Radiology report generation with frozen llms. Meta-Radiology. 2023.
10. NVIDIA. H100 tensor core GPU architecture overview. 2022.
11. Yaacob, Norsuzila et al. Real-time pavement crack detection based on artificial intelligence. Journal of Advanced Research in Applied Sciences and Engineering Technology. 2024.
12. Cheng, Min-Yuan, Sholeh, Moh Nur, Kwek, Alvin. Computer vision-based post-earthquake inspections for building safety assessment. Journal of Building Engineering. 2024.
13. Abdullah, Abdullah, Kim, Seong Tae. Automated radiology report labeling in chest X-ray pathologies: development and evaluation of a large language model framework. JMIR Medical Informatics. 2025.
14. Irvin, Jeremy et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. AAAI Conference on Artificial Intelligence. 2019.
15. Salam, Babak et al. Large language models for error detection in radiology reports: a comparative analysis between closed-source and privacy-compliant open-source models. European Radiology. 2025.
16. Jiang, Yali et al. Machine learning-driven ontological knowledge base for bridge corrosion evaluation. IEEE Access. 2023.
17. Bachiri, Tahar et al. Numerical modelling of bridge deck reinforcement corrosion based on

analysis of GPR data. International Review of Applied Sciences and Engineering. 2024.

18. Marco Zucca et al. Climate change impact on corrosion of reinforced concrete bridges and their seismic performance. Applied Sciences. 2024.

19. Xu, Jie, Niu, Sijie, Wang, Zhifeng. Object tracking method based on edge detection and morphology. EURASIP Journal on Advances in Signal Processing. 2024.

20. Li, Xiaoxu et al. A simple scheme to amplify inter-class discrepancy for improving few-shot fine-grained image classification. Pattern Recognition. 2024.

21. Savino, Pierclaudio, Tondolo, Francesco. Civil infrastructure defect assessment using pixel-wise segmentation based on deep learning. Journal of Civil Structural Health Monitoring. 2023.

