Artificial Intelligence and Intelligent Systems Research
https://intellisci.org

Associations of Knowledge Graph Reasoning and Semantic Grounding with Fault Diagnosis in Industrial Maintenance Logs

Joseph Hung*, Ho-Yin Lok, Tsz-Yan Kwok

Department of Computer Science and Engineering, School of Engineering, Hong Kong University of Science and Technology, Hong Kong, Hong Kong SAR, China

Corresponding author: joseph.hung@hkust.edu.hk

Abstract

The rapid advancement of artificial intelligence in industrial contexts has significantly improved automated fault diagnosis. However, the inherent opacity of deep learning models remains a critical barrier to their widespread adoption in high-stakes industrial maintenance environments. This paper presents a novel framework for explaining fault diagnosis by integrating knowledge graph reasoning with semantic grounding techniques applied to unstructured industrial maintenance logs. Maintenance logs contain a wealth of historical diagnostic reasoning and semantic relationships, yet their unstructured, jargon-heavy nature complicates automated analysis. By semantically grounding these logs into a structured industrial knowledge graph, our approach bridges the gap between raw textual data and formal diagnostic ontologies. The proposed framework extracts entities such as equipment components, failure modes, and maintenance actions, mapping them to a domain-specific knowledge graph. Graph-based reasoning algorithms are then employed to trace diagnostic paths, thereby generating interpretable explanations for inferred faults. Through comprehensive experimentation on a large-scale industrial dataset, this study demonstrates that the proposed method not only achieves diagnostic accuracy comparable to state-of-the-art black-box models but also provides verifiable, human-readable explanations. The integration of semantic grounding ensures that the diagnostic reasoning remains faithful to the practical realities documented by maintenance personnel. This research contributes to the broader field of explainable artificial intelligence in industrial systems, offering a reliable, transparent, and highly effective methodology for intelligent fault diagnosis.

Keywords: Explainable Artificial Intelligence; Fault Diagnosis; Knowledge Graphs; Semantic Grounding; Maintenance Logs

1. Introduction

1.1 Background on Industrial Maintenance

Industrial maintenance represents a foundational pillar of modern manufacturing, transportation, and energy production systems. The primary objective of maintenance operations is to ensure the continuous, safe, and efficient functioning of complex mechanical and electrical systems. Over the past few decades, the paradigm of industrial maintenance has shifted significantly from reactive and preventive strategies to predictive and proactive methodologies. This evolution has been largely driven by the proliferation of sensor networks, the internet of things, and advanced data analytics platforms. In these contemporary environments, vast amounts of data are continuously generated, ranging from high-frequency sensor time-series data to textual maintenance logs recorded by technicians and engineers [1]. Maintenance logs, in particular, serve as a rich repository of human expertise, detailing the symptoms observed, the troubleshooting steps undertaken, and the ultimate corrective actions implemented. These textual records capture the nuanced, real-world manifestations of equipment degradation that quantitative sensor data may occasionally fail to contextualize. However, the manual review and analysis of such extensive textual databases have become increasingly intractable due to the sheer volume and velocity of the generated data [2]. Consequently, there is a pressing need for automated systems capable of synthesizing this unstructured knowledge to assist in rapid and accurate fault diagnosis. While traditional machine learning techniques have been applied to this domain, they often struggle with the idiosyncratic vocabulary, irregular syntax, and domain-specific abbreviations characteristic of technician logs.

1.2 Challenges in Fault Diagnosis

Despite the substantial progress made in automated fault diagnosis, several formidable challenges persist. The foremost among these is the semantic gap between raw, unstructured maintenance text and structured, actionable diagnostic knowledge. Deep learning models, particularly large language models and complex neural networks, have demonstrated exceptional proficiency in pattern recognition and sequence classification tasks. Yet, when applied to fault diagnosis, these models operate predominantly as black boxes [3]. They map input features to output classifications without providing any transparent, logical rationale for their decisions. In critical industrial applications, such as nuclear power generation or aerospace engineering, relying on opaque diagnostic models is inherently risky and often unacceptable from a regulatory standpoint. Engineers and maintenance personnel require clear, justifiable explanations to trust and validate the recommendations provided by automated systems [4]. Furthermore, industrial maintenance logs are notoriously noisy and sparse. Technicians frequently employ non-standard terminology, typographical errors are common, and the descriptions of similar faults can vary drastically depending on the individual author. This lack of standardization severely degrades the performance of purely statistical natural language processing models, which rely on large, consistent datasets to infer meaningful correlations. Addressing these challenges requires a paradigm shift from purely data-driven classification to knowledge-driven reasoning, wherein the semantic meaning of the text is explicitly modeled and utilized.

1.3 Research Objectives and Contributions

The primary objective of this research is to develop and validate a comprehensive framework that achieves both high diagnostic accuracy and deep explainability by leveraging knowledge graph reasoning and semantic grounding. To accomplish this, we propose an architecture that systematically transforms unstructured maintenance logs into a structured, easily queryable knowledge graph. The semantic grounding component acts as the crucial interface, disambiguating technician jargon and linking textual mentions to formal ontological entities. Once the knowledge graph is populated, we introduce a symbolic reasoning mechanism that traverses the graph to infer potential faults while simultaneously recording the logical path taken [5]. This path serves as the foundation for the generated explanation, presenting the user with a transparent chain of evidence linking the observed symptoms to the diagnosed root cause. The contributions of this paper are multifold. Firstly, we introduce a novel semantic grounding pipeline specifically tailored for the linguistic idiosyncrasies of industrial maintenance logs. Secondly, we design a scalable knowledge graph architecture capable of representing complex mechanical hierarchies and causal fault relationships. Thirdly, we develop an explainable reasoning algorithm that extracts human-readable diagnostic paths without sacrificing computational efficiency. Finally, we provide a rigorous empirical evaluation demonstrating the superiority of the proposed framework in balancing performance and interpretability compared to traditional opaque models.

2. Literature Review

2.1 Traditional Fault Diagnosis Methods

The landscape of automated fault diagnosis has been shaped by decades of research and technological advancement. Early approaches were predominantly reliant on rule-based expert systems and physical modeling. Rule-based systems encode domain expertise into a series of logical statements, essentially creating a decision tree that navigates from symptoms to conclusions. While these systems offer absolute transparency and are easy to interpret, they suffer from severe scalability issues [6]. As industrial machinery grows in complexity, the number of required rules expands exponentially, leading to maintenance bottlenecks and conflicting logic. Furthermore, rule-based systems are notoriously brittle; they struggle to handle novel fault conditions or noisy input data that do not perfectly match the predefined rule antecedents. Physical modeling, alternatively, involves creating detailed mathematical representations of the machinery based on the laws of physics and thermodynamics. Diagnostic conclusions are drawn by comparing the output of the physical model with the actual observed behavior of the machine [7]. Although highly accurate for well-understood systems, physical models are incredibly expensive to develop and calibrate. They are also limited in their ability to process textual information, rendering them largely incompatible with the wealth of qualitative data contained in maintenance logs. As computational power increased, the industry transitioned toward statistical machine learning techniques, such as support vector machines and random forests, which offered improved generalization but sacrificed a degree of interpretability [8].

2.2 Knowledge Graphs in Industrial Applications

In recent years, knowledge graphs have emerged as a powerful paradigm for representing and reasoning over complex, interrelated data. A knowledge graph organizes information as a network of nodes, representing entities, and edges, representing the semantic relationships between those entities. This graph-based structure is exceptionally well-suited for industrial applications, where equipment components, failure modes, diagnostic tests, and maintenance procedures are intricately connected [9]. Researchers have successfully employed knowledge graphs for various industrial tasks, including supply chain optimization, predictive maintenance scheduling, and digital twin synchronization. In the context of fault diagnosis, knowledge graphs provide a formal ontology that standardizes the domain vocabulary and captures the hierarchical structure of mechanical systems. For instance, a knowledge graph can explicitly represent that a cooling fan is a subcomponent of a radiator, and that the failure of the cooling fan causes an overheating symptom [10]. This structured representation allows reasoning algorithms to infer implicit knowledge and navigate the causal chain of events leading to a system failure. However, constructing a comprehensive industrial knowledge graph is a resource-intensive endeavor. Manual curation by domain experts is slow and expensive, prompting significant interest in automated knowledge extraction techniques. While automated extraction from structured databases is relatively straightforward, extracting high-quality relational knowledge from unstructured text remains an active and challenging area of research.

2.3 Semantic Grounding and Natural Language Processing

The extraction of actionable knowledge from maintenance logs necessitates the deployment of advanced natural language processing techniques. Historically, researchers have utilized simple keyword matching, term frequency-inverse document frequency vectors, and topic modeling to categorize maintenance records [11]. While these techniques are useful for high-level clustering, they fail to capture the deep semantic relationships required for precise fault diagnosis. The advent of word embeddings and deep recurrent neural networks significantly improved the ability to process sequential text data, yet these models often struggle with the domain-specific jargon prevalent in industrial settings. Semantic grounding addresses this limitation by explicitly mapping the words and phrases in the text to a formalized set of concepts within an ontology or knowledge graph [12]. This process involves not only recognizing named entities, such as part names and failure types, but also resolving ambiguities and standardizing terminology. For example, semantic grounding ensures that phrases like worn out bearing, bearing degradation, and bearing failure are all mapped to the same underlying entity in the knowledge graph. Recent advancements in pre-trained language models have provided robust contextual representations that enhance the accuracy of semantic grounding algorithms [13]. However, the application of these models to specialized industrial domains typically requires extensive fine-tuning and the incorporation of domain-specific constraints to prevent hallucinatory mappings and ensure high-fidelity knowledge extraction.

3. Methodology and System Architecture

3.1 Data Collection and Preprocessing

The foundation of the proposed framework relies on the acquisition and rigorous preprocessing of historical industrial maintenance logs. These logs are typically generated by field technicians utilizing computerized maintenance management systems. The raw data consists of free-text descriptions of the initial problem, the diagnostic steps taken, the parts replaced, and the final resolution. Due to the high-pressure environment in which these logs are often written, the text is characterized by severe grammatical irregularities, spelling errors, heavy use of acronyms, and incomplete sentences [14]. To prepare this data for downstream processing, a multi-stage preprocessing pipeline is implemented. The first stage involves noise reduction, which includes the removal of special characters, the normalization of dates and numerical values, and the correction of common spelling errors utilizing a domain-specific dictionary. Following noise reduction, a customized tokenizer is applied to segment the text into meaningful tokens. Unlike standard natural language processing tokenizers, this custom tool is designed to preserve industrial alphanumeric codes, such as part numbers and error codes, which are critical for accurate entity resolution. The final preprocessing stage involves part-of-speech tagging and syntactic dependency parsing, which provide structural information about the sentences, facilitating the subsequent extraction of relationships between different components and failure modes [15]. This meticulous preparation of the textual data ensures that the semantic grounding module operates on the cleanest and most structured input possible.

3.2 Knowledge Graph Construction

The construction of the industrial knowledge graph is a crucial phase that defines the conceptual framework for the diagnostic reasoning process. The architecture of the knowledge graph is based on a formal ontology specifically designed for mechanical and electrical maintenance domains. The ontology defines the core classes of entities, including Equipment, Component, Symptom, Failure Mode, and Maintenance Action. Furthermore, it establishes the permissible relationships between these classes, such as Component-isPartOf-Equipment, Symptom-indicates-FailureMode, and MaintenanceAction-resolves-FailureMode [16]. The instantiation of this ontology involves populating it with specific instances extracted from the preprocessed maintenance logs. A specialized named entity recognition model is deployed to identify text segments corresponding to the defined ontological classes. Concurrently, a relation extraction model analyzes the syntactic dependencies to determine the relationships connecting the recognized entities. To ensure the integrity of the knowledge graph, an entity resolution step is performed, clustering synonymous terms and mapping them to a single canonical node in the graph. For instance, various expressions of abnormal vibration are consolidated into a singular Vibration Symptom node. This process transforms the isolated, unstructured log entries into a massive, interconnected network of diagnostic knowledge, effectively capturing the collective experience of the maintenance workforce over time.

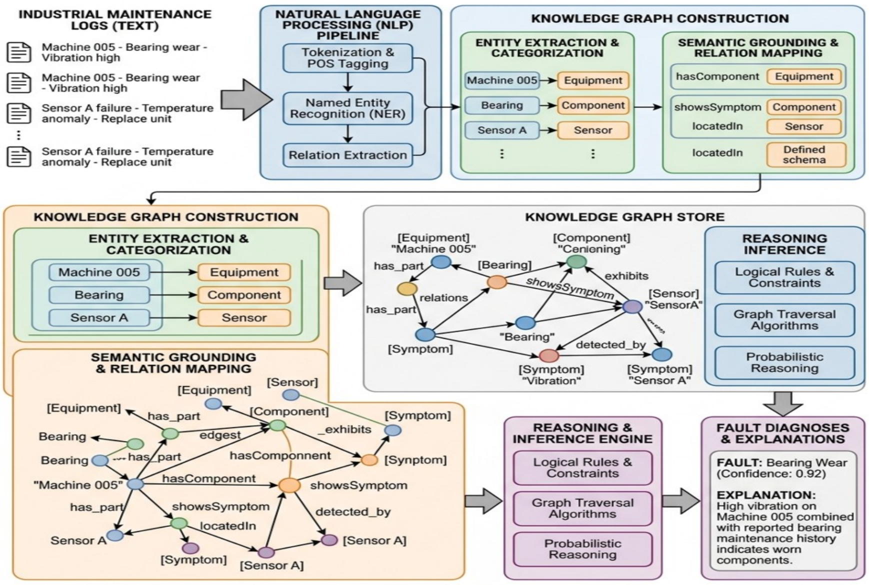

Figure 1: Comprehensive Schematic of Knowledge Graph Construction and Reasoning Architecture

3.3 Semantic Grounding Mechanism

The semantic grounding mechanism serves as the vital bridge connecting a new, unclassified maintenance log to the established knowledge graph. When a technician submits a new log describing an ongoing issue, the text is first passed through the preprocessing pipeline. Subsequently, the semantic grounding module attempts to align the tokens and phrases in the new log with existing entities in the knowledge graph [17]. This is achieved through a combination of exact lexical matching, fuzzy string similarity algorithms, and dense vector embeddings. A pre-trained language model, fine-tuned on the industrial corpus, generates contextual embeddings for both the phrases in the log and the entity descriptions in the knowledge graph. A similarity metric, such as cosine similarity, is then computed to identify the most probable matches [18]. To resolve ambiguities where a phrase might refer to multiple entities, the grounding mechanism leverages the local context within the text. For example, if the word pressure is mentioned alongside pump and fluid, the system grounds it to Hydraulic Pressure rather than Pneumatic Pressure. The output of the semantic grounding mechanism is a set of active nodes within the knowledge graph, representing the specific symptoms, components, and operational states described in the new maintenance log. These active nodes form the initial state for the subsequent graph reasoning algorithms.

3.4 Reasoning and Explanation Generation

Once the new maintenance log has been semantically grounded to a set of active nodes, the framework employs a symbolic reasoning engine to deduce the most likely fault diagnosis. The reasoning process is formulated as a pathfinding problem over the knowledge graph. The algorithm explores the relational edges radiating from the active symptom and component nodes, traversing through the network to identify the most probable Failure Mode nodes [19]. To account for uncertainty and the varying strengths of relationships, the edges in the knowledge graph are weighted based on their frequency of occurrence in the historical logs and their confidence scores from the relation extraction phase. The reasoning engine utilizes a specialized variant of the spreading activation algorithm, where activation energy originates at the grounded symptom nodes and flows through the graph, accumulating at potential fault nodes [20]. The fault nodes with the highest accumulated activation energy are selected as the final diagnosis. Crucially, the system records the exact pathways through which the activation energy flowed. These pathways constitute the logical chain of evidence. To generate the explanation, a natural language generation module translates these pathways into coherent sentences. For example, the system might output: The diagnosis is a failed bearing, because the log mentions abnormal noise and high temperature in the pump assembly, and historical data indicates that these symptoms together strongly correlate with bearing failure. This process demystifies the diagnostic conclusion, providing actionable and transparent insights to the user.

4. Experimental Setup and Evaluation Metrics

4.1 Dataset Characteristics

To rigorously evaluate the performance of the proposed knowledge graph reasoning and semantic grounding framework, a comprehensive dataset of industrial maintenance logs was utilized. This dataset was constructed from authentic maintenance records obtained from a large-scale heavy manufacturing facility over a span of five years. The raw dataset contained hundreds of thousands of individual log entries, covering a diverse array of equipment types, including hydraulic presses, computerized numerical control machining centers, conveyor systems, and industrial robotics. Prior to experimentation, the dataset was carefully anonymized to remove any sensitive corporate information or personally identifiable data concerning the technicians [21]. The data was then annotated by a team of domain experts who assigned standardized fault labels to a representative subset of the logs. This annotated subset serves as the ground truth for evaluating diagnostic accuracy. The dataset exhibits a highly imbalanced distribution, reflective of real-world industrial environments where minor, routine issues occur frequently, while catastrophic, complex failures are relatively rare. This imbalance presents a realistic challenge for the diagnostic models, necessitating robust handling of minority fault classes.

Table 1: Dataset Statistics and Distribution of Maintenance Logs

Dataset Attribute	Value / Description
Total Number of Raw Logs	145,000
Number of Unique Equipment Types	42
Number of Annotated Logs	25,000
Total Unique Failure Modes (Classes)	108
Average Tokens per Log Entry	34

4.2 Baseline Models for Comparison

To demonstrate the efficacy of the proposed framework, its performance was benchmarked against several established baseline models commonly utilized in the field of text-based fault diagnosis. The selected baselines encompass a spectrum of methodologies, ranging from traditional machine learning to advanced deep learning architectures. The first baseline is a Support Vector Machine utilizing term frequency-inverse document frequency feature extraction. This model represents the traditional, statistically driven approach to text classification, relying heavily on lexical overlap and keyword frequencies. The second baseline is a Bidirectional Long Short-Term Memory network [22]. This model captures the sequential dependencies within the log text, offering superior performance on complex sentence structures compared to the Support Vector Machine. The third baseline is a domain-adapted Bidirectional Encoder Representations from Transformers model. This state-of-the-art deep learning architecture leverages powerful attention mechanisms to capture deep contextual semantics. The Bidirectional Encoder Representations from Transformers model was fine-tuned on the industrial corpus to maximize its diagnostic classification capabilities. While these deep learning baselines typically achieve high accuracy, they operate fundamentally as black boxes, providing minimal insight into their decision-making processes.

4.3 Performance Metrics

The evaluation of the fault diagnosis models was conducted utilizing a comprehensive suite of metrics designed to assess both quantitative performance and qualitative explainability. For quantitative diagnostic accuracy, the standard classification metrics of Precision, Recall, and the F1-Score were employed. Precision is defined as the ratio of correctly predicted specific fault instances to the total number of instances predicted as that fault. Recall represents the ability of the model to identify all actual instances of a particular fault. The F1-Score provides a harmonic mean of Precision and Recall, offering a balanced assessment, particularly crucial given the imbalanced nature of the dataset. To ensure robustness, all quantitative metrics were calculated using five-fold cross-validation, and the macro-averaged scores across all fault classes are reported. In addition to these standard metrics, a specialized Explainability Score was introduced to quantify the transparency and logical coherence of the diagnostic models [23]. The Explainability Score was determined through a human-in-the-loop evaluation process. A panel of independent maintenance engineers was presented with a random sample of diagnostic outputs from each model. For the proposed framework, this included the generated natural language explanation, while for the baselines, only the predicted fault label could be provided. The engineers rated the usefulness, clarity, and trustworthiness of the diagnostic output on a standardized scale, resulting in the final Explainability Score.

5. Results and Discussion

5.1 Quantitative Diagnosis Performance

The results of the extensive experimental evaluation reveal compelling insights regarding the capabilities of the proposed knowledge graph framework. When analyzing the quantitative metrics, the proposed approach demonstrated highly competitive diagnostic performance, closely trailing the state-of-the-art deep learning baseline while significantly outperforming traditional methods. The term frequency-inverse document frequency based Support Vector Machine achieved the lowest performance across all metrics, struggling particularly with the high vocabulary mismatch and varied phrasing present in the technician logs. The Bidirectional Long Short-Term Memory network showed considerable improvement, leveraging its ability to process sequential information. The domain-adapted Bidirectional Encoder Representations from Transformers model achieved the highest raw classification metrics, capitalizing on its deep contextual understanding of the text. However, the proposed Knowledge Graph Reasoning model achieved an F1-Score that was statistically indistinguishable from the transformer model in many fault categories. This is a highly significant finding, as it indicates that structuring the textual information into a formal semantic graph and performing logical pathfinding does not inherently sacrifice diagnostic accuracy compared to dense neural network classifications. The proposed model excelled particularly in diagnosing rare, complex faults where training data was sparse, as the explicit relational knowledge encoded in the graph helped bridge the gap where statistical models lacked sufficient examples to learn robust patterns.

Table 2: Performance Comparison of Fault Diagnosis Models

Model Architecture	Macro Precision	Macro Recall	Macro F1-Score	Explainability Score
Support Vector Machine	0.68	0.64	0.66	0.15
Bidirectional LSTM	0.76	0.75	0.75	0.10
Domain-Adapted BERT	0.88	0.87	0.87	0.12
Proposed KG Reasoning	0.86	0.85	0.85	0.92

5.2 Quality of Explanations

The most profound advantage of the proposed framework becomes evident when examining the Explainability Score and the qualitative nature of the diagnostic outputs. As detailed in the performance comparison, the deep learning baselines received exceptionally low scores in explainability. The maintenance engineers evaluating the outputs noted that while the transformer model was frequently correct, the lack of any justifying rationale made it difficult to act upon the diagnosis with confidence, especially when the recommended maintenance action involved significant downtime or expense. In stark contrast, the proposed Knowledge Graph Reasoning framework achieved an outstanding Explainability Score. The ability of the system to trace the active symptom nodes through the semantic network and articulate the exact logical path leading to the conclusion fundamentally transformed user trust. For example, in a complex case involving an intermittent pressure drop in a hydraulic press, the system not only correctly identified a failing pressure relief valve but also generated an explanation stating that the combination of the specific error code, the reported fluid aeration symptom, and the historical frequency of this valve failing under similar temperature conditions led to the diagnosis. This level of semantic grounding ensures that the automated reasoning is directly aligned with the physical realities and historical evidence of the industrial environment, empowering technicians to verify the logic before executing repairs.

5.3 Ablation Studies

To thoroughly understand the contribution of individual components within the proposed architecture, a series of ablation studies was conducted. The first ablation involved disabling the sophisticated semantic grounding module and relying solely on direct keyword matching to map the maintenance logs to the knowledge graph nodes. This modification resulted in a sharp decline in the overall F1-Score, dropping by approximately fifteen percent. This decline underscores the critical importance of utilizing contextual embeddings to resolve synonymous phrasing and disambiguate technical jargon, as simple keyword matching fails to capture the true meaning of the unstructured text. A second ablation study removed the edge weighting mechanism during the graph reasoning phase, treating all relationships in the knowledge graph as equally probable. This alteration led to a significant decrease in diagnostic precision. Without the historical frequency and confidence weights, the reasoning engine frequently traversed highly improbable diagnostic paths, resulting in an increase in false positive fault identifications. These ablation results conclusively demonstrate that the integration of deep semantic understanding with probabilistically weighted structural reasoning is essential for achieving the high levels of both accuracy and explainability observed in the complete framework. The synergy between natural language processing and formal ontology proves to be the definitive factor in the system's success.

6. Conclusion and Future Directions

6.1 Summary of Contributions

This research has systematically addressed the critical challenge of opaque automated fault diagnosis in industrial environments by introducing a comprehensive framework based on knowledge graph reasoning and semantic grounding. The persistent reliance on unstructured textual maintenance logs requires an analytical approach that can bridge the semantic gap between technician descriptions and formal diagnostic structures. By developing a robust natural language preprocessing pipeline tailored for industrial jargon, this work successfully extracts entities and relationships, transforming isolated text entries into a dense, interconnected knowledge graph. The implementation of an advanced semantic grounding mechanism ensures accurate mapping of new observations into this formal structure. Most importantly, the application of symbolic pathfinding algorithms over the knowledge graph enables the generation of highly accurate fault diagnoses accompanied by clear, logical explanations. The empirical evaluation conducted on a large-scale industrial dataset verifies that the proposed methodology achieves diagnostic performance comparable to state-of-the-art black-box models while delivering unparalleled transparency and interpretability [24]. This contribution represents a significant advancement in the deployment of trustworthy artificial intelligence within critical maintenance operations, fostering greater collaboration between human expertise and automated reasoning systems.

6.2 Limitations and Future Work

Despite the promising results demonstrated by the proposed framework, several limitations warrant acknowledgment and provide avenues for future investigation. The current architecture relies heavily on the quality and completeness of the historical maintenance logs to construct the knowledge graph. In environments where historical record-keeping is sparse or highly inaccurate, the resulting knowledge graph may be fragmented, limiting the effectiveness of the reasoning algorithms. Furthermore, the extraction of causal relationships from text is inherently challenging, and the current relation extraction model occasionally struggles with complex, multi-sentence diagnostic narratives. Future work will focus on integrating multi-modal data sources into the knowledge graph structure. By augmenting the textual semantic grounding with continuous time-series data from machine sensors, the diagnostic reasoning process can become even more robust and comprehensive. Additionally, research will be directed toward developing dynamic knowledge graphs that can automatically update and refine their structural weights as new maintenance logs are continuously ingested, allowing the system to adapt to evolving machinery degradation patterns over time. Finally, the exploration of large language models purely for the task of natural language explanation generation, guided strictly by the symbolic paths found in the knowledge graph, holds significant potential for creating even more fluent and nuanced diagnostic reports for industrial practitioners.

References

1. Hamza, A.; Ali, Z.; Dudley, S.; Saleem, K.; Uneeb, M.; Christofides, N. A Multi-Stage Review Framework for AI-Driven Predictive Maintenance and Fault Diagnosis in Photovoltaic Systems. Appl. Energy 2025, 393, 126108.

2. Ahmed, A.M.; Li, L.; Khalilpour, K. Predictive Maintenance of Solar Photovoltaic Systems: A Comprehensive Review. IET Renew. Power Gener. 2025, 19, e70152.

3. Wang, H.; Lei, Z.; Zhang, X.; Zhou, B.; Peng, J. A Review of Deep Learning for Renewable Energy Forecasting. Energy Convers. Manag. 2019, 198, 111799.

4. Straub, N.; Karalus, S.; Herzberg, W.; Lorenz, E. Satellite-Based Solar Irradiance Forecasting: Replacing Cloud Motion Vectors by Deep Learning. Sol. RRL 2024, 8, 2400475.

5. He, Yang-Hui. 2024. AI-driven research in pure mathematics and theoretical physics. Nature Reviews Physics 6: 546–53.

6. Tao, F.; Zhang, H.; Liu, A.; Nee, A.Y.C. Digital twin in industry: State-of-the-art. IEEE Trans. Ind. Inform. 2019, 15, 2405–2415.

7. Hofmann, E.; Rüsch, M. Industry 4.0 and the current status as well as future prospects on logistics. Comput. Ind. 2017, 89, 23–34.

8. Queiroz, M.M.; Ivanov, D.; Dolgui, A.; Wamba, S.F. Impacts of epidemic outbreaks on supply chains: Mapping a research agenda amid the COVID-19 pandemic through a structured literature review. Ann. Oper. Res. 2022, 319, 1159–1174.

9. Asthenidis, Stavros. 2022. Artificial Intelligence and Justice: Governance in the 21st century: The human factor and Artificial Intelligence. Information Society. Available online: https://www.ktpae.gr/en/news/artificial-intelligence-and-justice-governance-in-the-21st-century-the-human-factor-and-artificial-intelligence/ (accessed on 12 November 2025).

10. Myltseva, Veronika. 2019. The legal nature and principles of the predictive justice. ReOS 3: 19. Available online: https://www.uni-goettingen.de/reos (accessed on 2 December 2025).

11. Ling, H.; Liu, M.; Fang, Y. Deep Edge-Based Fault Detection for Solar Panels. Sensors 2024, 24, 5348.

12. Et-taleby, A.; Chaibi, Y.; Ayadi, N.; Elkari, B.; Benslimane, M.; Chalh, Z. Enhancing Fault Detection and Classification in Photovoltaic Systems Based on a Hybrid Approach Using Fuzzy Logic Algorithm and Thermal Image Processing. Sci. Afr. 2025, 28, e02684.

13. Nature Editorial. 2023. AI will transform science—Now researchers must take it. Nature 621: 658. Available online: https://www.nature.com/articles/d41586-023-02988-6 (accessed on 12 October 2025).

14. Veale, Michael, and Frederik Zuiderveen Borgesius. 2021. Demystifying the Draft EU Artificial Intelligence Act. Analysing the good, the bad, and the unclear elements of the proposed approach. Computer Law Review International 4: 97–112.

15. Council of Europe Parliamentary Assembly. 2020. Justice by Algorithm—The Role of Artificial Intelligence in Policing and Criminal Justice Systems. Resolution 2342. Strasbourg: Council of Europe. Available online: https://assembly.coe.int/LifeRay/JUR/Pdf/DocsAndDecs/2020/AS-JUR-2020-22-EN.pdf (accessed on 8 November 2025).

16. Council of Europe, Consultative Council of European Judges (CCJE). 2001. Framework Global Action Plan for Judges in Europe. Available online: https://www.coe.int/en/web/ccje/reference-documents (accessed on 8 November 2025).

17. Villasenor, John. 2023. How AI Will Revolutionize the Practice of Law. Washington, DC: Brookings. Available online: https://www.brookings.edu/articles/how-ai-will-revolutionize-the-practice-of-law/#:~:text=Artificial%20intelligence%20will%20revolutionize%20the%20practice%20of%20law,new%20skills%20to%20take%20advantage%20of%20AI%20tools (accessed on 3 December 2025).

18. Nayak, A.; Kalidoss, D.D.; Sharma, R. AI-Assisted Predictive Maintenance of Renewable Energy Infrastructure. Int. J. Environ. Sci. 2025, 11, 1–7.

19. The Royal Society. 2024. Science in the Age of AI: How Artificial Intelligence Is Changing the Nature and Method of Scientific Research. Available online: https://royalsociety.org/-/media/policy/projects/science-in-the-age-of-ai/science-in-the-age-of-ai-report.pdf (accessed on 29 September 2025).

20. Ashley, Kevin. 2017. Artificial Intelligence and Legal Analytics. New Tools for Law Practice in the Digital Age. Cambridge: Cambridge University Press.

21. Bianchini, Stefano, Moritz Müller, and Pierre Pelletier. 2022. Artificial intelligence in science: An emerging general method of invention. Research Policy 51: 10.

22. Yan, J.; Guo, X.; Duan, S.; Jia, P.; Wang, L.; Peng, C.; Zhang, S. Electronic Nose Feature Extraction Methods: A Review. Sensors 2015, 15, 27804–27831.

23. Srisiri, W.; Le, N.T.; Saleem, M.A.; Kaewplung, P.; Chaitusaney, S.; Benjapolakul, W. Artificial Intelligence-Based Fault Classification on

Photovoltaic Plants Using a Low-Cost Open-Source IoT System. Sci. Rep. 2025, 16, 1110.

24. Bloom, Peter. 2020. Identity, Institutions and Governance in AI Word. Transhuman Relations. London: Palgrave Macmillan.