

Article

An Explainable Time-Series Knowledge Graph Framework with Dynamic Temporal Segmentation for Industrial Spindle Health Monitoring

Chun-Shih Cheng ^{1,2,*}  and Guan-Ju Peng ³ 

¹ Doctoral Program in Big Data Analytics for Industrial Applications, National Chung Hsing University, Taichung 402202, Taiwan

² Department of Industrial Engineering and Management, Chaoyang University of Technology, Taichung 413310, Taiwan

³ Department of Applied Mathematics, National Chung Hsing University, Taichung 402202, Taiwan; gjpeng@email.nchu.edu.tw

* Correspondence: t2025017@cyut.edu.tw; Tel.: +886-4-2332-3000 (ext. 4645)

Abstract

This study presents an explainable knowledge graph (KG) framework that transforms continuous spindle monitoring time-series data into transparent, reasoning-ready diagnostic structures. Existing data-driven approaches, while accurate, often lack the interpretability required for high-stakes industrial decision-making and are sensitive to operating condition drifts. To address these limitations, we propose a two-level temporal segmentation method combining label transition detection and statistical drift analysis to identify meaningful state boundaries. Furthermore, a percentile-based discretization mechanism converts statistical features into interpretable semantic tags. A Neo4j-based state–event–feature schema captures lifecycle evolution and evidence relations, enabling attribution path reasoning that links failure events to salient precursor features. Experiments on real industrial spindle data demonstrate a fault detection accuracy of 84.97% and a false alarm rate of 3.43%, effectively capturing stable baselines and intermittent abnormal bursts. The proposed framework provides a distinct novelty in bridging the gap between numerical time-series and symbolic reasoning, offering a practical pathway for deploying explainable and maintainable spindle health analytics.

Keywords: knowledge graph; explainable AI; temporal segmentation; spindle health monitoring



Academic Editor: Ibrahim Tansel

Received: 5 January 2026

Revised: 20 February 2026

Accepted: 23 February 2026

Published: 4 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Since the concept of Industry 4.0 was proposed at the Hannover Messe in 2011, digitalization and data-driven decision-making have evolved from emerging trends into essential enablers of modern manufacturing competitiveness [1]. In smart manufacturing environments, the health condition of production equipment directly affects the overall equipment effectiveness (OEE), machining quality, and delivery reliability. For high-precision production lines, any unplanned downtime may lead to substantial economic loss and irreversible damage to critical components [2–4].

Among such components, the high-speed motor spindle is a high-value and mission-critical module in advanced equipment such as five-axis machining centers. Operating under long-term high rotational speed and variable load conditions, spindle bearing systems are prone to gradual degradation phenomena including wear, spalling, looseness, and

lubrication deterioration, which are well-documented in the rotating machinery diagnostics literature [3–5]. These degradations are typically manifested through increasing vibration energy, intensified impulsive signals, chatter, and surface quality deterioration, which may eventually evolve into severe failures or unexpected shutdowns [4,5].

Consequently, prognostics and health management (PHM) and predictive maintenance (PdM) have become fundamental capabilities in advanced manufacturing [3,4,6]. Compared with corrective maintenance and preventive maintenance, PdM aims to leverage continuous multivariate sensor streams to detect early degradation before failures occur, thereby enabling proactive maintenance scheduling [3,4,6,7]. However, in real production environments, spindle health monitoring often faces a critical dilemma: although signals are measurable, robust, deployable, and maintainable diagnostic solutions remain insufficient. Continuous data streams, varying operating conditions, and weak or non-stationary fault precursors significantly challenge the reliability and long-term operability of monitoring systems across different operating regimes, batches, and time horizons [6,8].

1.1. Motivation and Problem Statement

In recent years, data-driven approaches have achieved remarkable success in PHM applications. Deep learning models such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs/LSTMs) are capable of automatically learning discriminative representations from large-scale multivariate time-series data and have demonstrated high diagnostic accuracy for machinery fault diagnosis on benchmark datasets [9–12]. Nevertheless, the increasing complexity of such models introduces critical obstacles to industrial deployment.

First, most deep models operate as black-box predictors that provide limited causal or semantic interpretability, which restricts their deployment in high-stakes industrial scenarios [13–15]. Second, even when post hoc explanation tools are applied, the resulting explanations may be unstable or inconsistent under high-risk manufacturing scenarios. As emphasized by Rudin [16], inherently interpretable models should be preferred over black-box models with post hoc explanations when high-stakes decisions are involved.

More specifically, in spindle monitoring applications, simply outputting an “abnormal probability” is insufficient to support maintenance actions. What practitioners truly require is a traceable evidence chain: when a machine state begins to deteriorate, which statistical indicators deviate significantly and what are the mechanical causes these deviations may correspond to (e.g., increased kurtosis caused by bearing impacts or looseness) [5]? Moreover, most industrial AI systems treat continuous sensor data (numeric) and maintenance knowledge (symbolic) as separate entities, creating a semantic gap between numerical time-series representations and reasoning-oriented symbolic knowledge structures [17–19]. From a time-series mining perspective, such a gap can be addressed through symbolic temporal abstraction techniques that convert continuous signals into discrete semantic states.

Based on the above observations, this study focuses on the following three verifiable research gaps:

1. Lack of systematic and theoretically grounded transformation from continuous time-series signals to reasoning-ready symbolic structures, as most existing approaches directly classify them without intermediate state modeling [10–12].
2. Lack of statistically robust semantic baselines across varying operating conditions, since fixed thresholds are sensitive to distribution shifts and concept drift [6,20,21].
3. Lack of formalized fault attribution mechanisms that provide explicit causal paths between precursor states, salient features, and fault events, which limits explainable decision-making [16,22,23].

While deep learning has advanced significantly, a gap remains between theoretical models and industrial deployment. As highlighted by Lipton [15], existing black-box models often struggle in complex environments typical of shop floors. They lack methodological formalism to provide traceable evidence chains, which are essential for maintenance personnel to trust and act upon diagnostic outputs.

1.2. Literature Review

This section reviews the major technological paradigms in machine fault diagnosis and health monitoring and clarifies the research gaps addressed by this study. The evolution of related research can be broadly categorized into three stages: signal-processing-based physical feature methods, data-driven and deep learning approaches, and the integration of knowledge graphs with explainable AI.

1.2.1. Signal-Processing-Based and Physical Feature Methods

Early fault diagnosis methods heavily relied on physical models and expert rules. Lei et al. [11] reported that vibration-based decomposition and time–frequency analysis have long been the mainstream approaches in rotating machinery health monitoring. Jardine et al. [3] systematically reviewed condition-based maintenance and diagnostics/prognostics practices for machinery systems. In addition, Randall and Antoni [5] provided a tutorial-style discussion for rolling element bearing diagnostics and the interpretation of vibration signatures. However, Widodo and Yang [20] pointed out that such approaches are highly dependent on prior knowledge and tend to degrade under non-stationary signals and variable operating conditions. Moreover, operating condition changes and lifecycle evolution often lead to threshold drift and reduced robustness, highlighting the need for more maintainable monitoring strategies [6,21].

1.2.2. Data-Driven and Deep Learning Approaches

With the advancement of computational resources and data availability, deep learning has driven the development of end-to-end diagnostic models. Jia et al. [10] and Zhao et al. [12] demonstrated the effectiveness of deep feature learning for rotating machinery diagnosis and machine health monitoring. Zhang et al. [24] verified the capability of RNN/LSTM models in capturing long-term dependencies for remaining useful life (RUL) prediction, while Li et al. [25] employed domain adaptation to improve generalization under variable operating conditions. Despite these performance gains, many high-accuracy models still suffer from limited interpretability and decision transparency, which remain major barriers to industrial deployment [16,26].

1.2.3. Knowledge Graphs and Explainable AI

To bridge the gap between data-driven and knowledge-driven paradigms, knowledge graphs (KGs) have been introduced as a structured paradigm for representing entities, relations, and causal reasoning in complex systems [18,22,23]. However, existing research often fails to meet industrial requirements because most methods are sensitive to operating condition drifts and lack semantic context. While KGs enable the structured representation of equipment entities and fault modes, conventional approaches remain largely disconnected from raw, dynamic time-series data [23]. Most industrial KG studies focus on static knowledge, such as converting manuals into graph form or using simplified threshold rules, which are insufficient to address the non-stationary nature of spindle monitoring. These approaches cannot consistently transform continuous data into reasoning-ready structures while preserving temporal evolution and explicit attribution paths. Accordingly, this study proposes a framework centered on statistical feature semanticization to

bridge the gap between numeric signals and symbolic reasoning, delivering traceable and explainable diagnostics.

1.3. Research Objectives

Based on the background and research gaps, this study aims to propose an explainable time-series knowledge graph construction framework based on statistical feature extraction and to validate its effectiveness using real-world spindle monitoring data from an industrial production line. The specific research objectives are as follows:

1. To establish a rigorous two-level dynamic temporal segmentation mechanism (output: StateSegment): By integrating supervised label change detection with unsupervised statistical drift analysis, continuous time-series data are functionally transformed into semantically meaningful state segments characterized by explicit temporal boundaries.
2. To formulate a robust, distribution-aware feature semanticization rule (output: MetricFeature): Global percentiles (P_{25} , P_{75}) derived from normal operation base-lines are used to discretize continuous statistical features into semantic categories (high/low/stable), reducing sensitivity to noise and minor drifts while preserving engineering interpretability.
3. To define a semantically formalized state–event–feature knowledge graph schema (output: KG Schema): A structured schema comprising StateSegment, TransitionEvent, and MetricFeature nodes and their relations (e.g., HAS_FEATURE, ENDED_BY) is designed to enable scalable and automated graph construction from time-series data.
4. To enable quantifiable and traceable diagnostic attribution (output: reasoning path): The system implements backward reasoning to generate structured explanation paths that link abnormal events to their salient precursor features, providing auditable and decision-oriented diagnostics to address the “black-box” limitations of existing models.

2. Materials and Methods

2.1. Materials and Experimental Setup

2.1.1. Experimental Scenario and Target Equipment

The experimental data used in this study were collected from a real production line in a precision machine factory. The monitored equipment is a five-axis vertical machining center, which is widely used for high-precision milling operations of aerospace aluminum alloy components. The operational stability of such equipment is critical to machining accuracy, surface quality, and production reliability.

The primary diagnostic target is the high-speed motor spindle with a maximum rotational speed of 24,000 RPM. The angular contact ball bearings inside the spindle are among the most failure-prone components and therefore constitute the focus of this research, as they are particularly failure-prone under variable load conditions. To address the “black-box” limitations of conventional models, the experimental framework adopts a four-layer architecture for real-time, explainable spindle health monitoring.

As illustrated in Figure 1, vibration signals are acquired through NI-based hardware and processed on an AMD Ryzen™-based industrial edge AI PC. This edge computing layer is specifically configured to handle high-frequency data (25.6 kS/s) and perform statistical feature extraction within strict 1 s intervals, ensuring that the subsequently constructed knowledge graph provides a traceable and timely evidence chain for maintenance decisions.




Figure 1. Experiment environment.

2.1.2. Data Acquisition Architecture

To enable reliable spindle health monitoring under real machining conditions, we deployed a three-layer data acquisition architecture comprising a sensor layer, a data acquisition layer, and an edge computing layer. The sensor layer captures spindle vibration and motor load variations using tri-axial IEPE accelerometers and non-contact Hall-effect current sensors, respectively. High-resolution vibration data are digitized by an NI cDAQ-9174 platform with an NI-9234 module (24-bit ADC, 25.6 kS/s per channel), providing a sufficient dynamic range and bandwidth to represent both steady-state and impulsive degradation signatures.

To ensure real-time processing capabilities at the high sampling rate of 25.6 kS/s, the edge computing layer utilizes an AMD Ryzen™-based industrial edge AI PC. We implemented a sliding window size of 12,800 samples (0.5 s) with a 50% overlap (see Table 1). This specific configuration balances frequency resolution and computational load, ensuring that the Kolmogorov–Smirnov (KS) test and feature extraction are completed well within the 1 s aggregation interval.

Table 1. Data acquisition and edge computing system configuration.

Layer	Component	Model Platform	Key Specifications
Sensor Layer	IEPE accelerometers (tri-axial)	PCB Piezotronics 356A15	Sensitivity: 100 mV/g; frequency response: 0.5–5 kHz
	Hall-effect current sensors	(Hall-effect, non-contact)	Non-contact measurement; three-phase load monitoring
Data Acquisition Layer	DAQ chassis	NI cDAQ-9174	Modular chassis for multi-channel acquisition
	Vibration acquisition module	NI-9234	ADC: 24-bit; sampling rate: 25.6 kS/s per channel
Edge Computing Layer	Industrial edge AI PC	AMD Ryzen™-based	Real-time computation; 0.5 s window with 50% overlap

2.1.3. Dataset Description

After edge-side preprocessing, the raw high-frequency vibration signals were aggregated into statistical feature records at a rate of one sample per second. The dataset consists of time-indexed statistical features and corresponding state labels (Table 2).


Table 2. Description of the experimental dataset fields.

Name	Symbol	Data Type	Physical Meaning
Timestamp	Time	Time series	Index of data sampling time
State Label	Label/HRC	Categorical	Label = 1 (normal), Label = 3 (abnormal)
Mean Energy	Z_{mean}	Numerical	Overall vibration energy level
Maximum Peak	Z_{max}	Numerical	Instantaneous impact intensity
Minimum Value	Z_{min}	Numerical	Lower bound of waveform amplitude
Skewness	Z_{skew}	Numerical	Asymmetry of signal distribution
Kurtosis	$Z_{kurtosis}$	Numerical	Impulsiveness and shock severity

2.2. Proposed Methodology

This study proposes an explainable knowledge graph construction framework composed of four main stages (Figure 2):

1. Dynamic temporal segmentation;
2. Statistical feature discretization;
3. Knowledge graph construction;
4. Explainable attribution reasoning stage.

**Figure 2.** Overview of the proposed framework.

The overall objective is to transform a continuous sensor time series into semantically meaningful graph entities and relations, enabling both state traceability and interpretable fault attribution.

2.2.1. Dynamic Temporal Segmentation

Because raw sensor features are continuous time-series signals, they cannot be directly represented as discrete and reasoning-ready nodes in a knowledge graph. We therefore adopt a two-level segmentation strategy that integrates supervised label transitions (coarse boundaries) and unsupervised statistical drift detection (fine boundaries within normal operation).

Let the normalized feature sequence be $X = \{x_t\}_{t=1}^T$, $x_t \in \mathbb{R}^d$, where $x_t \in \mathbb{R}^d$ is the d -dimensional feature vector at time t .

Let the corresponding label sequence be $L = \{l_t\}_{t=1}^T$, $l_t \in \{1, 3\}$ (1: normal; 3: abnormal), where 1 indicates normal (good) operation and 3 indicates abnormal (bad) operation.

Level 1 segmentation via label transition

A coarse boundary is detected whenever the label changes between consecutive timestamps:

$$l_t \neq l_{(t-1)} \quad (1)$$

All detected transition points partition the sequence into coarse segments, $S_{coarse} = \{s^1, s^2, \dots, s^N\}$, which separate macro-level normal and abnormal operational intervals.

Level 2 segmentation (statistical drift within normal segments). For each normal segment (Label = 1), a sliding window Kolmogorov–Smirnov (KS) test is employed to detect internal distributional shifts. Let W_{ref} and W_{test} denote the reference and test windows, respectively. For a scalar feature z , with empirical cumulative distribution functions (ECDFs) $F_{ref}(z)$ and $F_{test}(z)$, the KS statistics are:

$$D_{KS} = \sup_z |F_{ref}(z) - F_{test}(z)| \quad (2)$$

A drift boundary is inserted if:

$$D_{KS} > D_{critical} \text{ and } p_value < 0.05 \quad (3)$$

where the significance level α is set to 0.05. $D_{critical}$ is calculated as $\frac{1.36}{\sqrt{n}}$ for large sample sizes, where n is the effective window sample size. This statistical rigor ensures that segmentation is driven by significant distributional changes rather than random noise.

Graph instantiation of segments

Each final segment s_j is instantiated as a graph node (: *StateSegment*) with properties $Start_{Idx}$, End_{Idx} , and Label.

2.2.2. Statistical Feature Discretization

To bridge the gap between continuous numerical signals and high-level semantics, this study transforms raw statistical features (e.g., a Kurtosis value of 2.5) into discrete, human-readable semantic tags (e.g., *Kurtosis_{High}*). This process ensures that the subsequent knowledge graph construction captures physically meaningful states rather than noisy fluctuations. We employ a Global Baseline Profiling strategy to define machine-specific healthy boundaries based on empirical data.

Global Baseline Profiling (normal reference)

Using all samples labeled as normal (Label = 1), for each feature f we compute, $P(25)^{f(25th\ percentile)}$ and $P(75)^{f(75th\ percentile)}$. The choice of the Interquartile Range (IQR) as the robust baseline boundary is justified by the non-stationary and non-Gaussian nature of industrial spindle data. Unlike mean-variance thresholds ($\mu \pm 3\sigma$), which are sensitive to outliers and assume a normal distribution, percentile-based thresholds satisfy the requirement for robust statistics, effectively distinguishing nominal variability from genuine anomalies in complex machining environments.

Segment-wise aggregation

Following Level 1 and Level 2 segmentation, each resulting segment S_i is summarized by the mean value of its constituent features. For a specific feature f , the segment-wise mean $\mu_{i,f}$ is calculated as:

$$\mu_{i,f} = \frac{1}{|S_i|} \sum_{t \in S_i} z_{t,f} \quad (4)$$

where $|S_i|$ denotes the segment length and $z_{t,f}$ represents the feature value at time t .

Semantic mapping rules

A discrete feature node is generated only when a segment deviates beyond the healthy percentile band:

High deviation : If $\mu_{i,f} > P(75)^f$, a semantic tag $Tag_{i,f} = \text{High}(f)$ is generated. (5)

Low deviation : If $\mu_{i,f} < P(25)^f$, a semantic tag $Tag_{i,f} = \text{Low}(f)$ is generated. (6)

- Stable range:

If $P(25)^f \leq \mu_{i,f} \leq P(75)^f$, no feature node is created, as the variation is treated as negligible background fluctuation. (7)

This design suppresses non-informative variability and retains only diagnostically salient deviations.

2.2.3. Knowledge Graph Construction

A knowledge graph is constructed in Neo4j by defining machine lifecycle entities and their relations as a semantic network.

Node types

- StateSegment: Operational segment node representing a time interval with a stable state.
- TransitionEvent: Event node indicating state switching (e.g., failure, repair) with Property Event Type.
- MetricFeature: Discretized semantic feature node produced by Section 2.2.2.

Edge definitions.

- Temporal lifecycle links:

These relations preserve ordered machine evolution:

$$(: \text{StateSegment}) - [:\text{ENDED}_{BY}] \rightarrow (: \text{TransitionEvent}) \quad (8)$$

$$(: \text{TransitionEvent}) - [:\text{STARTS}] \rightarrow (: \text{StateSegment}) \quad (9)$$

- Feature association links:

These relations bind a segment to its salient semantic deviations:

$$(: \text{StateSegment}) - [:\text{HAS}_{FEATURE}] \rightarrow (: \text{MetricFeature}) \quad (10)$$

- Feature co-occurrence links:

If two semantic features frequently appear in the same abnormal segments, a co-occurrence edge is added:

$$(: \text{MetricFeature}) - [:\text{CO}_{OCCURS}_{WITH}] \rightarrow (: \text{MetricFeature}) \quad (11)$$

2.2.4. Attribution Path Reasoning

The knowledge graph's primary value is explainable inference. To answer "why did the machine fail?" we compute an attribution (reasoning) path linking a failure event to its immediately preceding operational evidence.

Objective: Given a target failure event node E_{fail} , identify the most informative evidence in the predecessor segment S_{prev} and its semantic features.

Reasoning procedure:

5. Select a target failure event node E_{fail}
6. Trace back to the predecessor segment S_{prev} , satisfying:

$$S_{prev} - [\text{ENDED}_{BY}] \rightarrow E_{fail} \quad (12)$$

7. Retrieve all features connected to S_{prev} :

$$F = f|(S_{prev} - [HAS_{FEATURE}] \rightarrow f) \quad (13)$$

8. Map feature set F to a failure mode M using an external expert rule base or domain ontology.

Explanation output format (example):

$$\begin{aligned} & [Event]abnormalshutdown \\ & \leftarrow [Predecessorstate]segment \\ & \leftarrow [Salientevidence]Kurtosis_{High}(\text{exceeds healthy baseline}) \end{aligned}$$

By construction, each explanation is grounded in explicit segment boundaries, baseline-derived thresholds, and graph-traceable relations, ensuring mathematical clarity and reproducibility from raw signals to final interpretability.

3. Results

This section reports the empirical results of the proposed machine health monitoring framework. Following the hierarchical logic of microscopic signal dynamics $\rightarrow$ macroscopic statistical distributions $\rightarrow$ semantic knowledge graph reasoning, we validate the effectiveness of the approach through both quantitative metrics and qualitative reasoning paths. In response to the reviewers' suggestions, we first establish the framework's reliability using a Quantitative Performance Matrix, followed by an analysis of temporal transition detection, discriminative feature behaviors, and explainable attribution paths.

3.1. Temporal Dynamics and State Transition Analysis

We first examine the machine's operational evolution from normal to failure in the time domain. A continuous monitoring record of 35,000 s is analyzed to characterize both global and local transition behaviors. To validate the effectiveness of the proposed two-level temporal segmentation and diagnostic framework, we first establish its reliability using a Quantitative Performance Matrix (Table 3). These metrics demonstrate the system's high precision in identifying state boundaries and distinguishing between stable baselines and abnormal events.

Table 3. Quantitative Performance Matrix.

Metric	Value	Description
Fault Detection Accuracy	84.97%	Correctly identified normal/abnormal states
False Alarm Rate (FAR)	3.43%	Low rate of false positives in stable regions
Segmentation Accuracy	>95%	Boundaries aligned with labeled transitions
Avg. Detection Delay	~0.5s	Processing within 12,800 sample window

The results in Table 3 indicate that the framework achieves a high degree of robustness, with detection delays well within operational safety margins (approximately 0.5 s). This quantitative foundation supports the qualitative observations presented in the following temporal and spectral analyses.

3.1.1. Global Trend Observation

Figure 3 shows the global temporal profile of the mean vibration energy indicator, (Z_{mean}), over the entire monitoring record (x-axis in record count). Overall, the signal remains near a stable baseline (approximately $Z_{mean} \approx 1.1$) for most of the observation period, indicating that the spindle operates under predominantly steady conditions.




Figure 3. Global time-series trajectory of mean vibration energy, Z_{mean} .

Notably, multiple short-duration bursts are observed, where Z_{mean} exhibits sharp impulsive peaks (up to approximately 1.8). These bursts occur intermittently across the record (e.g., around 2–4k, 9–11k, 13–16k, and 20–21k in record count), suggesting transient high-energy events rather than a single sustained mean shift. The red-shaded windows highlight the segments identified as abnormal/transition intervals by the proposed segmentation scheme, which coincides with elevated peak density and increased local variability.

In summary, the global trend indicates a “stable baseline with intermittent abnormal bursts” pattern, supporting the necessity of event-oriented temporal segmentation to localize fault onset and extract reasoning-ready state segments for subsequent knowledge graph construction.

3.1.2. Micro-Level Zoom-In of Transition Dynamics

To validate the precision of the proposed segmentation strategy, a zoomed view of the “normal $\rightarrow$ failure” boundary is examined (Figure 4). The vertical dashed line denotes the automatically detected cut point.




Figure 4. Zoomed transition region and automatic temporal segmentation result.

- **Segmentation performance:** Within the 0.5 s window preceding the transition, Z_{mean} shows a slight rise but remains within short-term fluctuation. A new segment boundary is only inserted when the signal exhibits persistent distributional drift, indicating that the method avoids reacting to transient noise and instead captures the true onset of failure-relevant instability.

3.2. Statistical Feature Distribution and Correlation Analysis

After segmentation, we analyze the distributions of five key statistics (mean, max, min, skewness, and kurtosis) under different health states to identify features with strong diagnostic relevance.

3.2.1. Boxplot-Based Distribution Comparison

Boxplots are used to compare the normal state (Label = 1) and abnormal state (Label = 3).

1. Kurtosis indicator ($Z_{kurtosis}$):

- Observation: The abnormal group exhibits a substantially higher median and a pronounced long-tail distribution (Figure 5).
- Physical interpretation: Kurtosis reflects the degree of impulsiveness. Rolling-element bearing defects (e.g., spalling) often generate periodic impulse-like components, leading to a rapid increase in kurtosis. This result indicates that $Z_{kurtosis}$ is among the most discriminative indicators in the considered scenario.




Figure 5. Zoomed transit Boxplot comparison of $Z_{kurtosis}$ under normal vs. abnormal states.

2. Peak/impact indicator (Z_{max}):

- Observation: Z_{max} also shows a clear separation between normal and abnormal regimes (Figure 6), suggesting the presence of instantaneous high-energy releases during abnormal operation.




Figure 6. Boxplot comparison of Z_{max} under normal vs. abnormal states.

3.2.2. Feature Clustering Scatter Analysis

To study multivariate relationships, we plot the two-dimensional scatter distribution using Z_{mean} (x-axis) and $Z_{kurtosis}$ (y-axis) (Figure 7).




Figure 7. Scatter distribution of Z_{mean} vs. $Z_{kurtosis}$ for normal and abnormal samples.

1. Normal cluster (green points): Normal samples are concentrated around ($Z_{mean} \approx 0$, $Z_{kurtosis} \approx 0$), indicating high consistency under healthy operation.
2. Abnormal dispersion (red points): Abnormal samples spread across multiple regions, revealing heterogeneous failure mechanisms, including:
 - Type A: High Z_{mean} with low $Z_{kurtosis}$ (lower right). This pattern is consistent with unbalance, which increases the vibration energy but does not necessarily produce impulsive shocks.
 - Type B: High $Z_{kurtosis}$ with low Z_{mean} (upper left). This pattern matches early bearing damage, where impulsive shocks appear before the total energy level increases.

This finding implies that a single “abnormal” label may contain multiple underlying physical causes, motivating the use of a knowledge graph for fine-grained semantic attribution.

3.3. Results of Explainable Knowledge Graph Construction

By integrating the segmented states and discretized semantic features into the ontology-driven graph generation engine, the first version of the machine health knowledge graph is constructed.

3.3.1. Graph Statistics

The graph generated contains:

- Total nodes: 43 (including operational state segments and extracted semantic feature nodes).
- Total edges: 76 (including temporal relations and feature association relations).

These statistics indicate that the framework compactly represents machine lifecycle evolution while preserving interpretable feature evidence. As illustrated in the growth analysis in Figure 8 over the 35,000-sample lifecycle, the accumulation of nodes and edges exhibits a distinct phased behavior driven by the machine’s health state. During stable operational phases, the graph expands linearly at a gradual pace, reflecting the routine

addition of sequential state segments and basic temporal relations. However, within anomalous intervals (highlighted by the shaded regions), sharp bursts in structural growth occur. Notably, the cumulative edge count increases at a significantly steeper rate than the node count during these critical periods. This divergence highlights the system’s dynamic capability to densify semantic interconnectedness, efficiently linking multiple diagnostic features to sudden health degradations without unnecessarily inflating the overall graph size during nominal operations.




Figure 8. Knowledge graph growth analysis.

3.3.2. Representative Nodes and Relations

Tables 4 and 5 provide representative examples illustrating how numerical patterns are transformed into semantic entities and relations.

Table 4. Representative knowledge graph nodes.

ID	Node Type	Key Attributes	Semantic Interpretation
Seg_{15}	StateSegment	Label = 1, Mean = 0.2	Stable operation segment
Seg_{16}	StateSegment	Label = 3, Mean = 2.5	Severe abnormal segment
Feat_K_Hi	MetricFeature	Metric = Kurtosis, Level = High	High-impulse feature node
Evt_Fail	TransitionEvent	EventType = Breakdown	Failure (shutdown) event

Table 5. Representative knowledge graph relations.

ID	Node Type	Key Attributes	Semantic Interpretation
Seg_{15}	ENDED_BY	Evt_Fail	Normal segment terminates at failure
Seg_{16}	HAS_FEATURE	Feat_K_Hi	Segment exhibits high kurtosis evidence

A notable observation is that two consecutive fault-related segments may be annotated differently (e.g., one segment with $Peak_{High}$ and another with $Energy_{High}$), reflecting a plausible physical evolution in which impulsive shocks occur before sustained energy escalation.

3.3.3. Graph Visualization

Figure 9 shows the global Neo4j visualization. A time-ordered backbone (green/red segment chain) is clearly observed, while semantic feature nodes radiate from each segment. This topology preserves the temporal structure and explicitly exposes state–evidence links, providing a direct interface for explainable reasoning.




Figure 9. Overview visualization of the machine health knowledge graph.

3.4. Automated Attribution Case Studies

To demonstrate explainability, we query two representative reasoning cases from the constructed graph: one for abnormal shutdown attribution and another for interpretable variability under nominal operation.

3.4.1. Case 1: Failure Attribution

Scenario: The machine experiences an unplanned shutdown at $t = 1024$ s.

Queried reasoning path:

Event(Failure) $\leftarrow$ $Seg_{16}(Bad) \rightarrow$ HAS_FEATURE $\rightarrow$ Kurtosis_High $\rightarrow$ SUGGESTS $\rightarrow$ Bearing_Damage

Generated diagnostic explanation:

The system detects an abnormal shutdown event. Tracing back to the predecessor segment (Seg_{16}), the dominant evidence is high kurtosis ($Kurtosis_{High}$), which indicates strong impulsive behavior. According to the domain knowledge base, this pattern is consistent with early rolling-element bearing damage. Therefore, inspection of the front spindle bearing is recommended as the primary maintenance action.

3.4.2. Case 2: Explainable Variability Within Normal Operation

Scenario: The machine remains labeled as normal (Label = 1), but a secondary warning is triggered.

Queried reasoning path (Figure 10):

Seg_8 (Normal) $\rightarrow$ HAS_FEATURE $\rightarrow$ Z_mean_Rising (Sub-threshold)




Figure 10. Zoom-in view of an attribution path for a representative failure event.

Generated diagnostic explanation:

Although the current segment remains within the acceptable normal band, the mean vibration energy shows an increasing trend. This behavior may be induced by operational factors such as increased feed rate or mild tool wear. Immediate shutdown is not required; however, the segment is recommended for enhanced monitoring.

Overall, the results validate three key points:

1. The proposed segmentation accurately localizes transition onsets while suppressing transient noise.
2. Discriminative statistics such as $Z_{kurtosis}$ and Z_{max} capture physically meaningful fault signatures.
3. The knowledge graph provides structured, queryable, and human-interpretable attribution paths beyond binary classification, supporting actionable diagnostics and maintenance recommendations.

4. Discussion and Conclusions

4.1. Discussion

This study addresses a core challenge in smart manufacturing and Industrial Internet of Things (IIoT): how to transform large-scale machine sensing data into decision-relevant and trustworthy diagnostic knowledge. While deep learning approaches have achieved strong fault classification accuracy, their black-box nature limits interpretability, reduces shop-floor trust, and constrains practical adoption. To overcome these limitations, we developed an explainable knowledge graph construction framework based on temporal statistical feature extraction and validated it using real-world machine operation data.

To further demonstrate the competitive advantages of our approach, we compared the proposed Explainable KG framework with conventional thresholding and state-of-the-art deep learning models (see Table 6).

Table 6. Comparison of diagnostic approaches in spindle health monitoring.

Feature	Traditional Thresholding	Deep Learning (CNN/RNN)	Proposed Explainable KG
Interpretability	Medium (Simple Logic)	Low (Black-box)	High (Semantic Graph Paths)
Adaptability	Low (Fixed Limits)	Medium (Retraining Needed)	High (Dynamic Drift Detection)
Context Awareness	None	Implicit	Explicit (State–Event–Evidence)
False Alarm Rate	High (Sensitive to Noise)	Low	Low (3.43%)

- Semantic structuring of continuous time series:

A key methodological contribution is the proposed dynamic temporal segmentation strategy, which converts continuous sensor streams into discrete and semantically meaningful state entities (state segments). Unlike fixed-length sliding windows that may mix heterogeneous states or fragment the fault evolution process, the proposed two-level mechanism integrates (i) label transition detection and (ii) statistical drift monitoring. As shown in the results, the method accurately localizes the onset of abnormality (e.g., the structural break around $t = 1000$ s), while also capturing subtle pattern variations within nominal operation. In addition, the global percentile-based discretization maps raw statistical values into human-readable semantic symbols, bridging the gap between numeric measurements and symbolic engineering knowledge. For a feature f , the baseline band defined by $P(25)^f$ and $P(75)^f$ enables the system to highlight deviations of operational significance rather than normal fluctuations.

- From classification outputs to explainable attribution paths:

Beyond binary decisions or probabilistic outputs, the proposed framework provides structured reasoning paths that explicitly connect an event to its predecessor segment and salient feature evidence and further to a hypothesized fault mode. The case studies demonstrate that the system can generate actionable explanations (e.g., linking high kurtosis and high peak impact to bearing damage), aligning with expert diagnostic logic: phenomenon (data) $\rightarrow$ evidence (semantic features) $\rightarrow$ cause (knowledge). This form of transparency supports targeted maintenance actions and reduces the risk of blind part replacement.

- Flexibility and maintainability via a two-layer design:

The separation between a data layer (objective historical operational traces) and a knowledge layer (abstract fault patterns and causal rules) improves system extensibility. When new fault types or updated expert knowledge become available, the knowledge layer can be revised without retraining the entire data-driven pipeline. This decoupling is beneficial for long-term maintenance, cross-machine transfer, and iterative deployment in industrial environments.

4.2. Conclusions

This work proposes and validates an explainable knowledge graph framework for machine health monitoring. The main conclusions are:

1. The proposed temporal segmentation effectively transforms continuous sensor streams into discrete state segments with meaningful boundaries, enabling the graph-based modeling of machine lifecycle transitions.
2. Percentile-based discretization successfully converts statistical indicators into interpretable semantic features, retaining diagnostically salient deviations while suppressing non-informative fluctuations.
3. The constructed knowledge graph supports explainable attribution through reasoning paths, allowing the system to provide diagnostic evidence and maintenance-oriented recommendations beyond conventional classifiers.
4. The two-layer (data/knowledge) architecture offers practical extensibility and maintainability, supporting updates to domain knowledge without requiring re-training of the entire pipeline.

Overall, the proposed approach provides a feasible pathway to deploy explainable, traceable, and decision-supportive machine health analytics in smart manufacturing settings.

4.3. Limitations

Despite the validated feasibility, several limitations remain:

1. Static thresholds: The current discretization relies on global percentiles, which may become suboptimal under aging drift, potentially increasing false alarms over long-term operation.
2. Single-modality sensing: The present evaluation primarily uses vibration-derived statistics; performance may be limited for electrical or thermal faults that are weakly reflected in vibration features.
3. Knowledge base coverage: The current inference relies on predefined rules (e.g., high kurtosis implies bearing damage). For unseen composite failure modes, explanations may be incomplete or misleading.

4.4. Future Perspectives

Future work will focus on three directions to strengthen robustness and applicability:

1. Multi-modal data fusion: Extend the graph to incorporate current signatures, acoustic emission, and CNC controller context variables (e.g., spindle load, feed override) to distinguish operating condition changes from genuine faults and to improve coverage across fault categories.
2. Adaptive learning and dynamic thresholds: Introduce online learning or adaptive filtering to update baseline bands over time, enabling the dynamic adjustment of $P(25)^f$ and $P(75)^f$ under aging drift. Unsupervised discovery (e.g., DBSCAN) can further identify emerging patterns beyond predefined statistics.
3. LLM-assisted knowledge acquisition: Use large language models to extract entities and relations from maintenance manuals, technical reports, and historical work orders, enabling semi-automatic expansion of the knowledge layer and supporting self-evolving machine health management.

These extensions are expected to improve long-term stability, broaden fault coverage, and reduce the knowledge engineering burden, thereby enabling scalable and trustworthy deployment in real production environments.

Author Contributions: Conceptualization, C.-S.C.; Methodology, C.-S.C.; Software, C.-S.C.; Validation, C.-S.C.; Writing—original draft, C.-S.C.; Visualization, C.-S.C.; Supervision, G.-J.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Due to confidentiality considerations related to the experimental setup and industrial data sources, the datasets generated and analyzed during the current study are not publicly available but can be made available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kagermann, H.; Wahlster, W.; Helbig, J. *Recommendations for Implementing the Strategic Initiative INDUSTRIE 4.0: Securing the Future of German Manufacturing Industry*; Final Report of the Industrie 4.0 Working Group; acatech—National Academy of Science and Engineering: Munich, Germany, 2013; pp. 1–84.
2. Lee, J.; Bagheri, B.; Kao, H.-A. A Cyber-Physical Systems Architecture for Industry 4.0-Based Manufacturing Systems. *Manuf. Lett.* **2015**, *3*, 18–23. [\[CrossRef\]](#)
3. Jardine, A.K.S.; Lin, D.; Banjevic, D. A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance. *Mech. Syst. Signal Process.* **2006**, *20*, 1483–1510. [\[CrossRef\]](#)
4. Lee, J.; Wu, F.; Zhao, W.; Ghaffari, M.; Liao, L.; Siegel, D. Prognostics and Health Management Design for Rotary Machinery Systems—Reviews, Methodology and Applications. *Mech. Syst. Signal Process.* **2014**, *42*, 314–334. [\[CrossRef\]](#)
5. Randall, R.B.; Antoni, J. Rolling Element Bearing Diagnostics—A Tutorial. *Mech. Syst. Signal Process.* **2011**, *25*, 485–520. [\[CrossRef\]](#)
6. Lei, Y.; Yang, B.; Jiang, X.; Jia, F.; Li, N.; Nandi, A.K. Applications of Machine Learning to Machine Fault Diagnosis: A Review and Roadmap. *Mech. Syst. Signal Process.* **2020**, *138*, 106587. [\[CrossRef\]](#)

7. Mobley, R.K. *An Introduction to Predictive Maintenance*; Butterworth-Heinemann: Oxford, UK, 2002.
8. Li, W.; Huang, R.; Li, J.; Liao, Y.; Chen, Z.; He, G.; Yan, R.; Gryllias, K. A Perspective Survey on Deep Transfer Learning for Fault Diagnosis in Industrial Scenarios: Theories, Applications and Challenges. *Mech. Syst. Signal Process.* **2022**, *167*, 108487. [[CrossRef](#)]
9. LeCun, Y.; Bengio, Y.; Hinton, G. Deep Learning. *Nature* **2015**, *521*, 436–444. [[CrossRef](#)]
10. Jia, F.; Lei, Y.; Lin, J.; Zhou, X.; Lu, N. Deep Neural Networks: A Promising Tool for Fault Characteristic Mining and Intelligent Diagnosis of Rotating Machinery with Massive Data. *Mech. Syst. Signal Process.* **2016**, *72–73*, 303–315. [[CrossRef](#)]
11. Lei, Y.; Lin, J.; He, Z.; Zuo, M.J. A Review on Empirical Mode Decomposition in Fault Diagnosis of Rotating Machinery. *Mech. Syst. Signal Process.* **2013**, *35*, 108–126. [[CrossRef](#)]
12. Zhao, R.; Yan, R.; Chen, Z.; Mao, K.; Wang, P.; Gao, R.X. Deep Learning and Its Applications to Machine Health Monitoring: A Survey. *Mech. Syst. Signal Process.* **2019**, *115*, 213–237. [[CrossRef](#)]
13. Samek, W.; Wiegand, T.; Müller, K.-R. Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models. *ITU J. ICT Discov.* **2017**, *1*, 39–48.
14. Adadi, A.; Berrada, M. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access* **2018**, *6*, 52138–52160. [[CrossRef](#)]
15. Lipton, Z.C. The Mythos of Model Interpretability. *Queue* **2018**, *16*, 31–57. [[CrossRef](#)]
16. Rudin, C. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. [[CrossRef](#)]
17. Ehrlinger, L.; Wöß, W. Towards a Definition of Knowledge Graphs. In *Proceedings of the SEMANTiCS (Posters, Demos, SuCCESS)*; RWTH Aachen University: Aachen, Germany, 2016; pp. 1–4.
18. Paulheim, H. Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods. *Semant. Web* **2017**, *8*, 489–508. [[CrossRef](#)]
19. Ji, S.; Pan, S.; Cambria, E.; Marttinen, P.; Yu, P.S. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Trans. Neural Netw. Learn. Syst.* **2022**, *33*, 494–514. [[CrossRef](#)]
20. Widodo, A.; Yang, B.-S. Support Vector Machine in Machine Condition Monitoring and Fault Diagnosis. *Mech. Syst. Signal Process.* **2007**, *21*, 2560–2574. [[CrossRef](#)]
21. Si, X.-S.; Wang, W.; Hu, C.-H.; Zhou, D.-H. Remaining Useful Life Estimation—A Review on the Statistical Data Driven Approaches. *Eur. J. Oper. Res.* **2011**, *213*, 1–14. [[CrossRef](#)]
22. Wang, Q.; Mao, Z.; Wang, B.; Guo, L. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Trans. Knowl. Data Eng.* **2017**, *29*, 2724–2743. [[CrossRef](#)]
23. Noy, N.F.; Gao, Y.; Jain, A.; Narayanan, A.; Patterson, A.; Taylor, J. Industry-Scale Knowledge Graphs: Lessons and Challenges. *Commun. ACM* **2019**, *62*, 36–43. [[CrossRef](#)]
24. Zhang, Y.; Xiong, R.; He, H.; Pecht, M.G. Long Short-Term Memory Recurrent Neural Network for Remaining Useful Life Prediction of Lithium-Ion Batteries. *IEEE Trans. Veh. Technol.* **2018**, *67*, 5695–5705. [[CrossRef](#)]
25. Li, X.; Zhang, W.; Ding, Q.; Sun, J.-Q. Multi-Layer Domain Adaptation Method for Rolling Bearing Fault Diagnosis. *Signal Process.* **2019**, *157*, 180–197. [[CrossRef](#)]
26. Arrieta, A.B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; Garcia, S.; Gil-Lopez, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.