Semantic Web 1 (2025) 1–4
IOS Press

Scalable Knowledge Representation for Fault Diagnosis of Cyber-Physical Systems: A Systematic Literature Review

Ameneh Naghdipour [a,*], Benno Kruit [a], Stefan Schlobach [a] and Jieying Chen [a]
[a] *Department of Computer Science, Vrije Universiteit Amsterdam, De Boelelaan 1111, 1081 HV Amsterdam, The Netherlands*
E-mails: a.naghdipour@vu.nl, b.b.kruit@vu.nl, k.s.schlobach@vu.nl, j.y.chen@vu.nl

Editors: First Editor, University or Company name, Country; Second Editor, University or Company name, Country
Solicited reviews: First Solicited Reviewer, University or Company name, Country; Second Solicited Reviewer, University or Company name, Country
Open reviews: First Open Reviewer, University or Company name, Country; Second Open Reviewer, University or Company name, Country

Abstract. Fault diagnosis in Cyber-Physical Systems (CPS) is essential for minimizing downtime, ensuring operational safety, and improving system resilience. As CPSs become increasingly interconnected and complex, traditional diagnostic methods struggle to capture their dynamic interactions and dependencies. Semantic technologies, including knowledge graphs and ontologies, provide mechanisms for structured representation, integration, and reasoning over diverse sources of diagnostic knowledge. This paper provides a comprehensive review of semantic approaches for fault diagnosis in CPS through a Systematic Literature Review (SLR). It covers key stages such as knowledge acquisition from domain experts, knowledge extraction from documents, semantic modeling of domain knowledge and data, and model enhancement. To the best of our knowledge, no prior systematic literature review has covered all these critical aspects. Unlike previous reviews, we systematically analyze and categorize the findings related to each stage. Additionally, we explore the role of available manufacturing data sources and their integration with semantic models. By bridging the gap between fault diagnosis and semantic technologies, this work highlights the potential of semantic representations to enhance interpretability, interoperability, and automation in CPS fault detection. We further discuss open challenges and outline future research directions, emphasizing the role of semantic frameworks in advancing intelligent fault diagnosis. Our findings aim to guide researchers and practitioners in leveraging semantic technologies for more robust and explainable fault diagnosis in CPS, while also highlighting existing research gaps and outlining directions for future work.

Keywords: Fault diagnosis, Semantic representation, Knowledge graph, Ontology, Information extraction, Knowledge acquisition, Model enhancement, Cyber-physical system, Literature review

*Corresponding author. E-mail: a.naghdipour@vu.nl.

1570-0844/$35.00 © 2025 – IOS Press. All rights reserved.

1. Introduction

Over the past decades, many industrial systems have become fully integrated with *Cyber-Physical Systems* (CPS) [57]. Lee [51] defines CPS as the integration of calculation and physical process, which involves embedded computer and networks monitoring and controlling the physical processes. Physical processes affect computations by feedback loops and vice versa. Reliability and correct functioning pose particular challenges in this context [50]. Faults and degradations can cause systems to go down or underperform for extended periods of time, resulting in failures to provide service and unscheduled maintenance. The significant costs associated with such events are not only direct (missed performance, spare parts, and labor costs) but also indirect, such as schedule disruption and reputation loss. In one estimate from 2022, unplanned downtime now costs Fortune Global 500 companies 11% of their yearly turnover: almost $1.5tn, up from $864bn (8% of turnover) two years earlier [72]. Therefore, it is imperative that maintenance engineers discover and understand the nature of these faults and degradation quickly and efficiently.

However, providing effective support for fault diagnosis in modern CPS remains challenging for several reasons. First, the existence of large volumes of fragmented, lengthy, and unstructured data that can not be efficiently used to guide maintenance engineers in performing*Fault Diagnosis* (FD). Second, hidden knowledge that has not been documented. Addressing these challenges can be supported by *Knowledge Representation* (KR): the encoding of information in a form that allows computer systems to perform complex problem solving, such as through logical or statistical reasoning [84]. However, the scale and complexity of modern CPS have significant implications for the KR techniques suitable for FD in practice [18]. It is no longer feasible to manually create formalized knowledge structures, and established diagnosis techniques may not be able to handle the complexity of modern systems. Moreover, automating and scaling up the creation of knowledge repositories invariably introduces noise, which demands a higher robustness of diagnosis methods. Additionally, in the application of knowledge-based diagnosis methods, changes and improvements to the system are common, which means the diagnosis must adapt to these changes.

These challenges have motivated the development of a wide range of fault diagnosis approaches, evolving alongside advances in artificial intelligence, modeling techniques, and data-driven methods. Computer-aided diagnosis was among the first applications investigated when digital computers became available, and its development has closely followed the scientific progress in the field of artificial intelligence. From the start, one of the main techniques has been symbolic diagnosis (see [58] for a historical overview). Many diagnosis techniques were developed upon the assumption that the entire system could be described and modeled or simulated with high fidelity, which have come to be known as *model-based* or *first principles* systems [5]. Other diagnosis techniques have leveraged high volumes of (sensor) data for diagnosis using statistical (machine learning) techniques [68, 70]. Additionally, many techniques combine data-driven, (physics) model-based, and knowledge-based methods [92]. However, while such techniques can be used successfully to diagnose component-level faults, modern CPS are too complex to model completely with sufficient fidelity. For example, in Industrial Internet-of-Things applications [18], the decentralized interaction of many components produces highly complex behavior, which makes system-level models unfeasible. Therefore, it is important to consider the fundamental requirements for the development of an efficient system-level FD method for CPS, which include scalability, robustness, and evolution. In the following, we elaborate on these three aspects to justify the exploration of semantic representation methods within this paper.

- *Scalability.* This is a vital consideration in fault diagnosis for large and complex systems as their increasing size and complexity lead to increased interdependencies and interconnected components [2]. Therefore, fault diagnosis methods must efficiently scale to address intricate relationships. Rule-based fault diagnosis methods, exemplified by if-then rules, represent some of the earliest efforts to capture and model fault data for diagnostic purposes. These methods require extensive rule definitions and maintenance [24], making them inefficient for systems with multiple interacting factors where a single fault may result from multiple interacting factors. Additionally, any changes demand rule modifications, further limiting scalability. Other methods such as first-principles modeling rely on

precise mathematical equations to diagnose faults. In this method, an explicit mathematical model of the system is developed by laws of physics that govern the process behavior that enable to diagnose the fault by looking into the difference between observed and expected behaviors [38]. However, constructing accurate models for complex systems is difficult and unrealistic due to the need for deep understanding of system behavior and detailed knowledge. [73]. Consequently, researchers have been working towards developing methods that can handle the inherent complexity of CPS. *Knowledge graphs* (KGs) have emerged as a promising solution in this context. They can efficiently model relationships, dependencies, and hierarchies within CPS, providing a scalable framework for managing large and intricate systems. For instance, the Dimension KG [23] exemplifies this capability, containing 32 billion statements by integrating data from 143 million publications, 160 million patents, 30 million datasets, etc. Such resources cannot be built manually, and must therefore be constructed (semi-)automatically.

- *Robustness.* This is a challenge that arises due to the inherent inaccuracies and inconsistencies in real-world data collected from sensors, logs, or manual reports. Noisy or incomplete data hinder the creation of reliable representations. For example, during KG construction, erroneous relationships may be extracted, leading to incorrect connections between faults, causes, and actions. This highlights the need for robust strategies to mitigate these effects. Researchers have been addressing these challenges using different techniques for denoising, de-duplication, and entity resolution [13, 91] that will be discussed in Section 4.2. Additionally, establishing a well-structured schema for the knowledge graph helps prevent the extraction of incorrect entities and relationships.
- *Evolution.* This refers to the adaptability and maintainability of fault diagnosis frameworks to incorporate new data, rules, relationships, or components over time without the need for complete restructuring. The evolution of fault diagnosis systems is a critical aspect in the context of CPS as they encounter changes such as the introduction of new components, sensors, or subsystems, changes in operational conditions or environments, shifts in fault patterns due to aging or upgrading of components, and the adoption of new standards and technologies. Rigid approaches, such as rule-based methods, struggle with adaptability because any system change necessitates extensive rule revisions, a process that becomes increasingly resource-intensive as the system grows. KGs address the challenge of evolution by providing an inherently flexible structure. They can be expanded by adding new nodes (entities) and edges (relationships) without disrupting the existing structure. Moreover, machine learning and reasoning techniques allow KGs to be refined by inferring new and implicit information based on explicit knowledge [36, 94].

In this study, we conduct a *Systematic Literature Review* (SLR) to investigate semantic-based fault diagnosis methodologies. The SLR provides a structured and transparent process for systematically analyzing existing research. The review involves several key steps: formulating research questions to guide the review, constructing a search string to identify relevant studies, and establishing inclusion and exclusion criteria to filter the most relevant literature. This is followed by article identification, selection, and assessment, after which the final selected articles are systematically synthesized.

This paper investigates five research questions covering the main stages of semantic-based fault diagnosis in CPSs: (1) fault knowledge acquisition from maintenance experts, (2) *Information Extraction* (IE) from maintenance documents, (3) fault knowledge representation, (4) diagnosis using the represented knowledge, and (5) enhancement of the constructed knowledge base. These questions were derived from the identified challenges of capturing undocumented expert knowledge and transforming large volumes of fragmented and unstructured fault-related information into scalable knowledge representations that can support fault diagnosis. Accordingly, the review examines the methods and tools employed at each of these stages.

The main contributions of this survey are as follows. First, it provides a systematic review of semantic-based fault diagnosis in CPSs, covering the five aforementioned research questions and develops a structured taxonomy of the methods, techniques, and tools reported across these research questions, supplemented by figures and tables that facilitate easy comparison, as well as a summary of the input sources (fault data) used by existing studies to identify commonly used and underexplored sources. Second, it

identifies methodological gaps in the current literature, highlights emerging research directions and open challenges for advancing semantic-based fault diagnosis in CPSs.

The scope of this review covers studies published between 1995 and 2024 and focuses specifically on semantic-based fault diagnosis in industrial CPSs. We consider approaches that leverage domain knowledge and human-generated maintenance-related information, rather than relying solely on sensor data. Other maintenance tasks, such as predictive maintenance and *Remaining Useful Life* (RUL) estimation are outside the scope of this review. Predictive maintenance aims to anticipate when a failure is likely to occur in order to prevent it, while RUL estimation focuses on determining how long a component can continue to operate before it requires replacement. In contrast, fault diagnosis seeks to identify the causes of observed faults or abnormal conditions after they have occurred.

The rest of the paper is organized as follows: Section 2 presents the related work, covering previous surveys and reviews in the field of fault diagnosis, and highlighting how this review differs from previous studies. Section 3 details the research methodology, including *Research Questions* (RQs), search strategy, database selection, paper selection criteria, and evaluation methods. Section 4 describes the results corresponding to each research question. In Section 5, a detailed discussion is presented that highlights the findings and provides recommendations and future directions. Section 6 provides a brief overview of the development and emerging trends identified in studies published in 2025 and 2026. The limitations of this study are presented in Section 7. Finally, Section 8 concludes the paper.

2. Related Work

Existing literature reviews on fault diagnosis and predictive maintenance have surveyed a range of methods, techniques, and frameworks aimed at improving the identification, reasoning, and handling of faults in complex systems. For example, Xia et al. [95] and Chi et al. [18] explored specific aspects such as graph-based approaches and knowledge-based fault diagnosis in *Industrial Internet of Things* (IIoT) settings. These surveys primarily focus on the capabilities of particular methodologies, such as graph-based semantic inference or the use of ontologies for interoperability and reasoning. Similarly, Chen et al. [17] and Ji et al. [42] delved into knowledge graph reasoning and representation, providing valuable insights into technical applications, but do not address broader challenges like scalability or the evolving nature of semantic models.

A key distinguishing feature of this study is its focus on semantic knowledge representation techniques, particularly in the domain of fault diagnosis for CPS. Unlike previous surveys, which often emphasize specific tools or approaches, this work systematically reviews the entire life-cycle of semantic representation for fault diagnosis. This includes knowledge acquisition, extraction, representation, and enhancement, with a particular emphasis on the practical implications of applying these techniques in highly interconnected and dynamically evolving CPS environments.

Notably, Jieyang et al. [44] and Ucar et al. [83] reviewed AI-based and data-driven predictive maintenance techniques, but lack a comprehensive analysis of semantic representation methods and their unique challenges in the diagnosis of CPS faults. While Liu et al. [56] highlights ontology-based fault diagnosis, it primarily addresses historical developments and integration with machine learning, without discussing the scalability or dynamic adaptation of ontologies over time. Similarly, Gao et al. [27] and Turner et al. [82] emphasize traditional model-based and signal-based fault diagnosis but do not address the emerging challenges posed by semantic representations in the context of CPS.

In contrast, this paper systematically categorizes and analyzes methods for knowledge acquisition from domain experts, information extraction from text, the scalable representation of knowledge, and the enhancement of the KB. Additionally, the study examines methods for fault diagnosis under incomplete knowledge conditions to support maintenance in CPS. By addressing these related aspects, the review provides an overview of current research directions and identifies areas for further investigation in semantic fault diagnosis.

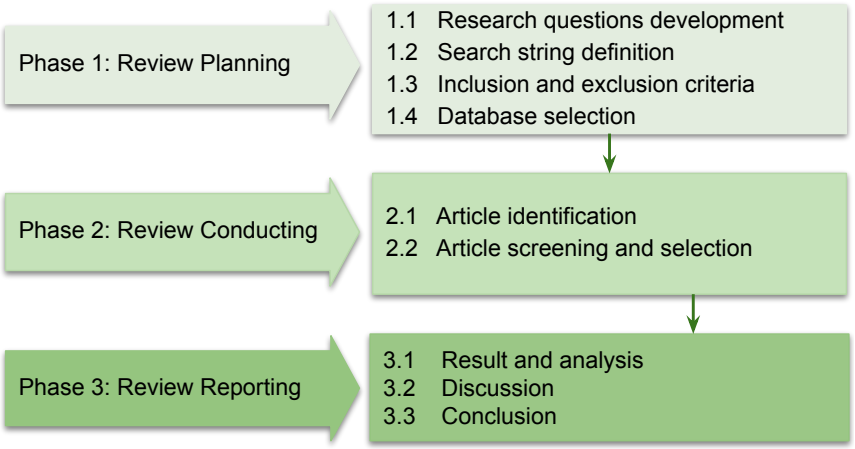


Fig. 1. The steps of the systematic literature review process

3. Systematic Literature Review Methodology

Systematic research methods, such as *Systematic Mapping Studies* (SMS) [65] and SLR [46], are widely used to synthesize and organize existing research. While SMS primarily focuses on identifying research trends and categorizing topics, SLR aims to identify relevant primary studies and systematically extract and analyze data from them. Since the objective of this study is to address specific research questions through analysis of existing literature, we adopt the SLR methodology. Following established SLR guidelines, the review process consists of three main phases: planning, conducting, and reporting, as illustrated in Fig. 1. The following subsections describe the planning and conducting phases of the review.

3.1. ***Systematic Review Planning***

The planning phase defines the protocol and review strategy for conducting the systematic literature review. This includes the formulation of *Research Questions* (RQs), the development of the search string, the establishment of inclusion and exclusion criteria, and the selection of databases used to identify relevant studies. The following subsections describe each of these steps in detail.

3.1.1. ***Research Question Development***

Research questions serve as a guide for conducting a literature review. In this study, the formulation of the research questions is motivated by two key challenges in the field of CPS fault diagnosis. The first challenge concerns the large volume of fragmented, lengthy, and unstructured industrial data, which is difficult to use efficiently for fault diagnosis. The second challenge relates to the domain knowledge that remains

Table 1
Research questions addressed in this systematic literature review

#	Research Question	Keyword
RQ 1	What methods for knowledge acquisition from domain experts are specifically applied to support semantic-based fault diagnosis in CPSs?	Acquisition
RQ 2	What are different methods for extracting knowledge, based on Information Extraction from text documents to support semantic-based fault diagnosis in CPSs?	Extraction
RQ 3	What are different semantic methods for scalable knowledge representation in fault diagnosis of CPSs?	Representation
RQ 4	How can faults be diagnosed based on incomplete knowledge?	Diagnosis
RQ 5	What are different methods for enhancing represented knowledge by domain experts to support semantic-based fault diagnosis in CPSs?	Enhancement

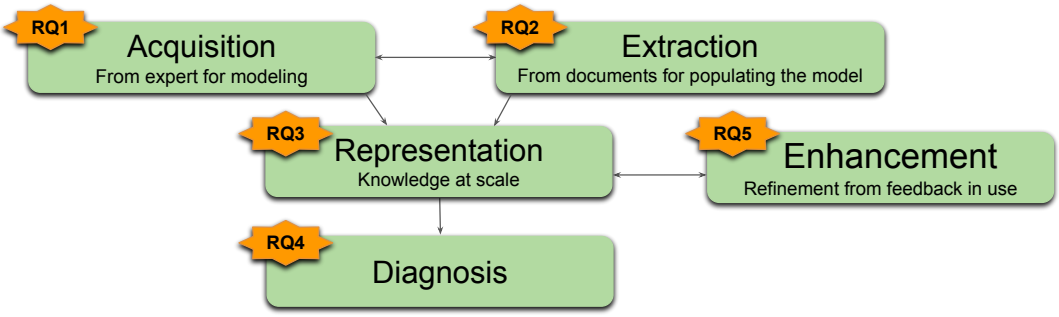


Fig. 2. Visualization of relationships among keywords associated with the research questions

undocumented and is often retained as tacit knowledge held by human operators and engineers. To address these challenges, knowledge graphs were identified as a promising approach for integrating heterogeneous data sources and formalizing expert knowledge within a unified and structured framework. Accordingly, knowledge acquisition, information extraction, and knowledge graph-based representation were considered as complementary processes for transforming heterogeneous and undocumented fault-related knowledge into structured knowledge that can support fault diagnosis. Therefore, the research questions were formulated to investigate the main processes involved in acquiring, extracting, representing, applying, and enhancing knowledge for fault diagnosis in CPS. Specifically, we define five RQs, as presented in Table 1. RQ1 seeks to investigate how fault-related knowledge is acquired from domain experts. RQ2 focuses on methods for extracting information from documented sources. RQ3 examines semantic knowledge representation methods for fault related[1] knowledge. RQ4 explores how structured knowledge representations are utilized for diagnostic purposes. Finally, RQ5 investigates approaches for enhancing and updating the constructed representation.

Fig. 2 illustrates the relationships among keywords associated with the research questions. The process begins with two complementary inputs: knowledge acquisition from experts and information extraction from documents. These inputs jointly support the construction of a structured knowledge representation, which forms the basis for diagnostic reasoning. The knowledge base can subsequently be enhanced and updated as new knowledge and information become available.

3.1.2. ***Search String Definition***

The search string consists of keywords and phrases designed to retrieve relevant literature. We identified key terms and their synonyms by analyzing related work and consulting industry experts.

As illustrated in Fig. 3, the search string is structured around four conceptual elements: research tasks, information sources, semantic representation methods, and application goals. The first element includes keywords related to the four tasks addressed by the research questions: information extraction, knowledge acquisition, knowledge representation, and knowledge enhancement. *Natural Language Processing* (NLP) was also included because it is a broad field encompassing information extraction and is widely used in the literature when referring to techniques for extracting information from text. The task-related keywords were combined with the sources-related keywords such as repair manuals, logs, technical reports, procedures, tacit knowledge, and so on. These sources-related keywords were included because knowledge acquisition, information extraction, and knowledge representation methods are often studied across multiple domains. Restricting the search to maintenance-related sources therefore improved the relevance of

[1] In this review, a fault denotes an abnormal condition of a machine, component, or subsystem that can degrade function, reduce output quality, trigger unsafe operation, or lead to failure. Fault diagnosis refers to identifying the likely fault type, cause, location, or corrective action after abnormal behavior has been observed. This differs from predictive maintenance and RUL estimation, which primarily aim to anticipate future failures or estimate remaining operating time.

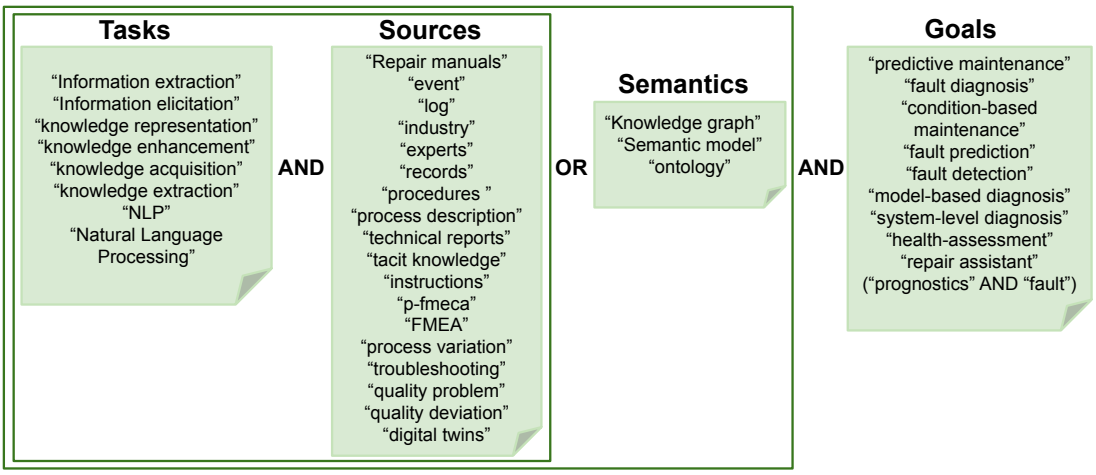


Fig. 3. Structured formulation of the search string used for retrieving relevant studies in the systematic review

the retrieved studies to fault diagnosis. In parallel, semantic representation keywords, including knowledge graph, semantic model, and ontology, formed a separate branch in the search string to capture studies focusing directly on scalable semantic modeling approaches, which corresponds to the focus of this study on semantic knowledge representation techniques.

The final element of the search string corresponds to the application goal of the review and includes keywords related to fault diagnosis. The presence of other maintenance-related terms, such as prediction and condition-based monitoring reflects the fact that researchers sometimes use these terms while performing fault diagnosis.Moreover, the term `prognostics` is intentionally combined with `fault` to reduce ambiguity in the search results, as `prognostics` is used across multiple domains, including medicine and clinical prediction, which are outside the scope of this study. Combining it with `fault` increases the relevance of the retrieved studies to fault diagnosis.

The search design involves a deliberate trade-off. Broad methodological terms such as information extraction, knowledge elicitation, knowledge representation, enhancement, and reasoning were used to capture relevant work across the five RQs. At the same time, the search was narrowed at the application level to fault diagnosis and at the representation level to scalable semantic knowledge representation. This design was intended to broaden the coverage of relevant methods while maintaining a focused scope. The search is fully reproducible by entering the exact keywords and Boolean logic: *((("Information extraction" OR "Information elicitation" OR "knowledge representation" OR "knowledge enhancement" OR "knowledge acquisition" OR "knowledge extraction" OR "NLP" OR "Natural Language Processing") AND ("Repair manuals" OR "event" OR "log" OR "industry" OR "experts" OR "records" OR "procedures" OR "process description" OR "technical reports" OR "tacit knowledge" OR "instructions" OR "p-fmeca" OR "FMEA" OR "process variation" OR "trouble shooting" OR "quality problem" OR "quality deviation" OR "digital twins")) OR ("knowledge graph" OR "semantic model" OR "ontology")) AND ("predictive maintenance" OR "fault diagnosis" OR "condition-based maintenance" OR "fault prediction" OR "fault detection" OR "model-based diagnosis" OR "system-level diagnosis" OR "health-assessment" OR "repair assistant" OR ("prognostics" AND "fault"))* .

3.1.3. ***Inclusion and Exclusion Criteria***

This step defines the criteria for filtering and selecting relevant studies. Inclusion criteria define studies that are aligned with research objectives and fall within the scope of this review, while exclusion criteria eliminate studies that are outside of this scope. Given that fault diagnosis has been extensively studied and a substantial body of literature exists, our focus is specifically on fault diagnosis in industrial equipment and aviation. Accordingly, the inclusion and exclusion criteria are defined in accordance with this objective and are summarized in Table 2.

Table 2
Inclusion and exclusion criteria for paper assessment

Inclusion Criteria	Exclusion Criteria
• Semantic knowledge representation methods for supporting fault diagnosis in industrial systems and aviation • Studies proposing any tool,technique or method for acquiring, extracting, representing, enhancing knowledge to support fault diagnosis • Papers published between 1995 and 2024	• Studies based on complete physical system models • Studies using only sensor data as input source • Medical diagnosis studies • Papers other than the English language • Technical reports, thesis, keynote, abstracts, posters, doctoral symposiums, whitepapers. • Duplicate or overlapping studies for which a more comprehensive version is available • Papers focusing on other aspects of maintenance instead of fault diagnosis • Papers with other domain than industrial systems and aviation

The inclusion criteria cover studies published between 1995 and 2024 that propose tools, techniques, or methods for acquiring, extracting, representing, and enhancing knowledge pertinent to CPS fault diagnosis. In addition, semantic knowledge representation approaches are explicitly included in line with the objectives of this review.

As for exclusion, studies relying solely on sensor data are excluded, as this review focuses on approaches that leverage domain knowledge and human-generated maintenance-related information rather than solely machine-generated measurements. Studies within the medical field are also excluded, as this review focuses on machine fault diagnosis. Articles focusing on other aspects of maintenance, such as failure prediction and remaining useful life estimation, are not considered, since their goal is to predict failures or estimate part replacement times, whereas our focus is on identifying the causes of observed faults after they have occurred. Studies that rely on complete system-level models for fault diagnosis are excluded, as such approaches are often difficult to scale, extend, and maintain in complex CPS environments. Studies focusing on machines and systems outside the industrial and aviation domains are excluded, as this review specifically focuses on semantic-based fault diagnosis within these domains. Finally, duplicate or overlapping studies are excluded for which a more comprehensive version is available.

3.1.4. ***Database Selection***

To identify relevant studies, the search string was applied in the Scopus database. Scopus was selected due to its extensive coverage of peer-reviewed literature and its advanced filtering capabilities. To mitigate the limitations associated with relying on a single database, both forward and backward snowballing were also applied. Forward snowballing identifies studies that cite the selected papers, while backward snowballing examines the references of the selected studies to identify earlier relevant work. This complementary strategy helps broaden coverage and reduce the risk of missing relevant studies.

3.2. ***Review Conducting***

This phase covers the identification, screening and selection of papers. The tasks are explained in the following sub-sections.

3.2.1. ***Article Identification***

The article identification stage was conducted using the search string defined during the planning phase. By entering the search string into Scopus, a total of 1,273 articles were retrieved based on title, abstract, and keyword search. These articles formed the initial pool for the subsequent selection process.

3.2.2. ***Article Screening and Selection***

We started the selection process by screening the abstract of the retrieved articles according to the established inclusion and exclusion criteria, resulting in 304 articles being retained for further assessment. During this stage, a conservative approach was adopted, whereby articles were retained whenever abstracts

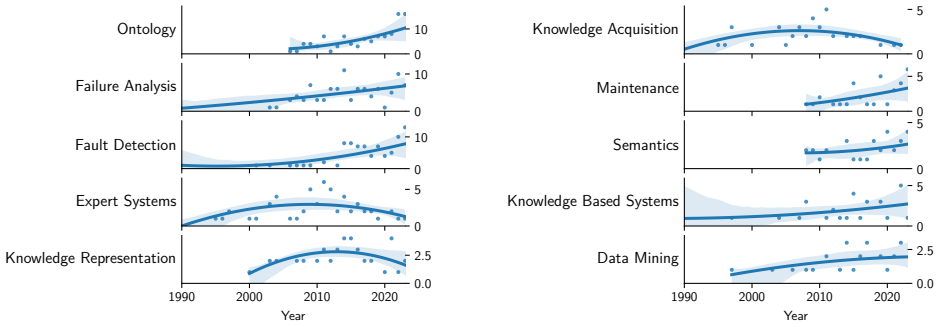


Fig. 4. Top 10 most frequent index keywords of papers after abstract screening, by year. Lines and shading show a quadratic regression fit and its confidence interval. Most keywords show an upwards trend, except *Expert Systems*, *Knowledge Representation*, and *Knowledge Acquisition*.

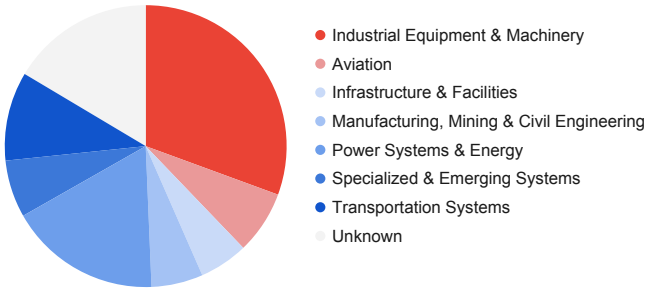


Fig. 5. Application domains of papers after abstract screening. Papers with applications in the *Industrial Equipment & Machinery* and *Aviation* domain were selected in the domain screening step.

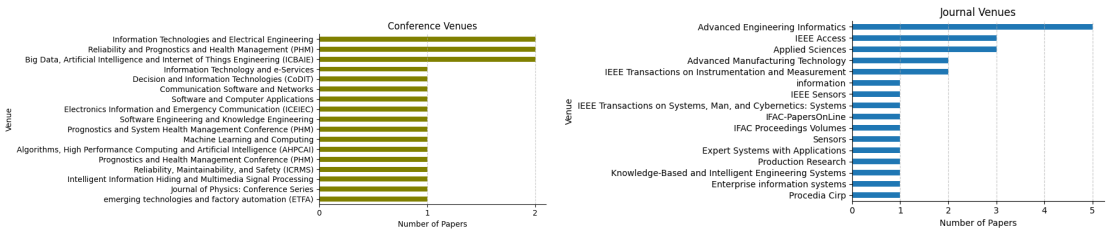


Fig. 6. Venue distributions

did not provide sufficient information to make a definitive inclusion or exclusion decision. Fig. 4 shows temporal trends of the ten most frequent keywords associated with these articles.

The 304 studies retained after abstract screening were further assessed according to the established inclusion and exclusion criteria. Papers were excluded if they were unavailable in English, relied solely on sensor data, lacked sufficient methodological detail, focused exclusively on proposals or reviews, or presented non-scalable approaches. Studies addressing predictive maintenance, prognostics, or Remaining Useful Life estimation were also excluded. This step resulted in 141 papers being retained for full-text scope-assessment.

In the final selection step, the full texts of the remaining studies were screened, resulting in the selection of 39 papers. Although CPS can encompass a wide range of application domains such as infrastructure, power systems, transportation, and civil engineering systems, this review adopts a narrower scope, focusing on industrial equipment and machinery and aviation domain such as aircraft systems. Fig. 5 shows the application domains of the 304 papers after abstract screening and highlights the selected domains retained in this study.

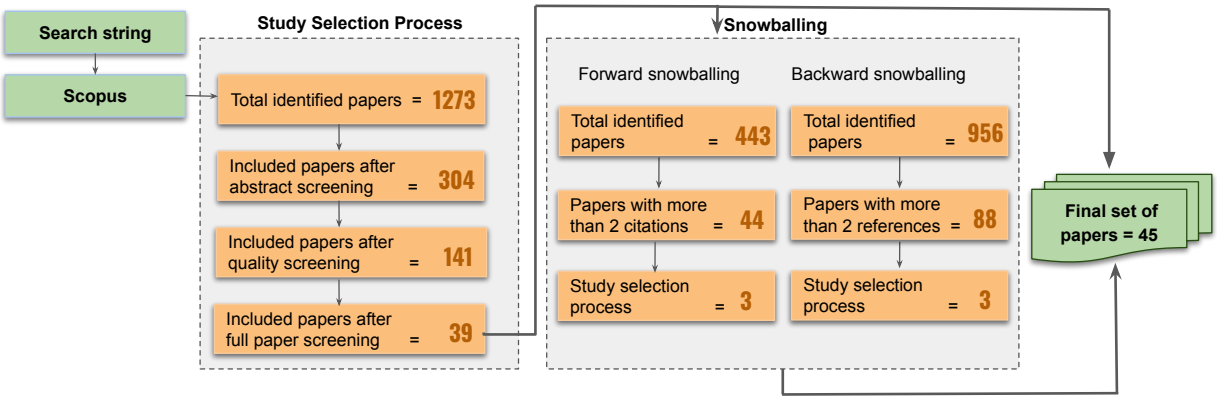


Fig. 7. Study identification, screening, and selection using inclusion and exclusion criteria, with snowballing [4, 6, 10, 13–16, 20, 21, 25, 26, 28–30, 35, 36, 43, 52, 53, 55, 59, 60, 69, 71, 74–76, 79–81, 87, 89–91, 93, 94, 96, 97, 99–105]

As a complementary strategy, both backward and forward snowballing were subsequently applied to identify potentially relevant papers that may not have been retrieved by the database search, for example, because the relevant terminology was not present in their title, abstract, or keywords. In backward snowballing, the reference lists of selected papers are reviewed, while forward snowballing examines studies that cited them. This process identified 956 papers through backward snowballing and 443 through forward snowballing. Papers were considered for further screening if they had at least two references or had been cited at least twice. This criterion was applied to reduce the inclusion of marginally relevant studies while still allowing papers beyond the initial database search to be considered. These studies then underwent the same screening and assessment process, resulting in the inclusion of an additional six studies. In total, 45 papers were included in this literature review.

Notably, the search and selection process was conducted independently by two reviewers. Following this, both reviewers convened to discuss the selected papers, ensuring alignment with the SLR methodology. Any disagreements were resolved through discussion and reassessment until consensus was reached.

4. Results and Analysis

This section presents the results derived from the final set of selected papers, structured according to the five research questions outlined in Section 3.1.1: knowledge acquisition, information extraction, knowledge representation, fault diagnosis, and knowledge bases enhancement. Each subsection corresponds to a specific research question and summarizes the identified methods, their strengths and limitations, and the research gaps.

4.1. ***Knowledge Acquisition***

Knowledge elicitation refers to the process of obtaining knowledge from human sources, such as domain experts. It is part of the larger knowledge acquisition process, which involves not only the elicitation of knowledge, but also the explication and formulization of that knowledge [19].

A review of the included papers indicates that none specifically addresses the development of dedicated methods to elicit knowledge from domain experts in semantic-based fault diagnosis, nor does it employ established elicitation frameworks beyond traditional interviews. In fact, traditional interviews appear to be the only method used for knowledge elicitation. Moreover, there is no reporting on the effort, time, or resources involved to conduct these interviews. This lack of reporting on the elicitation process makes it difficult to assess the scalability, reproducibility, and practical effectiveness of the current approach.

Table 3
Expert involvement in semantic-based fault diagnosis methodology development (*Knowledge Acquisition*)

Expert Involvement Category	References	Total Number
Constructing or validating the ontology /schema	[6, 10, 15, 20, 21, 28, 53, 79, 81, 99, 103]	11
Constructing or validating rules	[28, 69, 96]	3
Validating the proposed methodology	[26, 60, 81]	3

Consequently, future work could investigate more scalable and interactive approaches for capturing expert knowledge for fault diagnosis, particularly AI-assisted techniques such as *Large Language Models* (LLMs) for conversational agent or chatbot-driven elicitation systems [9] that can support structured knowledge elicitation from domain experts.

Given the limited attention to elicitation methods, this study instead focuses on identifying the stages of the semantic-based fault diagnosis pipeline to which domain experts contribute, as shown in Table 3. These categories are elaborated as follows:

- **Ontology/schema construction or validation.** Domain experts play a key role in defining the structure of knowledge by contributing to the construction of ontologies (defined in section 4.3.1) or schema and validating their accuracy and relevance. Notably, ontology or schema construction, with 11 papers [6, 10, 15, 20, 21, 28, 53, 79, 81, 99, 103], represents the most frequently reported area of expert involvement.
- **Rule construction or validation.** This category encompasses the development and validation of rules based on domain-specific knowledge to support reasoning and knowledge extraction in semantic-based fault diagnosis systems. Domain experts contributed by defining, refining, or validating these rules, as reported in three studies [28, 69, 96].
- **Methodology validation.** This category refers to the validation of the proposed fault diagnosis methodology by domain experts, particularly service engineers, to assess its practical applicability and usefulness in supporting them to identify fault causes and determine appropriate solutions. This form of expert validation was reported in three studies [26, 60, 81].

Besides the above categories that explicitly reported domain expert involvement, four studies [13, 29, 52, 104] followed ontology engineering methodologies without explicitly mentioning consultation with domain experts. In addition, 11 papers [14, 30, 59, 71, 75, 87, 90, 94, 97, 102, 105] defined an ontology or schema without providing details on the construction process, leaving expert involvement unclear. Four other studies [25, 26, 55, 74] reused or were inspired by existing ontologies. Finally, several works [4, 35, 36, 76, 93, 100, 101] adopted fully automatic bottom-up or data-driven approaches for ontology or knowledge graph construction, where the schema is derived from data.

Overall, this distribution indicates a mismatch between the reliance on domain expertise and the limited reporting on how such knowledge is elicited. Notably, the frequent use of knowledge graphs and ontologies as the knowledge bases in the reviewed literature may partly explains the emphasis on domain knowledge.

4.2. ***Information Extraction***

Information extraction (IE) is the process of converting unstructured text into structured representations containing selected information from the text [32]. Among the reviewed studies, 26 out of 45 papers either proposed or employed an IE method. As shown in Fig. 8, IE methods can be grouped into five main categories: data mining–based, rule-based, linguistic, *Deep Learning* (DL)-based, and hybrid approaches. Since DL-based methods constitute the majority of the reviewed studies, a dedicated table summarizes and compares them, while the remaining categories are discussed in the following subsections.

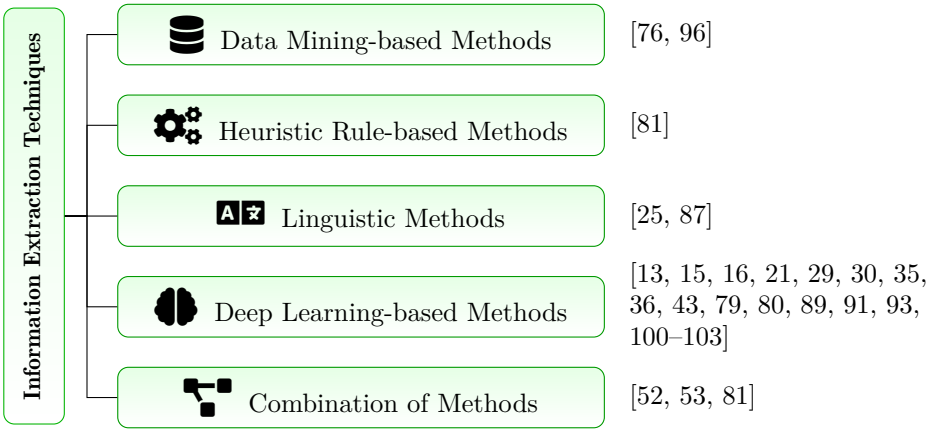


Fig. 8. A categorization of information extraction methods identified in the reviewed studies

4.2.1. ***Data Mining–based Information Extraction***

Data mining is the process of automatically discovering useful information from large amount of data [78]. Among the reviewed studies, Xiao et al. [96] employed *Association Rule Mining* (ARM), a data mining technique used to discover frequently co-occurring event patterns, to extract associations between fault phenomena and fault solutions from historical maintenance data. ARM evaluates associations using measures such as support and confidence. Support quantifies how frequently an event combination occurs, while confidence estimates the likelihood of event B occurring given event A. The authors leveraged the Apriori algorithm [1], one of the most widely used algorithms for ARM. However, ARM-based approaches identify associations based on measures such as support and confidence. Consequently, infrequent fault patterns may have low support, while weakly associated patterns may have low confidence, causing them to be excluded when minimum thresholds are applied.

Sun and Wang [76] employed Pearson Correlation Coefficients [66], a statistical measure of linear dependence between two variables, to identify relationships among heterogeneous data collected from different levels, including production, energy, material, and process data. These relationships were represented as weighted edges between entities in the knowledge graph. To integrate information from multiple sources, the authors further applied entity alignment and inconsistency-checking procedures to construct a unified graph. The resulting correlation-based relationships were then analyzed using edge weights to identify strongly correlated interaction paths within the system. However, their method is limited to identifying linear statistical relationships and may fail to capture non-linear dependencies among variables.

4.2.2. ***Heuristic Rule-based Method for Information Extraction***

Heuristic rule-based information extraction relies on manually designed patterns, templates, and predefined rules to identify and extract information from semi-structured or unstructured sources. These rules are often implemented using regular expressions, keyword lists, or hand-crafted templates. Because the extraction logic is explicitly defined, this approach offers interpretability and fine-grained control over the extracted results. For example, Tang et al. [81] noted that in equipment technical documentation, a sentence containing keywords such as `cause` or `lead to` between two fault entities often indicates a causal relation. Such domain-specific patterns can therefore be used to formulate rules for extracting relevant relationships. However, the performance of rule-based approaches depends heavily on the quality and coverage of the manually constructed rules. Designing and maintaining these rules can require substantial effort, and their applicability may be limited when the data contain substantial linguistic variation or originate from heterogeneous sources.

4.2.3. ***Linguistic-based Method for Information Extraction***

Linguistic approaches exploit the syntactic and grammatical properties of natural language to extract entities and relationships from unstructured text. These approaches analyze the linguistic structure of

sentences to identify relevant entities and relational patterns that can subsequently be transformed into structured knowledge. A common technique is dependency parsing [62], which represents grammatical relationships between words in the form of a dependency tree. Wang et al. [87] employed dependency parsing to identify relational structures in technical text. In the sentence `Expert A diagnosed the fault`, the word `diagnosed` functions as the head word, while `Expert A` and `fault` correspond to the subject and object, respectively. Such grammatical dependencies can then be transformed into knowledge graph triples. This method can reduce the need to manually define extraction patterns. However, its performance is highly dependent on parsing accuracy. Complex sentence structures, ambiguous expressions, and domain-specific language can lead to parsing errors and inconsistent relation extraction. In addition, different syntactic formulations expressing the same semantic meaning may produce different dependency structures, reducing extraction consistency across documents. Ebrahimpour et al. [25] combined parsing analysis with semantic interpretation to extract structured knowledge from unstructured maintenance logs. Their approach first identifies and normalizes key terms, and then filters them using domain standards including ISO 14224 [41] and OREDA [63] to reduce ambiguity and noise. The extracted terms are then organized into hierarchical structures representing equipment-categories and equipment-failure relationships. Semantic interpretation is subsequently applied to identify relationships between these elements, particularly cause-effect relationships. For example, in the maintenance record `Valve failed to open due to diaphragm wear`, the method identifies `failed to open` as the failure mode, `diaphragm` as the affected component, and `wear` as the underlying cause. These elements are then linked into a causal structure representing fault propagation (`wear → diaphragm damage → valve fails to open`). This approach transforms raw textual descriptions into structured knowledge suitable for KG construction. However, this method relies on predefined causality logic, lexical matching and cue words, making its effectiveness dependent on their coverage and quality.

4.2.4. ***Deep Learning-based Methods for Information Extraction***

Deep learning-based methods employ neural network models to automatically learn feature representations from data. These methods have become increasingly popular for information extraction tasks due to their ability to automatically learn useful features [98]. Since a substantial portion of the included papers in this review adopted deep learning techniques for information extraction, we categorize them into two groups: schema-based IE and Schema-free IE. Table 4 summarizes methods that perform extraction under a predefined schema, where the objective is to identify entities and relations according to predefined categories. In contrast, Table 5 summarizes approaches that do not rely on a rigid ontology/schema and instead aim to extract knowledge in a more open and flexible manner. This categorization facilitates a systematic comparison of existing methods and highlights differences in extraction extraction strategies.

Among the reviewed studies, *Bidirectional Encoder Representations from Transformers* (BERT) [22], *A Lite BERT* (ALBERT) [49], *Bidirectional Long Short-Term Memory* (BiLSTM) [31], and *Conditional Random Fields* (CRF) [48] are the most frequently used architectures and models. Therefore, a brief overview of them is provided below.

- **BERT**. It is a transformer-based language representation model that learns contextual token representations by jointly considering the left and right context of each token. It is pre-trained on large-scale unlabeled text using masked language modeling and next sentence prediction, and can subsequently be fine-tuned for a wide range of downstream natural language processing tasks [22].
- **ALBERT**. It is a lightweight, parameter-efficient variant of BERT designed to reduce the number of parameters and memory requirements of large-scale pre-trained language models. It achieves this through cross-layer parameter sharing and factorized embedding parameterization, and introduces *Sentence Order Prediction* (SOP) as an additional self-supervised pre-training objective [49].
- **BiLSTM**. It is a neural network architecture that processes sequential data in both forward and backward directions, allowing the model to learn from both past and future contextual information [31].
- **CRF**. It is a probabilistic sequence labeling model that defines a conditional probability distribution over label sequences given an input sequence, and can be used to predicts the most likely label sequence for a given input [48].

The reviewed studies used these architectures in different combinations depending on the extraction task and data characteristics.

Table 4: Deep learning-based methods for information extraction under a predefined schema

Information Extraction Method	Short Description	Pros (+)/Cons (–)	Reference
Joint Entity and Relation Extraction: SBALGN model	*Self-attention-based stacked BiLSTM with Label weight embedding and Graph Convolution Network* (SBALGN) is proposed for event argument entity and relation joint extraction. First, the input text is converted into character-level embeddings using a Word2Vec model. These embeddings are then fed into a stacked multi-layer BiLSTM. Next, a self-attention mechanism is applied to model dependencies between all character pairs, generating character dependency representations. Using these representations, the model performs the first-stage joint extraction: a CRF layer predicts entity labels, while simultaneously the dependency features, combined with label-weight embeddings, are used to compute relation probabilities between character pairs. After this initial extraction, a BiGCN layer constructs a graph based on the predicted relations, propagates relational information between nodes (characters), and updates the contextual features accordingly. Finally, with these refined features, the model performs a second-stage joint extraction of entities and relations, improving both entity boundary detection and relation classification accuracy.	+ It outperformed baselines including Baseline-stacked-BiLSTM and Baseline-stacked-BiLSTM+self-attention. + By operating at the character level, it avoids the need for error-prone word segmentation in Chinese text. – Compared to simpler joint extraction models, SBALGN requires more training time. – SBALGN achieves slightly lower overall joint extraction performance (−0.23% F1) compared to the BERT pre-trained model of Huang et al. [40], although it improves relation extraction accuracy.	[21]
Entity Extraction: BiLSTM-CRF **Relation Extraction:** BERT model **Joint Entity and Relation Extraction:** Joint Extraction Model Based on BERT	Entities are identified using a BiLSTM-CRF model, where BiLSTM captures bidirectional contextual dependencies in text and CRF ensures consistent sequence labeling. For relation extraction, the model leverages BERT. In addition, the paper introduces a joint BERT-based model that performs entity recognition and relation extraction simultaneously by using a CRF layer after the BERT-based vectorization of text. Then Soft label embeddings encode entity type information from logits, helping maintain consistency in predictions. A multi-sigmoid layer is used to handle cases where an entity belongs to multiple relations. Relation extraction task is transformed into multi-head selection problem to handle cases where a token participates in multiple relations.	+ The BERT joint extraction model achieves better F-measure compared to the BERT-BiLSTM-CRF. + The model supports multi-label relation extraction, where one entity can participate in multiple relationships simultaneously. – While recall improves in the Joint model, precision is relatively low.	[30]
Entity Extraction: BERT-BiLSTM-CRF **Relation Extraction:** CNN	First, entity recognition is performed by converting the text into word embeddings using a pre trained-BERT model, which are then passed into a BiLSTM model to capture contextual information from both past and future tokens. The sequence is then decoded using a CRF layer with *Beginning, Inside, Outside* (BIO) tagging to identify fault-related entities. For relationship extraction, a *Convolutional Neural Network* (CNN)-based classifier is used to detect relationships between entity pairs by learning sentence-level features and classifying whether a valid relation exists.	+ Leverages BERT-based contextual modeling combined with BiLSTM and CRF decoding, enabling robust representation of domain-specific terminology while improving entity boundary consistency through constrained sequence labeling. – CNN treats relation extraction as classification per sentence that may miss long-distance dependencies or cross-sentence relations. – Entity recognition and relation extraction are handled as two separate steps, which can propagate errors from entity extraction module to the relation extraction. – Lack of comparison with other methods.	[79]
Entity Extraction: LR-CNN **Relation Extraction:** BiGRU-Attention model	First, fault-related texts are preprocessed using word segmentation, lexical annotation, and feature selection with domain-specific dictionaries, and words are converted into embeddings using Word2Vec. Named entity recognition is performed using a *Lexicon Rethinking* (LR-CNN) model enhanced with a lexicon mechanism to resolve word conflicts. The model further integrates entity embeddings, contextual information, and graph-based representations through a multi-channel CNN architecture to capture both local and contextual semantic features. For relation extraction, the BiGRU captures contextual information in both forward and backward directions, while the attention mechanism learns to assign higher weights to the most important words in a sentence that contribute to determining the relationship.	+ The BiGRU-Attention model emphasizes informative relational terms by assigning higher attention weights to important words during relation classification. – The paper does not explicitly evaluate the performance of entity recognition or relation extraction as standalone tasks.	[29]

Table 4 (continued)

Information Extraction Method	Short Description	Pros (+)/Cons (–)	Reference
Entity Extraction: BERT-BiLSTM-CRF	First, BERT is used to generate deep contextual word representations by encoding the entire sentence in a bidirectional manner. These embeddings are then fed into a BiLSTM network. Finally, a CRF layer is applied for sequence labeling and decoding, ensuring globally optimal entity predictions by considering dependencies between output labels.	+ BERT-BiLSTM-CRF achieves better F-measure compared to other models like HMM, BiLSTM, and BiLSTM-CRF. – Model only handles entity recognition. – It relies on manual annotation and expert-defined labeling guidelines, making corpus construction time-consuming and labor-intensive.	[103]
Entity Extraction: Pre-existing NER tool **Relation Extraction:** PCNN-ATTRA-RL	It proposes a *Piecewise Convolutional Neural Network with Reinforcement Learning* (PCNN-ATTRA-RL) framework for relation extraction in noisy, distantly supervised datasets. It first uses a PCNN model with word and position embeddings to extract sentence-level features (ATTRA), and applies sentence-level attention to assign higher weights to more relevant sentences for each relation type. As distant supervision automatically labels sentences based on known entity relations defined in the ontology, it often introduces noise because not all sentences mentioning the entities actually express the intended relation. To reduce noise, the RL method is used, where an agent learns to remove irrelevant or falsely labeled sentences based on improvements in F1 score. A CNN-based policy network is used to decide whether to keep or discard sentences.	+ The combination of sentence-level attention and RL helped mitigate the impact of noisy sentences. + PCNN-ATTRA-RL outperforms several baseline DSRE methods including PCNN, PCNN-RL, and PCNN-ATTRA. – It operates at sentence/bag level, so it cannot capture long-range dependencies.	[13]
Entity Extraction: MF-GCN **Relation Extraction:** MF-GCN	The *Multi-Feature Graph Convolutional Network* (MF-GCN) consists of three layers: a feature fusion layer, a stacked BiLSTM layer, and a BiGCN entity relation inference layer. In the *feature fusion layer*, the input text is converted into character embeddings using Word2Vec, and three types of features are extracted: context features using a BiLSTM over the embeddings, context salient features using a CNN over the embeddings, and local context features by combining character embeddings with radical-level embeddings and applying CNN. These three feature vectors are then fused. Next, in the *stacked BiLSTM layer*, the fused features are passed through a two-layer BiLSTM to capture long-range dependencies, and a CRF layer produces an initial entity recognition result. Finally, in the *BiGCN relation inference layer*, a graph is constructed based on entity relations, and a BiGCN layer updates node representations using adjacency matrices to model relationships between entities; this step refines entity boundaries and incorporates relation information, and the output is again passed through a CRF to generate the final entity predictions.	+ Utilizing Bi-GCN module in MF-GCN enables global relational reasoning. + MF-GCN achieves better F-measure in entity recognition, compared to Lattice LSTM, WC-LSTM, and BERT-CRF. + The use of character-level and radical-level embeddings, along with contextual features, helps handle technical terms and sparse entities. + The two-stage CRF prediction improves entity boundary detection – The model relies on radical-level features, which are specific to languages like Chinese and may not generalize easily to other languages. – The model's performance depends on the availability of entity–relation information, as removing the BiGCN layer results in a drop in F1-measure from 93.58% to 90.30%.	[15]
Joint Entity and Relation Extraction: ALBERT-BiLSTM-CRF **Table Detection:** Modified Mask R-CNN	For textual information, it applies ALBERT-BiLSTM-CRF that first generates contextual word embeddings using ALBERT, then captures sequence dependencies with a BiLSTM, and an attention layer assigns higher importance to relevant words in a sentence. Finally, it uses a CRF layer to jointly perform named entity recognition and relation extraction. For tabular information, the method first detects tables in documents using a modified Mask R-CNN model, then applies OCR to extract the table content. After that, ALBERT-BiLSTM-CRF is used to identify entities within the extracted text. Zhou et al. [102] used the same IE method in [101].	+ ALBERT-BiLSTM-CRF performs better than BiLSTM and BiLSTM-CRF in entity recognition in terms of accuracy. + Table detection based on modified Mask R-CNN achieves better F-measure compared to Faster R-CNN and standard Mask R-CNN. – Hyperparameters were selected empirically and may not be universally optimal. – It is tailored for regular tables with clear cause-and-effect relationships and may struggle with irregular layouts.	[101, 102]
Joint Event Extraction: BERT+ Cascade Binary Tagging + Entity Alignment	It uses event extraction to identify fault events, including triggers, event types, and faulty components (arguments). The framework uses a shared BERT encoder with three sequential decoders for trigger extraction, type detection, and argument extraction. Self-attention is used to capture contextual dependencies, while *Conditional Layer Normalization* (CLN) integrates trigger and type information into downstream subtasks. After extraction, entity alignment is performed using the Dedupe library and semantic similarity matching to merge inconsistent entity names before storing triples.	+ The proposed method achieves a 1.5% improvement over CasEE on the Airbus dataset and 7.9% for ARJ21 dataset. + Entity alignment reduces redundancy caused by abbreviations, spelling variations, and inconsistent maintenance terminology. – Entity alignment relies mainly on text semantic similarity and clustering, which the authors themselves state is still limited and needs improvement. – Multiple sequential decoding stages increase model complexity and error propagation risk between trigger, type, and argument extraction.	[91]

Table 5: Deep learning-based methods for information extraction with non-rigid schema

Information Extraction Method	Short Description	Pros (+)/Cons (–)	Reference
Joint Entity and Relation Extraction: ROMGCJE model	The *Relation Oriented Model with Global Context information for Jointly Extracting overlapping relations and entities* (ROMGCJE) is proposed to deal with the overlapping relation problem. First, the input text is transformed into word embeddings through the *Primary Task-Shared Representations Layer* (PTSRL). Next, these embeddings are encoded by the *Contextual Word Representations Layer* (CWRL) to capture contextual semantic information. Then, a *Relation-Based Attention with Gating Mechanism* (RBAGM) is applied to enhance entity recognition by incorporating detected relation information into the contextual representations. In the decoder part, the relations are extracted by a *Multiple Relation Classifier* (MRC), and the entities are recognized by an improved LSTM. In addition, the *Attention for Joint Extraction* (AJE) module is designed to match entity pairs with their corresponding relations. Finally, the *Minimum Risk Training* (MRT) mechanism is introduced to jointly train the model to extract triples from text.	+ ROMGCJE exhibited faster convergence during training compared to HRL, RSAN, and WDec. + ROMGCJE can jointly extract overlapping entities and relations from unstructured texts. + Relation-aware attention improves entity recognition. – It requires many sensitive hyperparameters, making tuning difficult and performance unstable if not carefully optimized. – The multiple attention layers, gating mechanisms, and global embeddings make the model difficult to interpret, as it is unclear why a relation is extracted or which component contributes most.	[36]
Fault Event Identification: BERT **Trigger Word Extraction:** BERT-BiLSTM-CRF **Argument Extraction:** CLN	The fault event extraction is formulated as a three-step pipeline. First, it identifies whether a sentence describes a fault event using a BERT-based binary classification model. If a fault event is detected, trigger words are extracted by treating the task as a named entity recognition problem, using a BERT-BiLSTM-CRF model to capture contextual dependencies and ensure consistent labeling. Finally, argument extraction is performed by incorporating trigger information into the sentence representation through *Conditional Layer Normalization* (CLN), enabling the model to extract relevant event arguments (e.g., time, location).	+ It achieves strong performance with high F1-measure (0.95 for fault detection, 0.93 for trigger extraction, and up to 0.91 for argument extraction). – While the argument extraction remains relatively simplified, it may struggle with overlapping arguments and long-distance dependencies between multiple events. – Lack of comparison with other methods.	[80]
Text Classification: Fine-tuned BERT	For structured data source (troubleshooting manuals), each task number corresponds to a specific fault type and its associated causes, which are directly extracted. For unstructured maintenance records, where formatting is inconsistent, the method applies a BERT-based binary text matching model: event descriptions are paired and labeled according to whether they belong to the same fault type. This allows new maintenance texts to be mapped to existing fault categories, indirectly associating them with their known causes. Before modeling, data is cleaned and invalid records are removed, and the final BERT model is fine-tuned on labeled fault pairs to perform fault-type classification.	+ BERT model has a better precision score compared to MATCHPayramid and MVLSTM. – Pairwise text matching can become computationally expensive when comparing large numbers of maintenance records	[100]
Entity Extraction: BERT-Decoder-CRF	First, a pr-trained-BERT model encodes the input sequence into deep bidirectional contextual representations. These representations are then fed into a Transformer-based Decoder, where masked multi-head attention is applied to model interactions between tokens and their corresponding label sequences, allowing the model to better align textual features with entity tags during training and prediction. Afterward, the Decoder output passes through Feed Forward layers together with residual Add&Norm operations to produce refined sentence feature vectors. Finally, these vectors are input into a CRF layer, which models dependencies between adjacent labels and predicts the optimal global label sequence for entity recognition.	+ It improves entity recognition compared to models like IDCNN-CRF and BiLSTM-CRF. + It shows good generalization by performing well on both industrial and public datasets. – F1-measure of the *Faulty Time* entity is relatively low due to difficulty handling heterogeneous patterns.	[89]
Joint Entity and Relation Extraction: JERE	*Joint Entity and Relation Extraction* (JERE) model is proposed to automatically extract knowledge in the form of triples from text. First, a BERT encoder generates contextual token embeddings, which are then passed through a BiLSTM layer to capture bidirectional contextual dependencies. A CRF layer is applied on top for accurate named entity recognition using a sequence labeling scheme *Beginning, Inside, Last, Outside, or Unit* (BILOU). To enable interaction between entity recognition and relation extraction, the predicted entity labels are transformed into soft label embeddings. Finally, a multi-head mechanism is used to classify relationships between entity pairs, producing scores for candidate triples.	+ JERE achieves the highest scores across P, R, F1 for triple extraction compared to MultiR, HRL, CoType, and LSTM-CRF methods. + The joint extraction architecture reduces error propagation between entity recognition and relation extraction by optimizing both tasks simultaneously through a shared objective function. + The multi-head relation extraction mechanism allows the model to identify multiple relations associated with the same entity – Requires large labeled dataset with high-quality entity/relation annotations. – Relation extraction is performed mainly at the sentence/token level, making the approach less suitable for cross-sentence relation extraction.	[35]

Table 5 (continued)

Information Extraction Method	Short Description	Pros (+)/Cons (–)	Reference
Entity Extraction: BERT-BiLSTM-CRF	First, a pre trained-BERT model encodes the input text into contextualized token representations. These representations are then passed to a BiLSTM layer, which further models sequential dependencies by incorporating both past and future context in the sequence. Finally, a CRF layer performs joint decoding to assign the most probable BIO labels to each token while considering dependencies between adjacent labels.	+ Superior performance in extracting 5 kinds of entities compared to Lattice LSTM, BERT NER, BiLSTM-CRF, ZEN, and ERNIE models. – Inability to handle overlapping entities which forced authors to split the dataset into separate files (e.g., equipment entities vs. other sequences) – Performance varies with sequence length, the results show that the model performs better on shorter sequences.	[16]
Joint Entity and Relation Extraction: CASREL+BERT-wwm-ext+Att+Decoder	*Cascaded Adversarial Relation Extraction* (CASREL) framework is used to perform joint entity and relation extraction in a cascade structure. The model first uses a pre-trained encoder, specifically BERT-wwm-ext, to convert the input text into contextualized representations. On top of this, a Head Entity Tagger identifies subject entities by predicting their start and end positions using a binary classification scheme. Then, conditioned on the detected subject representation, a Tail Entity Tagger extracts object entities and their corresponding relations, effectively modeling triples. To better capture local contextual dependencies, a scaled dot-product self-attention module is introduced, and the decoder component is enhanced by replacing a single linear layer with a multi-layer perceptron to improve nonlinear feature modeling. It addresses overlapping relations through its cascaded, subject-to-object prediction strategy.	+ BERT-wwm-ext shows a better F1-measure compared to BERT and ALBERT. + CASREL+BERT-wwm-ext+Att+Decoder shows a better F1-measure compared to based models like CASREL+BERT-wwm-ext. + It is effective at handling overlapping relation extraction. – Requires careful hyperparameter tuning (e.g., attention heads, hidden dimensions significantly affect performance).	[43]
Object Detection in Image: Faster R-CNN	First, entity extraction happens at the visual level using an improved Faster R-CNN model. This model detects objects in images (e.g., transformers, radiators, insulators) and outputs bounding boxes along with feature vectors representing each detected object. These detected objects effectively serve as the entities. The feature vectors obtained from visual detection are compared with pre-defined entity, attribute, and relation stored in the KG, using cosine similarity, to align a detected object with corresponding nodes and attributes in the graph.This is used for reasoning and will be explained in the Reasoning section.	+ The improved Faster R-CNN model achieves a higher mean Average Precision compared to Fast R-CNN, Faster R-CNN, YOLOv3, YOLOv5, and EFGRnet. – Detection of key objects (like transformers in fault scenes) can be unreliable. This is attributed to insufficient training data and low-quality or unclear targets.	[93]

Table 4 provides a comparative overview of schema-based DL methods, summarizing the models used for information extraction, the proposed methodologies, and their main advantages and limitations. Two observations can be made from this comparison. First, most methods combine pretrained encoders or sequence models with task-specific decoders or auxiliary modules such as CRF, attention, *Convolutional Neural Network* (CNN) [54], or *Graph Convolutional Network* (GCN) [45]. These combinations are used to incorporate complementary capabilities, such as contextual representation, sequence dependency modeling, attention-based feature weighting, and relational reasoning. This suggests that combining complementary modules can improve performance compared with individual baseline models. Second, joint extraction is used when the task requires tighter coupling between entities and relations, as separate pipelines can be vulnerable to cascading errors. However, joint modeling also involves trade-offs; particularly in precision, and complexity. Overall, the reviewed studies suggest that recent developments in this area involve the design of task-specific model combinations and the incorporation of domain-specific information.

Table 5 presents a comparative overview of deep learning-based information extraction approaches that operate without relying on a strictly predefined schema. However, most are only partially schema-free in practice because these approaches may still depend on task-specific output spaces, predefined entity categories, or relation labels. The table summarizes the models used for information extraction, the proposed methodologies, and their main advantages and limitations. The comparison highlights several important aspects of how these approaches address information extraction for semantic-based fault diagnosis of CPSs.

The main methodological trend is toward joint and interaction-aware extraction, where entities and relations are coupled through cascading decoding, attention, graph refinement, or label-aware mechanisms.

Table 6
Information extraction using hybrid methods

Information Extraction Method	Short Description	Pros(+)/Cons(–)	Reference
Table Segment Detection: Mask R-CNN **Entity Extraction:** ALBERT-BiLSTM-CRF **Relation Extraction:** Dependency Parsing	The method treats PDF tables as images and applies a Mask R-CNN–based model to detect and segment table cells. These detected cells are then passed through *Optical Character Recognition* (OCR) to extract raw text. Next, ALBERT-BiLSTM-CRF is used to extract entities, converting entity extraction into a sequence labeling task. For relation extraction, the method goes beyond simple pairwise relations by modeling multi-nary relationships in two ways. First, it uses the table structure (rows and columns) to infer schema-level relations like Symptom–Cause–Action. Second, it applies dependency parsing with rule-based conditions inside table cells to detect coordinated phrases and build hyperedges, linking multiple related entities (e.g., one symptom connected to several causes) at once.	+ Combines visual table structure, OCR, sequence labeling, and dependency parsing, enabling extraction of both standard entity relations and more complex multi-nary relations. + Exploits row/column structure to infer implicit semantic relations that are difficult to capture using sentence-based extraction alone. – Assumes well-structured tables with consistent templates, performance drops if tables are messy or layouts vary.	[53]
Entity Extraction: SP-LEAR **Relation Extraction:** Rule-based and Deep learning based	The *Start Position and Label-knowledge Enhanced Representation* (SP-LEAR) is proposed for entity recognition which treats entity extraction as a span prediction problem, instead of BIO tagging. It starts with BERT embeddings, then enhances them by injecting label knowledge (entity type descriptions) using an attention mechanism. A key improvement is that it explicitly models the dependency between start and end positions: the predicted start position is fed back into the representation to help predict the end position more accurately. For relation extraction, Rule-based extraction is used for structured or semi-structured data (e.g., detecting causal relations via keywords like `cause`) and Deep learning classification for unstructured text, where relation extraction is framed as a binary classification problem: given two entities, predict whether a specific relation exists. The model builds features at two levels: Sentence-level features (via BERT-CNN-pooling) and Paragraph-level context features (via CNN-pooling). These are concatenated and fed into a classifier to determine the relationship. BERT converts each sentence into contextual embeddings, CNN (convolution layer) then scans these embeddings to detect important local patterns (e.g., phrases like `caused by, leads to`), and Pooling compresses this into a fixed-size vector.	+ Entity recognition using SP-LEAR outperforms other models like BERT+Softmax, BERT+CRF, LEAR, and BERT+Span in terms of F1-measure. – SP-LEAR improves F1-measure only slightly compared to LEAR (0.814 → 0.820). – The model performs worse on entity types with longer or less clearly defined spans, such as `Fault` and `Test` entities, suggesting difficulty in handling entities with vague semantic boundaries.	[81]
Entity Extraction: Rule-based and Deep learning based (Transformer encoder-CRF) **Relation Extraction:** BiLSTM + self-attention model	For entity extraction, a rule-based approach (regular expressions and POS tagging) is used specifically for the `Power` entity class. It detects noun phrases and applies pattern rules to identify entities with high precision. For all other entity types, a sequence labeling model called Transformer encoder + CRF is used to build rich token embeddings by combining multiple features including word embeddings (Word2Vec), part-of-speech tags, domain-specific phrase indicators, and frequency information. These are passed through the Transformer to capture context, and the CRF layer ensures coherent label sequences. For relation extraction, a BiLSTM and self-attention model is used. Each candidate entity pair in a sentence is evaluated by encoding the sentence with embeddings that include word features, POS tags, and relative position features (distance to head and tail entities). Feeding these into a BiLSTM to capture sequential context, followed by self-attention to focus on the most relevant parts of the sentence. A Softmax classifier predicts the relation. To handle class imbalance (many entity pairs have no relation), the paper introduces a two-stage classification strategy. First, the binary classifier decides whether any relation exists between two entities. Then, the same model architecture determines the specific relation type.	+ The DL-based entity extraction achieved an F1-measure of 0.82, while the rule-based method for `Power` entities scored 1.00. Combining both methods resulted in F1-measure 0.83. + The two-stage relation extraction approach showed improvement over the one-stage. – The lower performance of the `Function` entity suggests a need for more training data or advanced models.	[52]

While pre-trained language models, particularly BERT and its variants, are used as the main encoding layer in many of the reviewed approaches, the main innovation lies in architectural strategies such as cascading extraction, graph-based refinement, and label-aware interactions. However, despite these advances, many methods still rely on multi-stage pipelines and task-specific heuristics, indicating that fully unified and generalizable information extraction for complex technical domains remains an open challenge.

Table 6 provides an overview of the articles that combine different methodologies for entity and relation extraction. The reviewed methods use different techniques for different components of the extraction process. While entity extraction is performed using approaches such as CRF-based sequence labeling or span-based models, relation extraction incorporates additional mechanisms, including multi-stage decision, self-attention, or two-stage classification strategies. This reflects a broader insight: identifying `what exists` in text is generally easier than determining `how things are connected`, especially when relations span multiple sentences or depend on broader context.

The comparison also shows that these methods are rarely fully unified. They rely on hand-crafted rules, task-specific decomposition, or structural assumptions about the source data such as applying separate rules for `Power` entities, using sentence vs paragraph level features, or splitting relation detection into relation detection and relation type classification. These design choices allow the methods to address specific characteristics of the data and extraction tasks but also make the resulting approaches dependent on task and domain-specific configurations. Overall, while hybrid methods can combine the strengths of different techniques, they also introduce task-specific components and dependencies that may limit their general applicability across different technical domains.

4.3. ***Knowledge Representation***

Knowledge Representation concerns the encoding of knowledge in a form that enables computer systems to reason automatically and use that knowledge to solve problems [37]. In this review, we focus on scalable KR methods used for CPS fault diagnosis. The included studies used different forms of knowledge representation, including ontologies, knowledge graphs, Bayesian networks, hypergraphs, and xSGD.

In addition, Table 7 summarizes the input sources used in these studies, providing an overview of the types of information and datasets used to construct and populate knowledge representations. We first examine these input sources before discussing the representation methods.

In all reviewed studies, input datasets consist mainly of textual documents of various types with differences in their frequency of use. Among these sources, `maintenance records` stand out as the most frequently utilized, appearing in 31 studies. Their widespread use may be attributed to the rich historical information they provide regarding failures and maintenance actions. `Sensor data` is the second most frequently used source, appearing in 12 studies, and provide information on operational conditions and system parameters. It should be noted that, although sensor data are extensively employed in the broader fault diagnosis literature, studies relying exclusively on sensor data without incorporating other human-generated knowledge sources were excluded from this review. `technical manuals` were used in nine papers and provide descriptions and instructions on finding the cause of the failure and solving it. *Failure Mode and Effects Analysis* (FMEA) appeared in five studies, indicating its importance as a source of knowledge about failure mechanisms and their consequences. `Online sources`, including websites, books, and scientific articles, were used in six studies. Furthermore, eight studies did not explicitly specify their input sources, referring only generally to domain knowledge or textual data. *Bill of Materials* (BoM), which contains information about the physical structure of the system, was used in three studies, indicating comparatively limited exploitation to detailed system structural information. In addition to the frequently used sources, several input types appeared only once in the reviewed literature and are therefore grouped under the category `Others`. Examples include design knowledge, network knowledge, and development and production databases, etc.

This overview of input sources allows for a clearer understanding of the data and knowledge landscape in semantic-based fault diagnosis of CPSs, helping to identify commonly used and underexplored sources.

The following sub sections provide a detailed discussion of the identified KR formalisms, together with an analysis of their respective strengths and limitations.

4.3.1. ***Ontologies***

An ontology is a formal and explicit specification of shared conceptualization [33]. In this context, a conceptualization describes the concepts and relationships within a universe of discourse, while a specification defines these concepts and relationships precisely using a formal language. The term shared indicates that the conceptualization is intended to be commonly understood and reused by different users or systems. Ontologies are typically created using the *Web Ontology Language* (OWL) [106], a standard semantic web language for ontology modeling. Ontologies have been widely used in the domain of fault diagnosis with entities such as components, functions, faults, symptoms, causes, maintenance actions, and their interdependencies. Their widespread use can primarily be attributed to three factors: the ability to integrate heterogeneous data and knowledge sources [10, 13, 15, 21, 29], the provision of structured semantic representations of domain [25, 26, 28], and support for reasoning and inference [14].

Table 7
Input data source for each research paper

References	Maintenance records	FMEA	Sensor data	Technical Manuals	Bill of Material	Online data	Not mentioned precisely (domain knowledge/ text/ etc.)	Others
Chen et al. [14]	✓	✓	✓		✓			
Dendani et al. [20]							✓	
Steinegger et al. [74]	✓		✓		✓			
Su et al. [75]	✓			✓				
Xiao et al. [96]	✓							
Naqvi et al. [60]	✓							
Wang et al. [89]	✓							
Chen et al. [16]	✓							
Tang et al. [79]	✓	✓						Quality control data
Wang et al. [87]	✓							Organization information
Han et al. [36]	✓					✓		
Gong and Wang [30]							✓	
Emmanouilidis et al. [26]	✓		✓					
Zhou et al. [103]				✓		✓		
Arroyo et al. [4]			✓					System structure behavior
Sarazin et al. [69]	✓		✓	✓				
Li et al. [53]	✓							
Tang et al. [81]	✓	✓		✓				
Geng et al. [29]						✓	✓	
Tang et al. [80]	✓							
Lim et al. [55]	✓				✓			
Wang et al. [91]	✓							
Yang et al. [99]	✓		✓					Design knowledge
Wu et al. [93]							✓	
Zhang et al. [100]	✓			✓				
Han et al. [35]	✓					✓		
Chen et al. [15]	✓							
Sun and Wang. [76]	✓							textTE process data
Ebrahimipour and Soumaya. [25]	✓							
Chen et al. [13]	✓					✓		
Zhou et al. [104]	✓	✓						
Geng and Fu [28]	✓			✓				Network knowledge
Yang et al. [97]	✓		✓					
Cai et al. [10]	✓		✓					
Jia and Yang [43]							✓	
Beden et al. [6]			✓					Relational database
Wang et al. [90]				✓				
Nagy et al. [59]							✓	
Xia et al. [94]	✓		✓					
Zhou et al. [101]							✓	
Shi et al. [71]		✓		✓				Development and production database, experimental tests, in-orbit operation database, design reliability data, production process data, in-orbit management data
Zhou et al. [102]	✓							fault investigation sheets
Zhou et al. [105]			✓				✓	
Deng et al. [21]	✓					✓		
Li et al. [52]			✓	✓				
Total number	31	5	12	9	3	6	8	9

We categorized the studies that employed ontologies (32 out of 45 papers) according to their primary objectives. Table 8 summarizes these categories and provides a brief description of each objective, together with the advantages and limitations.

As can be seen, 11 papers [14, 25, 26, 28, 59, 74, 75, 79, 90, 104, 105] employed ontologies as the primary framework for structuring and representing domain knowledge and data. In all of these studies, the upper-level ontology (including the main classes, subclasses, relationships, and attributes) was manually designed with the involvement of domain experts. Following the definition of the ontology schema, the relevant data were automatically extracted, organized, and mapped to the corresponding ontology elements. Although these studies employed ontologies as the primary knowledge representation framework, the diag-

nostic approaches integrated with the ontology varied considerably. Several studies used SPARQL query language [85] to retrieve ontology-based knowledge, whereas others combined ontology with probabilistic or analytical techniques, such as Bayesian theory and fault tree analysis [67], to infer fault causes and assess system behavior under uncertainty. The specific diagnostic techniques integrated with ontology-based representations are discussed further in Section 4.4.

Sixteen studies [10, 13, 15, 21, 29, 30, 52, 53, 55, 71, 81, 87, 93, 94, 97, 102] employed ontologies primarily as a semantic schema for knowledge graph construction. In these approaches, the ontology defines the core domain concepts and relationships that provide the structural basis for the KG and support the integration of heterogeneous fault-related information. In these studies, the ontology mainly serves as a lightweight semantic backbone that supports downstream tasks such as entity organization, data integration, and information retrieval within the KG representation.

Table 8: Ontology representation for semantic-based fault diagnosis of CPSs and the pros and cons

The Purpose of Using Ontology Representation	Short Description	Pros(+)/Cons(−)	References
Ontology as the primary knowledge representation	In twelve papers ontology served as the main framework for organizing and formalizing heterogeneous fault diagnosis knowledge. These ontologies were typically developed by defining classes, subclasses, individuals (instances), attributes, and relations, all encoded using the OWL and developed with the involvement of domain experts. Several studies also incorporated Semantic *Semantic Web Rule Language* (SWRL)-based reasoning, where the SWRL is used to support rule-based inference over ontology knowledge. In addition, some studies reused or extended established ontologies and standards, such as ISO 14224, ISO 15926, and the SOSA/SSN ontology.	+ This structured approach ensures that diverse and complex knowledge is represented in a unified, extensible, and machine-interpretable way. + Facilitates integration of data and knowledge from multiple sources. + Supports semantic reasoning and inference of implicit knowledge. − Requires domain expertise and careful structuring to ensure consistency and completeness. − It can become difficult to maintain and extend as the ontology grows, leading to scalability issues.	[14, 25, 26, 28, 59, 74, 75, 79, 90, 104, 105]
Ontology-guided knowledge graph construction	sixteen studies employed ontology primarily as a semantic schema for structuring knowledge graphs or enabling semantic inference over knowledge graph representations. In these approaches, the ontology defines the domain concepts, entity types, relationships, constraints, and attributes relevant to fault diagnosis. The ontology guides entity extraction, relation extraction, data integration, and KG construction from heterogeneous data sources. Some studies additionally used ontology-driven reasoning to discover hidden relationships and implicit knowledge.	+ Supports reasoning over the KG and enhances downstream querying. + Provides a formal schema that improves the consistency of entity and relation extraction. + Enables semantic integration of heterogeneous data sources into a structured KG. − Designing the domain-specific ontology is time-consuming and resource-intensive. − Highly structured ontologies may reduce flexibility and make it harder to represent novel or unexpected entities and relations.	[10, 13, 15, 21, 29, 30, 52, 53, 55, 71, 81, 87, 93, 94, 97, 102]
Other ontology applications	Several studies used ontologies to support specific functions in conjunction with other methods or techniques. For example, Beden et al. [6] built an ontology and reused Time ontology to model domain knowledge as a semantic layer over distributed relational tables, enabling unified access and querying without physical data integration. Dendani et al. [20] formalized domain knowledge using ontology to organize cases in the CBR, thereby enabling semantic similarity matching between fault cases. Zhou et al. [103] employed ontology as a corpus annotation schema to define a structured taxonomy of fault-related entities for consistent corpus construction and supervised entity labeling.	+ Enables semantic integration across heterogeneous databases without physical data fusion. + Supports structured case organization and semantic similarity matching in CBR. + Facilitates consistent corpus annotation and supervised learning tasks. − Depends on carefully designed schemas to preserve semantic consistency. − Ontology-based access to relational databases introduces additional mapping and query translation overhead.	[6, 20, 99, 103]

A smaller group of studies used ontologies to support specific functions in conjunction with other methods, including *Case-Based Reasoning* (CBR) [20, 99], corpus annotation for supervised learning [103] and querying heterogeneous relational databases [6]. In CBR-based fault diagnosis, the ontology was used to represent the domain knowledge and to integrate this knowledge with the case model. The cases model was represented using the ontology, allowing the case base to be described using the concepts and relationships defined in the domain model. The ontology was also used for calculating case similarity, supporting semantic similarity assessment during retrieval. Zhou et al. [103] employed an ontology to define entities of the industrial robot fault diagnosis domain and used it as a basis for corpus annotation. The ontology defined a hierarchical taxonomy of fault-related entities and was then used to standardize corpus annotation and label industrial fault records. In this context, the ontology provided a consistent semantic structure for organizing fault terminology and generating annotated training data for deep learning models such as BERT-BiLSTM-CRF. Beden et al. [6] used ontology as a semantic layer for integrating and accessing data from heterogeneous industrial databases. Through mappings implemented in the Ontop framework [11], a virtual ontology-based data access system, relational database tables were linked to ontology classes and properties, allowing the relational data to be virtually represented as RDF triples without physically integrating the databases. This enables users to query heterogeneous industrial data through the ontology using SPARQL at the virtual level, without physically integrating the underlying databases.

These studies demonstrate the flexibility of ontologies in serving different roles within semantic-based fault diagnosis of CPSs. Depending on the objective, ontologies were used as primary knowledge representation frameworks, semantic schemas for knowledge graph construction, or supporting components for case-based reasoning, corpus annotation, and semantic access to relational database. Across these roles, ontologies provided structured representations of domain concepts and relationships, supporting tasks such as knowledge organization, semantic retrieval, data integration, and interoperability. This is particularly important in CPS environments, where fault-related information originates from multiple heterogeneous systems and formats.

However, the analysis also reveals several limitations. Most ontology construction processes rely on expert involvement and manual engineering, making them time-consuming, although this approach ensures that the resulting ontology remains closely aligned with the specific domain. As CPSs become increasingly complex and dynamic, maintaining large ontologies and ensuring semantic consistency becomes a significant challenge. In addition, highly structured ontological models may reduce flexibility when dealing with evolving systems, as changes in system components or relationships may require corresponding modifications to the ontology.

4.3.2. ***Knowledge Graphs***

A knowledge graph is a graph of data intended to accumulate and convey knowledge of the real world, whose nodes represent entities of interest and whose edges represent potentially different relations between these entities [39], organized in a standardized format known as triples [86]. These triples, structured as ⟨`subject, predicate, object`⟩, express facts such as ⟨`leaking, has cause, crack`⟩. This structure is well-suited for capturing and modeling in formation from large-scale datasets.

Among the reviewed studies, twenty-two papers employed KGs as the primary knowledge base for integrating and structuring fault-related information. Their widespread use is mainly attributed to several key advantages, including the ability to integrate heterogeneous data sources [29, 35, 55, 87], support efficient knowledge retrieval [87], and facilitate updating and evolution of data and knowledge [100].

However, our analysis reveals that, despite sharing a common objective of constructing knowledge graphs as a semantic knowledge base for fault diagnosis, the approaches used to construct them differ. Accordingly, we categorize these works based on their KG construction methods, as summarized in Table 9.

Sixteen papers [10, 13, 15, 21, 29, 30, 52, 55, 71, 81, 87, 93, 94, 97, 101, 102] adopted a schema or ontology-driven approach, in which an ontology is manually designed to guide KG construction, and domain data are subsequently extracted and mapped to the predefined structure. In contrast, six studies [35, 36, 69, 76, 91, 100] constructed knowledge graphs without relying on a predefined schema. While this schema-free approach offers greater flexibility in handling heterogeneous data, it typically increases the complexity

of the extraction process and introduces additional challenges related to redundancy, inconsistency, and ambiguity resolution [36, 91].

In addition, we observed differences in how the KGs are structured and deployed across studies. Seventeen articles [10, 13, 15, 21, 29, 30, 35, 36, 52, 69, 71, 81, 87, 91, 93, 94, 97] used a single knowledge graph architecture to represent the entire domain, typically formalized using standard ⟨subject, predicate, object⟩ triples. In contrast, a smaller subset of studies employed decomposed or multi-graph architectures. Three papers constructed multiple knowledge graphs. For example, Lim et al. [55] built separate knowledge graphs for individual assets, rather than introducing a single centralized KG. Zhou et al. [101, 102] built a two-graph architecture that distinguished between a knowledge graph capturing structured equipment information and a causal graph modeling cause–effect relations among fault phenomena, causes, and maintenance actions, with both graphs subsequently integrated for fault diagnosis. Other studies introduced structural variants within a single KG. For example, the multi-level knowledge graph proposed by Sun et al. [76] represents different abstraction layers of industrial systems including production, process, energy, and materials, capturing cross-level dependencies. Similarly, Zhang et al. [100] introduced edge weights to represent fault frequency, enabling the KG to distinguish between high and low frequency failure patterns.

Once knowledge graphs are constructed as the primary framework for integrating and organizing domain knowledge, various approaches have been employed to support fault diagnosis. A detailed discussion of these approaches is provided in Section 4.4. The most prevalent approach, reported in 11 papers [13, 21, 30, 36, 52, 71, 81, 87, 91, 100, 101], involved directly querying the knowledge graph using query languages such as SPARQL, Cypher [61], and AQL [3], depending on the underlying graph database. The second most common strategy is knowledge graph embedding, applied in six studies [29, 35, 53, 93, 94, 102], although the specific embedding models and optimization objectives differ across implementations and are further discussed in the reasoning section. Two studies [76, 97] utilized the constructed knowledge graphs to build Bayesian Networks for probabilistic reasoning. In addition, three papers [10, 55, 69] explored other approaches, including graph-based algorithms and rule-based inference, to support fault diagnosis.

It is important to note that while many papers aimed to construct KGs for FD, some primarily focused on techniques for entity and relation extraction from input data, a critical step in KG construction. These studies often did not discuss how the resulting knowledge graphs were subsequently utilized for fault diagnosis. Therefore, we do not include them in this section as part of the KG representation.

Overall, the reviewed studies demonstrate that knowledge graphs have been widely used to structure domain knowledge in semantic-based fault diagnosis, particularly due to their ability to integrate heterogeneous data, support efficient knowledge retrieval, enable continuous updates, and scale to large and complex industrial systems. However, differences in construction strategies (ontology-guided vs. data-driven), KG architecture (single vs. multiple graphs), and reasoning mechanisms (query-based, embedding-based, probabilistic, and rule-based) reflect different design priorities and trade-offs across the reviewed studies.

4.3.3. ***Other Representations: xSDG, Hypergraph, and Bayesian Network Representations***

By reviewing the papers, we identified three other representations, namely *extended Signed Digraph* (xSDG), Hypergraph, and Bayesian network, each presented in a single paper. Accordingly, these approaches are discussed jointly in this subsection due to their limited occurrence in the set of reviewed papers of this study. Notably, Bayesian networks are widely used as a diagnostic reasoning method in the fault diagnosis domain; however, they are less frequently adopted as a primary knowledge representation formalism than knowledge graphs or ontologies, which are more commonly used for integrating and representing knowledge.

Arroyo et al. [4] proposed a model-based diagnostic approach by modeling knowledge using xSDG which enhances causal modeling by introducing additional node and arc values, offering a more detailed qualitative description than standard SDGs, which uses simple "+" or "-" signs to indicate the direction of influence between nodes. The xSDG uses five node values–`HighHigh, High, Normal, Low, and LowLow`–to express different states of deviation, and seven arc values–`Positive strengthening, Positive constant`, and etc, to characterize the strength and direction of causal influences. However, determining the appropriate

Table 9
Knowledge graph representation for fault diagnosis of CPSs and the pros and cons

KG Construction method	Short Description	Pros (+)/Cons (–)	References
KG construction by adhering to a schema	In sixteen papers, the construction of the KG began with the definition of domain knowledge through an ontology. This process involved specifying the key concepts, entities, relationships, attributes, and constraints. After establishing the ontology, various information extraction techniques, described in Section 4.2, were applied to collect and extract data from different sources, resulting in the construction and population of the knowledge graph. Once the KG was established, different reasoning techniques were employed to support fault diagnosis. The specific reasoning and diagnostic approaches adopted in these studies are discussed in detail in Section 4.4.	+ Ontology provides a well-defined semantic schema that ensures consistent representation of entities and relations across the KG. + Supports explicit reasoning and inference due to structured ontology constraints. – Requires manual effort for ontology design and maintenance. – Ontology definition relies on domain experts, making construction time-consuming.	[10, 13, 15, 21, 29, 30, 52, 55, 71, 81, 87, 93, 94, 97, 101, 102]
KG construction without having a schema	Six papers adopted a fully-automatic approach to knowledge graph construction without relying on a formally defined ontology at the initial stage. In these works, the graph structure is derived directly from textual industrial data through information extraction techniques. So the resulting representation closely reflects the scope, granularity, and reliability of the underlying data sources, as well as the extraction process used to generate them.	+ Reduces upfront effort by avoiding manual ontology engineering. + Enables faster adaptation to new or evolving data sources. – Often results in weaker semantic consistency across the graph. – Limited interoperability due to lack of shared conceptual schema. – More prone to noisy or ambiguous relation structures from extraction errors.	[35, 36, 69, 76, 91, 100]

arc value can be complex. Fuzzy set theory is employed to address this complexity by providing a systematic way to assign these values based on observed process data. Moreover, arc consistency (using product operator) is checked to validate that the relationships between cause nodes, connecting arcs, and effect nodes are consistent. This consistency ensures that deviations or disturbances can be accurately propagated through the graph. Formally, an arc is said to be consistent if the product of the cause node, the directed arc, and the effect node is positive. Finally, to extend the model, causal relationships are added based on first-physical principles, such as mass and energy conservation laws. Although the presented method provides a structured procedure to derive causal diagnostic models, it still involves several manual steps, including the definition of arc values. The authors note that the required modeling effort is lower than that of full first-principles modeling.

Li et al. [53] proposed a *Knowledge Hypergraph* (KHyG) as an alternative to conventional triple-based knowledge graph for representing multi-nary relations in fault diagnosis. At the schema level, hyperedges capture structural constraints among multiple entity types; for example, the `Only Works When` hyperedge preserves the relationship among a symptom, its cause, and the corresponding troubleshooting action. At the instance level, hyperedges connect specific entity instances involved in multi-nary relations. For instance, an instance-level hyperedge may link the symptom `Hanging Cable Break` to two causes, `Moving to Inappropriate Positions` and `Not-enough Reserved Torque`, under the relation `Has Joint Causes`. However, constructing KHyGs involves substantial manual effort from domain experts, particularly for extracting schema-level hyperedges and ensuring the quality of the extracted multi-nary relations. In addition, from an implementation perspective, mature toolkits for storage, querying, and visualization of KHyGs are still lacking compared compared with those available for conventional KGs such as GraphDB. This limits their integration into real-world industrial knowledge-based systems.

Xiao et al. [96] proposed a knowledge representation approach that maps association rules mined from maintenance data into a BN to support diagnostic decision-making. The association rules, extracted using a modified Apriori algorithm focused on 2-itemsets, are characterized by support and confidence measures, which are subsequently reused to parameterize the BN. In this formulation, rules of the form $X \rightarrow Y$ are represented as directed edges, where antecedent and consequent correspond to connected nodes, while support represent the joint occurrence probability, confidence represent conditional probabilities. This allows association-based knowledge extracted from historical maintenance records to be represented as a probabilistic graphical model that supports inference about fault–solution relationships. However, the approach remains tightly coupled to the quality and coverage of the resulting association rules, which is itself dependent on the completeness and consistency of the maintenance data. Moreover, because the method generates rules from 2-itemsets and constructs a two-layer BN, more complex dependencies involving multiple variables are not explicitly represented, and the resulting structure may become less expressive in scenarios involving complex fault interactions.

4.4. ***Diagnosis***

In previous sections, we examined knowledge extraction, acquisition, and representation. The ultimate objective of these three tasks is to support effective fault diagnosis. This section presents our findings on various fault diagnosis methods identified across the selected studies. Fig. 9 summarizes these methods, which are grouped into four main categories, including semantic/knowledge-driven methods, representation learning methods, similarity and structure-based methods, and hybrid methods. Each category is discussed in the following subsections.

4.4.1. ***Semantic-based Methods***

Semantic-based fault diagnosis methods exploit the structured semantics encoded in knowledge representations. Rather than relying on large-scale training data, these approaches typically operate over explicitly defined entities and relationships, enabling diagnosis through structured queries or logical inference mechanisms. To provide a clearer understanding, we categorize these methods into four groups: query-based approaches, Bayesian frameworks, rule-based inference methods, and hybrid semantic approaches, as summarized in Table 10.

Querying. This method has been used to retrieve diagnostic knowledge from semantic knowledge bases, such as knowledge graphs and ontologies, where domain knowledge is explicitly modeled through entities and relationships. In total, we included 16 papers in this category, however, the implementation details vary across papers. Some studies [36, 74, 104] explicitly implement query mechanisms, such as SPARQL or Cypher, to retrieve relevant knowledge from the underlying knowledge base. Others [52, 71, 81] provide a search interface for accessing diagnostic information without specifying the underlying query language. In some studies [13, 14, 25, 26, 30], querying is only identified as a supported capability, with limited details on its implementation. By contrast, several studies [21, 91, 100] stored the constructed KG in Neo4j and visualize it without reporting a query mechanism for fault diagnosis. Several studies [14, 79, 104, 105] incorporated SWRL rules with the ontology to enhance inference capabilities.

Wang et al. [87] stored the KG in a Neo4j graph database and developed an application system supporting knowledge retrieval and question answering. Natural-language fault queries were analyzed using word segmentation, dependency analysis, entity linking, semantic rules, and entity similarity to generate structured query statements. These queries were then executed against the KG using Cypher to retrieve relevant fault-diagnosis knowledge and provide answers to users.

Beden et al. [6] mapped the constructed ontology to relational databases through the Ontop framework. SPARQL queries were executed over the virtual knowledge graph, with Ontop translating these queries to retrieve relevant information directly from the underlying relational databases. This enabled semantic querying without physically transforming the relational data into a separate knowledge graph.

Overall, the reviewed studies indicate that semantic knowledge bases can support structured access to diagnostic knowledge. However, not all studies reported the extent to which this capability is implemented

or evaluated, with some providing limited details on the underlying retrieval process. Direct querying may require users to have knowledge of both the query language and the structure of the underlying knowledge base, while approaches that provide search interfaces or query-generation mechanisms reduce this requirement but still generally rely on predefined entities, relationships, or semantic patterns. As a result, these approaches primarily support targeted retrieval of diagnostic knowledge rather than exploratory search, limiting the ability to discover potentially relevant information. A short description, advantages, and limitations of query-based methods are summarized in Table 10.

Bayesian framework. Bayesian methods leverage probabilistic reasoning to handle uncertainty but typically require prior knowledge about relationships between variables. In the reviewed studies, this prior structural knowledge is derived from knowledge graphs [10, 76, 97], ontologies [75], or association rules mined from historical data [96]. Although all approaches combine structured knowledge with Bayesian-based reasoning, they differ in how the Bayesian component is modeled and applied.

In [10, 76, 97] a knowledge graph was used to structurally organize fault-related information and narrow candidate faults, after which Bayesian reasoning is introduced to handle uncertainty and rank or infer fault causes. However, the role of Bayesian methods differs across the three works. Sun et al. [76] used Bayesian theory mainly as a posterior-probability ranking mechanism for candidate faults extracted from the KG, without constructing a full probabilistic graphical model. Cai et al. [10] used Bayesian inference to update priors and compute posterior probabilities over candidate faults identified via graph matching.

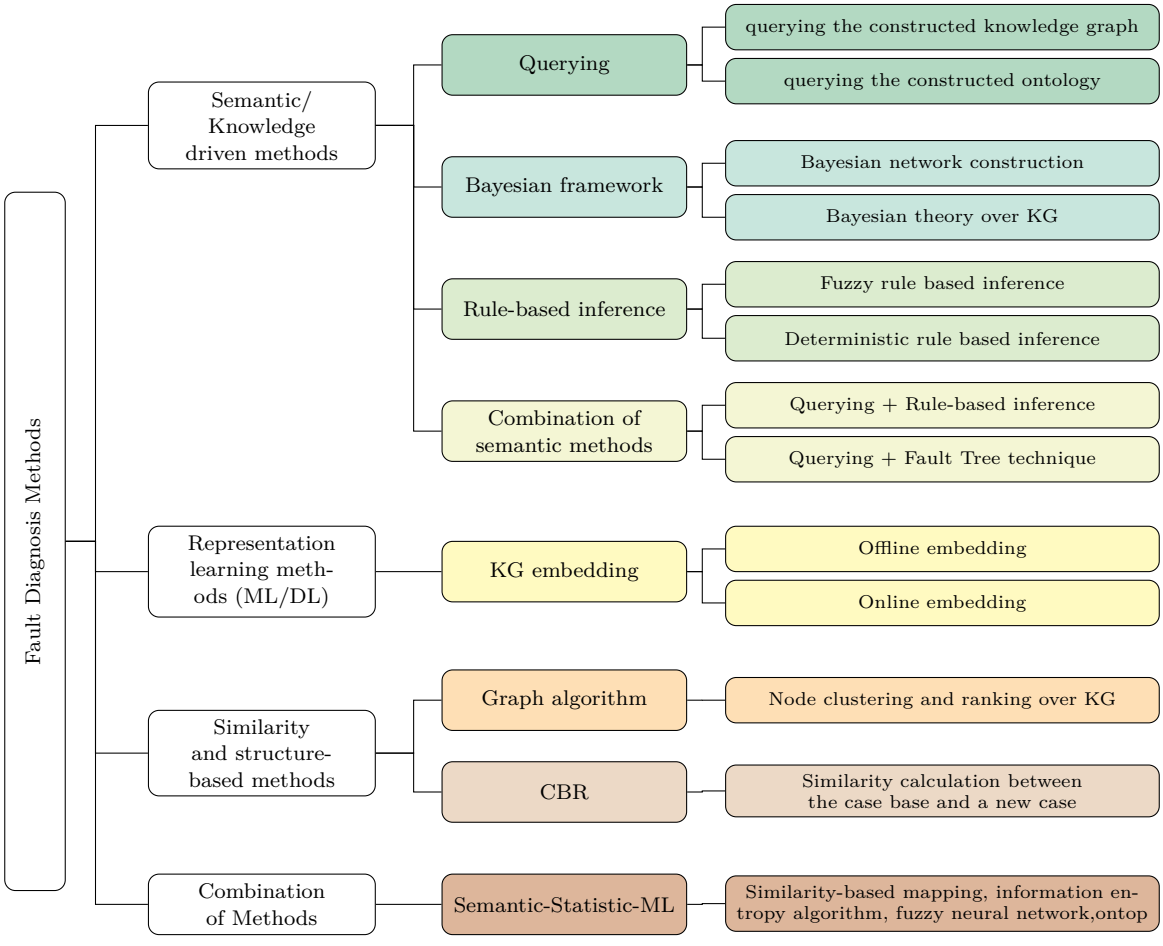


Fig. 9. A categorization of the fault diagnosis methods

In contrast, Yang et al. [97] explicitly transformed the KG into a full three-layer BN, enabling iterative probabilistic reasoning over symptoms, fault modes, and fault causes.

Su et al. [75] mapped observed test information to ontology nodes, and used *Test–Fault* (TF) correlation matrix to identify reachable fault states through information flow analysis. This step reduced the set of possible fault causes based on structural and functional propagation relationships. In parallel, a *Fault–Maintenance* (FM) mapping layer was constructed by combining ontology knowledge with a Bayesian network, where the network is primarily used to encode relationships between concepts and support similarity propagation for ontology alignment. The Bayesian network does not perform explicit probabilistic fault inference; instead, it facilitates iterative refinement of fault–maintenance mappings by leveraging indirect dependencies in the ontology structure. The final diagnosis was obtained by combining deterministic graph-based propagation with ontology-driven FM mapping, providing both fault identification and corresponding maintenance guidance.

Xiao et al. [96] proposed a diagnosis framework in which association rules extracted from historical maintenance data are represented as a Bayesian network. Frequent relationships between faults and solutions are first mined using an Apriori-based association rule method, and the resulting support and confidence values are used to define probabilistic dependencies in the BN. During diagnosis, observed fault information is used as evidence in Bayesian reasoning to compute posterior probabilities of candidate solutions. The BN is simplified by retaining only nodes relevant to the observed conditions, and the final maintenance solution is selected based on inferred probabilities and maintenance cost considerations.

Rule-based inference. This method relies on explicitly defined inference rules. Two types of rule-based inference were used : fuzzy rules and deterministic rules. Fuzzy inference leverages fuzzy logic to handle imprecise information, whereas deterministic inference relies on predefined if-then rules that produce exact outcomes without uncertainty.

Arroyo et al. [4] performed diagnosis using a rule-based reasoning process on an xSDG model. Fault symptoms are propagated through the network according to consistency rules that only allow causally compatible node–arc–node relations, while inconsistent paths are blocked to avoid invalid fault propagation. A fuzzy layer is then used to manage uncertainty by assigning graded influence to partially consistent relations. The final fault candidates are obtained from reachable root-cause nodes and ranked based on aggregated fuzzy consistency and causal strength.

Sarazin et al. [69] constructed a *Heterogeneous Information Network* (HIN) as the base of facts and a set of if–then rules as the base of rules for reasoning. Diagnosis was performed by first retrieving relevant nodes and edges from the HIN based on the observed symptoms and then extracting the corresponding rules associated with these subgraphs. If multiple fault candidates are retrieved, rule-based reasoning is used to filter and rank possible solutions. In cases where multiple candidates are still valid, probabilistic scoring (e.g., based on historical occurrence data) was introduced to prioritize the most likely fault.

Combination of semantic methods. Geng et al. [28] combined SPARQL queries and fuzzy logic for fault diagnosis. SPARQL queries are used to retrieve relevant entities and relationships from the ontology for the observed symptoms, while Prolog [47], a logic programming language, was used to apply logical inference rules to derive fault causes from the observed symptoms. For fuzzy or uncertain fault information, the method extends this process using fuzzy membership functions. Numerical sensor values are first converted into fuzzy linguistic values through fuzzification, then mapped into ontology attributes. These fuzzy descriptors are matched against fuzzy inference rules, allowing the system to infer fault causes even when symptoms are imprecise.

Nagy et al. [59] applied ontology reasoning mechanisms, mainly SWRL rules for logical inference and SPARQL queries for retrieving and filtering relevant entities and relationships. The fault tree analysis was then used as a supporting structural knowledge to guide the identification of root causes and structure the reasoning process. The resulting fault hypotheses are then ranked based on rule satisfaction and semantic relationships, and can be linked to recommended maintenance or response actions stored in the ontology.

Wang et al. [90] transformed the equipment structure and functional dependencies extracted from technical manuals into a semantic diagnostic network. The resulting directed graph was represented using

adjacency, path, and reachability matrices. The hierarchical diagnosis algorithm uses these matrices to analyze functional dependencies, starting from an observed fault and tracing related functions through the directed graph. The reachability matrix identifies functions that can influence or be affected by the abnormal function, while path length indicate their potential influence. This enables the identification of suspicious features and possible fault locations.

4.4.2. ***Machine Learning and Deep Learning Methods***

Machine learning and deep learning have been applied to exploit semantic representations for fault diagnosis. KG embedding was identified as ML/DL technique in the reviewed studies, applied in either an online or offline setting. KG embedding methods transform entities and relationships in a knowledge graph into low-dimensional vector representations. These vector representations enable similarity computation, link prediction, classification, and reasoning tasks using numerical operations [88]. Table 11 provides a short description of each method and their pros and cons. A detailed explanation is provided below.

Offline KG embedding. Offline methods precompute embeddings for the entire KG in advance, avoiding the need to learn embeddings at query time. In the reviewed studies [35, 53, 94, 101, 102], the primary purpose of offline KG embedding was to support answering diagnostic queries by representing the KG and the query in a vector space, enabling similarity-based matching to identify and rank potential answers. Although the studies shared this overall purpose, they differed in how the embeddings were used. In

Table 10
Semantic-based fault diagnosis methods

Semantic Method	Short Description	Pros (+)/Cons (–)	References
Querying	Retrieves diagnostic knowledge from knowledge bases, such as knowledge graphs and ontologies, by accessing entities and relationships through structured query mechanisms or search interfaces. Across studies, implementations range from explicit SPARQL/Cypher queries to search, while some studies construct and visualize knowledge graphs without reporting a diagnostic query mechanism, others mention its support for querying.	+ Enables direct retrieval of relevant diagnostic knowledge without requiring model training. + The same query always returns the same results, improving reproducibility and reliability. + Works well with established semantic web technologies (e.g., SPARQL, OWL ontologies, RDF stores). – Explicit query mechanisms require knowledge of query languages such as SPARQL or Cypher and the structure or the knowledge base, unless a custom interface is built. – Search or query results depend on the exact pattern match.	[6, 13, 14, 21, 25, 26, 30, 36, 52, 71, 74, 81, 87, 91, 100, 104]
Bayesian framework	Bayesian-based diagnosis methods combine semantic representations such as knowledge graphs, ontologies, or association rules with probabilistic reasoning to handle uncertainty in fault diagnosis. Semantic structures are first used to organize relationships between entities such as symptoms, faults, causes, and maintenance actions, while Bayesian inference or Bayesian networks are then applied to compute posterior probabilities, rank candidate faults, or iteratively update diagnoses based on observed evidence.	+ Supports ranking of candidate faults instead of only deterministic matching. + Can iteratively refine diagnoses when new evidence becomes available. – Requires prior knowledge or probabilistic dependencies to define the Bayesian model.	[10, 75, 76, 96, 97]
Rule-based inference	Rule-based diagnostic methods infer faults by applying predefined if–then rules. Some approaches define deterministic rules while others extend this process with fuzzy or scoring mechanisms to handle uncertainty or imprecise cases.	+ Fuzzy rules can handle imprecision. + Deterministic rules make the reasoning process straightforward, interpretable, and transparent for users. – The quality and accuracy of the system depends on the completeness, accuracy, and relevance of the predefined rules. – Predefined rules may struggle to adapt to novel situations, faults, or data patterns not explicitly covered in the KB.	[4, 69]

particular, two studies [35, 101] used KG embedding to support multi-hop question-answering. Han et al. [35] used ComplEx to learn representations of the entities and relationships in the KG, which are then used to match query entities and relationships with relevant KG facts. An RL-based reasoning module subsequently traverses the KG to identify candidate answers along multi-hop paths. Zhou et al. [101] identified and transformed the entities and relationships in the question into a question embedding. The central entity in the query is then used as the starting node, and a relevant multi-hop subgraph is processed to obtain relationship-aware node embeddings. The approach subsequently formulates multi-hop reasoning as a link-prediction task, adopting the idea of TransE [8] to score and rank potential answer nodes or paths. Xia et al. [94] leveraged KG embedding for link prediction to enrich the KG, which further aids in the computation of similarities between the embedded KG and query vectors.

Online KG embedding. In contrast to offline methods, online KG embedding methods generate vector representations during query processing. Geng et al. [29] utilized KG embedding for fault classification. When a new issue is observed, the system extracts relevant information by identifying feature words, connecting them to entities, and creating a subgraph. The TransD model is then used to learn a representation of this subgraph, which is combined with textual features in a CNN for fault classification. However, they did not implement online embedding in the sense of training on the graph to derive answers; rather, they aimed to identify a more accurate class for the given fault.

4.4.3. ***Similarity and Structure-based Methods***

Similarity and structure-based fault diagnosis methods rely on similarity measures and structural properties of knowledge representations to identify relationships between newly observed fault cases and previously observed data stored in a knowledge base. Unlike semantic reasoning or learning-based approaches, these methods do not depend on explicit logical inference or predictive model training, but instead exploit case similarity or graph structural organization to support fault identification and reasoning. Among the reviewed studies, *Case-Based Reasoning* (CBR) and graph algorithm were the two main categories of similarity and structure-based methods. These approaches are discussed in the following.

CBR. CBR is an approach for fault diagnosis that solves new problems by retrieving and adapting solutions from previously encountered cases stored in a knowledge base. Two studies [20, 60] used a CBR framework that includes a knowledge base storing past cases. In these approaches, cases are structured using an ontology-based schema. Cosine similarity [12] is used to compute the similarity between a query case and previously stored cases. This approach can adapt to new problems by retrieving and modifying similar past cases. However, scaling CBR systems can be challenging, as efficient retrieval from large case bases often requires advanced indexing techniques.

Table 11
ML/DL methods for fault diagnosis

ML/DL Method	Short Description	Pros (+)/ Cons (–)	References
Offline KG embedding	The entire KG is encoded into a fixed-dimensional vector space where entities and relations are represented as continuous vectors. During diagnosis, query vector representations are generated and matched against precomputed KG embeddings using similarity measures or scoring functions. The reviewed studies use this method for similarity-based retrieval, multi-hop reasoning, and link prediction.	+ Efficient at query time since embeddings are precomputed. – The resulting representations provide limited interpretability of the reasoning process.	[35, 53, 94, 101, 102]
Online KG embedding	Online KG embedding methods generate vector representations during query processing. In the reviewed study, a query-related subgraph is constructed and embedded using TransD, with the resulting representation combined with textual features for fault classification.	+ Allows the representation to focus on information relevant to the current fault case. – Requires additional computation during diagnosis to construct and encode the relevant subgraph.	[29]

Graph algorithm. Graph based method in fault diagnosis aims to group similar entities in a knowledge graph based on structural or semantic proximity, so that faults, symptoms, and related components that are closely connected can be analyzed together. These methods help reduce the search space for potential fault causes by identifying densely connected subgraphs or communities. Lim et al. [55] applied the Louvain algorithm [7], a greedy optimization method for community detection, in the knowledge graph to identify clusters of nodes closely associated with the observed fault. The algorithm iteratively moves nodes between communities to maximize network modularity. Modularity measures how effectively the graph is partitioned into distinct clusters, with higher values indicating stronger separation between communities. This clustering process helps narrow down potential fault causes. The PageRank algorithm [64], an iterative graph-based ranking algorithm that assigns importance scores to nodes based on the number and quality of edges, is then used to rank nodes within the identified communities according to their importance and influence in the network. Candidate solutions are subsequently validated using simulation capabilities integrated into the individual asset digital twins.

The method provides no explicit probabilistic confidence measure for diagnostic outcomes, meaning that different plausible fault propagation paths cannot be quantitatively compared in terms of likelihood. As a result, the final diagnosis may be sensitive to small structural changes in the graph and requires external validation

4.4.4. ***Hybrid Diagnostic Approaches***

Hybrid diagnostic approaches combine data-driven processing with knowledge-based and/or model-based reasoning.

Yang et al. [99] proposed an integrated fault diagnosis approach that explicitly separates and coordinates shallow (case-based) and deep (model-based) knowledge within two parallel knowledge bases, each with its own inference mechanism. The shallow layer is an ontology-based fault case system that stores structured maintenance experience and supports semantic retrieval of similar historical cases using case-based reasoning to quickly reuse prior solutions. The deep layer is a model-based diagnostic system that represents fault detection and fault propagation models, enabling causal reasoning over sensor-derived symptoms to infer root causes and propagation paths. In application, the system dynamically selects or combines these two reasoning modes: case-based retrieval is preferred when sufficient historical fault data exists, while model-based causal reasoning is emphasized for new or complex failures where past cases are unavailable or insufficient. A coordination mechanism links both layers, allowing fault symptoms and diagnostic outputs to be exchanged between them.

Wu et al. [93] proposed a multimodal knowledge graph-based fault diagnostic model that combines visual object detection, knowledge graph reasoning, and text generation. An improved Faster R-CNN is used to detect primary and nonprimary objects in equipment images and extract feature vectors representing the detected visual information. These feature vectors are mapped to entities, attributes, and relationships in the KG using cosine similarity. Based on the matched KG elements, a triple is formed from the main device entity, secondary device entity, and fault-phenomenon attribute value. Semantic matching templates, path queries, and multihop queries are then used to determine the corresponding fault entity. Finally, the inferred information is converted into an ordered text feature representation and processed by BART to generate an interpretable fault diagnosis text.

Zhou et al. [105] proposed a hybrid diagnosis method where raw signals are first processed in the fault identification stage using selectable or plugin-based algorithms (such as statistical feature extraction combined with machine learning classifiers), which convert sensor data into labeled fault indicators. These outputs are then mapped to an ontology-based KB. In the fault diagnosis stage, reasoning engines apply OWL axioms and SWRL rules (and when uncertainty is present in sensor data or fault interpretation probabilistic or fuzzy methods are used) to infer fault causes, affected components, and likely failure mechanisms. Finally, in the fault solving stage, the system retrieves maintenance strategies and reports the results, while also allowing feedback-based knowledge updates to refine future diagnosis cycles.

It should be noted that several studies [15, 16, 43, 80, 89, 103] focused on information extraction or knowledge enhancement and do not explicitly specify a diagnostic reasoning mechanism. Nevertheless,

their methods contribute indirectly to diagnosis by improving the completeness, quality, and structure of the knowledge graph used for downstream retrieval and reasoning. For example, entity-relation extraction in [15] achieves a higher F1-score than competing extraction models, suggesting a more accurate knowledge representation. Taken together, these works appear to prioritize the construction of scalable and robust knowledge graphs rather than the development or using an explicit diagnosis method.

4.5. ***Knowledge Enhancement***

Knowledge enhancement is the process of enriching or updating an existing knowledge representation to improve its completeness and quality. The reviewed studies show that only a limited subset explicitly addresses knowledge enhancement, and the corresponding methods are often described only briefly. Among these studies, two main approaches can be identified: Explicit Information Addition and Implicit Knowledge Inference. The following subsections describe each category in detail, and Table 12 provides a short description of the methods, along with their advantages and limitations.

4.5.1. ***Explicit Information Addition***

This approach focuses on directly enriching the knowledge representation by inserting new explicit information such as cases, facts, entities, or relationships, or by updating existing knowledge elements. Although the reviewed studies share the general goal of improving diagnostic capability through enriched knowledge representations, they differ in how this enrichment is achieved.

Dendani et al. [20] and Naqvi et al. [60] acknowledged the crucial role of case base maintenance and expansion in improving the CBR system's overall performance. Their CBR cycle explicitly includes the *Retain* stage, highlighting the process of adding new cases to the case base. This incremental addition increases the coverage of previously unseen scenarios, improves retrieval accuracy through richer precedent availability, and supports adaptation to evolving operational conditions. In this setting, knowledge enhancement is achieved through continuous accumulation of experience.

Sarazin et al. [69] extend a *Heterogeneous Information Network* (HIN)-based expert system by adding new nodes and updating existing ones through dedicated operators. The associate operator integrates newly acquired knowledge into the network structure, while update operators modify node attributes in response to new observations or calculations. For example, the state of a `Maintainable Item` is dynamically updated over time, reflecting operational usage such as flight cycles and maintenance events, where its lifetime is incrementally adjusted and reset following repair or replacement actions.

Sun et al. [76] enhanced the knowledge graph by continuously refining it through mining deep-level associations using weighted critical paths, which capture cross-layer dependencies that are not directly observable in the raw data. During operation, the graph structure was also updated implicitly through changes in variable relationships and weights, which reflect evolving system states and enable more accurate fault detection and diagnosis.

Zhou et al. [105] proposed an architecture that supports the ongoing enrichment and updating of multiple knowledge bases through the knowledge updating module, which synchronizes field data and maintenance feedback with a centralized knowledge platform. In particular, newly observed faults, sensor data, and diagnosis outcomes generated during system operation are mapped to ontological representations and incorporated into existing knowledge structures. In addition, the system allows extension through configurable fault identification plugins and reasoning strategies, enabling new diagnostic knowledge to be added without modifying the core system.

Zhang et al. [100], in contrast, did not propose an explicit knowledge enhancement method; instead, the knowledge graph is used as a continuously updatable representation. Historical and operational data can be incorporated into the graph over time, allowing maintenance engineers to access updated statistical information and past fault records to support diagnosis and repair tasks. In this sense, the knowledge representation is maintained through usage-driven updates rather than a formally defined enhancement mechanism.

Table 12
Knowledge enhancement methods for fault diagnosis

Enhancement Category	Short Description	Pros(+)/Cons(–)	References
Explicit Information Addition	Explicit Information Addition refers to knowledge enhancement methods that directly enrich the knowledge representation by inserting new cases, entities, relationships, or attribute updates derived from operational data, expert input, or system feedback. The knowledge base is expanded or modified in a structured manner, allowing the system to incorporate newly observed fault instances, evolving system states, and updated relational information to improve diagnostic coverage and accuracy over time.	+ Improves diagnostic coverage by continuously incorporating new cases and observed fault scenarios. + Enables straightforward integration of new knowledge without redesigning the entire system (e.g., plugins, retain stages, or update operators). – Risk of inconsistency or redundancy if updates are not properly validated or managed. – In some implementations, updates are implicit or weakly formalized, reducing transparency of how knowledge evolves.	[20, 60, 69, 76, 100, 105]
Implicit Knowledge Inference	Implicit Knowledge Inference enhances the knowledge representation by deriving new relationships or semantic associations from existing graph structures rather than explicitly inserting new knowledge. These methods typically rely on link prediction, graph embeddings, or reinforcement learning to infer missing relations and improve the completeness and connectivity of the knowledge graph.	+ Discovers hidden or previously unknown relationships automatically. + Handles sparse or incomplete graphs through representation learning. – Inferred relationships may introduce uncertainty or incorrect associations. – Learned representations are often difficult to interpret compared to explicit rules.	[36, 71, 94]

Although these studies differ in implementation, they all contribute to explicit knowledge enrichment by adding new cases, updating existing entities and relationships, or synchronizing newly acquired diagnostic information with the knowledge base.

4.5.2. ***Implicit Knowledge Inference***

This approach leverages the existing knowledge within the model to infer new information or relationships that were not previously explicitly stated. It often involves reasoning over the knowledge base to discover hidden patterns, uncover latent relationships, or derive conclusions that were not initially apparent from the explicit data.

Han et al. [36] addressed knowledge enhancement in a fault diagnosis knowledge graph using *Reinforcement Learning* (RL) [77], where an agent learns optimal actions through reward feedback, to infer and complete missing relationships based on the existing graph structure. Moreover, to handle new or unseen entities, such as previously unrecorded faults or equipment types, inductive *Graph Neural Networks* (GNNs) [34] can generate embedding vectors for these new entities based on their available features and neighborhood information, allowing the KG to adapt and address novel scenarios without requiring complete retraining.

Xia et al. [94] addressed the sparsity problem in maintenance-oriented knowledge graphs, where many entities have limited connections. To improve graph completeness, they proposed the ACRGCN model, a GNN model, which predicts missing links and integrates new entities through graph representation learning techniques. By learning latent structural relationships within the graph, the model uncovers implicit relationships between entities and supports the inference of feasible maintenance solutions.

Shi et al. [71] enriched an existing knowledge graph using an embedding translation algorithm combined with semantic constraints to predict missing relationships between entities. Their approach extends the TransE link prediction method to infer implicit semantic relations and improve the completeness of the knowledge graph.

Overall, only nine out of 45 papers explicitly discussed knowledge enhancement for semantic-based fault diagnosis systems. Nevertheless, one of the inherent advantages of knowledge graphs and ontologies is their ability to evolve by incorporating new knowledge over time. The explicit information addition to

KGs and ontologies differs from that of CBR systems. While CBR focuses on retaining and incorporating new cases through a process of retrieve, reuse, revise, and then retain, KGs and ontologies are enhanced by expanding and refining their knowledge structures through the addition of new entities, relationships, attributes, rules, or axioms. This structural flexibility allows them to integrate evolving domain knowledge with minimal modification to the existing representation, making them scalable, adaptable, and suitable for dynamic fault diagnosis environments.

5. Discussion

This section highlights the main findings for all research questions. In addition, it outlines the opportunities and challenges for further developments in semantic-based fault diagnosis of CPSs.

5.1. ***Knowledge Acquisition***

This review highlights a clear gap in the literature regarding knowledge acquisition for semantic-based fault diagnosis in industrial equipment and machinery and aviation CPSs. Although domain experts are involved in several stages of system development, particularly ontology/schema construction, rule definition, and methodology validation, the process of acquiring and formalizing this knowledge has not been discussed. Most studies rely on traditional approaches such as interviews or manual knowledge encoding, while providing no information about the elicitation process itself, including the required effort, scalability, or reproducibility.

The findings reveal a strong dependence on expert knowledge despite limited methodological attention to how such knowledge is captured and maintained. The reviewed literature does not treat expert knowledge acquisition as a core research challenge.

These findings suggest that future research should place greater emphasis on scalable and structured knowledge acquisition methodologies. In particular, AI-assisted approaches, such as LLM-based conversational elicitation systems, interactive annotation platforms, or human-in-the-loop knowledge engineering frameworks, could help reduce the manual burden on experts while improving maintainability of semantic fault diagnosis systems.

5.2. ***Information Extraction***

The reviewed studies on entity–relation extraction (Tables 4, 5, 6) show important trends and innovations. Among the selected studies that addressed entity-relation extraction, several recent papers used transformer-based models such as BERT or ALBERT, often combined with sequence-labeling architectures such as BiLSTM or CRF, which have improved extraction accuracy by leveraging deep contextual embeddings. Moreover, several hybrid architectures, such as MF-GCN-Bi-GCN [15], MPCNN-ATTRA-RL [13], and SBALGN [21] reported better performance than the baselines used in their respective studies. Because the datasets, extraction targets, and baseline models differ across papers, these results support a qualitative observation rather than a direct ranking of extraction methods. However, these observations suggest that well-designed combinations of methods have the potential to push the boundaries of information extraction performance.

The reviewed literature demonstrates a clear methodological evolution from traditional NLP approaches toward deep learning and transformer-based architectures such as BERT and ALBERT. Nevertheless, most works still formulate information extraction as discriminative tasks such as entity or relation extraction, while recent generative AI and LLM-based approaches remain underrepresented in the reviewed corpus as it was conducted in 2024. A brief analysis of newer studies published in 2025 and 2026 (section 6) indicates growing interest in generative AI for fault diagnosis and knowledge extraction, including zero-shot and few-shot learning to reduce dependency on large labeled datasets. However, these newer studies were not examined through full-text analysis and therefore fall outside the scope of the systematic review; consequently, their challenges regarding hallucination, explainability, domain adaptation, and reliability in safety-critical industrial environments cannot be assessed here.

5.3. Knowledge Representation

Knowledge representation for semantic-based fault diagnosis in industrial equipment and machinery and aviation CPSs is dominated by ontologies and knowledge graphs. Ontologies serve both as standalone knowledge base and as semantic layers supporting knowledge graphs and other diagnostic methods. Their use as semantic schemas for KGs highlights their role in providing common semantic structure and interoperability. However, the reliance on manual ontology engineering and domain expertise can limit scalability and adaptation as industrial systems evolve, highlighting the need for more automated and maintainable approaches to ontology development and evolution.

Knowledge graphs, in contrast, are increasingly being used as operational knowledge bases that can be queried, updated, and integrated with different diagnostic methods. Their flexible structure allows them to integrate heterogeneous fault-related knowledge and support querying, learning, and reasoning approaches. Moreover, a brief analysis of studies published in 2025 and 2026 indicates that knowledge graphs can support LLM-based fault diagnosis by providing structured, domain-specific knowledge for information retrieval, entity linking, causal reasoning, and grounding of diagnostic responses. However, the flexibility of KGs results in divers design choices across studies, particularly in how graphs are structured, constructed, and integrated with diagnostic methods.

Alternative representations such as rule-based systems, hypergraphs, Bayesian networks, and signed digraph provide more specialized modeling capabilities but are less widely adopted in the reviewed literature. Rule-based systems, while interpretable, do not scale well with increasing system complexity. More expressive structures such as hypergraphs and extended signed digraph models improve representational power but introduce additional construction and maintenance overhead, often requiring detailed expert specification.

5.4. Diagnosis

The reviewed studies show that knowledge representation serves as a foundation for a diverse range of diagnostic approaches. While ontologies and knowledge graphs provide the underlying semantic structure, the represented knowledge can be queried directly or transformed into other forms, such as vector embeddings or Bayesian networks, to support different types of inference, knowledge retrieval, and diagnostic reasoning. Despite this flexibility, the reviewed studies often make limited use of the broader knowledge available within these representations. In some reviewed studies, diagnosis is performed by matching observed symptoms with directly connected entities or relationships in the knowledge representation, with limited exploitation of broader graph context. This may limits the ability to exploit broader contextual information embedded in the graph structure. In contrast, richer reasoning could be achieved by leveraging subgraph-level context around the observed symptoms, which would allow diagnostic systems to capture indirect dependencies and fault-propagation relationships, thereby providing more informative guidance to service engineers even when direct links are insufficient.

Another important limitation concerns the handling of uncertainty and incompleteness, which are inherent in real industrial CPS environments. Although such conditions are common in practice, they are not consistently reflected in either the representation models or the downstream reasoning processes. As a result, many approaches operate under assumptions of relatively complete and reliable knowledge, which limits their robustness when dealing with noisy, missing, or ambiguous diagnostic evidence.

Overall, the literature shows a clear reliance on semantic representations as the foundation for fault diagnosis, while also revealing a gap between the richness of these representations and the depth of their use in diagnostic reasoning. While knowledge structures are often flexible and reusable across multiple diagnostic methods, their current use in fault diagnosis often remains relatively limited and does not fully exploit their potential for exploratory knowledge retrieval and uncertainty-aware reasoning.

5.5. *Knowledge Enhancement*

Within the scope of this review, most studies primarily focus on constructing the initial knowledge base, while comparatively little attention is given to how the knowledge base should continuously evolve after deployment. We identified two forms of knowledge enhancement: explicit information addition, where new entities, relations, attributes, or rules are incorporated into the knowledge base, and implicit knowledge inference, where new information or relationships are derived from existing knowledge. In ontology and knowledge graph-based systems, explicit mechanisms for knowledge updating and maintenance are less frequently defined as part of the methodology. This may partly reflect the extensibility of these representations, in which new entities, relations, and attributes, can be added without fundamentally restructuring the existing model. This contrasts with other knowledge-based approaches such as case-based reasoning and expert systems, where knowledge enhancement is typically implemented as an explicit stage within the methodology itself, for example through dedicated case-retention or rule-update procedures.

Another important observation is that most studies focus primarily on expanding the knowledge base, while giving comparatively little attention to the consistency and validity of the evolving representation. Although semantic models are often described as flexible and continuously updatable, no work explicitly addresses how newly added knowledge is validated, reconciled with existing structures, or protected from introducing contradictions and redundancies. As knowledge bases grow in size and heterogeneity, maintaining semantic coherence and reliability becomes increasingly important, particularly in dynamic CPS environments where operational knowledge evolves continuously. This highlights an important direction for future research toward developing frameworks capable not only of accumulating knowledge but also of enhancing it reliably over time.

6. Developments and Emerging Trends in 2025–2026: A Brief Overview

This manuscript was submitted to the Semantic Web Journal in March 2025 and remained under review for approximately one year before acceptance in 2026. Since the systematic literature review was conducted based on the literature available at the time of submission, extending the review to cover the additional publication period would require substantial work and could significantly change the scope of the manuscript. Therefore, rather than repeating the complete systematic review, we added this section to provide a discussion of the main developments and emerging research trends in the field during the additional publication period from 2025 to 10th September 2026, when the manuscript was accepted.

To identify recent developments, we applied the same search string (in section 3.1.2) to Scopus, searching the title, abstract, and keyword fields. This search yielded 747 publications in English. We performed the title screening, based on the same inclusion and exclusion criteria defined in section 3.1.3. In this stage, publications were excluded for several reasons, including focus on medical applications, being review or survey papers, addressing domains outside the scope of this research, such as buildings or energy systems, and focusing on remaining useful life estimation rather than fault diagnosis. Following this screening process, 409 papers were included. It should be noted that the screening at this stage was based solely on the title, rather than the multi-stage screening process applied previously, which involved screening both the abstracts and the full texts. As a result, the relevance of some publications could not always be determined with certainty. Therefore, the result of the title screening should be interpreted as an initial filtering step rather than a definitive assessment of study eligibility.

The retained publications were subsequently used to examine the research concepts and methodological trends emerging during the additional publication period.

Figure 10 presents a word cloud generated from the titles of the 409 papers retained after title screening. While terms directly related to fault diagnosis, knowledge graphs, and industrial applications are prominent, particular attention can be drawn to the frequent occurrence of the terms *intelligent*, *large*, *language*, *model*, and *learning*. The prominence of these terms reflects the strong presence of intelligent and learning-based approaches, particularly those involving large language models, in the identified body of

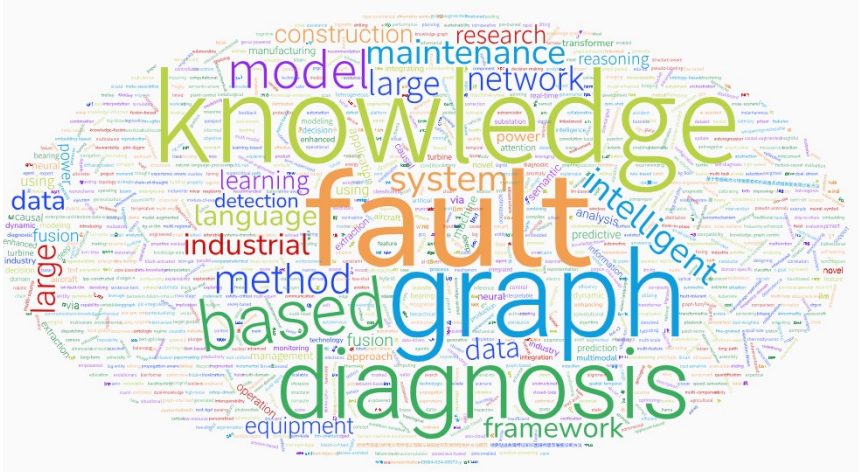


Fig. 10. Word cloud of key terms identified from the titles of papers after title screening

literature. To further investigate this LLM-related trend, we screened the abstracts of the 65 papers that include the keywords *large language model* or *LLM* in their titles, focusing on the role of LLMs in relation to knowledge graphs. Of these 65 papers, 15 papers focus on KG construction, using LLMs to extract and structure fault-related knowledge from unstructured or semi-structured text. Another 26 papers focus on LLM-based fault diagnosis supported by KGs. In these studies, KGs provide structured, domain-specific knowledge that supports different stages of diagnosis, such as retrieving relevant fault information, linking symptoms to KG entities, identifying causal relationships and fault propagation paths, and constraining or grounding the LLM's diagnostic reasoning and response generation. Five papers combine these two roles, using LLMs both to extract knowledge for KG construction and to leverage the resulting KGs for downstream LLM-based diagnosis. A smaller group of three papers investigates LLMs for KG completion and link prediction, where the models are used to infer or predict missing graph information. Two papers explore other KG-related applications of LLMs, with one focusing on graph rectification and the other on diagnostic report generation. The remaining 14 papers discuss LLM applications that are not directly centered on KGs.

7. Limitations

The present review is based on a structured and systematic methodology with a clearly defined scope. Within this scope, several limitations are identified and are outlined below.

First, this review deliberately focused on fault diagnosis in industrial machinery and aviation domains in order to maintain a clear and manageable scope. While this narrow focus strengthens thematic consistency, it may also exclude relevant insights from related diagnostic domains, such as medical diagnosis or other application areas, where transferable methodological ideas may exist. Future work could explore cross-domain connections in a broader comparative review.

Second, while the search strategy was systematic and restricted to peer-reviewed Scopus-indexed studies with well-defined eligibility criteria, regional distribution was not used as a filtering or analytical dimension. Consequently, the resulting corpus exhibits an uneven geographical distribution, with a substantial proportion of studies originating from China, which was not further examined in this study due to the predefined scope of focusing on methodological synthesis rather than regional comparison.

8. Conclusion

The importance of fault diagnosis has been recognized for more than fifty years, and with the emergence of CPS, it has become even more complex and crucial. This paper reviewed the existing literature on semantic-based fault diagnosis methods for industrial CPSs by investigating five main research questions addressing knowledge acquisition, information extraction, knowledge representation, diagnosis, and knowledge enhancement. These questions cover the main stages involved in developing, operating, and maintaining semantic-based fault diagnosis systems. For each research question, the reviewed studies were analyzed and categorized according to their methodological approaches. Where appropriate, categorization tables were provided to summarize the identified methods together with short descriptions, advantages, and limitations.

In addition, this review examined the industrial data and knowledge sources used to develop fault diagnosis systems.These sources were summarized in a dedicated table, providing an overview of commonly used inputs and highlighting potentially underexplored sources in the existing literature. The findings were further synthesized through a discussion of the main observations, limitations, and research gaps identified across the reviewed studies. This synthesis highlights several open challenges and provides directions for future research. In addition, a targeted and brief overview of the literature published in 2025 and 2026 was provided to identify recent developments and emerging trends, with particular attention to the growing role of large language models in semantic-based fault diagnosis.

As future work, broader cross-domain studies on fault diagnosis could be conducted, as well as an analysis of regional differences in research trends to obtain a more balanced global perspective.

Acknowledgements

This publication is part of the project ZORRO with project number KICH1.ST02.21.003 of the research programme Key Enabling Technologies (KIC) which is (partly) financed by the Dutch Research Council (NWO). The authors extend their heartfelt thanks to Dr. Godfried Webers from Philips Healthcare manufacture and Dr. Peter Kruizinga from Canon Producting Printing for their valuable support and feedback on this work.

References

[1] R. Agrawal and R. Srikant, Fast Algorithms for Mining Association Rules in Large Databases, in: *Very Large Data Bases Conference*, 1994. https://api.semanticscholar.org/CorpusID:3131928.

[2] C. Alippi, S. Ntalampiras and M. Roveri, Model-Free Fault Detection and Isolation in Large-Scale Cyber-Physical Systems, *IEEE Transactions on Emerging Topics in Computational Intelligence* **1** (2017), 61–71. https://api.semanticscholar.org/CorpusID:6536453.

[3] ArangoDB, ArangoDB - The Multi-Model Database, 2025. https://arangodb.com/.

[4] E. Arroyo, D.M. Schulze, L. Christiansen, A. Fay and N.F. Thornhill, Derivation of Diagnostic Models Based on Formalized Process Knowledge, *IFAC Proceedings Volumes* **47** (2014), 3456–3464. https://api.semanticscholar.org/CorpusID:26517658.

[5] R.V. Beard, Failure accomodation in linear systems through self-reorganization., PhD thesis, Massachusetts Institute of Technology, 1971.

[6] S. Beden, Q. Cao and A. Beckmann, SCRO: A Domain Ontology for Describing Steel Cold Rolling Processes towards Industry 4.0, *Inf.* **12** (2021), 304. https://api.semanticscholar.org/CorpusID:237458562.

[7] V.D. Blondel, J.-L. Guillaume, R. Lambiotte and E. Lefebvre, Fast unfolding of communities in large networks, *Journal of Statistical Mechanics: Theory and Experiment* **2008** (2008), P10008. https://api.semanticscholar.org/CorpusID:334423.

[8] A. Bordes, N. Usunier, A. García-Durán, J. Weston and O. Yakhnenko, Translating Embeddings for Modeling Multi-relational Data, in: *Neural Information Processing Systems*, 2013. https://api.semanticscholar.org/CorpusID:14941970.

[9] F. Caffaro and G. Rizzo, Knowledge-Enhanced Conversational Agents, *Journal of Computer Science and Technology* **39** (2024), 585–609. https://api.semanticscholar.org/CorpusID:271363083.

[10] C. Cai, Z. Jiang, H. Wu, J. Wang, J. Liu and L. Song, Research on knowledge graph-driven equipment fault diagnosis method for intelligent manufacturing, *The International Journal of Advanced Manufacturing Technology* (2024). https://api.semanticscholar.org/CorpusID:267049129.

[11] D. Calvanese, B. Cogrel, S. Komla-Ebri, R. Kontchakov, D. Lanti, M. Rezk, M. Rodriguez-Muro and G. Xiao, Ontop: Answering SPARQL queries over relational databases, *Semantic Web* **8** (2016), 471–487. https://api.semanticscholar.org/CorpusID:7721653.
[12] S. Ceri, A. Bozzon, M. Brambilla, E.D. Valle, P. Fraternali and S. Quarteroni, An Introduction to Information Retrieval, 2013. https://api.semanticscholar.org/CorpusID:262553219.
[13] C. Chen, T. Wang, Y. Zheng, Y. Liu, H. Xie, J. Deng and L. Cheng, Reinforcement learning-based distant supervision relation extraction for fault diagnosis knowledge graph construction under industry 4.0, *Adv. Eng. Informatics* **55** (2023), 101900. https://api.semanticscholar.org/CorpusID:256714877.
[14] J. Chen, Z. Hu, J. Lu, X. Zheng, H. Zhang and D. Kiritsis, A data-knowledge hybrid driven method for gas turbine gas path diagnosis, *Applied Sciences* **12**(12) (2022), 5961.
[15] S. Chen, L. Cheng, J. Deng and T. Wang, Multi-Feature Fusion Event Argument Entity Recognition Method for Industrial Robot Fault Diagnosis, *Applied Sciences* **12**(23) (2022), 12359.
[16] T. Chen, J. Zhu, Z. Zeng and X. Jia, Compressor fault diagnosis knowledge: A benchmark dataset for knowledge extraction from maintenance log sheets based on sequence labeling, *IEEE Access* **9** (2021), 59394–59405.
[17] X. Chen, S. Jia and Y. Xiang, A Review: Knowledge Reasoning over Knowledge Graph, *Expert systems with applications* **141** (2020), 112948. doi:10.1016/j.eswa.2019.112948.
[18] Y. Chi, Y. Dong, Z.J. Wang, F.R. Yu and V.C. Leung, Knowledge-based fault diagnosis in industrial internet of things: a survey, *IEEE Internet of Things Journal* **9**(15) (2022), 12886–12900.
[19] N.J. Cooke, Varieties of knowledge elicitation techniques, *Int. J. Hum. Comput. Stud.* **41** (1994), 801–849. https://api.semanticscholar.org/CorpusID:2401148.
[20] N. Dendani-Hadiby and M.T. Khadir, A fault diagnosis application based on a combination case-based reasoning and ontology approach, *International Journal of Knowledge-based and Intelligent Engineering Systems* **17**(4) (2013), 305–317.
[21] J. Deng, T. Wang, Z. Wang, J. Zhou and L. Cheng, Research on event logic knowledge graph construction method of robot transmission system fault diagnosis, *IEEE Access* **10** (2022), 17656–17673.
[22] J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: *North American Chapter of the Association for Computational Linguistics*, 2019. https://api.semanticscholar.org/CorpusID:52967399.
[23] Dimensions.ai, Dimensions Knowledge Graph, 2025, Accessed: 2025-01-17.
[24] D. Dou, J. Yang, J. Liu and Y. Zhao, A rule-based intelligent method for fault diagnosis of rotating machinery, *Knowledge-Based Systems* **36** (2012), 1–8.
[25] V. Ebrahimipour and S. Yacout, Ontology-based schema to support maintenance knowledge representation with a case study of a pneumatic valve, *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **45**(4) (2015), 702–712.
[26] C. Emmanouilidis, M. Gregori and A. Al-Shdifat, Context ontology development for connected maintenance services, *IFAC-PapersOnLine* **53**(2) (2020), 10923–10928.
[27] Z. Gao, C. Cecati and S.X. Ding, A Survey of Fault Diagnosis and Fault-Tolerant Techniques—Part I: Fault Diagnosis with Model-Based and Signal-Based Approaches, *IEEE transactions on industrial electronics* **62**(6) (2015), 3757–3767. doi:10.1109/TIE.2015.2417501.
[28] D. Geng and X. Fu, Research on Fault Diagnosis Mechanism of Production Line Equipment Based on Semantic, *2020 IEEE 10th International Conference on Electronics Information and Emergency Communication (ICEIEC)* (2020), 220–223. https://api.semanticscholar.org/CorpusID:220939614.
[29] D. Geng, L. Zhou, Y. Zhao and H. Li, Fault Diagnosis of Motor based on Knowledge Graph and CNN, in: *2022 3rd International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE)*, IEEE, 2022, pp. 742–745.
[30] R. Gong and T. Wang, Construction of Knowledge Graph of Pumping Wells Based on Improved BERT Model of Joint Information Extraction, in: *2022 2nd International Conference on Algorithms, High Performance Computing and Artificial Intelligence (AHPCAI)*, IEEE, 2022, pp. 555–558.
[31] A. Graves and J. Schmidhuber, 2005 Special Issue: Framewise phoneme classification with bidirectional LSTM and other neural network architectures, *Neural Networks* **18** (2005), 602–610. https://api.semanticscholar.org/CorpusID:1856462.
[32] R. Grishman, Twenty-five years of information extraction, *Natural Language Engineering* **25** (2019), 677–692. https://api.semanticscholar.org/CorpusID:203701085.
[33] T.R. Gruber, Toward principles for the design of ontologies used for knowledge sharing?, *Int. J. Hum. Comput. Stud.* **43** (1995), 907–928. https://api.semanticscholar.org/CorpusID:1652449.
[34] W.L. Hamilton, Z. Ying and J. Leskovec, Inductive Representation Learning on Large Graphs, *ArXiv* **abs/1706.02216** (2017). https://api.semanticscholar.org/CorpusID:4755450.
[35] H. Han, J. Wang and X. Wang, Leveraging Knowledge Graph Reasoning in a Multi-Hop Question Answering System for Hot Rolling Line Fault Diagnosis, *IEEE Transactions on Instrumentation and Measurement* (2023).
[36] H. Han, J. Wang, X. Wang and S. Chen, Construction and evolution of fault diagnosis knowledge graph in industrial process, *IEEE Transactions on Instrumentation and Measurement* **71** (2022), 1–12.
[37] F.V. Harmelen, V. Lifschitz and B.W. Porter, Handbook of Knowledge Representation, 2008. https://api.semanticscholar.org/CorpusID:11808268.

[38] H. Hassanpour, P. Mhaskar, J.M. House and T.I. Salsbury, A hybrid modeling approach integrating first-principles knowledge with statistical methods for fault detection in HVAC systems, *Comput. Chem. Eng.* **142** (2020), 107022. https://api.semanticscholar.org/CorpusID:224873621.

[39] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J.E.L. Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A.-C.N. Ngomo, S.M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab and A. Zimmermann, Knowledge Graphs, *ACM Computing Surveys (CSUR)* **54** (2020), 1–37. https://api.semanticscholar.org/CorpusID:212414739.

[40] W. Huáng, X. Cheng, T. Wang and W. Chu, BERT-Based Multi-Head Selection for Joint Entity-Relation Extraction, in: *Natural Language Processing and Chinese Computing*, 2019. https://api.semanticscholar.org/CorpusID:201058602.

[41] International Organization for Standardization, ISO 14224:2016 Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment, 2016, ISO 14224:2016.

[42] S. Ji, S. Pan, E. Cambria, P. Marttinen and S.Y. Philip, A Survey on Knowledge Graphs: Representation, Acquisition, and Applications, *IEEE transactions on neural networks and learning systems* **33**(2) (2021), 494–514. doi:10.1109/TNNLS.2021.3070843.

[43] L. Jia and F. Yang, Research on Relation Extraction in the Field of Aviation Electrical Fault Diagnosis, *2023 4th International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE)* (2023), 146–148. https://api.semanticscholar.org/CorpusID:264292533.

[44] P. Jieyang, A. Kimmig, W. Dongkun, Z. Niu, F. Zhi, W. Jiahai, X. Liu and J. Ovtcharova, A Systematic Review of Data-Driven Approaches to Fault Diagnosis and Early Warning, *Journal of Intelligent Manufacturing* **34**(8) (2023), 3277–3304. doi:10.1007/s10845-022-02020-0.

[45] T. Kipf and M. Welling, Semi-Supervised Classification with Graph Convolutional Networks, *ArXiv* **abs/1609.02907** (2016). https://api.semanticscholar.org/CorpusID:3144218.

[46] B. Kitchenham, O.P. Brereton, D. Budgen, M. Turner, J. Bailey and S. Linkman, Systematic literature reviews in software engineering–a systematic literature review, *Information and software technology* **51**(1) (2009), 7–15.

[47] P. Körner, M. Leuschel, J. Barbosa, V.S. Costa, V. Dahl, M.V. Hermenegildo, J.F. Morales, J. Wielemaker, D. Diaz, S. Abreu and G. Ciatto, Fifty Years of Prolog and Beyond, *Theory Pract. Log. Program.* **22** (2022), 776–858. https://api.semanticscholar.org/CorpusID:247159044.

[48] J.D. Lafferty, A. McCallum and F. Pereira, Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data, in: *International Conference on Machine Learning*, 2001. https://api.semanticscholar.org/CorpusID:219683473.

[49] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma and R. Soricut, ALBERT: A Lite BERT for Self-supervised Learning of Language Representations, *ArXiv* **abs/1909.11942** (2019). https://api.semanticscholar.org/CorpusID:202888986.

[50] S. Lazarova-Molnar and N. Mohamed, Reliability Analysis of Cyber-Physical Systems, *Simulation for Cyber-Physical Systems Engineering* (2020). https://api.semanticscholar.org/CorpusID:229247134.

[51] E.A. Lee, Computing Foundations and Practice for Cyber- Physical Systems: A Preliminary Report, 2007. https://api.semanticscholar.org/CorpusID:15989874.

[52] C. Li, X. Yang, S. Luo, M. Song and W. Li, Towards Domain-Specific Knowledge Graph Construction for Flight Control Aided Maintenance, *Applied Sciences* **12**(24) (2022), 12736.

[53] X. Li, F. Zhang, Q. Li, B. Zhou and J. Bao, Exploiting a knowledge hypergraph for modeling multi-nary relations in fault diagnosis reports, *Advanced Engineering Informatics* **57** (2023), 102084.

[54] Z. Li, F. Liu, W. Yang, S. Peng and J. Zhou, A Survey of Convolutional Neural Networks: Analysis, Applications, and Prospects, *IEEE Transactions on Neural Networks and Learning Systems* **33** (2020), 6999–7019. https://api.semanticscholar.org/CorpusID:214803074.

[55] K.Y.H. Lim, T.S. Yosal, C.-H. Chen, P. Zheng, L. Wang and X. Xu, Graph-enabled cognitive digital twins for causal inference in maintenance processes, *International Journal of Production Research* **62**(13) (2024), 4717–4734.

[56] B. Liu, J. Wu, L. Yao and Z. Ding, Ontology-Based Fault Diagnosis: A Decade in Review, in: *Proceedings of the 11th International Conference on Computer Modeling and Simulation*, 2019, pp. 112–116. doi:10.1145/3307363.3307381.

[57] Y. Liu, Y. Peng, B. Wang, S. Yao and Z. Liu, Review on cyber-physical systems, *IEEE/CAA Journal of Automatica Sinica* **4** (2017), 27–40. https://api.semanticscholar.org/CorpusID:146680.

[58] P. Lucas, Symbolic diagnosis and its formalisation, *The Knowledge Engineering Review* **12**(2) (1997), 109–146.

[59] L. Nagy, T. Ruppert and J. Abonyi, Towards an Ontology-Based Fault Detection and Diagnosis Framework - A Semantic Approach, *2023 9th International Conference on Control, Decision and Information Technologies (CoDIT)* (2023), 1267–1272. https://api.semanticscholar.org/CorpusID:264456234.

[60] S.M.R. Naqvi, M. Ghufran, S. Meraghni, C. Varnier, J.-M. Nicod and N. Zerhouni, Cbr-based decision support system for maintenance text using nlp for an aviation case study, in: *2022 Prognostics and Health Management Conference (PHM-2022 London)*, IEEE, 2022, pp. 344–349.

[61] Neo4j, Neo4j - The Graph Database Platform, 2025. https://neo4j.com/.

[62] J. Nivre, Dependency Grammar and Dependency Parsing, 2005. https://api.semanticscholar.org/CorpusID:16318436.

[63] OREDA Participants Group, *OREDA: Offshore and Onshore Reliability Data Handbook*, 6th edn, DNV GL, 2015.

[64] L. Page, S. Brin, R. Motwani and T. Winograd, The PageRank Citation Ranking : Bringing Order to the Web, in: *The Web Conference*, 1999. https://api.semanticscholar.org/CorpusID:1508503.

[65] K. Petersen, S. Vakkalanka and L. Kuzniarz, Guidelines for conducting systematic mapping studies in software engineering: An update, *Inf. Softw. Technol.* **64** (2015), 1–18. https://api.semanticscholar.org/CorpusID:13134455.
[66] J.L. Rodgers, W.A. Nicewander and D.C. Blouin, Thirteen ways to look at the correlation coefficient, *The American Statistician* **42** (1988), 59–66. https://api.semanticscholar.org/CorpusID:2478952.
[67] E. Ruijters and M. Stoelinga, Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools, *Comput. Sci. Rev.* **15** (2014), 29–62. https://api.semanticscholar.org/CorpusID:206990190.
[68] A.R. Sahu, S.K. Palei and A. Mishra, Data-driven fault diagnosis approaches for industrial equipment: A review, *Expert Systems* **41**(2) (2024), e13360.
[69] A. Sarazin, J. Bascans, J.-B. Sciau, J. Song, B. Supiot, A. Montarnal, X. Lorca and S. Truptil, Expert system dedicated to condition-based maintenance based on a knowledge graph approach: Application to an aeronautic system, *Expert Systems with Applications* **186** (2021), 115767.
[70] M. Schwabacher, A survey of data-driven prognostics, *Infotech@ Aerospace* (2005), 7002.
[71] H.-B. Shi, D. Huang, L. Wang, M.-Y. Wu, Y.-C. Xu, B.-E. Zeng and C. Pang, An information integration approach to spacecraft fault diagnosis, *Enterprise information systems* **15**(8) (2021), 1128–1161.
[72] Siemens Senseye Predictive Maintenance, The True Cost of Downtime 2022, https://assets.new.siemens.com/siemens/assets/api/uuid:3d606495-dbe0-43e4-80b1-d04e27ada920/dics-b10153-00-7600truecostofdowntime2022-144.pdf (accessed 10-01-2024).
[73] M. Soleimani, F. Campean and D. Neagu, Diagnostics and prognostics for complex systems: A review of methods and challenges, *Quality and Reliability Engineering International* **37**(8) (2021), 3746–3778.
[74] M. Steinegger, M. Melik-Merkumians, J. Zajc and G. Schitter, A framework for automatic knowledge-based fault detection in industrial conveyor systems, in: *2017 22nd IEEE international conference on emerging technologies and factory automation (ETFA)*, IEEE, 2017, pp. 1–6.
[75] Y. Su, H. Wang, J. Wang and X. Liang, A Three-Dimensional Information Flow Modeling Method for Integrated Fault Diagnosis and Maintenance of Complex System, in: *2019 Prognostics and System Health Management Conference (PHM-Paris)*, IEEE, 2019, pp. 51–56.
[76] T. Sun and Q. Wang, Multi-source fault detection and diagnosis based on multi-level Knowledge Graph and Bayesian theory reasoning (S)., in: *SEKE*, 2019, pp. 177–232.
[77] R.S. Sutton and A.G. Barto, Reinforcement Learning: An Introduction, *IEEE Trans. Neural Networks* **9** (1998), 1054–1054. https://api.semanticscholar.org/CorpusID:60035920.
[78] P.-N. Tan, M.S. Steinbach, A. Karpatne and V. Kumar, Introduction to Data Mining (2nd Edition), 2018. https://api.semanticscholar.org/CorpusID:196001763.
[79] X. Tang, J. Wang, C. Wu, B. Hu and S.M. Noman, Constructing Aircraft Fault Knowledge Graph for Intelligent Aided Diagnosis, in: *Proceedings of the 4th International Conference on Information Technologies and Electrical Engineering*, 2021, pp. 1–4.
[80] X. Tang, X. Xie, L. Cui, X. Xu and J. Wang, Fault Knowledge Acquisition of Aircraft Based on Event Extraction Technology, in: *Proceedings of the 5th International Conference on Information Technologies and Electrical Engineering*, 2022, pp. 566–571.
[81] X. Tang, G. Chi, L. Cui, A.W. Ip, K.L. Yung and X. Xie, Exploring research on the construction and application of knowledge graphs for aircraft fault diagnosis, *Sensors* **23**(11) (2023), 5295.
[82] C.J. Turner, C. Emmanouilidis, T. Tomiyama, A. Tiwari and R. Roy, Intelligent Decision Support for Maintenance: An Overview and Future Trends, *International Journal of Computer Integrated Manufacturing* **32**(10) (2019), 936–959. doi:10.1080/0951192X.2019.1667033.
[83] A. Ucar, M. Karakose and N. Kırımça, Artificial Intelligence for Predictive Maintenance Applications: Key Components, Trustworthiness, and Future Trends, *Applied Sciences* **14**(2) (2024), 898. doi:10.3390/app14020898.
[84] F. Van Harmelen, V. Lifschitz and B. Porter, *Handbook of knowledge representation*, Elsevier, 2008.
[85] W3C, SPARQL 1.1 Query Language, 2013, Accessed: 2025-03-24. https://www.w3.org/TR/sparql11-query/.
[86] W3C RDF and SPARQL Working Group, RDF 1.2 Concepts and Abstract Data Model, 2026, W3C Candidate Recommendation Snapshot, 7 April 2026.
[87] J. Wang, C. Yan, Y. Zhang, Y. Li and H. Wang, Construction and application of knowledge graph for fault diagnosis of turbine generator set based on ontology, in: *Journal of Physics: Conference Series*, Vol. 2184, IOP Publishing, 2022, p. 012015.
[88] Q. Wang, Z. Mao, B. Wang and L. Guo, Knowledge Graph Embedding: A Survey of Approaches and Applications, *IEEE Transactions on Knowledge and Data Engineering* **29** (2017), 2724–2743. https://api.semanticscholar.org/CorpusID:19135805.
[89] R. Wang, C. Yan, J. Lu, J. Wang and Y. Zhang, Chinese Named Entity Recognition Based BERT-Decoder-CRF in Turbine Generator Set Fault Diagnosis, in: *International Conference on Intelligent Information Hiding and Multimedia Signal Processing*, Springer, 2022, pp. 79–89.
[90] Z. Wang, Y. Qian, L. Wang, S. Zhang and X. Luo, The Extraction of Hidden Fault Diagnostic Knowledge in Equipment Technology Manual Based on Semantic Annotation, *Proceedings of the 2019 8th International Conference on Software and Computer Applications* (2019). https://api.semanticscholar.org/CorpusID:153314579.

[91] Z. Wang, X. Wang, Y. Zhang, Y. Yang and Y. Li, Information Extraction of Aircraft Maintenance Records for Knowledge Graph Construction, in: *2022 Global Reliability and Prognostics and Health Management (PHM-Yantai)*, IEEE, 2022, pp. 1–6.

[92] Y. Wilhelm, P. Reimann, W. Gauchel and B. Mitschang, Overview on hybrid approaches to fault detection and diagnosis: Combining data-driven, physics-based and knowledge-based models, *Procedia Cirp* **99** (2021), 278–283.

[93] Y. Wu, F. Liu, L. Wan and Z. Wang, Intelligent fault diagnostic model for industrial equipment based on multimodal knowledge graph, *IEEE Sensors Journal* (2023).

[94] L. Xia, P. Zheng, S. Li, P. Lyu, C.K.M. Lee, J. Zhou and K. Wang, A Knowledge Graph-based Link Prediction for Interpretable Maintenance Planning in Complex Equipment, *2022 13th International Conference on Reliability, Maintainability, and Safety (ICRMS)* (2022), 301–305. https://api.semanticscholar.org/CorpusID:253555036.

[95] L. Xia, P. Zheng, X. Li, R.X. Gao and L. Wang, Toward cognitive predictive maintenance: A survey of graph-based approaches, *Journal of Manufacturing Systems* **64** (2022), 107–120. doi:10.1016/j.jmsy.2022.06.002.

[96] S. Xiao, Y. Hu, J. Han, R. Zhou and J. Wen, Bayesian networks-based association rules and knowledge reuse in maintenance decision-making of industrial product-service systems, *Procedia Cirp* **47** (2016), 198–203.

[97] N. Yang, G. Zhang and J. Wang, Research on Knowledge Graph and Bayesian Network in Fault Diagnosis of Steam Turbine, *2020 Global Reliability and Prognostics and Health Management (PHM-Shanghai)* (2020), 1–6. https://api.semanticscholar.org/CorpusID:229357907.

[98] Y. Yang, Z. Wu, Y.-X. Yang, S. Lian, F. Guo and Z. Wang, A Survey of Information Extraction Based on Deep Learning, *Applied Sciences* (2022). https://api.semanticscholar.org/CorpusID:252605355.

[99] Z. Yang, L. Qing and P. Lu, Integration of deep and shallow aircraft fault knowledge, in: *2011 IEEE 3rd International Conference on Communication Software and Networks*, IEEE, 2011, pp. 320–324.

[100] S. Zhang, Y. Zhang, Y. Yang, W. Cheng, H. Zhao and Y. Li, Knowledge graph construction for fault diagnosis of aircraft environmental control system, in: *2021 Global Reliability and Prognostics and Health Management (PHM-Nanjing)*, IEEE, 2021, pp. 1–5.

[101] B. Zhou, B. Hua, X. Gu, Y. Lu, T. Peng, Y. Zheng, X. Shen and J. Bao, An end-to-end tabular information-oriented causality event evolutionary knowledge graph for manufacturing documents, *Advanced Engineering Informatics* **50** (2021), 101441.

[102] B. Zhou, J. Li, X. Li, B. Hua and J. Bao, Leveraging on causal knowledge for enhancing the root cause analysis of equipment spot inspection failures, *Advanced Engineering Informatics* **54** (2022), 101799.

[103] J. Zhou, T. Wang and J. Deng, Corpus construction and entity recognition for the field of industrial robot fault diagnosis, in: *Proceedings of the 2021 13th International Conference on Machine Learning and Computing*, 2021, pp. 410–416.

[104] Q. Zhou, P. Yan and Y. Xin, Research on a knowledge modelling methodology for fault diagnosis of machine tools based on formal semantics, *Adv. Eng. Informatics* **32** (2017), 92–112. https://api.semanticscholar.org/CorpusID:10612333.

[105] Q. Zhou, P. Yan, H. Liu, Y. Xin and Y. Chen, Research on a configurable method for fault diagnosis knowledge of machine tools and its application, *The International Journal of Advanced Manufacturing Technology* **95** (2018), 937–960.

[106] OWL 2 Web Ontology Language Document Overview, 2012. https://www.w3.org/TR/owl2-overview/.