Front. Eng. Manag. 2024, 11(1): 143–158
https://doi.org/10.1007/s42524-023-0273-1

RESEARCH ARTICLE

Zhulin HAN, Jian WANG

Knowledge enhanced graph inference network based entity-relation extraction and knowledge graph construction for industrial domain

Abstract With the escalating complexity in production scenarios, vast amounts of production information are retained within enterprises in the industrial domain. Probing questions of how to meticulously excavate value from complex document information and establish coherent information links arise. In this work, we present a framework for knowledge graph construction in the industrial domain, predicated on knowledge-enhanced document-level entity and relation extraction. This approach alleviates the shortage of annotated data in the industrial domain and models the interplay of industrial documents. To augment the accuracy of named entity recognition, domain-specific knowledge is incorporated into the initialization of the word embedding matrix within the bidirectional long short-term memory conditional random field (BiLSTM-CRF) framework. For relation extraction, this paper introduces the knowledge-enhanced graph inference (KEGI) network, a pioneering method designed for long paragraphs in the industrial domain. This method discerns intricate interactions among entities by constructing a document graph and innovatively integrates knowledge representation into both node construction and path inference through TransR. On the application stratum, BiLSTM-CRF and KEGI are utilized to craft a knowledge graph from a knowledge representation model and Chinese fault reports for a steel production line, specifically SPOnto and SPFRDoc. The $F1$ value for entity and relation extraction has been enhanced by 2% to 6%. The quality of the extracted knowledge graph complies with the requirements of real-world production environment applications. The results demonstrate that KEGI can profoundly delve into production reports, extracting a wealth of knowledge and patterns, thereby providing a comprehensive solution for production management.

Keywords knowledge graph construction, industrial, BiLSTM-CRF, document-level relation extraction, graph inference

1 Introduction

With the development of intelligent manufacturing technologies such as the "industrial Internet" and artificial intelligence, the traditional manufacturing industry is placing emphasis on implementing the "industrial upgrading" strategy. Intelligent manufacturing has thus emerged as a major trend and central aspect of enterprise development (Qi et al., 2017; Zhou et al., 2019).

Following years of industrialization and informatization, industrial enterprises have laid the groundwork necessary to foster enterprise data and digital transformation. Nonetheless, they confront substantial challenges related to core technology bottlenecks, transformation, and upgrading. The accumulation of information and knowledge through long-term production processes has led to an uneven level of enterprise informatization, giving rise to several challenges, such as information silos and dormant information (Kamble et al., 2018). Addressing how to enhance the information structure of enterprises, delve into high-density knowledge, and facilitate control and decision-making in complex environments constitutes the scientific inquiries to be approached during the process of digitization (Wang et al., 2018; 2021).

The process of industrial upgrading has resulted in the

Received Mar. 21, 2023; revised Jul. 17, 2023; accepted Aug. 8, 2023

Zhulin HAN (✉), Jian WANG
College of Electronic and Information Engineering, Tongji University, Shanghai 200092, China
E-mail: hanzhulin@tongji.edu.cn

This work was supported by the National Science and Technology Innovation 2030 New Generation Artificial Intelligence Major Project (Grant No. 2018AAA0101800), and the National Natural Science Foundation of China (Grant No. 72271188).

extensive integration of intelligent devices, creating an urgent requirement for efficient management of production data. This encompasses various forms of textual information, including equipment manuals, documentation, and fault reports, which have become indispensable (Wan et al., 2018). Human decision-making and control have consequently become integral to the production process, relying on prior knowledge in most situations. Therefore, the effective handling and application of empirical knowledge throughout the manufacturing process is pivotal to the success of intelligent manufacturing. However, the actual incorporation of knowledge within the contemporary industrial sector faces obstacles. This is partially due to a shortfall in focus on knowledge by enterprises and challenges related to knowledge management. Presently, industrial domain knowledge data management mainly depends on traditional relational databases, which display critical limitations such as fragmented distribution, increased redundancy, and restricted storage capacity. Additionally, standard text representation methods lack semantic connections and ordering between textual data, thereby hindering the efficiency of information and knowledge search and reasoning and significantly obstructing the reutilization of industrial knowledge.

In 2012, Google initially introduced the concept of the knowledge graph and its application in search engines, offering a fresh outlook on text knowledge representation. Unlike simple sequences of strings, the knowledge graph represents text as a network composed of entities and semantic links (Hu et al., 2020). By employing knowledge extraction and semantic mining of industrial data, knowledge graphs enable the deep mining of semantic information from extensive industrial data, the creation of structured information networks, and the support of production process control and decision-making (Zhang et al., 2021). Moreover, compared to traditional keyword-based search, the utilization of a knowledge graph in semantic searches facilitates a more thorough understanding of engineering problems, infers hidden relationships from existing connections, and provides effective solutions for intelligent applications. As a result, the incorporation of a knowledge graph can act as a valuable instrument in improving production efficiency and strengthening core competitiveness within the industrial sector. However, the deployment of natural language processing technology in constructing knowledge graphs within the industrial domain still faces specific challenges.

• Lack of industrial domain annotation data: The natural language processing model for knowledge extraction depends on large-scale corpus annotation. However, the scale of industrial domain data is limited due to manual annotation, and conventional algorithms fail to achieve the desired computational performance and predictive accuracy on such datasets. How to enhance the precision of knowledge extraction in the industrial domain remains an area for further research (Zhang et al., 2021).

• Insufficient depth of knowledge mining: Extraction models that focus on interdependence within individual sentences are inadequate for knowledge extraction in the industrial sector. This is particularly true given that fault reports in this domain are frequently archived in documents that display complex logical connections between sentences. Extracting these relationships requires a comprehensive understanding of the content and logical reasoning (Zhou et al., 2022).

To address the aforementioned challenges, this paper introduces the knowledge-enhanced graph inference (KEGI) network, a relation extraction method designed for long paragraphs in the industrial domain. KEGI constructs the document graph to capture the intricate interactions among entities and mentions, and it innovatively incorporates knowledge representation into node construction and path inference through TransR. On this basis, a semi-automatic framework for constructing a knowledge graph in the industrial domain has been developed.

This framework for knowledge graph construction can enhance the extraction efficiency of domain knowledge, even with a limited domain corpus size, and exhibits strong portability. Additionally, the study proposes the creation of a domain-specific knowledge graph to validate the effectiveness of the approach. The knowledge graph is formed by mapping the extracted relationships between entities and their attributes onto a graph structure. This structure enables efficient querying and visualization of the connections between entities, shedding light on the underlying relationships within the industrial domain.

The primary research content of this paper includes 1) named entity recognition (NER) based on bidirectional long short-term memory conditional random field (BiLSTM-CRF) and a knowledge dictionary; 2) document graph construction coupled with node representations fused with knowledge embedding; and 3) relation extraction and path inference utilizing knowledge augmentation.

The structure of this paper is organized as follows. Section 2 reviews the related research progress of knowledge graph construction, NER, and relation extraction. Section 3 delineates the framework for knowledge graph construction, and the KEGI approach is proposed and explained in detail. Section 4 presents a case to evaluate performance, and Section 5 provides the conclusion.

2 Related work

2.1 Knowledge management in the industrial domain

The significance of knowledge management has escalated as data have become increasingly critical and distributed. However, industrial production and services continue to face persistent challenges, including value chain instability, risk prediction, and information extraction. As a result, the knowledge and ontology generated may not satisfy the requirements for standardization, customization, and integration in particular manufacturing fields, leading to substantial barriers in sharing and reusing knowledge. To address these challenges, companies strive to integrate a variety of resources (Gui et al., 2020).

During the product design phase, Huet et al. (2021) advocated for the use of knowledge graphs to facilitate the combination of complex engineering design rules and computer-aided design software. This integration enables the provision of design rule recommendations via the knowledge graph, verification of design schemes, and automated reasoning of the design program.

In the production phase, Xu et al. (2020) introduced a clustering method to discover new insights about corrective actions from historical quality problem-solving data. The method involved a two-stage clustering approach based on verb and noun clustering. The resulting clusters were used to establish connections between historical problems and corrective action groups, subsequently forming a knowledge graph to organize corrective measure knowledge. Similarly, Zheng et al. (2021) presented an embedding algorithm founded on a graph neural network and a multi-agent reinforcement learning technique that used an industrial knowledge graph derived from the Self-X cognitive manufacturing network. Liu et al. (2022) offered a method for data representation in support of cognitive manufacturing within the industrial Internet using a knowledge graph-based framework created by integrating production process data across various levels.

During the maintenance phase, Shi et al. (2021) explored the use of a knowledge graph to assimilate information throughout the entire life cycle of spacecraft, applying this approach to the creation of a fault diagnosis question-and-answer system. Deng et al. (2022) proposed an event logic knowledge graph for diagnosing faults in robot transmission systems, which included fault diagnosis ontology and event parameters, and employed a bidirectional long-term short-term memory network with a conditional random field (CRF) for entity recognition. Ren et al. (2021) implemented a multi-level knowledge graph method for fault detection, amalgamating domain knowledge and monitoring data to build a multi-level graph with monitoring variables as nodes.

These cited research contributions present innovative concepts and technological support for the continued progression of the industrial sector. The role of knowledge management in industrial production environments is clearly demonstrated. Industrial knowledge graphs, with their ability to execute cognitive integration, analysis, and services, actively contribute to enhancing productivity, sustaining production lines, and assisting in fault diagnosis throughout the product life cycle. The establishment of domain-specific knowledge graphs emerges as a feasible strategy to bolster manufacturing knowledge management capabilities.

2.2 Named entity recognition (NER)

Named entities form the foundation of relation extraction tasks, and the quality of their extraction significantly influences the accuracy of subsequent operations (Li et al., 2022). Early NER tasks relied on manually created dictionaries, which were highly dependent on human judgment and exhibited limited scalability. The advent of machine learning transformed NER into a sequence labeling problem. Lin et al. (2004) introduced a dictionary-assisted maximum entropy algorithm that overcame domain restrictions but suffered from high training overhead and a tendency to fall into local optimization.

Lafferty et al. (2001) proposed the CRF model to calculate global probabilities, bypassing local normalization and resolving the issue of label position offset. The CRF model is regarded as a fundamental model for NER tasks due to its excellent capacity to capture contextual features. In recent years, end-to-end deep learning approaches have emerged, substantially reducing the manual labor involved in constructing features for NER tasks.

Huang et al. (2015) developed the BiLSTM-CRF model to recognize specific characters in sequences by combining CRF with forward/backward transfer. By relying solely on word vectors, without the need for additional feature engineering, the model can achieve high-precision entity recognition. The inclusion of domain-specific dictionaries further enhances the model's results. Currently, this model represents the most widely adopted approach among deep learning-based NER methods.

Strubell et al. (2017) suggested that expanded convolution could supplement its long-sequence feature extraction capability while benefiting from the parallel computing advantage of convolutional neural network (CNN). This was combined with CRF to determine the labeling outcomes. Cao et al. (2018) presented a method for Chinese NER based on adversarial transfer learning. They integrated the Chinese word segmentation task and concurrently trained it with two NER tasks, allowing for effective long-distance relationships between sentences. This method has proven effective in Chinese entity extraction tasks.

This paper will concentrate on the construction of knowledge graphs within the industrial domain for entity extraction tasks. By employing the classic BiLSTM-CRF method and integrating domain ontology, it aims to achieve high-precision recognition of named entities within domain-specific texts.

2.3 Document-level relation extraction

Relation extraction is central to knowledge graph construction. Traditional relation extraction is primarily a binary extraction process between two entities, and the task of identifying a specific relationship is commonly abstracted as a multi-categorization problem. However, the complexities of application scenarios mean that simple relationship extraction is often inadequate for handling intricate real-world applications. In response to this challenge, document-level relationship extraction models have emerged, aiming to capture inter-sentence relationships and address the problem of entity coreference. Such extraction can be classified into two main categories: Sequence-based extraction models and document-level entity graph-based extraction models (Zhou et al., 2022).

The sequence-based approach has evolved from recurrent neural network (RNN) architectures to transformer architectures, implicitly addressing long-distance dependencies (Wang et al., 2019; Xiao et al., 2020; Han et al., 2020). For example, Zhou et al. (2021) developed a method to tackle the multi-label and multi-entity problem using adaptive thresholding and local context pooling techniques. This approach applies learnable entity-related thresholds to replace non-sensible global thresholds in multi-label classification, thus addressing the issue of using the same entity embedding across multiple entities. Furthermore, the structured self-attention network (SSAN), proposed by Xu et al. (2021), integrates the structural dependencies of entities and mentions into the self-attention mechanism, synthesizing the structure through an encoding network to achieve a fusion of contextual reasoning and structural reasoning.

In contrast to sequence-based models, document graph-based models introduce a graph structure, with graph convolutional networks (GCNs) being the dominant approach in this field (Sahu et al., 2019; Christopoulou et al., 2019; Nan et al., 2020). The GCN model by Sahu et al. (2019) learns the representation of each node in the document graph, aggregating entity mentions and classifying relationships through multi-instance learning. However, this approach does not consider the structural information of the document graph during reasoning. Christopoulou et al. (2019) employed a heuristic method, Edge-oriented Graph (EoG), to construct a document graph and ascertain the information flow between nodes through different types of edges, achieving a fitting of heterogeneous interaction relationships among documents. Building on the EoG model, Zeng et al. (2020) introduced the GAIN (graph aggregation-and-inference network) model, which weights the nodes of the mention-level graph, compresses it into an entity-level graph, and classifies the relationships through a multilayer perceptron (MLP) network. Despite these advancements, a recurrent limitation in current research is the lack of incorporation of graph structure information among nodes, particularly in the context of attention mechanisms in path inference. This shortcoming highlights an area for further exploration and innovation within the field, paving the way for more comprehensive and nuanced relation extraction models that can more effectively navigate the complexities of modern application scenarios.

The effectiveness of the aforementioned relation extraction models has indeed been validated; however, their application typically requires a substantial and well-curated corpus for training. Within the industrial domain, the scarcity of such corpora imposes certain restrictions on their utilization in specialized fields. To circumvent this obstacle, various research efforts have investigated the incorporation of external knowledge bases — including domain ontology, concepts, and attributes — to bolster the learning capabilities of these models. A graph-embedded representation model may facilitate resolving the external knowledge base representation issue.

Zhang et al. (2019) advanced language representation by regarding entities in the knowledge atlas as external knowledge, thereby training the enhanced language representation model ERNIE (Enhanced language RepresentatioN with Informative Entities). However, this method overlooked the relationships within the knowledge graph. Liu et al. (2020) introduced Knowledge-enabled Bidirectional Encoder Representation from Transformers (K-BERT), which addresses the limitation of the traditional BERT model's ability to handle only sequence structures, not graph structures. Through soft position encoding, they achieved the fusion of a knowledge graph and BERT encoding. More recently, Lyu et al. (2023) proposed utilizing enhanced semantic embedding techniques to construct the initial user relationship within a recommendation system, incorporating external knowledge bases, and subsequently learning the user behavior graph via the graph model.

From the above research, it is clear that the primary focus of document-level relation extraction lies in the confluence of document graphs and graph models. However, due to the reliance on large-scale corpora, there is an evident need to more effectively integrate external knowledge bases into these models. Present knowledge enhancement models can merely overlay domain ontology onto text-embedded representations and fail to meld seamlessly with the document graph. This shortfall hampers the guidance provided by external knowledge bases and impairs the reasoning performance of models, especially when dealing with small-scale corpora.

Building on this body of research, the current study aims to delve deeper into knowledge extraction methodologies tailored to the characteristics of industrial domain knowledge. By addressing these limitations and infusing both domain ontology structure information and mention-level document graph structure information into the path inference mechanism, this work presents an innovative approach for relation extraction within the industrial domain. The proposed methodology demonstrates enhanced accuracy in relation extraction, and the development of a domain-specific knowledge graph serves as a potent tool for the effective querying and analysis of relationships within the industrial sphere. The findings of this investigation carry significant implications for future research in industrial domain relation extraction and lay the groundwork for ongoing enhancement and refinement of techniques within this ever-evolving field.

3 Methodology

3.1 Framework

The construction of a knowledge graph necessitates the melding of an existing ontology model with the process of knowledge extraction. The architecture of this integration is illustrated in Fig. 1, and the input for this process is bifurcated into two components: 1) enterprise production data, encompassing three primary facets — expert experience, production system information, and product information; and 2) existing ontology models within the industrial domain, represented as dictionary participant word embeddings in NER and forming part of the node information within the composition document graph for relation extraction. These ontology models also serve a pivotal function in the path inference mechanism.

The initial phase of the knowledge extraction process entails the identification of entities within a document via a BiLSTM-CRF model. To heighten the model's accuracy, domain-specific knowledge from the industrial sector is integrated into the initialization of the word embedding matrix. Subsequent to this step, a document graph is formulated to mirror the interactions among entities, and these entities are embedded utilizing the BERT model. Furthermore, the portrayal of industrial domain knowledge is innovatively amalgamated into the node representation. The document graph undergoes processing by a GCN to amalgamate the entity nodes. The concluding stage features the proposal of a knowledge-enhanced path inference mechanism, which entails the synthesis of the extant knowledge representation model with the path information.

In summation, the knowledge graph stands as a robust instrument that enables the systematic organization and depiction of information, especially within the industrial domain. The fusion of the existing ontology model with the knowledge extraction procedure furnishes a holistic framework for the fabrication of a knowledge graph that embeds domain-specific knowledge and augments the model's accuracy. The proposed knowledge-enhanced path inference mechanism represents an avant-garde approach that further amplifies the applicability and efficacy of the knowledge graph within the industrial context.

3.2 Named entity recognition module

The procedure for constructing a document graph is initiated with the identification of named entities and mentions within the document. The manner in which text is represented in industrial domains is conditioned by particular attributes, which may be harnessed to enhance the learning of entity labels via the inclusion of antecedent knowledge. To bolster the accuracy of NER within this sector, the conventional BiLSTM-CRF model is supplemented with an industrial domain dictionary. This enrichment amplifies the word features that are utilized in the recognition process. A schematic representation of the model framework is depicted in Fig. 2.

During the data preprocessing phase, each character of a Chinese sentence is ascribed a corresponding type label, and these character types are categorized into the set {O, B-TYPE, I-TYPE, ...}, in line with the BIO annotation method. These classifications correspond to untyped characters, characters that mark the beginning of an entity, and characters that mark the end of an entity, respectively. The domain dictionary is replete with specialized vocabulary drawn from an existing knowledge ontology model. This dictionary is input into the Word2Vec tool, which yields a fixed-dimension word vector dictionary. This output is subsequently employed to initialize the character embedding matrix via unsupervised learning. As illustrated in Fig. 2, the model framework can be dissected into three discrete layers:

• The first layer is the lookup layer, which solely engages in transmuting each character representation from a one-hot vector to character embedding. A regular embedding matrix is spawned for random initialization, and to refine the embedding representation, words from the dictionary are tagged, and pretrained word embeddings are deployed. Character embeddings are enhanced using the associated tagged word embeddings (Li et al., 2020).

• The second layer is the BiLSTM layer, endowed with the capability to efficiently extract the forward and backward information features of the character embedding.

• The third layer is the CRF layer, which is tasked with ascertaining the label of each character within a sentence. When contrasted with the Soft-Max layer, the CRF layer can leverage sentence-level label information and model the transformative behavior between every two disparate labels. This process averts the independent tagging of character positions and the inadvertent disregard of the dependencies between labels.

For a given input sequence $X=(x_1, x_2, \dots, x_n)$, the output matrix of BiLSTM is P, and $P_{i,j}$ is the score of word vector i on label j. Then, for the predicted sequence $y=(y_1, y_2, \dots, y_n)$, the score is:

$$s(X, y)=\sum_{i=0}^{n} A_{y_i,y_{i+1}}+\sum_{i=1}^{n} P_{i,y_i}, \tag{1}$$

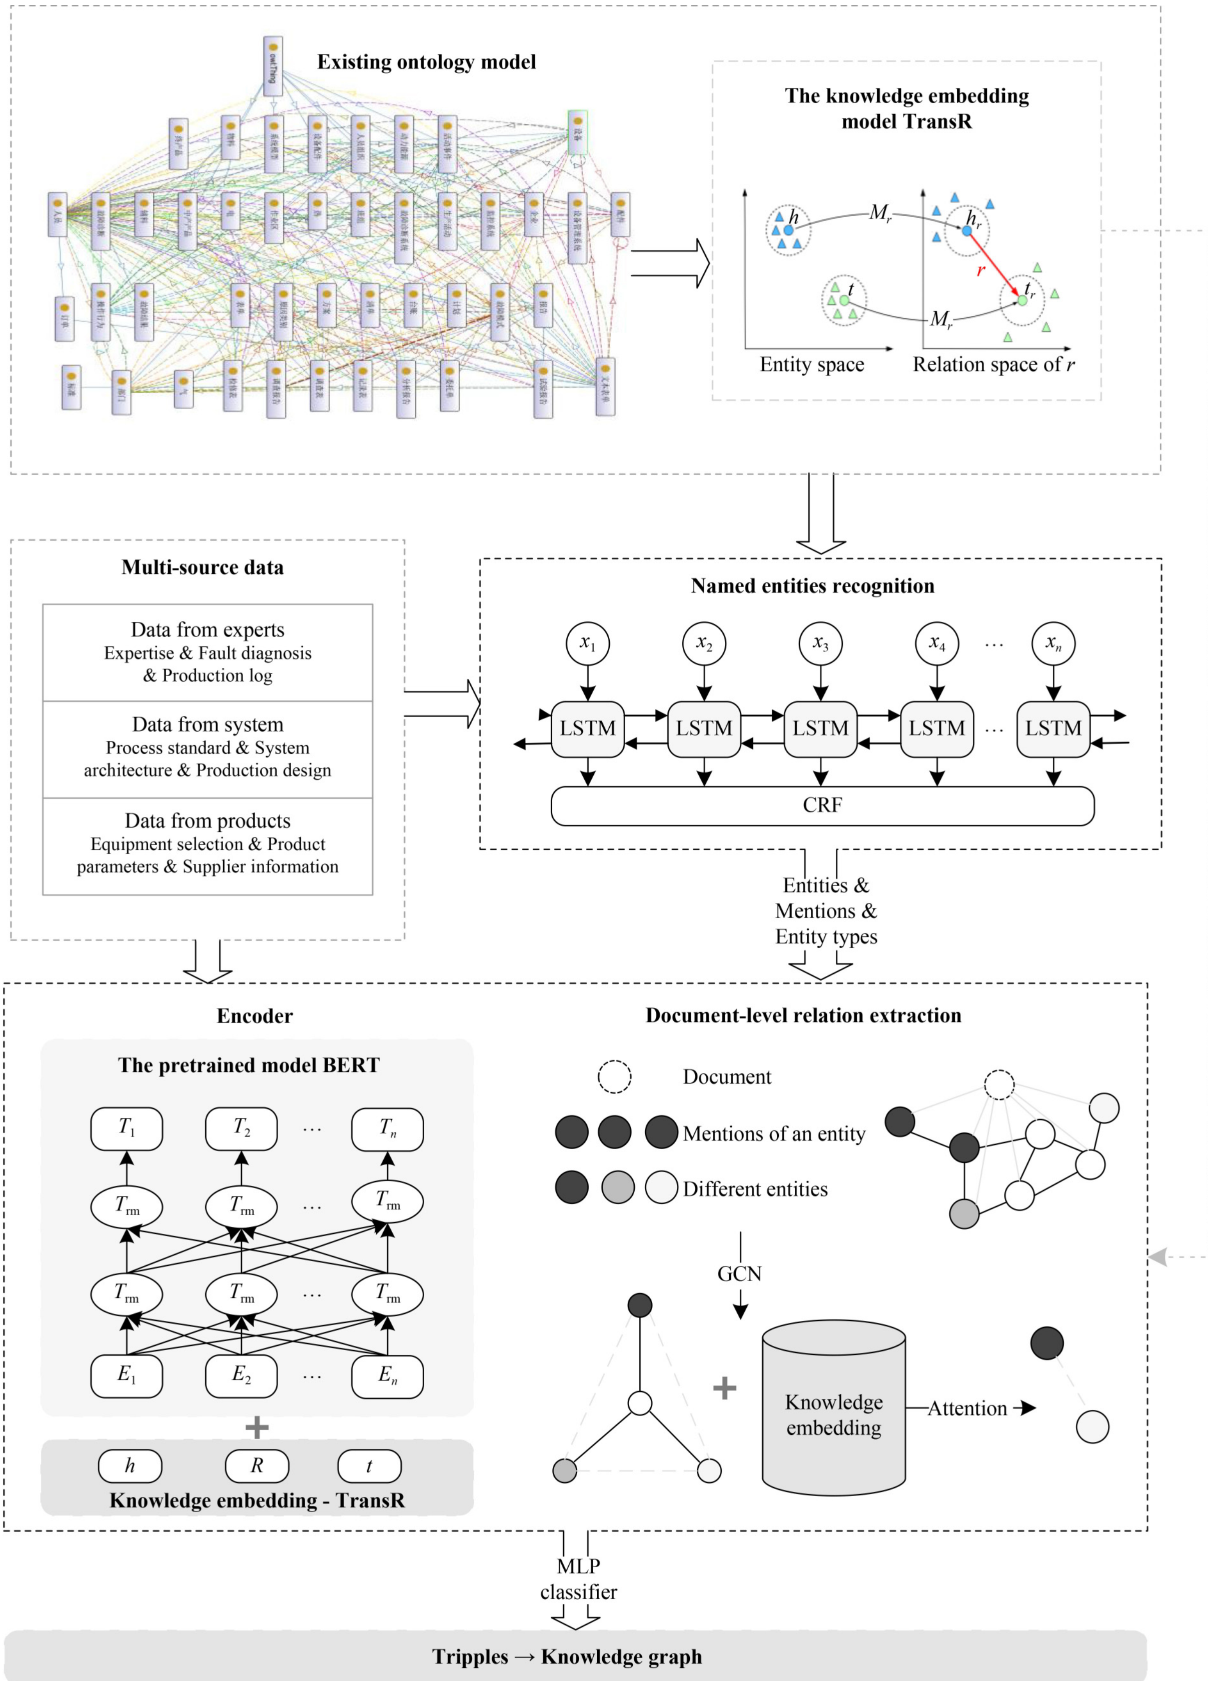

Fig. 1 Framework of the knowledge graph construction method based on document-level relation extraction.

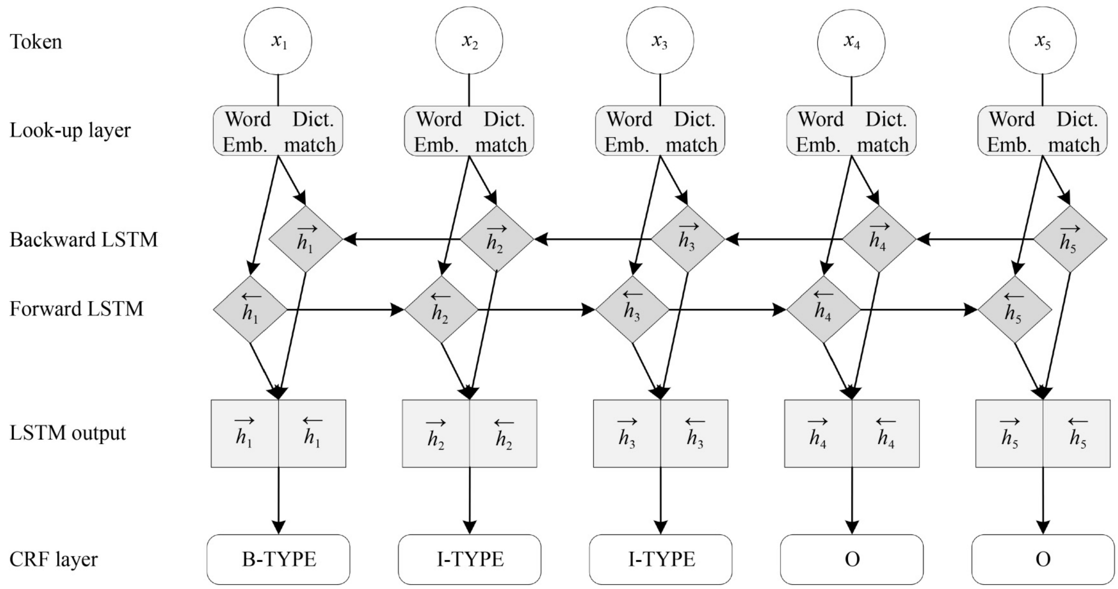

Fig. 2 BiLSTM-CRF network model combined with a dictionary.

where A is the state transfer matrix score. Maximizing the likelihood probability of a correctly labeled sequence is:

$$\log[p(y|X)] = P(X, y) - \log\Big(\sum_{\tilde{y}\in Y_X} e^{P(X, \tilde{y}y)}\Big), \tag{2}$$

$$y^* = \arg\max_{\tilde{y}\in Y_X} s(X, \tilde{y}), \tag{3}$$

where $x_1, x_2, \ldots, x_n$ is the set of possible labels of the input sequence X. The model parameters are searched by the Adam optimizer to minimize the loss, and the loss function is:

$$loss_{\mathrm{NER}} = -\log[p(y|X)]. \tag{4}$$

3.3 Document graph construction module

Upon the completion of entity extraction, the next phase involves the extraction of inter-entity relations within the document, a task accomplished by KEGI, as illustrated in Fig. 3. Document-level relation extraction is a process that necessitates the inference of inter-sentence entity relations, aiming to bolster the effectiveness of single-sentence extraction. To this end, entities that share identical names and classes are recognized as co-referential entities or mentions, a determination made based on both character features and precedence knowledge.

In this particular context, co-referential entities allude to entities utilized to reference the same real-world object within the document. The discernment of these co-referential entities constitutes a vital juncture in document-level relation extraction, since it sets the stage for the generation of accurate and semantically rich relationships between entities. Such information becomes the foundational bedrock for the construction of a graph-based representation of the relationships that exist among the entities, thereby paving the way for knowledge graph formulation.

KEGI’s method for document-level relation extraction hinges on the detection of co-referential entities coupled with the inference of inter-sentence entity relations. This strategy facilitates a faithful representation of the relationships that interconnect entities within a document and underpins the creation of a knowledge graph mirroring the structural architecture and thematic content of the document. Such a knowledge graph can subsequently be employed to enhance various industrial applications, encompassing areas such as information retrieval and question answering.

In sum, the extraction of inter-entity relations emerges as an indispensable constituent of knowledge graph construction. KEGI’s innovative approach, rooted in the pinpointing of co-referential entities and the derivation of inter-sentence entity relations, equips the process with the means to render a precise depiction of the relationships among entities. This approach culminates in the creation of a knowledge graph that authentically reproduces the intricate structure and substance of the document, thus contributing significantly to the field of knowledge extraction and representation.

The n Chinese characters in document $D = \{w_i\}_{i=1}^n$ are vectorized, and there are corresponding entities $\{E_1, E_2, \ldots, E_m\}$ in the existing knowledge base in the domain and relation R. In its triplet (h, r, t), $h, t \in R^k$, $r \in R^d$, there is a mapping matrix $M_r \in R^{k\times d}$ for relation r, which is then projected by the TransR model (Lin et al., 2017). Then, the vector representation can be obtained:

$$h_r = hM_r, \tag{5}$$

$$t_r = tM_r. \tag{6}$$

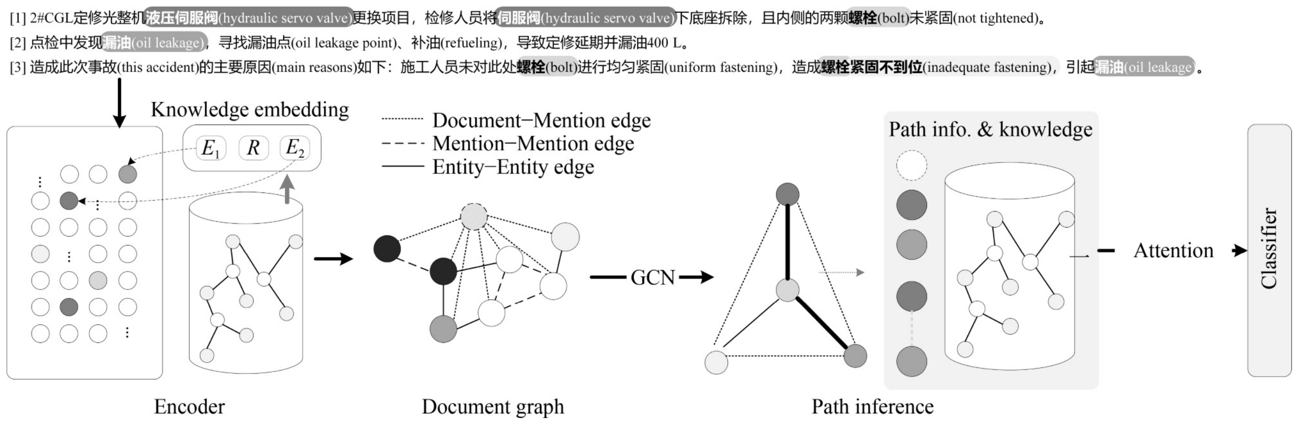

Fig. 3 Process of KEGI.

However, in the case of different vector representations of the same entity in the knowledge base (Dong et al., 2021), the average sum of *num* vector representations that the *i*th entity has is taken; then, the vector representation is $\{n_1, n_2, \ldots, n_m\}$:

$$n_i = \frac{\sum_{j=1}^{m} n_{ij}}{num}. \tag{7}$$

According to the encoder idea of Yao et al. (2019), each Chinese character w_i in the document is spliced into character embedding $E_w(w_i)$, entity type embedding $E_t(t_i)$ and coreference embedding $E_c(c_i)$. The *i*th Chinese character can be represented as:

$$x_i = [E_w(w_i);\ E_t(t_i);\ E_c(c_i)], \tag{8}$$

where t_i is the entity type and c_i is the entity ID.

The above two sets of vector representations $\{n_1, n_2, \ldots, n_m\}$ and $\{x_1, x_2, \ldots, x_n\}$ are superimposed through the Aggregator (Zhang et al., 2019):

$$\begin{aligned}&\{x'_1, x'_2, \ldots, x'_n\}\\&= \text{Aggregator}(\{x_1, x_2, \ldots, x_n\}, \{n_1, n_2, \ldots, n_m\}),\end{aligned} \tag{9}$$

where the Aggregator integrates the two separately for multi-headed attention with Eqs. (10) and (11):

$$\{\tilde{x}'_1, \tilde{x}'_2, \ldots, \tilde{x}'_n\} = MH-ATT(\{x_1, x_2, \ldots, x_n\}), \tag{10}$$

$$\{\tilde{n}'_1, \tilde{n}'_2, \ldots, \tilde{n}'_m\} = MH-ATT(\{n_1, n_2, \ldots, n_m\}). \tag{11}$$

Equation (12) integrates the above two sets of vectors:

$$v'_j = \sigma\left(W'_t\sigma\left(\tilde{W}'_t\tilde{x}'_j + \tilde{W}'_e\tilde{n}'_k + \tilde{b}'_t\right) + \tilde{b}'_t\right), \tag{12}$$

where W_t, W_e, and b_t are parameters to be trained, and σ is the activation function.

Thus far, the embedding vectors $\{v'_1, v'_2, \ldots, v'_n\}$ have integrated $\{n_1, n_2, \ldots, n_m\}$ and $\{x_1, x_2, \ldots, x_n\}$ by Eq. (12), which contains rich semantic knowledge. In the following, $\{v'_1, v'_2, \ldots, v'_n\}$ is referred to as $\{g_1, g_2, \ldots, g_n\}$.

Using existing character representations, a document graph can be constructed with entity mentions and documents serving as nodes to model the documents. In this architecture, each document graph connects all entity mentions with the document node at the center. The remaining characters are not added to the document graph, while the context information of the mentions is integrated into the vector representation through the encoder. Additionally, to facilitate the subsequent compression of the document graph, mentions referring to the same entity are connected. To ensure the efficacy of triplet extraction within a sentence, connections are established between different entities that appear in the same sentence.

In summary, the document graph comprises three types of edges: Document–mention, mention–mention, and entity–entity.

After that, the document graph is aggregated by GCN (Kipf and Welling, 2016). The convolution operation on node *u* is:

$$\rho_u^{l+1} = \sigma\Big(\sum_{k\in K}\sum_{v\in N_k(u)} W_k^l\rho_v^l + b_k^l\Big), \tag{13}$$

where K denotes different types of edges, and $N_k(u)$ is the neighboring nodes connected at node u through edge k. Similar to the GCN aggregation method for multi-relational graphs, here the weights are aggregated for different types of edges.

The final representation of node u needs to cover the features of all layers of the GCN and should consist of hidden state connections of the different layers:

$$m_u = [\rho_u^0;\ \rho_u^1;\ \ldots;\ \rho_u^N], \tag{14}$$

where the initialization of the entity mention node is expressed as $\rho_u^0 = \frac{1}{t-s+1}\sum_{j=s}^{t} g_j$, which represents a mention node from the *s*th to the *t*th word of the document, and the initial representation of a document node is the encoder output document representation.

3.4 Path inference and relation extraction module

As shown in Fig. 3, this module will extract the relation between entities based on the existing document graph and node representation and construct the entity interaction graph by compressing the coreference mention nodes into a single entity node. The entity node representation e_i is the average value of the N mentioned node representations of the entity:

$$e_i = \frac{\sum_{i=1}^{N} m_i}{N}. \tag{15}$$

Since there may be multiple connections between multiple mentions of different entities, these connections are compressed into edges between entities, and the edges between entities e_i and e_j are represented as:

$$e_{ij} = \sigma\left(W_q\left[e_i;\ e_j\right] + b_q\right). \tag{16}$$

After obtaining the representation of edges between entities, the inter-entity relations must be inferred. In this paper, we guide the path inference with domain knowledge, simultaneously combining the entity graph structure to optimize the inference ability.

The path representation of the head and tail entities is based on the common entity e_o through which they pass. The existing knowledge ontology of the industrial domain is utilized to identify the corresponding class of entity (e_h, e_o, e_t) and to extract the class, relationship, and related instance subgraph as external knowledge in the form of class-relationship-class. This can be expressed as a collection of (C_h, R, C_t) triples.

This module employs TransR as a knowledge embedding model. TransR can discern the differences between entities with various relations when embedding triples. When embedding word vectors, TransR will model both the entity space and the entity correspondence space so that the final entity vector contains the relation features between entities.

First, external knowledge is encoded. For the class-relationship-class triplet (C_h, R, C_t), there is a mapping matrix $M_r \in R^{k\times d}$ for the path between classes. The vector representations of C_h, C_t are:

$$r_h = C_h M_r, \tag{17}$$

$$r_t = C_t M_r. \tag{18}$$

The domain knowledge class $\{C_1,\ C_2,\ \ldots,\ C_n\}$ is represented by the vector obtained after TransR mapping as $\{r_1,\ r_2,\ \ldots,\ r_m\}$, where m denotes the number of knowledge class vector representations after embedding. Since the same knowledge class can obtain different vector representations in various relationships, m is slightly greater than n. Therefore, for the ith path between the head and the tail entities e_h, e_t of the relationship to be predicted, with the entity passing e_o, the set of representations of edges is $\{e_{ho},\ e_{ot},\ e_{to},\ e_{oh}\}$. The path between e_h and e_t is then represented as follows:

$$p^i_{h,t} = \left[e_{ho};\ e_{ot};\ e_{to};\ e_{oh};\ r_1;\ r_2;\ \ldots;\ r_m\right]. \tag{19}$$

Afterwards, the information about the different paths between the head and tail entities is fused into an attention mechanism (Bahdanau et al., 2014):

$$p_{h,t} = \sum_i \frac{e^{\sigma\left(\left[e_h;\ e_t\right]W_l p^i_{h,t}\right)}}{\sum_j e^{\sigma\left(\left[e_h;\ e_t\right]W_l p^j_{h,t}\right)}} p^i_{h,t}. \tag{20}$$

Through the process of path inference, an entity can be effectively represented by merging its node vectors with domain knowledge ontology. This representation manifests as informative entity interaction details across sentences within documents.

The ultimate stage in the process of relation classification can be conceptualized as a multi-label classification task (Han et al., 2020). This task entails the incorporation of information about entity pairs (e_h, e_t) acquired from preceding stages, ultimately deriving the category label through employment of a classifier. The input to this classifier encompasses the representations of both the head and tail entities, the representation of the document node within the document graph, and the path information. These three distinct types of vector representations are concatenated, culminating in the formation of the input for the multi-label classification layer:

$$U\left(r|e_h,\ e_t\right) = \mathrm{Sigmoid}\left(W_\beta \sigma\left(W_\alpha\left[e_h;\ e_t;\ e_{doc};\ p_{h,t}\right] + b_\alpha\right) + b_\beta\right), \tag{21}$$

where e_{doc} is the representation of the document node. The loss function is:

$$\begin{aligned} loss_{REL} = -\sum_{D\in S}\sum_{h\neq t}\sum_{r_i\in R} & I(r_i = 1)\log U(r_i|e_h,\ e_t) \\ & + I(r_i = 0)\log\left(1 - U\left(r_i|e_h,\ e_t\right)\right), \end{aligned} \tag{22}$$

where S is the set of triplets, R is the set of relationships, and $I(\cdot)$ is the indicator function.

The utilization of multi-label classification within the context of the relation classification task facilitates the depiction of numerous relationships between entities within a single document. This approach holds significant importance for the precise representation of relationships within the knowledge graph, as it permits the encompassment of multiple connections between entities present in a solitary document. The holistic training of the classification model guarantees the acquisition of knowledge concerning the interrelation between entities and their corresponding labels in a cohesive manner, thereby further enhancing the accuracy of the classification process.

Consequently, the ultimate step in the process of relation classification bears paramount significance in ensuring the meticulous representation of relationships among entities within the knowledge graph. The incorporation of multi-label classification and the comprehensive training of the classification model work in tandem to secure a precise portrayal of relationships within the knowledge graph. This approach supports the creation of an all-encompassing and meaningful representation of the document's content.

4 Case study

4.1 Problem description

The steel industry stands as a quintessential example of a traditional manufacturing sector, presently undergoing a transition toward intelligent manufacturing.

The ensuing inquiries are rooted in a thorough investigation of a steel company's continuous rolling line. The company's continuous rolling line exhibits intricate processes and a myriad of diverse production equipment types, encompassing both hot and cold rolling operations. The data generated during the production process comprise production plans, equipment data, production logs, and team member records. The management of these data predominantly relies on paper file archiving and information system storage, exhibiting a weak correlation between datasets. Production decisions often hinge on singular knowledge sources, resulting in a lack of synergy. The specific manifestations include inadequate supply chain coordination and obscured product management and statistics, culminating in inventory backlogs. Insufficient means to effectively manage production line equipment impedes efficient equipment maintenance. Dilemmas emerge with storing significant documents such as production plans and logs directly into databases; typically, they are saved as text files, thereby failing to achieve information sharing. The semantic information embedded within these documents remains underutilized. These challenges mirror the authentic needs of traditional industrial enterprises. Therefore, this paper adopts this enterprise as its focal point of research, employing BiLSTM-CRF and KEGI methodologies to extract knowledge from production reports originating from the hot-rolling production line. The objective is to construct a structured knowledge network. The experimental results show that the knowledge graph formulated in this study aligns with the precision requisites of domain knowledge. It successfully accomplishes the ensuing objectives: Thorough exploration of semantic knowledge within enterprise documents to foster enhanced knowledge utilization. An effective approach for enterprise data management has been provided, wherein the knowledge graph facilitates the management of production line equipment. This, in turn, enables equipment diagnosis and supports production decision-making through the graph's structural framework.

4.2 SPFRDoc dataset

To validate the method's efficacy within an industrial context, this paper establishes the SPFRDoc dataset for the fault diagnosis scenario in the steel manufacturing industry. This dataset encompasses 763 reports of equipment fault diagnosis, amounting to a total of 13158 triples. Among these triples, 2370 are deduced through inference from multiple sentences. Additionally, an ontology model, denoted as SPOnto, is formulated to delineate the categories of instances within the steel manufacturing sector and the attributes suitable for describing these instances. SPOnto comprises 1561 axioms, 352 class concepts, and 407 data attributes. It is constructed upon dependable and standardized knowledge, grounded in real-world manufacturing production scenarios of steel enterprises.

Following the relevant standards of the iron and steel manufacturing industry and drawing insights from the literature on hot-rolling production lines, the SPFRDoc dataset categorizes entities into eight distinct classes (Table 1) and relationships into six categories (Table 2). These categories of relationships encompass fault modes, reasons, operations, equipment, and more, including contextual relationships within the temporal, procedural, and spatial domains. Taking the iron and steel manufacturing industry as an illustrative case, SPFRDoc effectively

Table 1 Entity category information of SPFRDoc

ID	Label	Examples
1	Fault name	"The A1 probe hits the roll during the flaw detection of the 1# grinding machine of the comprehensive group"
2	Fault mode	"Short circuit", "Oil leakage"
3	Fault phenomenon	"Coil burnout", "Broken hinges"
4	Reason	"Oil pollution", "Load anomaly"
5	Operation	"Fasten", "Replace"
6	Equipment	"Motor", "Cylinder"
7	Fittings	"Bolt", "Bearing", "Wiring"
8	Person	"Wang Fei", "Peng Caihai"

Table 2 Relation category information of SPFRDoc

ID	Label	Examples of entity type pairs
1	Failure mode	(Fault name, Fault mode), (Equipment, Fault mode)
2	Failure position	(Equipment, Fittings)
3	Failure reason	(Fault name, Reason), (Fault phenomenon, Reason)
4	Solution	(Operation, Fittings)
5	Responsible person	(Fault name, Person), (Operation, Person)
6	Operation object	(Operation, Fittings)

encapsulates the relations intrinsic to production activities within the manufacturing sector. Consequently, the training set, validation set, and test set are randomly partitioned at a ratio of 8:1:1.

4.3 Baselines and evaluation metrics

In the comparative experiment, the following models were selected as baselines: The BERT-based model BERT-RE (relation extraction), SSAN, as well as the graph-based models like attention guided graph convolutional network (AGGCN), latent structure refinement (LSR), and GAIN.

BERT-RE (Wang et al., 2019) initially derives embedded representations through BERT and subsequently predicts the presence of a relationship between two entities. It then proceeds to predict the specific target relationship.

SSAN (Xu et al., 2021) employs hard-coded indicators to determine whether mentions exist within the same sentence and share the same entity. This approach integrates biaffine mechanisms into the scoring calculation segment of the transformer encoder.

AGGCN (Guo et al., 2019) harnesses an attention-guided GCN. This network inherently learns to selectively focus on pertinent substructures within dependency trees, which prove valuable for relation extraction tasks.

LSR (Nan et al., 2020) establishes graphs based on dependency trees and conducts relationship prediction through latent structure induction and GCN.

GAIN (Zeng et al., 2020) represents distant relations within documents through the creation of mention-level and entity-level graphs. It captures information interactions between diverse mentions via GCN.

The experiment utilizes widely accepted classification task evaluation metrics to assess model performance. The specific calculation formulas are as:

$$Precision = \frac{N_{\text{correct}}}{N_{\text{all}}}, \tag{23}$$

$$Recall = \frac{N_{\text{correct}}}{N_{\text{marked}}}, \tag{24}$$

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall}, \tag{25}$$

where N_{correct} is the number of labels that are predicted to be correct, N_{all} is the total number of predicted labels, and N_{marked} is the total number of marked labels. In addition, the comparison experiment uses the index N_{infer} to quantitatively compare the knowledge mining ability of the model on document information, in which N_{infer} represents the number of relationships inferred by the model excluding the common relation facts in the training and dev/test sets.

4.4 Experiments

Comparative experiment

The experiments detailed in this paper were conducted within a computational environment featuring 4 Tesla T4 GPUs, 256 GB of RAM, and an Intel Xeon 5238 CPU@2.1 Hz x 2.

To gauge the effectiveness of the entity extraction method, the paper evaluates the performance of two models: A standalone BiLSTM-CRF model and a BiLSTM-CRF model integrated with a domain-specific dictionary. These models are applied to the task of NER within the steel manufacturing domain. Prior to utilization, the text annotations in the dataset were transformed into BIO annotations. Here, B denotes the initiation of an entity, I indicates the continuation of an entity, and O signifies non-entities.

The outcomes of the entity recognition task are presented in Table 3, and the progression of loss in relation to iterations during model parameter learning is illustrated in Fig. 4(a). The findings suggest that the incorporation of a domain knowledge dictionary had a limited impact on tasks involving the prediction of labels within datasets boasting a considerable number of samples, such as “fault mode” and “fault name”. This is attributed to the capacity of deep learning models to glean features from abundant samples. Nonetheless, the enhanced method that amalgamated the domain knowledge dictionary displayed marked improvements in tasks associated with label prediction and was characterized by smaller sample sizes, such as “equipment” and “fittings”. These outcomes underscore the potential benefits of integrating domain knowledge dictionaries in scenarios constrained by limited training data.

Subsequent to the completion of the entity recognition phase, entity coreferences are consolidated based on the similarity of entity labels and textual expressions. Subsequent relationship extraction is carried out on the SPFRDoc dataset. Given the dataset’s Chinese documentation, the BERT Chinese pre-training model is employed for character encoding.

Table 3 The results of the entity recognition task

Label	BiLSTM-CRF			BiLSTIM-CRF + dictionary		
	Precision	Recall	*F*1	Precision	Recall	*F*1
Fault name	**89.77**	76.70	82.72	87.11	**86.90**	**87.00**
Fault mode	**83.57**	88.63	**86.02**	82.99	**89.05**	85.92
Fault phenomenon	**89.89**	77.67	83.33	85.29	**84.67**	**84.98**
Reason	77.92	**74.18**	**75.90**	**81.82**	69.23	75.00
Equipment	74.44	83.15	78.55	**85.11**	**87.59**	**86.33**
Fittings	**74.77**	70.34	72.49	70.34	**78.46**	**74.18**
Operation	85.29	**86.67**	83.57	**86.92**	81.67	**84.21**
Person	**83.80**	86.86	85.30	81.58	**90.85**	**85.96**

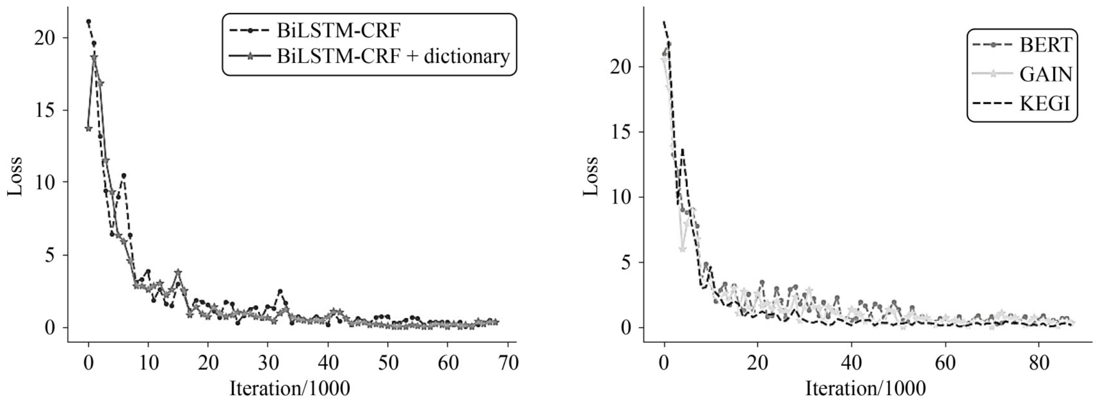

(a) NER models training on SPFRDoc (b) Relation extraction models training on SPFRDoc

Fig. 4 Model loss curves during training.

The experimental results for each method within the SPFRDoc dataset are shown in Table 4. Among the models utilizing various BERT versions or document graphs, KEGI consistently outperforms all robust baselines by a significant $F1$ score of 52.54 on the test set. Within the BERT-based models, KEGI exhibits notable $F1$ improvements of 4.80 and 3.53 for BERT-RE and SSAN, respectively. Further analysis of Table 4 indicates that in comparison with graph-based models, KEGI also achieves $F1$ score enhancements ranging from 0.99 to 5.19 on the test set. Additionally, KEGI ranks second best in the N_{infer} index, falling short only of the LSR model. Furthermore, compared to other graph-based models, KEGI showcases improvements of 62 and 13 for AGGCN and GAIN, respectively.

These results indicate that KEGI elevates prediction accuracy while preserving effective inferential capability. Moreover, it demonstrates superior effectiveness in document-level relation extraction tasks within the manufacturing domain.

To delve deeper into the classification efficacy of the analytical model across various relationships, representative BERT-RE and GAIN models were selected for comprehensive analysis of diverse relationship classification indicators. The resulting experimental outcomes are presented in Table 5 and Fig. 4(b).

Despite the dataset's robust domain characteristics, both the BERT-RE and GAIN models exhibit relatively lower accuracy in relation extraction tasks, primarily due to the scarcity of extensively annotated training datasets. Nonetheless, the incorporation of external knowledge guidance during the graph reasoning phase yields accuracy improvements. The graph reasoning model consistently exhibits superior performance, as evidenced by its higher $F1$ score.

Figure 5 provides a specific illustration of the capabilities of the BERT-RE model and the KEGI model, evaluating their respective competencies in recognizing and categorizing relationships, as well as accurately predicting failure reasons. This comparative analysis sheds light on the strengths and limitations of each model, thus guiding the development of future models aimed at enhancing relational reasoning capabilities. The figure highlights that the BERT-RE model adeptly identifies and categorizes three distinct relationship groups. Additionally, the KEGI model accurately predicts the "failure reason" of the "hydraulic servo valve" as "inadequate fastening". Notably, in a concise sentence containing only one entity, "oil leakage", the KEGI model adeptly incorporates this information into its relational reasoning by treating it as a mention. This underscores the model's capacity for logical reasoning that extends across sentences.

Table 4 The results of the comparative experiment

Model	Precision	Recall	$F1$	N_{infer}
BERT-RE	46.27	49.24	47.74	12
SSAN	47.91	50.06	49.01	59
AGGCN	45.72	49.10	47.35	31
LSR	47.95	52.93	50.32	**104**
GAIN	47.87	55.84	51.55	86
KEGI	**48.04**	**57.97**	**52.54**	93

To validate the efficacy of fusing concept and attribute information, this paper conducts ablation experiments on the SPFRDoc dataset. Training and prediction processes are executed after removing the TransR and path inference module, respectively. The outcomes are illustrated in Table 6.

Ablation experiment

Initially, the domain knowledge base is excluded from the embedded representation module by removing TransR^1. Node vectors are directly initialized using Eq. (8). This action results in a decline of 2.22 in the $F1$ score on the test set, underscoring the pivotal role played by the injection of the knowledge base in the embedding of node features.

Subsequently, TransR^2, which is responsible for representing the domain knowledge base in the path inference module, is removed. This entails eliminating $\{r_1, r_2, \ldots, r_m\}$ from Eq. (19). A marginal reduction in the $F1$ score

Table 5 The results of the relation extraction task

Label	BERT-RE			GAIN			KEGI		
	Precision	Recall	$F1$	Precision	Recall	$F1$	Precision	Recall	$F1$
Failure mode	**52.68**	40.14	45.56	49.79	54.94	52.24	47.83	**68.75**	**56.41**
Failure position	49.31	41.03	44.79	47.90	**55.59**	**51.45**	**49.61**	51.10	50.34
Failure reason	44.92	57.74	50.53	46.18	**58.55**	51.64	**49.00**	54.90	**51.79**
Solution	43.30	56.03	48.85	43.92	58.01	50.00	**45.80**	**58.87**	**51.52**
Responsible person	40.91	56.25	47.37	**49.05**	55.43	**52.04**	44.15	**60.45**	51.03
Operation object	46.52	44.22	45.34	50.39	52.53	51.44	**51.83**	**53.75**	**52.77**

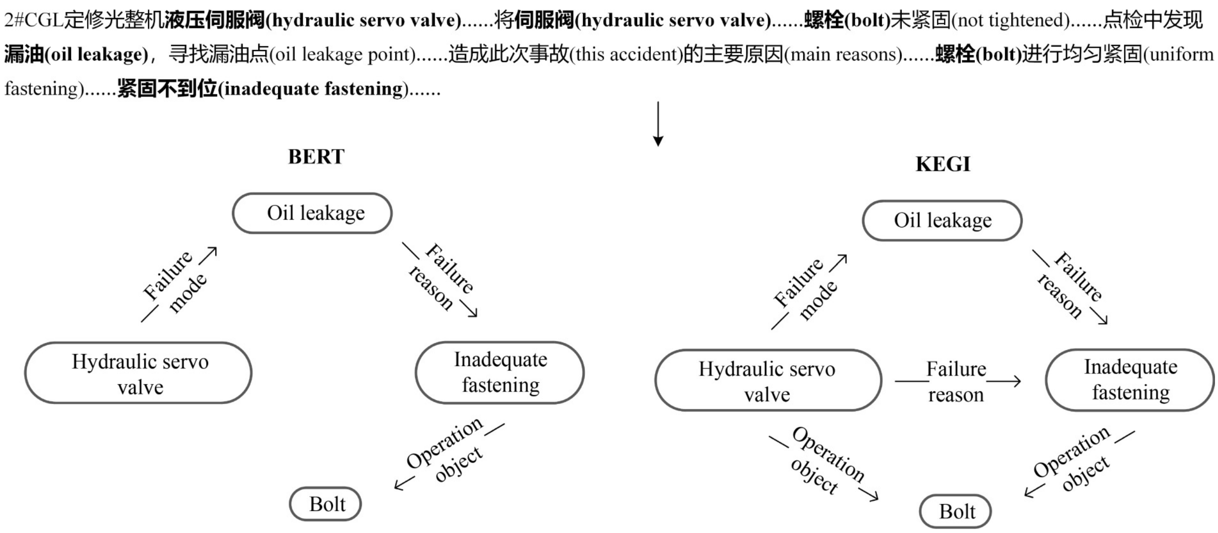

Fig. 5 Comparison of the BERT-RE model and the KEGI model in a specific case.

Table 6 The results of the ablation experiment

Model	Precision	Recall	$F1$
KEGI	**48.04**	**57.97**	**52.54**
−TransR[1]	46.96	57.33	50.32
−TransR[2]	47.30	55.86	51.22
−Inference	47.85	52.59	50.11

by 1.32 is observed, suggesting that the domain knowledge base contributes to path inference to a certain extent.

Following this, the entire path inference module is eliminated. The model then directly predicts relations based on entity representations, e_h and e_t. This omission leads to a noteworthy decline in the $F1$ score, notably a drop of 2.43 on the test dataset. This highlights the significance of the path inference module in capturing potential pathways for relation inference. Its presence effectively enhances the performance of document-level relation extraction tasks.

To ensure efficient querying and facilitate the implementation of advanced applications such as knowledge reasoning and computing, the establishment of an efficient knowledge graph storage mechanism holds paramount importance. In this context, graph databases, which are non-relational databases founded on graph theory principles, offer an apt solution. Graph databases enable the storage of entities and their inter-entity relationships, utilizing nodes, relationships, and attributes as their fundamental components. Among the well-regarded options, Neo4j stands as a commonly employed graph database. Within this study, the triplets were stored using Neo4j. The process involved transforming the mapped file into RDF (resource description framework) format using the open-source toolkit RDF2RDF, followed by importing it into the Neo4j graph database through utilization of the Neosemantics function plug-in. This approach seamlessly achieved the storage of the knowledge graph. Ultimately, the resultant triples find their place within the Neo4j graph database. Approximately 200 nodes were chosen for visualization, as depicted in Fig. 6.

5 Conclusions

The primary objective of this research is to tackle the knowledge management challenges encountered in the industrial domain and propose an innovative solution for extracting knowledge from lengthy paragraphs of text.

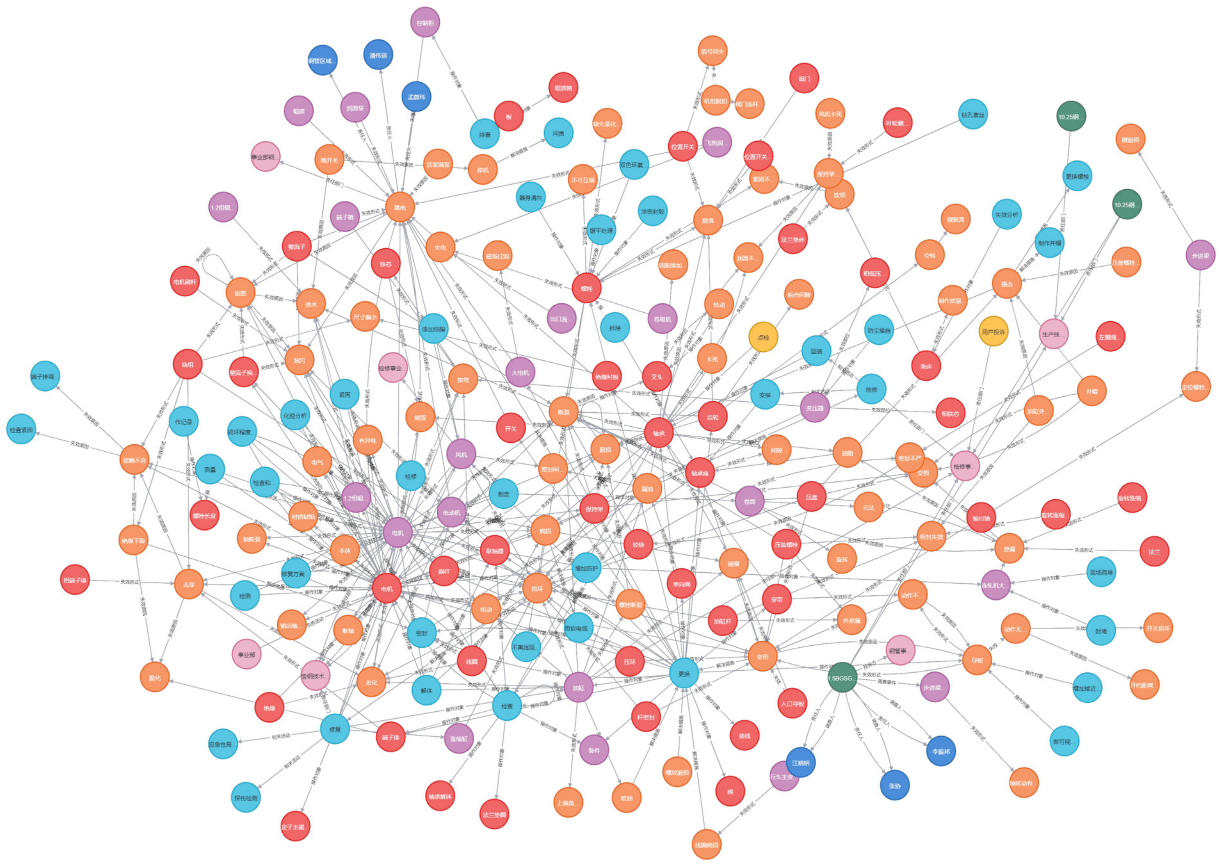

Fig. 6 The fault diagnosis knowledge graph of the Chinese steel manufacturing industry.

This goal is achieved through the development of KEGI, a relation extraction approach that amalgamates the capacity for constructing a document graph with the integration of knowledge representation into node formulation and path inference, all facilitated by the TransR model.

The framework introduced in this paper is semi-automatic in nature and designed to recognize named entities by leveraging the synergy between BiLSTM-CRF and a domain-specific dictionary. The integration of domain knowledge plays a pivotal role in guiding the processes of entity recognition and relation extraction. Furthermore, the optimization of the graph inference mechanism allows for the capture of intricate semantic details embedded within the text.

The proposed framework was practically applied to a real-world context, specifically the steel hot-rolling production line, resulting in the creation of the SPFRDoc and SPOnto datasets. Through this framework, triples were extracted, ultimately leading to the construction of a knowledge graph that embodies the semantic intricacies of the steel industry. This knowledge graph holds potential for diverse applications, including fault diagnosis and intelligent question answering within production processes.

This paper introduces two noteworthy contributions to the realm of relation extraction within the industrial domain:

• First, the KEGI model seamlessly integrates domain knowledge embeddings from the TransR model into the representation of document graph nodes. This integration, combined with the utilization of ontology models, enriches the reasoning mechanism between entities. This combined approach empowers node and path representations to not only encompass context information from the document but also incorporate structural insights from the domain ontology. This effectively addresses the limitations of manual annotation and entity interaction modeling within existing domain corpora.

• Second, the paper demonstrates the practical applicability of the KEGI framework by constructing a knowledge graph for the steel hot-rolling production line, leveraging the SPFRDoc and SPOnto datasets. The resulting knowledge graph encapsulates substantial semantic information and holds potential for tasks such as fault diagnosis and intelligent question answering in real-world production scenarios. Consequently, the KEGI model proficiently extracts information from enterprise documents and offers a robust solution for knowledge management within the industrial domain.

The trajectory of smart factories is poised to be influenced by knowledge-driven and collaborative operations, fostering the evolution of more sophisticated and intelligent decision-making systems. To advance this goal, future research should center around relational reasoning and knowledge fusion within the framework of production knowledge graphs. This approach would facilitate the exploration of knowledge-driven methodologies in decision-making and attribution contexts, thereby aiding businesses in addressing the challenges presented by ever-changing production environments. Such efforts are anticipated to propel the transformation of production management and control toward a more knowledge-driven paradigm, ultimately leading to heightened operational efficiency and effectiveness.

Competing Interests The authors declare that they have no competing interests.

Data Availability Statements The datasets generated and analyzed within the current study are available from China Baowu Steel Group Corporation Limited. Restrictions apply to the availability of steel production data, which were used under license for this study. Steel production data are available from the corresponding author with the permission of China Baowu Steel Group Corporation Limited.

References

Bahdanau D, Cho K, Bengio Y (2014). Neural machine translation by jointly learning to align and translate. arXiv preprint. arXiv: 1409.0473

Cao P, Chen Y, Liu K, Zhao J, Liu S (2018). Adversarial transfer learning for Chinese named entity recognition with self-attention mechanism. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. Brussels: ACL, 182–192

Christopoulou F, Miwa M, Ananiadou S (2019). Connecting the dots: Document-level neural relation extraction with edge-oriented graphs. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing / 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). Hong Kong: ACL, 4925–4936

Deng J, Wang T, Wang Z, Zhou J, Cheng L (2022). Research on event logic knowledge graph construction method of robot transmission system fault diagnosis. IEEE Access, 10: 17656–17673

Dong J, Wang J, Chen S (2021). Knowledge graph construction based on knowledge enhanced word embedding model in manufacturing domain. Journal of Intelligent & Fuzzy Systems, 41(2): 3603–3613

Gui W, Zeng Z, Chen X, Xie Y, Sun Y (2020). Knowledge-driven process industry smart manufacturing. Scientia Sinica Informationis, 50(9): 1345–1360 (in Chinese)

Guo Z, Zhang Y, Lu W (2019). Attention guided graph convolutional networks for relation extraction. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Florence: ACL, 241–251

Han X, Gao T, Lin Y, Peng H, Yang Y, Xiao C, Liu Z, Li P, Zhou J, Sun M (2020). More data, more relations, more context and more openness: A review and outlook for relation extraction. In: Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing. Suzhou: ACL, 745–758

Hu L, Wu G, Xing Y, Wang F (2020). Things2Vec: Semantic modeling in the Internet of Things with graph representation learning. IEEE Internet of Things Journal, 7(3): 1939–1948

Huang Z, Xu W, Yu K (2015). Bidirectional LSTM-CRF models for sequence tagging. arXiv preprint. arXiv:1508.01991

Huet A, Pinquié R, Véron P, Mallet A, Segonds F (2021). CACDA: A knowledge graph for a context-aware cognitive design assistant. Computers in Industry, 125: 103377

Kamble S, Gunasekaran A, Gawankar S (2018). Sustainable Industry 4.0 framework: A systematic literature review identifying the current trends and future perspectives. Process Safety and Environmental Protection, 117: 408–425

Kipf T N, Welling M (2016). Semi-supervised classification with graph convolutional networks. arXiv preprint. arXiv:1609.02907

Lafferty J, McCallum A, Pereira F (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the 18th International Conference on Machine Learning. San Francisco, CA: Morgan Kaufmann Publishers Inc., 282–289

Li G, Pan R, Mao J, Cao Y (2020). Entity recognition of Chinese electronic medical records based on BiLSTM-CRF network and dictionary resources. Journal of Modern Information, 40(4): 3–12, 58 (in Chinese)

Li J, Sun A, Han J, Li C (2022). A survey on deep learning for named entity recognition. IEEE Transactions on Knowledge and Data Engineering, 34(1): 50–70

Lin H, Liu Y, Wang W, Yue Y, Lin Z (2017). Learning entity and relation embeddings for knowledge resolution. Procedia Computer Science, 108: 345–354

Lin Y, Tsai T, Chou W, Wu K, Sung T, Hsu W (2004). A maximum entropy approach to biomedical named entity recognition. In: Proceedings of the 4th International Conference on Data Mining in Bioinformatics. Seattle WA: Springer-Verlag, 56–61

Liu M, Li X, Li J, Liu Y, Zhou B, Bao J (2022). A knowledge graph-based data representation approach for IIoT-enabled cognitive manufacturing. Advanced Engineering Informatics, 51: 101515

Liu W, Zhou P, Zhao Z, Wang Z, Ju Q, Deng H, Wang P (2020). K-BERT: Enabling language representation with knowledge graph. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence. New York, NY: AAAI, 2901–2908

Lyu Z, Wu Y, Lai J, Yang M, Li C, Zhou W (2023). Knowledge enhanced graph neural networks for explainable recommendation. IEEE Transactions on Knowledge and Data Engineering, 35(5): 4954–4968

Nan G, Guo Z, Sekulic I, Lu W (2020). Reasoning with latent structure refinement for document-level relation extraction. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL). Online: ACL, 1546–1557

Qi A, Sin T, Fathullah M, Lee C (2017). The impact of fit manufacturing on green manufacturing: A review. In: Proceedings of the 3rd Electronic and Green Materials International Conference (EGM). Krabi: AIP, 020083

Ren H, Chen Z, Jiang Z, Yang C, Gui W (2021). An industrial multilevel knowledge graph-based local–global monitoring for plant-wide processes. IEEE Transactions on Instrumentation and Measurement, 70: 1–15

Sahu S, Christopoulou F, Miwa M, Ananiadou S (2019). Inter-sentence relation extraction with document-level graph convolutional neural network. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL). Florence: ACL, 4309–4316

Shi H, Huang D, Wang L, Wu M, Xu Y, Zeng B, Pang C (2021). An information integration approach to spacecraft fault diagnosis. Enterprise Information Systems, 15(8): 1128–1161

Strubell E, Verga P, Belanger D, McCallum A (2017). Fast and accurate entity recognition with iterated dilated convolutions. arXiv preprint. arXiv:1702.02098

Wan Z, Ge P, Zhang X, Yin G (2018). Research on equipment manufacturing industry upgrading under intelligent manufacturing. World Sci-Tech R & D, 40(3): 316–327 (in Chinese)

Wang B, Yi B, Liu Z, Zhou Y, Zhou Y (2021). Evolution and state-of-the-art of intelligent manufacturing from HCPS perspective. Computer Integrated Manufacturing Systems, 27(10): 2749–2761 (in Chinese)

Wang B, Zang J, Qu X, Dong J, Zhou Y (2018). Research on new-generation intelligent manufacturing based on human-cyber-physical systems. Strategic Study of CAE, 20(4): 29–34

Wang H, Focke C, Sylvester R, Mishra N, Wang W (2019). Fine-tune BERT for DocRED with two-step process. arXiv preprint. arXiv: 1909.11898

Xiao C, Yao Y, Xie R, Han X, Liu Z, Sun M, Lin F, Lin L (2020). Denoising relation extraction from document-level distant supervision. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP). Online: ACL, 3683–3688

Xu B, Wang Q, Lyu Y, Zhu Y, Mao Z (2021). Entity structure within and throughout: Modeling mention dependencies for document-level relation extraction. In: Proceedings of the AAAI Conference on Artificial Intelligence. Online: AAAI, 14149–14157

Xu Z, Dang Y, Zhang Z, Chen J (2020). Typical short-term remedy knowledge mining for product quality problem-solving based on bipartite graph clustering. Computers in Industry, 122: 103277

Yao Y, Ye D, Li P, Han X, Lin Y, Liu Z, Liu Z, Huang L, Zhou J, Sun M (2019). DocRED: A large-scale document-level relation extraction dataset. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL). Florence: ACL, 764–777

Zeng S, Xu R, Chang B, Li L (2020). Double graph based reasoning for document-level relation extraction. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP). Online: ACL, 1630–1640

Zhang D, Liu Z, Jia W, Liu H, Tan J (2021). A review on knowledge graph and its application prospects to intelligent manufacturing. Journal of Mechanical Engineering, 57(5): 90–113 (in Chinese)

Zhang Z, Han X, Liu Z, Jiang X, Sun M, Liu Q (2019). ERNIE: Enhanced language representation with informative entities. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL). Florence: ACL, 1441–1451

Zheng P, Xia L, Li C, Li X, Liu B (2021). Towards Self-X cognitive manufacturing network: An industrial knowledge graph-based multi-agent reinforcement learning approach. Journal of Manufacturing Systems, 61: 16–26

Zhou J, Zhou Y, Wang B, Zang J (2019). Human-Cyber-Physical Systems (HCPSs) in the context of new-generation intelligent manufacturing. Engineering, 5(4): 624–636

Zhou W, Huang K, Ma T, Huang J (2021). Document-level relation extraction with adaptive thresholding and localized context pooling. In: Proceedings of the AAAI Conference on Artificial Intelligence. Online: AAAI, 14612–14620

Zhou Y, Huang H, Liu H, Hao Z (2022). Survey on document-level relation extraction. Journal of South China University of Technology (Natural Science Edition), 50(4): 10–25 (in Chinese)