Knowledge graph for blast furnace condition prediction and root cause tracing

Xiaojie Liu[1,2], Yanping Geng[1], Hongwei Li[1,2*], Yujie Zhang[1], Yifan Duan[1], Xin Li[3*]
1.North China University of Science and Technology, School of Metallurgy and Energy, Tangshan, China
2.North China University of Science and Technology, School of Science, Tangshan, China
3.North China University of Science and Technology, School of Civil Engineering, Tangshan, China

Xiaojie Liu: E-mail：liuxiaojie@ncst.edu.cn
Yanping Geng: E-mail：2791506833@qq.com
Hongwei Li: E-mail：lhwcoarising@ncst.edu.cn
Yujie Zhang: E-mail：yjzhang@stu.ncst.edu.cn
Yifan Duan: E-mail：yifanduan@stu.ncst.edu.cn
Xin Li: E-mail：lxxx@ncst.edu.cn
*Corresponding author: Hongwei Li, Xin Li
E-mail: lhwcoarising@ncst.edu.cn, lxxx@ncst.edu.cn

(Received 11 May 2026; Accepted 27 August 2026)

Abstract: Accurate prediction of blast furnace conditions and traceability of anomaly root causes are crucial for ensuring production stability and reducing energy consumption. However, existing methods struggle to simultaneously capture temporal dynamic evolution and causal structural relationships among parameters. To address this issue, this paper proposes a method for furnace condition prediction and root cause tracing that integrates a knowledge graph with dual-encoder deep learning. First, based on process mechanisms and expert rules, a knowledge graph containing anomaly types and their causal relationships is constructed, and historical cases are stored. Second, a temporal-graph dual encoder is designed: the temporal encoder employs a temporal convolutional network to extract dynamic evolution features from data, while the graph encoder uses a graph attention network to learn structural dependencies from the knowledge graph; a gated fusion mechanism then enables accurate prediction of six types of furnace conditions. Finally, a root cause tracing algorithm based on causal paths and case similarity is proposed, which leverages causal associations in the knowledge graph and historical cases to generate interpretable root cause chains. Experimental results show that the proposed method achieves an overall prediction accuracy of over 90% and an overall tracing accuracy of 91.7%, and the generated paths are highly consistent with process mechanisms.

Keywords: blast furnace condition prediction; knowledge graph; root cause tracing; fault diagnosis; deep learning; ironmaking

1 Introduction

The complexity of the blast furnace ironmaking process leads to furnace conditions being influenced by the coupling of multiple parameters. Accurately identifying anomalies and tracing their root causes is of great significance for production control. Blast furnace ironmaking is a core process in steel production, and its operating state directly determines hot metal quality, energy consumption, and production safety. The interior of a blast furnace involves complex physicochemical reactions and is characterized by high temperature, high pressure, and multiphase flow, making the in-furnace state difficult to observe directly; thus, it is known as a typical industrial "black box". Six typical anomalies, such as hanging, channeling, cold furnace, hot furnace, slip, and hearth buildup, occur frequently. These anomaly types are selected based on on-site statistical frequency and their impact on the process, consistent with industry standards and possessing good representativeness. Failure to detect and accurately locate the root cause in a timely manner will lead to reduced output, increased energy consumption, and even safety accidents[1].

Currently, the judgment of blast furnace conditions still primarily relies on the experience of operators, who make manual decisions by observing trends in key parameters such as blast pressure, permeability index, and molten iron temperature. However, manual experience has limitations, including strong subjectivity, lagging response, and difficulty in standardization, making it challenging to meet the urgent demands of the modern steel industry for intelligent operation and maintenance[2].

With the advancement of industrial big data and artificial intelligence technologies, data-driven furnace condition diagnosis methods have garnered widespread attention. Early research employed traditional machine learning models such as Support Vector Machines and Random Forests[3], achieving some success in classifying furnace conditions by manually extracting statistical features[4]. In recent years, deep learning models like Long Short-Term Memory networks and Convolutional Neural Networks have been applied to directly process time-series data, enabling automatic learning of temporal dependencies and further improving diagnostic accuracy[5]. However, most of these data-driven methods focus solely on changes in parameter values, overlooking the inherent causal relationships among blast furnace parameters. This leads to a lack of model interpretability, making it difficult to answer the critical question of "why an anomaly occurred" and failing to provide clear adjustment directions for field operators[1-2].

As a structured knowledge representation tool, knowledge graphs can explicitly model semantic relationships between entities, demonstrating unique advantages in the field of industrial fault diagnosis and root cause analysis. For instance, in the domain of Computer Numerical Control machine tools, researchers have constructed graphs incorporating equipment parameters, fault modes, and maintenance records to assist fault diagnosis through graph querying and reasoning[6]. However, existing research on industrial knowledge graphs is often limited to static knowledge and fails to integrate deeply with time-series data, making it difficult to capture causal relationships during the dynamic evolution of parameters. Recently, some cutting-edge studies have begun exploring the combination of time-series data and knowledge graphs[7-8]. For example, LMM-FD proposed a multimodal large model for fault diagnosis[9], and MM-TKGF constructed a multimodal temporal knowledge graph for reasoning[10-11]. These works offer new insights for blast furnace condition prediction, but how to effectively apply them to the complex operating conditions of a blast furnace and achieve accurate prediction still requires further investigation.

To address the aforementioned issues, this paper proposes a method for blast furnace condition prediction and root cause tracing based on a knowledge graph and dual-encoder deep learning. The main contributions include:

1) The construction of a blast furnace condition knowledge graph, which integrates process parameters, six types of anomalies and their causal relationships, and stores historical cases, providing a structured knowledge foundation for prediction and tracing.

2) The proposal of a TCN-GAT dual-encoder fusion model: the temporal encoder extracts dynamic evolution features from minute-level data, while the graph encoder learns structural dependencies among parameters. Gated fusion is employed to achieve complementary advantages and enhance the accuracy of furnace condition prediction.

3) The design of a root cause tracing algorithm based on causal paths and case similarity, which utilizes causal associations in the graph to generate candidate causes and combines temporal feature matching with historical case similarity to produce interpretable root cause chains.

4) The effectiveness of the proposed method is validated using industrial data. The knowledge graph and diagnostic model are integrated to develop an intelligent blast furnace condition monitoring system, which enables real-time monitoring, anomaly diagnosis, root cause visualization, and recommendation push.

2 Construction of the Blast Furnace Condition Knowledge Graph

2.1 Ontology Design

The ontology layer of a knowledge graph defines entity types, object properties, and data properties[12], providing a unified schema constraint for knowledge organization and reasoning[13]. Aiming at the tasks of blast furnace condition prediction and root cause tracing, this paper constructs an ontology structure covering multiple types of elements, with its core components as follows.

Entity Types: To comprehensively describe the blast furnace production process, the ontology defines various entity types, including process parameters, furnace condition states, historical cases, raw materials and fuels, operation regimes, production indicators, and blast furnace equipment. Among these, process parameters reflect the dynamic changes of furnace conditions; furnace condition states characterize the operational status; historical cases record instances of abnormal events; raw materials, fuels, and operation regimes describe production inputs and control conditions; production indicators measure operational performance; and blast furnace equipment is associated with physical entities. The aforementioned entities collectively constitute a knowledge description framework for the blast furnace process.

Object Properties: To characterize semantic associations between entities, the ontology defines several core object properties. Affects describes the direct influence between parameters or elements; causeBy is used to trace the causal path from an anomaly to its cause; manifestAs associates anomaly types with patterns of parameter changes; and instanceOf connects specific cases to their corresponding anomaly types. Furthermore, relationships such as hasFeature, hasParameter, and leadTo further enrich the semantic network among entities, providing a foundation for causal reasoning and path discovery.

Data Properties: To enable the quantitative representation of entity information, the ontology configures various data properties, such as name, description, unit, value, and time. These properties are attached to entities, endowing the knowledge graph with both symbolic semantics and numerical information, thereby supporting subsequent analytical tasks such as case feature matching and similarity calculation.

Fig. 1 illustrates a schematic diagram of the ontology structure of the knowledge graph constructed in this paper. This design provides a structured knowledge foundation for the subsequent temporal-graph dual-encoder prediction and root cause tracing.

Fig. 1 Ontology schema of the blast furnace condition knowledge graph

2.2 Knowledge Acquisition and Cleaning

The construction of a knowledge graph first relies on the acquisition and standardized processing of multi-source heterogeneous data[14]. Targeting the field of blast furnace condition diagnosis, this paper collects four types of knowledge sources: operation procedure documents, which cover standard operating procedures, parameter control ranges, and anomaly handling measures for each process step; fault case reports, which provide detailed descriptions, cause analyses, and handling processes of historical anomaly events; expert experience records, which store diagnostic rules and experiential knowledge accumulated by process experts over time in textual form; and a historical case database, which contains instances of anomaly events with annotated time periods and their corresponding parameter data. The sources mentioned above encompass various forms, ranging from normative knowledge to experiential knowledge, and from structured to unstructured data.

Significant differences exist among multi-source data in terms of format specifications, terminology expression, and level of detail, necessitating cleaning processes to ensure the effectiveness of subsequent knowledge extraction[15]. The cleaning steps include: standardizing formats by converting original documents such as PDF, Word, and plain text into UTF-8 encoded plain text to eliminate format heterogeneity; normalizing terminology by constructing a domain-specific synonym dictionary and performing consistent replacement for different expressions of the same concept, such as "blast pressure" and "hot blast pressure," or "hanging" and "slipping"; removing redundant information by deleting duplicate paragraphs, irrelevant annotations, and formatting marks; and correcting textual errors by automatically identifying and rectifying typos, punctuation misuse, and unit spelling errors based on rules and dictionaries. After the aforementioned cleaning process, a standardized domain-specific text corpus is formed, providing high-quality input data for subsequent knowledge extraction.

2.3 Knowledge Extraction and Fusion

Knowledge extraction aims to automatically identify domain entities and their semantic relations from the cleaned text corpus. Knowledge fusion is then used to integrate the extraction results from multiple sources, eliminating redundancy and conflicts to form a structured knowledge representation. This paper employs a joint model based on ALBERT-BiLSTM-CRF for knowledge extraction, integrates and disambiguates the extraction results through a knowledge fusion mechanism, and validates the effectiveness of the model through experiments [16].

2.3.1 Knowledge Extraction Model

This paper adopts the ALBERT-BiLSTM-CRF joint model as the knowledge extraction framework, as illustrated in Fig. 2. The model consists of four core layers:

1) ALBERT Layer: As a pre-trained language model, it maps the input text into dynamic, context-aware vector representations. It reduces model complexity through parameter sharing mechanisms, making it suitable for semantic encoding of domain-specific texts.

2) BiLSTM Layer: Building upon the output of ALBERT, a Bidirectional Long Short-Term Memory network is employed to capture contextual dependencies within the sequence, enhancing the model's ability to model long-distance semantic information.

3) CRF Layer: The Conditional Random Field is used to learn transition constraints between output labels. It ensures the rationality of the entity label sequence through global normalization, avoiding illegal label combinations.

4) Sigmoid Layer: After CRF decoding, a Sigmoid activation function outputs the probability distribution for each entity type and relation category. This converts the model's predictions into confidence scores, providing a screening basis for subsequent knowledge fusion.

Fig. 2 Architecture of the ALBERT-BiLSTM-CRF knowledge extraction model

2.3.2 Knowledge Fusion

As the extraction results originate from different sources, they may contain redundancies, conflicts, or semantic inconsistencies, necessitating integration through a knowledge fusion mechanism[17]. Specifically, an approach combining string similarity and semantic similarity is adopted for entity alignment, identifying different expressions referring to the same entity (e.g., "blast pressure" and "hot blast pressure") to achieve entity unification. For multiple relation descriptions that may exist between the same pair of entities, they are mapped to standard relation types (e.g., causeBy, affects) using a predefined relation dictionary, ensuring consistency in relational semantics. When conflicts arise among multi-source knowledge, such as different normal ranges assigned to the same parameter, a comprehensive judgment is made based on source authority, extraction confidence, and the majority principle, selecting the most credible knowledge as the final content of the graph. Additionally, a confidence score is assigned to each piece of fused knowledge, calculated by weighting the model's prediction probability, source authority, and multi-source consistency, which is used for quality control and application screening of the subsequent knowledge graph.

In case of multi-source knowledge conflicts, the following rules are applied: if multiple knowledge items describing the same fact are consistent (within a preset tolerance), the multi-source consistency indicator is increased; if conflicts exist, selection is made based on source authority ranking, and the conflict is recorded in the knowledge graph for subsequent manual review. The specific conflict determination logic is as follows: suppose there are n candidate knowledge items $\{K_1, K_2, ..., K_n\}$ describing the same fact, with the comprehensive confidence C_i calculated by the aforementioned formula. If $n=1$, the knowledge is directly adopted. If $n>1$, the difference $C_{\max} - C_{\mathrm{sec}ond}$ is computed. If the difference exceeds the threshold T (set to 0.15), the knowledge with the highest confidence is adopted. Otherwise, the conflict is marked for expert adjudication.

2.3.3 Analysis of Experimental Results

To validate the effectiveness of the ALBERT-BiLSTM-CRF model for knowledge extraction tasks in the blast furnace domain, this paper designs comparative experiments. CRF, BiLSTM-CRF, and BERT-BiLSTM-CRF are selected as baseline models, trained and tested on the same annotated dataset. The dataset is split into training, validation, and test sets in an 8:1:1 ratio, with precision, recall, and F1-score adopted as evaluation metrics.

The experimental results are shown in Table 1. The ALBERT-BiLSTM-CRF model outperforms all baseline models across all metrics, achieving an F1-score of 92.37%. This represents an improvement of 5.63 percentage points over BiLSTM-CRF and 1.33 percentage points over BERT-BiLSTM-CRF. The performance gain is primarily attributed to ALBERT's cross-layer parameter sharing and factorized embedding techniques, which reduce model complexity while retaining strong semantic representation capabilities. Combined with BiLSTM's contextual modeling and CRF's global constraints, the model can more accurately identify domain entities and their boundaries, effectively handling dependencies within label sequences.

Table 1 Comparison of performance among different knowledge extraction models (%)

Model	Precision	Recall	F1-score
CRF	78.64	76.21	77.41
BiLSTM-CRF	87.53	85.96	86.74
BERT-BiLSTM-CRF	90.85	91.23	91.04
ALBERT-BiLSTM-CRF	92.18	92.56	92.37

To further demonstrate the effectiveness of knowledge extraction, Table 2 lists examples of triples extracted from the text. These triples cover various relation types, including parameter-parameter, parameter-anomaly, and case-anomaly relationships, constituting the fundamental units of the knowledge graph. Here, “Case_0023” denotes a specific case identifier from the historical case database, stored as an instance node in the knowledge graph to associate with the corresponding anomaly type.

Table 2 Examples of knowledge extraction results

Head entity	Relation type	Tail entity	Confidence
Hot blast pressure	affects	Permeability index	0.95
Decrease in permeability index	manifestAs	Hanging	0.92
Hanging	causeBy	Deteriorated coke quality	0.78
Case_0023	instanceOf	Cool Furnace	0.87
Excessive burden descent speed	leadTo	Burden compaction	0.84
Hearth Buildup	hasParameter	Stave temperature	0.91

As shown in Table 2, the knowledge graph explicitly incorporates the influencing relationships between parameters through triples such as “Hot blast pressure → affects → Permeability index”. The graph encoder subsequently learns the structural dependencies among these parameters via the graph attention mechanism. For each anomaly type, key influencing parameters are extracted from expert knowledge and process mechanisms and represented as associative edges in the knowledge graph, enabling the model to perceive the influence weights of various parameters on each anomaly through the graph structure.

The confidence of each fused knowledge is calculated by weighted combination of three indicators: model prediction probability, source authority, and multi-source consistency, with the formula given as follows:

$$C = \alpha \cdot P_{\mathrm{model}} + \beta \cdot A_{source} + \gamma \cdot M_{consistency} \tag{1}$$

Where P_{model} is the prediction probability output from the Sigmoid layer of the ALBERT-BiLSTM-CRF model; A_{source} is the source authority, assigned based on the reliability level of the knowledge source; $M_{consistency}$ is the multi-source consistency, measured by the repetition frequency of the same triple across different sources; α, β, and γ are the corresponding weight coefficients, set to 0.5, 0.25, and 0.25, respectively.

The conflict resolution of the normal range for "Hot blast pressure" serves as a concrete example of the aforementioned procedure. Multiple candidate values were obtained from different sources during the knowledge extraction stage, derived from operating procedures (300–400 kPa), expert experience records (310–410 kPa), and historical literature (350–450 kPa). After similarity computation and semantic alignment, all three candidates were mapped to the same parameter entity. The source authority values were assigned as 0.9, 0.7, and 0.6, respectively; the model prediction probabilities were 0.88, 0.82, and 0.75, respectively; and the multi-source consistency metric was 0.8. The comprehensive confidences calculated by the formula were 0.86, 0.78, and 0.68, respectively. The difference between the maximum and the second maximum was 0.08, below the conflict determination threshold of 0.15, triggering a conflict flag for expert adjudication, which ultimately confirmed 300–400 kPa as the correct range.

The experimental results demonstrate that the ALBERT-BiLSTM-CRF model adopted in this paper can effectively extract high-quality knowledge triples from domain-specific texts, providing reliable data support for the subsequent construction of the knowledge graph.

2.4 Graph Storage and Visualization

Upon completion of knowledge extraction and fusion, the structured knowledge is stored in a graph database. This paper employs Neo4j as the underlying graph database, mapping entities such as process parameters, anomaly types, and cases as nodes, and the relationships between entities as labeled directed edges, with entity attributes stored in key-value pairs. An indexing mechanism is also established to improve query efficiency.

For visualization, the Neo4j Browser is utilized for interactive graph display. Fig. 3 presents a partial example of the constructed blast furnace condition knowledge graph, encompassing core parameters such as blast pressure and permeability index, as well as anomaly types like hanging and furnace cooling. The directed edges between nodes clearly depict relationships such as affects, causeBy, and manifestAs. Different colors denote different types of nodes, where pink represents process parameters, red represents anomaly types, and green represents judgement rules, orange represents operation recommendation, etc. This visualization scheme intuitively presents the complex associations among entities and provides support for highlighting causal paths in subsequent root cause tracing tasks.

Fig. 3 Architecture of the ALBERT-BiLSTM-CRF knowledge extraction model

3 Blast Furnace Condition Prediction Model Based on TCN-GAT Dual Encoder

To achieve accurate prediction of blast furnace conditions, this paper proposes a dual-encoder model integrating a Temporal Convolutional Network and a Graph Attention Network[18-19]. The model takes sequences of process parameters as input, while simultaneously incorporating the structural dependencies among parameters derived from the knowledge graph. The temporal dynamic features and graph structural features are extracted separately by the dual encoders, integrated via a gated fusion mechanism, and ultimately used for multi-class classification of furnace conditions[20]. The overall architecture of the proposed model is illustrated in Fig. 4.

Fig. 4 Overall architecture of the proposed model

Let the data collected from the blast furnace production process be denoted as $X = \{x_1, x_2, ..., x_T\}$, where $x_t \in \mathbb{R}^F$ represents the values of F process parameters at time step t. The objective of the task is to determine the furnace condition category $y_t \in \{C_0, C_1, ..., C_k\}$ at the current time step based on an observation window of length $L: \{x_{t-L+1}, ..., x_t\}$. Here, C_0 denotes the normal condition, while C_1 through C_k correspond to six types of anomalies: hanging, slipping, furnace cooling, furnace heating, channeling, and hearth accumulation. Additionally, the associative structure among parameters within the knowledge graph, represented as $G = (V, E)$, is utilized as auxiliary information. In this graph, the node set V corresponds to the process parameters, and the edge set E represents causal relationships or anomaly manifestation relationships between parameters, aiming to enhance the model's capability to capture interdependencies among parameters.

3.1 Temporal Encoder

The temporal encoder employs a Temporal Convolutional Network to extract dynamic evolution features from process parameter sequences. TCN is composed of multiple stacked residual blocks. Each residual block contains dilated convolution, weight normalization, ReLU activation, and Dropout layers, and incorporates residual connections to mitigate the vanishing gradient problem[21]. Dilated convolution expands the receptive field exponentially through increasing dilation factors, enabling the model to capture long-range temporal dependencies while maintaining the advantage of parallel computation.

Fig. 5 illustrates the internal structure of a TCN residual block. Each residual block consists of a dilated convolution layer, weight normalization, a ReLU activation function, and a Dropout layer. The input is added directly to the convolutional output via a residual connection, and then passed to the next layer after ReLU activation. Its computational process is described by Eq. (2).

$$Z^{(l)} = Dropout(\mathrm{Re}LU(WeightNorm(Conv1D(H^{(l-1)})))) \tag{2}$$

Here, Conv1D employs dilated convolution, with the dilation factor increasing exponentially with the layer depth. Given an input sequence $X \in \mathbb{R}^{L\times F}$, the TCN outputs temporal features $H_t \in \mathbb{R}^{L\times d_t}$ after multiple convolutional layers, where d_t is the hidden dimension of the temporal encoder. The output at the last time step is taken as the temporal encoding vector $h_t^{tcn} \in \mathbb{R}^{d_t}$, representing the temporal dynamic information at the current moment. In addition, the dilated convolutions in TCN enable the receptive field to expand exponentially with layer depth, effectively capturing parameter evolution trends over several hours, thereby mitigating the time lag effect of causal relationships in blast furnace operations to a certain extent.

Fig. 5 Schematic diagram of the TCN residual block structure

3.2 Graph Encoder

The graph encoder employs a Graph Attention Network to learn the structural dependencies among parameters. GAT assigns different weights to the neighbors of each node through a self-attention mechanism and aggregates neighbor information to update node representations, effectively capturing complex associations between nodes in the graph[22]. Based on the constructed blast furnace knowledge graph, this paper extracts a parameter subgraph relevant to the current moment, treating process parameters as nodes and relation types as edges, and utilizes GAT for node feature updating.

The basic features of the process parameters are used as the initial node features $h_v^0 \in \mathbb{R}^{d_g}$. Through a multi-head attention mechanism, the GAT layer computes the attention coefficient a_{vu} between node v and its neighbor u, and obtains the updated node representation through weighted aggregation. Fig. 6 illustrates the attention computation process of a single-head GAT layer: For a central node i and its neighbor node j, a shared linear transformation is first applied to obtain Wh_i and Wh_j. The attention coefficient e_{ij} is then calculated, normalized via softmax to obtain a_{ij}, and finally, the neighbor features are aggregated using the attention weights to produce the updated node representation. Its computational process is described by Eq. (3).

$$
\begin{gathered}
e_{vu} = LeakyReLU(a^{\top}[Wh_v \| Wh_u]) \\
\alpha_{vu} = \frac{\exp(e_{vu})}{\sum_{k \in N_v} \exp(e_{vk})} \\
h_v^{'} = \sigma(\sum_{u \in N_v} \alpha_{vu} W h_u)
\end{gathered}
\tag{3}
$$

Where W is a learnable weight matrix, a is the attention parameter vector, $\|$ denotes vector concatenation, and N_v is the set of neighbors of node v. Multi-head attention concatenates or averages the outputs of multiple independent heads as the final node representation. After multiple GAT layers, average pooling over all nodes is applied to obtain the global graph feature vector $h^{gat} \in \mathbb{R}^{d_g}$, representing the overall structural information among parameters.

Fig. 6 Schematic diagram of the graph attention mechanism computation

3.3 Feature Fusion and Classification

To effectively combine temporal dynamic features with graph structural features, this paper employs a gated fusion mechanism to integrate the outputs of the dual encoders. Gated fusion adaptively balances the contributions of the two types of features through learnable weights, with its computation formula shown in Eq. (4).

$$
\begin{gathered}
g = \sigma(W_g[h_t^{tcn}; h^{gat}] + b_g) \\
h^{fuse} = g \odot h_t^{tcn} + (1-g) \odot h^{gat}
\end{gathered}
\tag{4}
$$

Where $[\cdot\,;\,\cdot]$ denotes vector concatenation, W_g and b_g are learnable parameters, σ is the Sigmoid activation function, and $\odot$ represents element-wise multiplication.

The fused feature vector h^{fuse} is passed through a fully connected layer and a Softmax function to obtain the probability distribution for each category. The formula is as follows:

$$
\hat{y} = Soft\max(W_c h^{fuse} + b_c) \tag{5}
$$

The model is trained using the cross-entropy loss function, incorporating class weights to mitigate the issue of sample imbalance. The loss function is defined as:

$$\Lambda = -\frac{1}{N}\sum_{i=1}^{N} w_{y_i} \log(\hat{y}_i, y_i) \tag{6}$$

Where W_{y_i} is the weight for class y_i, set according to the inverse of the number of samples for each class in the training set.

Through the aforementioned gated fusion mechanism, the model can adaptively integrate temporal dynamic features and graph structural features, providing a high-quality fused representation for subsequent furnace condition classification.

3.4 Model Training and Optimization

The proposed model is implemented using the PyTorch framework and trained end-to-end with the Adam optimizer. The dataset is chronologically divided into training set (70%), validation set (10%), and test set (20%) to prevent future information leakage. The temporal encoder consists of four stacked TCN residual blocks with a hidden dimension of 64, kernel size of 3, and dilation factors of 1, 2, 4, and 8, respectively. The graph encoder comprises two multi-head GAT layers with four attention heads and a hidden dimension of 64. The fusion layer has a dimension of 128. The batch size is set to 64, and the initial learning rate is 0.001.

During training, a learning rate decay strategy is employed: if the validation loss does not decrease for five consecutive epochs, the learning rate is reduced to half of its current value. An early stopping mechanism (patience = 10) is also introduced to prevent overfitting. The loss function adopts weighted cross-entropy, with class weights set inversely proportional to the number of samples for each class in the training set to address the class imbalance problem. The experimental hardware environment is an NVIDIA RTX 3090 GPU. The model achieving the best performance on the validation set is selected for evaluation on the test set.

3.5 Performance Evaluation and Result Analysis

To validate the effectiveness of the proposed model, XGBoost, LSTM, TCN, and GAT are selected as baseline models for comparison. XGBoost takes manually constructed statistical features as input, while LSTM, TCN, and GAT use the raw time series or the constructed graph structure as input, respectively. To ensure fairness, all baseline models adopt the same training strategy as the proposed method.

To clarify the quantitative criteria for each anomaly type, this paper defines the corresponding measurable process parameters and their variation thresholds based on site operating procedures, as presented in Table 3. Representative parameters are listed for each anomaly, while comprehensive assessment of multiple parameters is required in practice. The "Normal range" column lists the operating intervals of each parameter under normal conditions, with data derived from site operating procedures. An anomaly is triggered when a parameter deviates from its normal operating range and the magnitude of change reaches the specified threshold.

Table 3 Quantitative criteria for each anomaly type

Anomaly type	Parameter	Trend	Threshold	Normal range
	Lower static pressure	Decrease	≥25 kPa/h	100–150 kPa
Hanging	Hot blast pressure	Increase	≥15 kPa/min	300–400 kPa
	...	...	...	...
	Stockline	Two consecutive drops	>0.2 m	1.5–2.5 m
Slip	Hot blast pressure	Instantaneous increase	≥1 kPa/10 s	300–400 kPa
	...	...	...	...
	Lower static pressure	Increase	≥25 kPa/h	100–150 kPa
Cool Furnace	Hot blast pressure	Decrease	≥10 kPa/30 min	300–400 kPa
	...	...	...	...
	Lower static pressure	Decrease	≥25 kPa/h	100–150 kPa
Hot Furnace	Hot blast pressure	Increase	≥10 kPa/30 min	300–400 kPa
	...	...	...	...

Channeling	Top temperature	Increase	≥50℃/min	150–250℃
	Top pressure	Increase	≥10 kPa/7.5 s	80–120 kPa
	...	...	...	...
Hearth Buildup	Hearth center temperature	Decrease	≥50℃	500–600℃
	Tapping volume (two tap holes)	Deviation from theoretical	≥10%	±5%
	...	...	...	...

The comparison results of Precision, Recall, Accuracy, and F1 for each model on the test set are shown in Fig. 7.

Fig. 7 Performance comparison of different models

As shown in Figure 7, the proposed TCN-GAT dual-encoder model outperforms all baseline models across all evaluation metrics, achieving an F1 score and an accuracy of 0.90, which are 6 percentage points and 11 percentage points higher than those of the suboptimal GAT model, respectively. Specifically, XGBoost relies on manually engineered statistical features and struggles to fully capture the temporal dynamics and complex dependencies among parameters, resulting in relatively low performance. Although LSTM can model temporal information, it makes insufficient use of the structural relationships among parameters. GAT only exploits the graph structure while ignoring temporal evolution, which limits its performance. TCN performs well in temporal modeling but lacks inter-parameter correlation information. By effectively integrating both types of features through a dual-encoder architecture, the proposed model significantly improves the accuracy and robustness of furnace condition prediction.

To further evaluate the performance of the proposed model in identifying various types of anomalies, Fig. 8 presents the confusion matrix on the test set. The rows of the matrix represent the true classes, while the columns represent the predicted classes. The diagonal elements indicate the classification accuracy for each class. In the figure, Class 0 denotes Normal, and Classes 1 to 6 correspond to six types of anomalies: Hanging, Slip, Cool Furnace, Hot Furnace, Channeling, and Hearth Buildup, respectively.

Fig. 8 Visual analysis of confusion matrix

As shown by the diagonal values, the model achieves a recognition accuracy of 93.2% for normal furnace conditions, while the recognition accuracy for various anomalies ranges from 87.7% to 92.0%. This indicates that the model can effectively distinguish between normal and abnormal states while maintaining stable discriminative ability across different anomaly types. The off-diagonal elements further reveal a few easily confused misclassification patterns. For example, 4.0% of Class 3 instances are misclassified as Class 4, and 3.2% of Class 5 instances are misclassified as Class 1. Overall, the misclassification distribution is relatively dispersed, with no significant systematic bias, verifying the model’s comprehensive robustness in identifying multiple anomaly categories.

3.6 Ablation Experiments

To evaluate the effectiveness of the dual-encoder architecture and the gated fusion mechanism, this paper designs four sets of ablation experiments, comparing the following model variants: T1 (TCN encoder only), T2 (GAT encoder only), T3 (simple concatenation of TCN and GAT), and T4 (the proposed TCN-GAT gated fusion model). The performance of each variant on the test set is presented in Table 4.

Table 4 Results of ablation experiments

Model Variant	Accuracy	Precision	Recall	F1
T1	0.82	0.81	0.84	0.82
T2	0.81	0.83	0.85	0.84
T3	0.89	0.84	0.86	0.85
T4	0.92	0.91	0.90	0.90

As shown in Table 4, the performance of single encoders T1 and T2 is limited. T3, which simply fuses temporal and graph structural information, achieves an accuracy of 0.89 and an F1 score of 0.85. T4 introduces gated fusion to further optimize feature integration, improving accuracy, precision, and recall to 0.90, 0.91, and 0.90, respectively—an increase of 1, 7, and 4 percentage points over T3. The results demonstrate that gated fusion enhances the synergy of the dual encoder and significantly improves diagnostic performance.

4 Root Cause Tracing Based on Causal Reasoning

To further locate the fundamental causes of anomalies after accurately identifying the anomaly types with the furnace condition prediction model[23], and to provide interpretable decision support for field operators, this paper proposes a root cause tracing method that integrates causal path reasoning with historical case similarity based on the constructed blast furnace knowledge graph. Starting from the anomaly node output by the prediction model, this method traces backward along the causal paths in the knowledge graph, combines similarity matching with historical case features, and ultimately generates a confidence-weighted root cause chain[24].

4.1 Tracing Path Reasoning

The knowledge graph defines direct influence relationships (affects) between parameters and causal associations (causeBy) between anomalies and parameters, providing structured prior knowledge for root cause tracing. Given the anomaly type $a \in A$ identified by the prediction model, the objective of the tracing task is to find the parameter change path $P = \{v_1, v_2, ..., v_k\}$ that may lead to the occurrence of this anomaly, where v_i is a parameter or state node, satisfying $v_1 \xrightarrow{causeBy} v_2 \xrightarrow{affects} ... \xrightarrow{affects} v_k = a$, the path consists of causal and influence edges all directed toward the anomaly.

This paper adopts a breadth-first reverse search strategy for path reasoning[25]. Starting from the anomaly node a, it searches backward along the causeBy relationship to find the direct parameter nodes that may cause the anomaly, and then continues tracing further upstream along the affects relationship to deeper cause parameters until no upstream nodes are found or the preset maximum depth D is reached. Let $G = (V, E)$ denote the knowledge graph, where each edge $e = (u, v, r, w)$ represents the relationship r from node u to node v and its confidence $w \in [0,1]$. For a path $P = (v_1, v_2, ... v_k)$, its initial confidence is defined as the product of the confidences of all edges along the path, as shown in Eq. (7).

$$conf(P) = \prod_{i=1}^{k-1} w(v_1, v_{i+1}) \qquad (7)$$

Where $w(v_i, v_{i+1})$ denotes the confidence of the relational edge from node v_i to node v_{i+1}. This product reflects the causal transmission credibility of the entire path.

4.2 Case Similarity Matching

The historical case database stores annotated anomaly events and their corresponding feature vectors. Each case c contains the anomaly type a_c, the occurrence time period $[t_{start}, t_{end}]$, and a feature vector $f_c \in \mathbb{R}^F$, where the feature vector is composed of the mean values of each process parameter during the 30 minutes prior to the anomaly, i.e., $f_c = [\overline{x_1}, \overline{x_2}, ..., \overline{x_F}]$, with F being the total number of process parameters. To enhance the reliability of the tracing results, the feature vector of the current anomaly is compared with the feature vectors of cases of the same anomaly type in the case database[26].

Let the feature vector of the current anomaly be $f_{cur} \in \mathrm{P}^{\mathcal{F}}$, and the feature vector of a historical case c be $f_c \in \mathrm{P}^{\mathcal{F}}$. Cosine similarity is used to measure the similarity between them, as shown in Eq. (8):

$$sim(c) = \frac{f_{cur} \cdot f_c}{\|f_{cur}\| \, \|f_c\|} \tag{8}$$

Where $f_{cur} \cdot f_c$ denotes the dot product of the two vectors, and $\|f_{cur}\|$ and $\|f_c\|$ represent the L2 norms of the vectors. A value closer to 1 indicates greater similarity between the two vectors.

For each anomaly type, the top k cases with the highest similarity are selected, and their corresponding root cause paths serve as supplementary references for candidate paths. If a path P obtained through path reasoning is consistent with the node sequence of a path P_c in the case database, the confidence of the path is adjusted as shown in Eq. (9).

$$conf_{new}(P) = conf(P) \times (1 + \beta \cdot sim(c)) \tag{9}$$

Where β is an adjustment factor used to control the influence of case similarity on the confidence, and $sim(c)$ is the similarity between the current case and the matched case. This adjustment allows paths that closely align with historical experience to achieve higher confidence.

4.3 Root Cause Chain Generation

By combining the results of causal path reasoning and case similarity matching, the root cause chain is generated. Algorithm 1 presents the specific procedure for root cause tracing.

Algorithm 1. Root Cause Tracing Algorithm

Input:

Anomaly node a, knowledge graph G, case base C, current feature vector f_{cur}, maximum depth D.

Output:

List of root cause chains R (each chain contains node sequence and confidence)

Initialize queue $Q \leftarrow \{(a, [a], 1.0)\}$, result set $R \leftarrow \varnothing$
while Q is not empty do
$(u, path, conf) \leftarrow pop(Q)$
if $len(path) > D$ then continue
for each incoming edge $(v \xrightarrow{r} u)$ in G do
If r is causeBy or affects then
$new\_path \leftarrow [v] + \mathrm{path}$
$new\_conf \leftarrow conf \times w_{(v,u)}$
If v has no incoming edge (i.e., root cause candidate) then
$R \leftarrow R \cup \{(new\_path, new\_conf)\}$
else
Push (Q, (v, new_path, new_conf))
end for
end while
for each path (path, conf) in R do
Compute case similarity $sim_{\max} = \max_{c \in C,\ a_c = path[0]} sim(c)$
$conf \leftarrow conf \times (1 + \beta \cdot sim_{\max})$

18. end for
19. Sort R in descending order of confidence and return R

The algorithm employs a breadth-first reverse search to traverse causal paths within the knowledge graph, adjusts path confidence by incorporating case similarity, and ultimately outputs a list of root cause chains ranked by confidence. The algorithm integrates the predefined causal structure of the knowledge graph with empirical knowledge from historical cases, ensuring that the generated root cause chains not only align with process mechanisms but also benefit from past experiences under similar operating conditions. This provides a quantitative basis for subsequent interpretable analysis and on-site decision-making[27].

4.4 Experimental Results and Analysis

To validate the effectiveness of the proposed root cause tracing method, 120 abnormal events from the test set (20 cases each of Hanging, Slip, Cool Furnace, Hot Furnace, Channeling, and Hearth Buildup) were selected for manual verification. Domain experts evaluated the root cause chains generated by the algorithm. If the key cause identified by the algorithm was consistent with the expert analysis, the trace was considered correct. The root cause tracing accuracy was defined as the ratio of correctly traced samples to the total number of samples. Table 5 presents the tracing results for each type of anomaly.

Table 5 Root cause tracing results for different types of anomalies

Anomaly Type	Test Samples	Correctly Traced	Accuracy (%)
Hanging	20	19	95.0
Slip	20	18	90.0
Cool Furnace	20	17	85.0
Hot Furnace	20	18	90.0
Channeling	20	18	90.0
Hearth Buildup	20	19	95.0
Total	120	110	91.7

As shown in Table 5, the overall root cause tracing accuracy reaches 91.7%, with the highest accuracy observed for hanging and hearth buildup, followed by slipping, furnace heating, and channeling, while furnace cooling exhibits a slightly lower accuracy. The higher accuracy for hanging, hearth buildup, and slipping can be attributed to their well-defined causal paths and the relatively complete depiction of their associated relationships in the knowledge graph. In contrast, furnace cooling and furnace heating have relatively lower accuracy due to their similar characteristics and frequent mutual transitions, which introduce ambiguity in the root cause chains for some samples. Nevertheless, the accuracy remains within an acceptable range.

The six types of anomalies described above are not entirely independent but exhibit certain interrelationships. For example, when hanging causes a sharp decline in permeability, the gas flow may pass through regions with lower resistance, potentially inducing channeling. Similarly, slip may disrupt burden distribution, leading to localized heat deficiency and subsequently causing a cool furnace. Such interrelationships increase diagnostic complexity. However, the knowledge graph constructed in this work defines causal edges between anomaly types, and the graph encoder captures cross-anomaly influence propagation through the attention mechanism, enabling the model to maintain high diagnostic accuracy under actual operating conditions.

Fig. 9 illustrates a typical root cause chain for the hanging anomaly. As shown in the figure, the root causes of hanging can be traced back to multiple upstream factors, such as improper operation, increased fines in raw materials and fuel, and deteriorated coke quality. These factors ultimately lead to hanging through intermediate parameters including burden descent speed and permeability. The confidence scores on the edges reflect the credibility of each causal relationship. This path is highly consistent with blast furnace process mechanisms, validating the interpretability and effectiveness of the root cause chains generated by the proposed method[28].

Factors such as scab formation, increased hot blast volume, and excessive blast pressure may also induce hanging. The modeling of root cause nodes and causal paths in the knowledge graph follows the principle of combining data-driven approaches with expert knowledge, prioritizing factors with higher frequency and clearer causal chains in historical cases. The current knowledge graph primarily incorporates three high-frequency factors: improper operation, increased fines in raw materials and fuel, and deteriorated coke quality. Additional factors can be gradually incorporated in the future through the dynamic update mechanism of the knowledge graph.

Fig. 9 Example of a root cause chain for hanging

The above experimental results indicate that the root cause tracing method proposed in this paper can accurately locate the fundamental causes across various anomaly scenarios, and the generated root cause chains are consistent with process mechanisms, providing interpretable decision-making support for field operators.

Compared with existing methods, the proposed method shows advantages in both diagnostic accuracy and interpretability. Traditional data-driven methods (e.g., SVM, random forest) can achieve anomaly classification but fail to fully utilize causal structural information among parameters, resulting in limited interpretability. Purely data-driven deep learning models (e.g., LSTM, TCN) can capture temporal dynamic features but lack root cause tracing capability. By explicitly modeling causal associations among parameters through the knowledge graph and combining it with the TCN-GAT dual encoder for temporal and graph structural feature fusion, the proposed method achieves favorable performance in prediction accuracy (>90%) and tracing accuracy (91.7%) while providing interpretable root cause reasoning.

5 System Development and Industrial Application

5.1 System Architecture Design

Based on the aforementioned furnace condition prediction and root cause tracing methods, this paper designs and develops an intelligent blast furnace condition diagnosis system, which has been deployed and applied in an actual industrial production environment[29]. The system adopts a layered modular architecture, consisting of a data layer, a knowledge graph layer, a prediction and tracing layer, and an application layer, with data interaction between layers facilitated through standard interfaces. The overall system architecture is shown in Fig. 10.

Fig. 10 System architecture diagram

The data layer is responsible for the acquisition and preprocessing of multi-source data, generating standardized time-series data through cleaning, alignment, missing value imputation, outlier detection and correction, and normalization. Missing values are handled using forward filling or backward filling strategies; outliers are detected and corrected based on the interquartile range (IQR) method; normalization is performed using Min-Max normalization to transform all parameters to a unified scale. The knowledge graph layer constructs a blast furnace knowledge graph based on the ontology, extracts entities and relationships from unstructured text, and stores process parameters, anomaly types, cases, and their causal associations after expert validation, providing prior knowledge for root cause tracing. The prediction and tracing layer serves as the core algorithmic layer: time-series data are processed by TCN to extract dynamic features, graph data are processed by GAT to learn structural dependencies, and multi-class furnace condition prediction is achieved through gated fusion. The prediction results trigger the root cause tracing algorithm, which performs reverse reasoning along causal paths and integrates case similarity to generate root cause chains, outputting anomaly types, confidence levels, and root cause paths. The application layer provides a visual interactive interface, integrating functions such as real-time monitoring, prediction display, graph visualization, historical query, and root cause retrieval, converting algorithmic outputs into charts and operational recommendations to assist rapid decision-making.

The above layered architecture achieves the organic integration of data, knowledge, algorithms, and applications. Each layer has clear responsibilities, and the modules are decoupled, facilitating system expansion and maintenance.

5.2 Industrial Application Effectiveness

To evaluate the system's performance in an actual production environment, a four-month pilot trial was conducted on the No.2 blast furnace of a steel enterprise. Fig. 11 shows the main interface of the system, which adopts a modular design and integrates core functions including furnace condition monitoring[30], parameter visualization, event tracing, and knowledge graph query. The top area of the interface displays the current furnace condition[32] (e.g., Normal, Hanging) in real time. The middle section dynamically presents the trends of key process parameters (such as permeability index, hot blast pressure, and molten iron temperature) through line charts. The left panel lists abnormal events in chronological order along with their processing status. The right panel displays the root cause chain and corresponding operational recommendations for each abnormal event. The bottom area provides an interactive visualization window for querying causal paths within the knowledge graph. This interface balances information density with operational convenience, enabling field personnel to quickly grasp the furnace status and locate the root causes of issues.

Fig. 11 Example of system interface

Based on the four-month field trial data, the system's positive impact on key process indicators became increasingly evident, as shown in Table 6. Each indicator improved continuously from the first month to the fourth month of the trial, confirming the system's effectiveness in stabilizing furnace conditions, enhancing energy efficiency, and promoting smooth operation.

Table 6 Key production indicators during trial operation

Indicator	Pre-application (baseline)	Month 1	Month 2	Month 3	Month 4
Hot metal output/tons/day	6230	6275	6330	6365	6385
Fuel ratio/($kg\cdot t^{-1}$)	526	524	523	523	522
Utilization factor/($t\cdot m^{-3}\cdot d$)	2.45	2.49	2.52	2.54	2.56
Gas utilization rate/%	48.2	48.5	48.7	48.8	49.2
Proportion of favorable hearth condition/%	78.45	79.79	83.52	85.11	87.03
Number of furnace fluctuations/times	3	2	2	1	1

Note: All data in the table are daily averages for each month.

As can be seen from Table 6, after the system was put into operation, all key production indicators exhibited a sustained improving trend. Hot metal output increased from 6275t in the first month to 6385t in the fourth month, a cumulative increase of 1.75%, indicating a steady improvement in blast furnace productivity. The utilization factor rose correspondingly from $2.49t\cdot m^{-3}\cdot d^{-1}$ to $2.56t\cdot m^{-3}\cdot d^{-1}$, reflecting a continuous enhancement in smelting intensity. The fuel ratio decreased from $524kg\cdot t^{-1}$ to $522kg\cdot t^{-1}$, a cumulative reduction of $2kg\cdot t^{-1}$, demonstrating effective energy consumption control. The gas utilization rate increased from 48.5% to 49.1%, a cumulative improvement of 0.6 percentage points, indicating a gradual optimization of gas energy efficiency and an improvement in the reducing atmosphere within the furnace. The proportion of favorable hearth condition increased from 79.79%

to 87.03%, a cumulative improvement of 7.24 percentage points—the most significant gain among all indicators—suggesting that hearth activity and thermal stability were notably enhanced. The number of furnace fluctuations remained steadily at two or fewer occurrences per month starting from the second month, decreasing further to one occurrence in the fourth month, representing a reduction of two occurrences compared to the baseline period. This indicates that the system's early warning and intervention capabilities effectively improved the stability of blast furnace operation.

It should be noted that the prediction accuracy exceeding 90% is evaluated on the offline test set, while the improvement of production indicators presented in Table 6 is constrained by multiple factors such as raw material conditions, equipment status, operational adjustment timeliness, and blast furnace thermal inertia. Consequently, the impact is unlikely to be fully reflected in macro indicators within a short period. Nevertheless, the system has demonstrated a positive trend in stabilizing furnace conditions and reducing abnormal fluctuation frequency, indicating its practical application value[32].

In addition, operator feedback indicated that the root cause paths recommended by the system were highly consistent with on-site experience, and the average troubleshooting time was reduced from 50 minutes before application to 15 minutes, significantly shortening the duration of abnormal event handling. During the trial period, the completeness of knowledge coverage was continuously improved. Overall, the system has demonstrated strong effectiveness and practicality in an actual industrial setting, providing robust technical support for intelligent blast furnace operation and maintenance.

6 Conclusions

This paper addresses the challenges of furnace condition prediction and abnormal root cause tracing in the blast furnace ironmaking process by proposing an intelligent diagnosis method that integrates a knowledge graph with a dual-encoder deep learning framework, and develops an industrial application system. The main work and conclusions are as follows:

1) A knowledge graph tailored to the blast furnace process was constructed. By combining process mechanisms with expert experience, key entity and relation types were defined. A natural language processing model was employed to automatically extract knowledge from multi-source texts, and after fusion and expert validation, the knowledge was stored in a graph database, providing a structured knowledge foundation for furnace condition prediction and root cause tracing.

2) A TCN-GAT dual-encoder furnace condition prediction model was proposed. The temporal encoder extracts dynamic evolution features from minute-level data, while the graph encoder learns structural dependencies among parameters. Through gated fusion, the model achieves accurate identification of six types of furnace conditions. Experimental results show that the proposed model achieves high accuracy on the test set, outperforming multiple baseline models.

3) A root cause tracing algorithm based on causal path reasoning and case similarity was designed. Starting from the prediction results, the algorithm performs reverse causal search along the knowledge graph and incorporates historical case similarity matching to generate confidence-weighted root cause chains. Field validation demonstrates that the tracing accuracy is high and the generated paths are highly consistent with process mechanisms, providing interpretable decision support for operators.

4) An intelligent blast furnace condition diagnosis system was developed and deployed in an industrial setting. The system integrates functions such as data acquisition, knowledge management, furnace condition prediction, root cause tracing, and visual interaction. Four months of trial operation data show continuous improvement in key production indicators and significantly enhanced furnace stability, validating the system's practicality and effectiveness.

References:

[1] J. Q. An, Y. P. Guo, X. M. Zhang, S. Du, Y. F. Huang, M. Wu. 高炉炼铁过程智能感知、诊断与控制方法的研究现状与展望 (Research status and prospects of intelligent sensing, diagnosis and control methods for blast furnace ironmaking processes), Metallurgical Industry Automation, 48 (02) (2024) 2-23. https://doi.org/10.3969/j.issn.1000-7059.2024.02.001 (in Chinese)

[2] X. J. Liu, T. S. Li, X. Li, Y. F. Duan, H. W. Li, Q. Lyu. 高炉炼铁智能化发展的探讨与展望 (Discussion and prospect of intelligent development of blast furnace ironmaking), China Metallurgy, 35 (01) (2025) 1-14+31. https://doi.org/10.13228/j.boyuan.issn1006-9356.20240391 (in Chinese)
[3] L. M. Liu, A. N. Wang, M. Sha, X. Y. Sun, Y. L. Li. Optional SVM for fault diagnosis of blast furnace with imbalanced data, ISIJ international, 51 (9) (2011) 1474-1479. https://doi.org/10.2355/isijinternational.51.1474
[4] S. W. Lou, X. J. Zhang, Y. L. Yang, Y. Li, Y. C. Zhao, C. J. Yang. 高炉炼铁过程故障检测与诊断综述:回顾,现状与展望 (Review of Fault Detection and Diagnosis Research in Blast Furnace Ironmaking Process : Retrospective), Status, and Prospects, Acta Automatica Sinica, 51 (08) (2025) 1739-1759. https://doi.org/10.16383/j.aas.c240660 (in Chinese)
[5] F. M. Li, C. H. Li, S. Liu, X. J. Liu, H. Xiao, J. Zhao, Q. Lyu. Diagnostic model for abnormal furnace conditions in blast furnace based on friendly adversarial training, Journal of Iron and Steel Research International, 32 (6) (2025) 1-14. https://doi.org/10.1007/s42243-024-01361-9
[6] Y. H. Liu, Y. Zhou, Y. F. Liu, X. Xu, Y. X. He. Intelligent Fault Diagnosis for CNC Through the Integration of Large Language Models and Domain Knowledge Graphs, Engineering, 53 (2025) 311-322. https://doi.org/10.1016/J.ENG.2025.04.003
[7] Y. J. Chen, M. Zhou, M. Z. Zhang, M. Zhao. Knowledge-Graph-Driven Fault Diagnosis Methods for Intelligent Production Lines, Sensors, 25 (13) (2025) 3912-3912. https://doi.org/10.3390/S25133912
[8] X. K. Huang, C. J. Yang, H. W. Zhang, et S. W. Lou, D. L. Gao, L. Y. Kong. Data and knowledge collaborative-driven fault identification and self-healing control action inference framework for blast furnace, Expert Systems with Application, 245 (2024) 123040-123040. https://doi.org/10.1016/j.eswa.2023.123040
[9] S. P. Yu, X. Li, Y. G. Lei, B. Yang, N. P. Li, K. Feng. Multimodal Data-enabled Large Model for Machine Fault Diagnosis Towards Intelligent Operation and Maintenance, Journal of Industrial Information Integration, 50 (2026) 101061-101061. https://doi.org/10.1016/j.jii.2026.101061
[10] G. J. Han, W. Chen, X. F. Zhang, A. Liu, L. Zhao. When Temporal Knowledge Graph meets multi-modality: A new perspective for Temporal Knowledge Graph Forecasting, Information Processing and Management, 63 (2PB) (2026) 104463-104463. https://doi.org/10.1016/J.IPM.2025.104463
[11] S. F. Chen, X. J. Liu, H. Y. Li, X. P. Bu, Q. Lyu, F. L. Liu. 高炉炼铁数据缺失处理研究初探 (Preliminary study on missing data processing of blast furnace ironmaking), China Metallurgy, 31 (02) (2021) 17. https://doi.org/10.13228/j.boyuan.issn1006-9356.20200358 (in Chinese)
[12] X. K. Huang, C. J. Yang, Y. Y. Zhang, S. W. Lou, L. Y. Kong, H. Zhou. Ontology guided multi-level knowledge graph construction and its applications in blast furnace ironmaking process, Advanced Engineering Informatics, 62 (PD) (2024) 102927-102927. https://doi.org/10.1016/J.AEI.2024.102927
[13] S. W. Lou, C. J. Yang, P. Wu, Y. L. Yang, L. Y. Kong, X. J. Zhang. Data-driven joint fault diagnosis based on RMK-ASSA and DBSKNet for blast furnace iron-making process, IEEE Transactions on Automation Science and Engineering, 21 (3) (2023) 3826-3841. https://doi.org/10.1109/TASE.2023.3287578
[14] J. N. Bian, Z. H. Mao, B. Jiang, Y. J. Ma, W. J. Liu. 基于知识图谱和多任务学习的工业生产关键设备故障诊断方法 (Fault diagnosis method of critical industrial equipment based on knowledge graphs and multi-task learning), Scientia Sinica (Informationis), 53 (04) (2023) 699-714. https://doi.org/10.1360/SSI-2022-0060 (in Chinese)
[15] X. R. Zhu, Z. X. Li, X. D. Wang, X. Y. Jiang, P. L. Sun, X. W. Wang. Multi-modal knowledge graph construction and application: A survey, IEEE Transactions on Knowledge and Data Engineering, 36 (2) (2022) 715-735. https://doi.org/10.1109/TKDE.2022.3224228
[16] J. Shobana, M. K. UmaMaheswari. ALBERT-BiLSTM-CRF with pointer network based model for named entity recognition, Multimedia Tools and Applications, 84 (39) (2025) 1-23. https://doi.org/10.1007/S11042-025-20970-4
[17] C. J. Qi, C. M. Wang, L. Z. Shao, D. M. Fu, K. Zhou, Z. Y. Zhao. 面向少样本的知识与数据的跨模态特征融合模型 (Integrating knowledge and data: a cross-modal feature fusion model for few-shot problems), Chinese

Journal of Engineering, 47 (08) (2025) 1662-1671. https://doi.org/10.13374/j.issn2095-9389.2024.10.30.004 (in Chinese)
[18] C. B. Li, H. W. Zhai, W. Wu, K. Dong, X. F. Zhang. 基于TCN-GAT与混合神经网络的汽车涂装烘干系统能耗异常检测 (Energy Consumption Anomaly Detection of Automobile Painting Drying System Based on TCN-GAT and Hybrid Neural Network), China Mechanical Engineering, 36 (08) (2025) 1864-1874. https://doi.org/10.3969/j.issn.1004-132X.2025.08.021 (in Chinese)
[19] Y. C. Ni, Q. Y. Jin, R. Hu. A Novel Unsupervised Structural Damage Detection Method Based on TCN-GAT Autoencoder, Sensors, 25 (21) (20205) 6724-6724. https://doi.org/10.3390/S25216724
[20] Q. Y. Zheng, J. Y. Zheng, F. Mei, A. Gao, X. Zhang,Y. Xie. TCN-GAT multivariate load forecasting model based on SHAP value selection strategy in integrated energy system, Frontiers in Energy Research, 11 (2023) 1208502-1208502. https://doi.org/10.3389/FENRG.2023.1208502
[21] D. Gad, S. Abdelfattah, G. Abdelrahman. Temporal graph memory networks for knowledge tracing, Knowledge and Information Systems, 68 (1) (2026) 112-112. https://doi.org/10.1007/S10115-026-02738-W
[22] Y. H. Han, Q. Li, C. Wang, Q. Zhao. A novel knowledge enhanced graph neural networks for fault diagnosis with application to blast furnace process safety, Process Safety and Environmental Protection, 166 (2022) 143-157. https://doi.org/10.1016/J.PSEP.2022.08.014
[23] J. Martinez-Gil, G. Buchgeher, D. Gabauer, B. Freudenthaler, D. Filipiak, A. Fensel. Root cause analysis in the industrial domain using knowledge graphs: a case study on power transformers, Procedia Computer Science, 200 (2022) 944-953. https://doi.org/10.1016/j.procs.2022.01.292
[24] J. Dong, K. R. Cao, K. X. Peng. Hierarchical causal graph-based fault root cause diagnosis and propagation path identification for complex industrial process monitoring, IEEE Transactions on Instrumentation and Measurement, 72 (2023) 1-11. https://doi.org/10.1109/TIM.2023.3268464
[25] K. Li, G. Y. Zhang, F. Han. 基于多模态谱学特征的古代黑曜石指纹识别与矿源溯源方法研究 (A Study on Fingerprint Recognition and Provenance Tracing of Ancient Obsidian Based on Multi-Modal Spectroscopic Features), Experimental Science and Technology, 24 (2025) 1-8. https://doi.org/10.12179/1672-4550.20250391 (in Chinese)
[26] L. G. Jia, B. L. Shao, C. Yang, G. Q. Bian. A Review of Research on Information Traceability Based on Blockchain Technology, Electronics, 13 (20) (2024) 4140-4140. https://doi.org/10.3390/ELECTRONICS13204140
[27] Z. H. Mao, H. Wang, B. Jiang, J. Xu, H. F. Guo. Graph convolutional neural network for intelligent fault diagnosis of machines via knowledge graph, IEEE Transactions on Industrial Informatics, 20 (5) (2024) 7862-7870. https://doi.org/10.1109/TII.2024.3367010
[28] C. S. Kim, H. B. Kim, J. M. Lee. Self-explanatory fault diagnosis framework for industrial processes using graph attention, IEEE Transactions on Industrial Informatics, 21 (4) (2025) 3396-3405. https://doi.org/10.1109/TII.2025.3526708
[29] R. Liu, W. G. Zhao, S. Liu, X. J. Liu, X. Li, Z. F. Zhang, Q. Lyu. Intelligent Prediction and Real-time Monitoring System for Gas Flow Distribution at the Top of Blast Furnace:Regular Article, ISIJ International, 63 (10) (2023) 1714-1726. https://doi.org/10.2355/ISIJINTERNATIONAL.ISIJINT-2023-099
[30] F. M. Li, L. R. Meng, X. J. Liu, X. Li, H. Y. Li, J. J. Mi. Blast Furnace Hanging Diagnosis Model Based On ReliefF-Decision Tree: Regular Article, ISIJ International, 64 (1) (2023) 96-104. https://doi.org/10.2355/ISIJINTERNATIONAL.ISIJINT-2023-350
[31] J. L. Zhu, W. H. Gui, Z. H. Jiang, Z. P. Chen, Y. J. Fang. 基于料面视频图像分析的高炉异常状态智能感知与识别 (Intelligent Perception and Recognition of Blast Furnace Anomalies viaBurden Surface Video Image Analysis), Acta Automatica Sinica, 50 (7) (2024) 1345-1362. https://doi.org/10.16383/j.aas.c230674
[32] J. L. Zheng, P. Wu, R. R. Zuo, X. Su, Y. Z. Liu, N. Kandel. TBDDQN: Imbalanced Fault Diagnosis for Blast Furnace Ironmaking Process via Transformer-BiLSTM Double Deep Q-Networks, Machines, 14 (3) (2026) 276. https://doi.org/10.3390/machines14030276

Acknowledgement

Thanks are given to the financial supports from the Yanzhao Iron and Steel Laboratory Regional Innovation Ability Promotion Project [grant number YZISL2024013, YZISL2024037]; Jiangxi Key Research and Development Program Project [grant number 20261BCE310008]; and the Hebei Provincial Natural Science Foundation Project [grant number E202409101, E202509007,E202509119].

Author Contributions

All authors participated in the conception and design of the study. Material preparation, data collection, and analysis were performed by Xiaojie Liu, Yifan Duan, and Haotian Xu. The first draft of the manuscript was written by Yanping Geng, and all authors commented on previous versions. All authors read and approved the final manuscript.

Data availability

The data used to support the findings of this study are available from the corresponding author upon request.

Conflicts of interests

The authors declare that they have no conflicts of interest.

Caption List

Table 1 Comparison of performance among different knowledge extraction models (%)
Table 2 Examples of knowledge extraction results
Table 3 Quantitative criteria for each anomaly type
Table 4 Results of ablation experiments
Table 5 Root cause tracing results for different types of anomalies
Table 6 Key production indicators during trial operation
Fig. 1 Ontology schema of the blast furnace condition knowledge graph.
Fig. 2 Architecture of the ALBERT-BiLSTM-CRF knowledge extraction model.
Fig. 3 Architecture of the ALBERT-BiLSTM-CRF knowledge extraction model.
Fig. 4 Overall architecture of the proposed model.
Fig. 5 Schematic diagram of the TCN residual block structure.
Fig. 6 Schematic diagram of the graph attention mechanism computation.
Fig. 7 Performance comparison of different models.
Fig. 8 Visual analysis of confusion matrix.
Fig. 9 Example of a root cause chain for hanging.
Fig. 10 System architecture diagram.
Fig. 11 Example of system interface.