INDERSCIENCE PUBLISHERS
Linking academia, business and industry through research

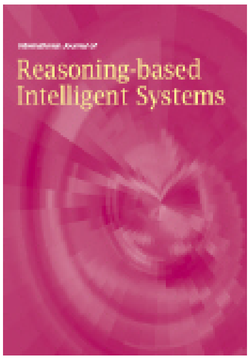

International Journal of Reasoning-based Intelligent Systems

ISSN online: 1755-0564 - ISSN print: 1755-0556
https://www.inderscience.com/ijris

Knowledge graph construction and GCN prediction model for tort liability elements in the Civil Code

Binjing Li

DOI: 10.1504/IJRIS.2025.10075258

Article History:

Received:	13 October 2025
Last revised:	05 November 2025
Accepted:	05 November 2025
Published online:	28 January 2026

Knowledge graph construction and GCN prediction model for tort liability elements in the Civil Code

Binjing Li

Department of Law and Literature,
Harbin Finance University,
Heilongjiang, 150030, China
Email: liyixuan0318@163.com

Abstract: This paper addresses the need for intelligent analysis of tort liability elements under the Civil Code by proposing a legal judgement prediction method that integrates knowledge graphs with graph neural networks. By constructing a knowledge graph of tort liability elements, this study proposes an element alignment method combining Laplace coding with attention mechanisms to precisely link factual circumstances with legal elements. Building upon this foundation, an end-to-end multi-task graph convolutional network prediction model was designed to simultaneously perform liability determination and identification of specific element statuses. Experiments on public datasets such as CAIL2018-Small demonstrate that this method achieves an accuracy of 89.7% and a Macro-F1 score of 88.5%, significantly enhancing both predictive performance and interpretability. This research provides a reliable technical pathway for intelligent judicial assistance systems and holds positive implications for advancing judicial intelligence.

Keywords: knowledge graph; GCN; tort liability elements; the Civil Code.

Reference to this paper should be made as follows: Li, B. (2026) 'Knowledge graph construction and GCN prediction model for tort liability elements in the Civil Code', *Int. J. Reasoning-based Intelligent Systems*, Vol. 18, No. 7, pp.1–11.

Biographical notes: Binjing Li is a Lecturer in the Department of Law and Literature at Harbin Finance University, China. She obtained her Bachelor's degree (2006) and Master's degree (2009) from the Heilongjiang University, China. She has published over ten papers. Her research interests are included civil law, economic law and knowledge graph.

1 Introduction

With the deep integration of artificial intelligence technology across various fields, legal AI is undergoing a paradigm shift from traditional knowledge engineering approaches based on logical rules to data-driven intelligent analysis methods. Against this backdrop, legal judgement prediction – a core task of legal intelligence – aims to automatically forecast case outcomes through computational models. This provides judicial assistance to judges, litigation strategy analysis to attorneys and predictable judicial guidance to the public, making its research both theoretically significant and practically valuable (Gao et al., 2024; Branting et al., 2021).

The promulgation and implementation of China's Civil Code have provided a more systematic and standardised framework for the field of tort liability (Li et al., 2024). However, the determination of tort liability itself constitutes a highly complex legal reasoning process, whose core lies in the rigorous examination and judgement of a series of constituent elements. These elements do not exist in isolation but are intertwined with strict logical connections and complex evidentiary dependencies. While existing traditional machine learning approaches can learn statistical patterns from case data, their model designs struggle to effectively capture and leverage the inherent structural relationships among legal elements (Bi et al., 2024; Tang et al., 2024). Meanwhile, the recent rise of deep learning-based pure text models, while excelling in text feature extraction, suffers from a black-box nature that renders decision-making processes non-traceable (Zhu et al., 2022). This prevents the clear presentation of the complete chain of reasoning – from case facts to legal elements to the final liability determination – which to some extent limits their in-depth application in judicial practice. In the digital age, infringement manifests new characteristics such as mass micro-infringement and mass aggregated infringement, posing significant challenges to traditional tort law in terms of damage assessment, fault determination, and causation analysis (Wu et al., 2024).

Current legal AI research mainly follows two paths: supervised classification models like BERT for legal prediction tasks, and large language models for legal reasoning (Yang et al., 2024). However, both mainstream approaches exhibit significant limitations when analysing the elements of tort liability. While supervision-based classification models can capture local textual features, they struggle to effectively model the complex topological relationships between legal elements (Benedetto et al., 2025). Large language models, meanwhile, often perform poorly on legal prediction tasks due to the abstract nature and unique characteristics of legal provisions (Cui et al., 2023; Shang, 2022). More critically, existing approaches generally overlook the deep semantic correspondence between case facts and legal elements, resulting in significant shortcomings in logical reasoning about legal requirements (He et al., 2023). Furthermore, the pervasive 'black box problem' in judicial AI remains unresolved, with the opacity of decision-making processes severely undermining its credibility and practical value in judicial practice (Tong et al., 2024).

To overcome the aforementioned limitations, this study adopts a technical approach combining knowledge graphs with graph convolutional networks. Taking the constructed fusion graph as input, the model performs information propagation and feature learning on the essential element correlation network through graph convolution operations. This design not only enhances the prediction performance of the primary task but, more importantly, enables the model to attribute the final liability determination to the fulfilment status of specific elements. Consequently, it significantly improves the model's interpretability and judicial practicality.

The main innovations and contributions of this work include:

1 The first fine-grained knowledge graph targeting the elements of tort liability under the Civil Code has been constructed. This graph systematically models multiple entities including elements, provisions, evidence, and defences along with their intricate logical relationships. It transforms legal norms into structured knowledge that is machine-understandable and inferable, laying a solid foundation for in-depth legal analysis. This study represents the first fine-grained knowledge graph dedicated to the elements of tort liability under the Civil Code, achieving sophisticated modelling of the complex logical relationships among these elements.

2 A case-element alignment method integrating Laplace position encoding with attention mechanisms is proposed. This approach precisely anchors case facts and legal elements through semantic attention, while leveraging Laplace encoding to embed global structural information of elements within the graph. It effectively bridges the semantic gap between case texts and structured knowledge, achieving precise mapping from facts to legal principles. This method pioneers the integration of Laplacian global structural encoding with attention-based local semantic alignment, resolving the mapping challenge between case texts and legal elements at the level of information representation.

3 An end-to-end GCN multi-task prediction model is designed. This model performs inference on a constructed fusion graph, simultaneously outputting both the final conclusion on whether tort liability is established and the fulfilment status of each specific element. This design not only enhances prediction performance but also significantly improves model interpretability by attributing decisions to specific elements, aligning with judicial practice's demand for transparent decision-making. This model coupled liability prediction with element state prediction through end-to-end multi-task learning. Its innovation lies in designing an intrinsically interpretable inference mechanism.

2 Relevant technologies

Knowledge graphs, as a technology that represents and stores knowledge through graph structures, are fundamentally designed to depict real-world concepts and their intrinsic relationships through entities, relationships, and attributes (Abu-Salih and Alotaibi, 2024; Zhang et al., 2024). In the legal domain, which demands extreme rigor and logical precision, constructing knowledge graphs presents unique challenges. Its core component – ontology design – requires not only identifying and classifying legal concept entities but also precisely depicting their hierarchically structured, logically coherent relationships. For instance, in tort law, fault-based liability and strict liability form mutually exclusive categories, with fault-based liability contingent upon fulfilling prerequisites such as unlawful acts and actual damages (Zhong et al., 2024). This precise, machine-readable logical constraint distinguishes legal knowledge graphs from general-purpose ones, laying a solid foundation for subsequent automated reasoning.

GCNs are deep learning models specifically designed to process graph-structured data (Bin, 2022). Its core concept originates from traditional convolutional neural networks, aiming to extend the applicability of convolutional operations from regular Euclidean spaces like images to non-Euclidean domains such as graph data structures (Liu et al., 2024; Yang et al., 2023). The fundamental operation of GCN follows a message-passing framework that learns node representations by aggregating neighbourhood information. Within this framework, each node receives information from its immediate neighbours and iteratively updates its own state accordingly. This characteristic makes GCN highly suitable for legal knowledge graphs. In such graphs, the semantics of a legal element do not exist in isolation but are jointly shaped by its own attributes and the associated elements it connects with.

To achieve neighbourhood aggregation, GCN employs an efficient layer-wise propagation rule based on spectral graph theory. Their core inter-layer propagation mechanism can be expressed by the following formula:

$$\mathbf{H}^{(l+1)} = \sigma\left(\hat{\mathbf{D}}^{-\frac{1}{2}}\hat{\mathbf{A}}\hat{\mathbf{D}}^{-\frac{1}{2}}\mathbf{H}^{(l)}\mathbf{W}^{(l)}\right) \tag{1}$$

where $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ denotes the adjacency matrix augmented with self-connections, precisely encoding the topological connections between nodes in the graph. In the legal knowledge graph constructed for this study, each non-zero element of matrix $\mathbf{A}$ specifically represents logical relationships between legal elements $\hat{\mathbf{D}}$ is the corresponding degree matrix, used to normalise the adjacency matrix for numerical stability. Using $\hat{\mathbf{D}}^{-\frac{1}{2}}\hat{\mathbf{A}}\hat{\mathbf{D}}^{-\frac{1}{2}}$ to perform symmetric normalisation on the adjacency matrix is a crucial step aimed at mitigating numerical instability issues that may arise from uneven node degree distributions in the graph, ensuring a stable training process. $\mathbf{H}^{(l)}$ is the node feature matrix for layer l. $\mathbf{W}^{(l)}$ is the trainable weight matrix for that layer, σ is the nonlinear activation function.

The core of the attention mechanism lies in its ability to enable models to focus differentially on different parts of the input information based on their importance. This is achieved by assigning dynamically computed weights to input elements, where higher weights indicate greater criticality for the current task. Unlike fixed weights in traditional pooling operations, these weights are dynamically generated for each input data point.

The core operation of this mechanism is commonly referred to as scaled dot-product attention. Their fundamental computational process can be expressed as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{QK}^{\top}}{\sqrt{d_k}}\right)\mathbf{V} \tag{2}$$

This formula can be decomposed into three steps: first, by computing the dot product between the query matrix $\mathbf{Q}$ and the key matrix $\mathbf{K}$, a similarity score matrix is obtained. This matrix reflects the matching degree between each query and all keys. Second, these scores are scaled by dividing them by the square root of the key vector's dimension $\sqrt{d_k}$. This prevents excessively large dot product results from causing the softmax function to get stuck in a region of minimal gradients. Subsequently, the scaled score matrix undergoes softmax transformation, normalising it into a probability distribution – the attention weights. Finally, these weights are applied to the value matrix $\mathbf{V}$, performing a weighted sum to yield the final output. Each position in the output represents a weighted combination of all value vectors, with the weight determined by the relevance of the query at that position to all keys.

In this study, the attention mechanism serves as a 'semantic bridge'. In the case-element alignment task, case fact texts serve as queries, while element nodes in the knowledge graph function as keys and values. Through attention calculations, the model automatically computes the association strength between each semantic unit in the case description and corresponding legal elements. This achieves precise, soft alignment from unstructured text to structured knowledge – far more flexible and robust than rule-based hard matching.

Knowledge graph technology provides a structured framework for precisely expressing complex logical relationships in the legal domain. Graph convolutional networks empower models to perform deep feature learning and relational reasoning within this framework, while attention mechanisms effectively bridge the gap between unstructured text and structured knowledge. The organic integration of these three components collectively forms the foundational technological pillars of this research methodology, providing the theoretical basis for the implementation of specific models in subsequent sections.

3 Knowledge graph construction

The ontology design of a knowledge graph constitutes the soul and framework of its knowledge system. This chapter follows the classic technical approach in knowledge engineering: ontology design, knowledge extraction, graph augmentation, and quality assessment. Ontology design is grounded in legal ontology, aiming to construct a machine-understandable legal conceptual framework. Knowledge extraction employs a hybrid model designed to balance the representational power of deep learning models with the global constraints of sequence labelling models. Graph augmentation incorporates Laplace coding, theoretically motivated to endow graph neural networks with global structural awareness. The quality assessment framework adheres to the classic three dimensions of software and data quality standards to ensure the reliability of the constructed outcomes. By precisely defining core concepts and their interrelationships, it provides a fundamental paradigmatic foundation for the deep structural representation of knowledge in the field of tort liability. Based on a systematic analysis of the legislative spirit and normative framework of the tort liability section of the Civil Code, this study constructs an ontology model comprising five categories of core entities and five dimensions of foundational relationships. This model aims to transcend the limitations of traditional knowledge bases that merely list concepts. Through a refined relational network design, it profoundly reveals and reproduces the complex and rigorous logical structure among legal elements and judicial reasoning pathways. The technical approach to knowledge graph construction is illustrated in Figure 1.

3.1 *Ontology design*

The ontology's entity structure design fully considers the characteristics of knowledge in the field of tort liability, establishing five core entities: legal elements, liability principles, defences, legal provisions, and types of evidence. These entities do not exist in isolation within the knowledge graph but are tightly interconnected through a multi-layered, multi-dimensional, mutually corroborating network of relationships, forming a knowledge ecosystem that is both clearly hierarchical and seamlessly integrated.

Figure 1 Knowledge graph technology roadmap

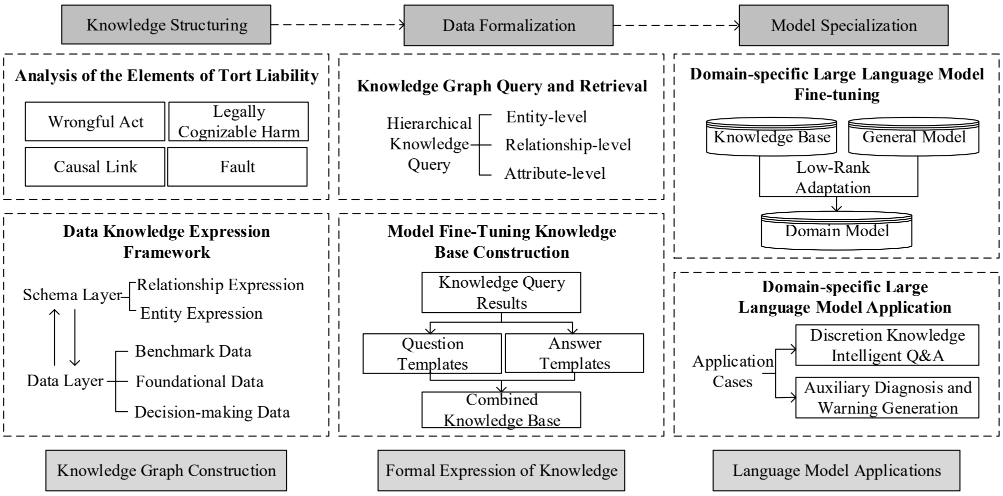

Figure 2 Model architecture diagram (see online version for colours)

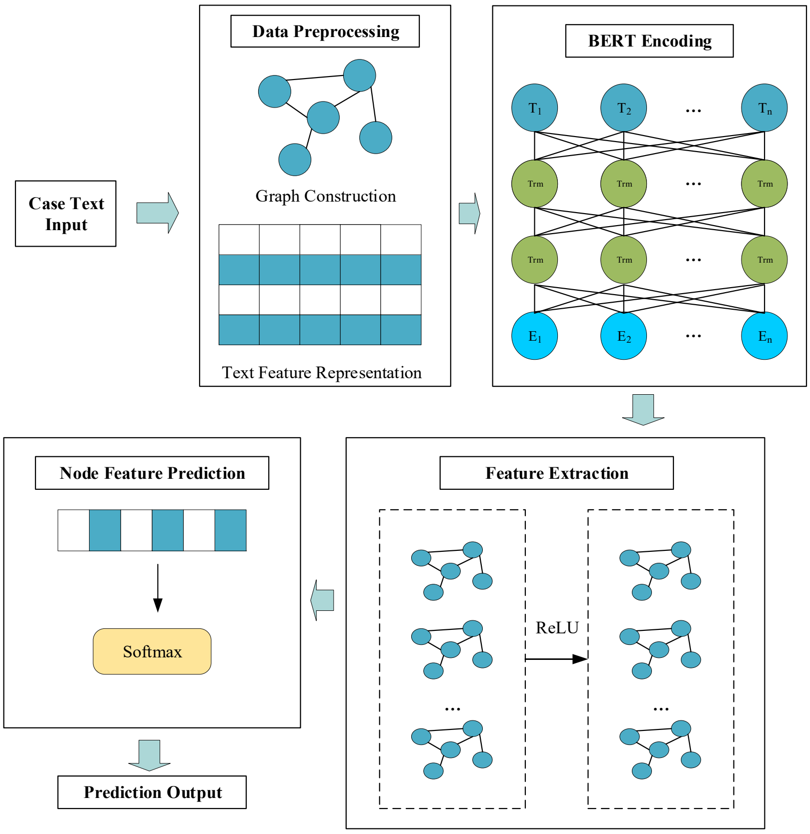

Five relational models collectively form an organic whole that mutually supports and corroborates each other. Hierarchical classification establishes vertical layers of knowledge, logical dependencies construct horizontal chains of reasoning, mutually exclusive negation ensures comprehensive logical completeness, proof associations enable practical application of norms, and legal provision tracing safeguards the system's normative credibility. Working in concert, they systematically recreate the complete cognitive process of tort liability.

3.2 *Knowledge extraction*

After completing the core design, the next critical task is knowledge extraction – specifically, the automatic extraction of entities and relationships from unstructured legal texts. This study employs a hybrid model based on bidirectional encoder representations from transformers (BERT)-bidirectional long short-term memory (BiLSTM)-conditional random field (CRF) for entity recognition. This model fully leverages the powerful representational capabilities of pre-trained language models alongside the advantages of sequence labelling models.

Specifically, this model first utilises BERT to obtain character-level context-aware representations, capturing rich semantic information in legal texts through a multi-layer Transformer encoder. It then employs a BiLSTM network to further capture sequence dependency features, learning long-range dependencies within legal texts. Finally, a CRF layer ensures the global optimality of the label sequence by considering transition constraints between labels, thereby guaranteeing the accuracy and consistency of entity recognition. The loss function for this entity recognition process employs the cross-entropy loss function, whose mathematical expression is as follows:

$$L_{entity} = -\sum_{i=1}^{N}\left[y_i \log\left(\hat{y}_i\right) + \left(1 - y_i\right)\log\left(1 - \hat{y}_i\right)\right] \tag{3}$$

where y_i represents the true label of the i^{th} character, $\hat{y}_i$ denotes the corresponding predicted label probability by the model, N represents the length of the input sequence. This loss function guides the learning process of model parameters by minimising the discrepancy between the predicted probability distribution and the true label distribution. Notably, due to the imbalanced distribution of entities in legal texts, we introduced a category weighting mechanism in practice. This mechanism assigns higher loss weights to entities in minority categories, thereby enhancing the model's ability to recognise low-frequency entities.

For relation extraction tasks, we employ a joint learning approach based on semantic similarity, determining relation types by calculating representation similarity between entity pairs. This method first obtains contextual representations of entity pairs through a pre-trained language model, then computes their vector similarity, and finally classifies them based on similarity scores. The formula for calculating relation similarity is as follows:

$$\operatorname{sim}\left(e_i, e_j\right) = \frac{\mathbf{h}_i \cdot \mathbf{h}_j}{\left\|\mathbf{h}_i\right\| \cdot \left\|\mathbf{h}_j\right\|} \tag{4}$$

where $\mathbf{h}_i$ and $\mathbf{h}_j$ are the representation vectors of entities e_i and e_j respectively, obtained through pre-trained language models. This formula measures the semantic similarity between entity vectors by calculating their cosine similarity with values ranging from –1 to 1. The closer the value is to 1, the greater the semantic similarity between entities and the higher the probability of a relationship existing. This method effectively addresses the challenges of relation extraction in legal texts arising from their high complexity and diversity, demonstrating particular strength in identifying implicit relationships. Additionally, we designed a multi-task learning mechanism enabling entity recognition and relation extraction to share the underlying text. This collaborative optimisation enhances overall performance while improving the model's training efficiency.

3.3 Graph enhancement

This study employs Laplace position encoding to enhance the structural representation of knowledge graphs, thereby serving downstream graph neural network models. The encoding leverages spectral decomposition of the graph to capture global structural information of nodes, generating a unique identity identifier for each node. In graph neural networks, traditional position encoding methods typically capture only local neighbourhood information. In contrast, Laplacian position encoding characterises a node's structural position within the graph from a global perspective, which is crucial for distinguishing core |legal elements from peripheral auxiliary elements. The computation of Laplacian position encoding is based on the eigenvalue decomposition of the graph's Laplacian matrix, with the following mathematical foundation:

$$\Delta = \mathbf{U}^{\top}\Lambda\mathbf{U} \tag{5}$$

where Δ is the Laplacian matrix of the graph, defined as the degree matrix minus the adjacency matrix, reflecting the graph's topological structure; $\mathbf{U}$ is the eigenvector matrix, whose column vectors correspond to the eigenvectors of the Laplacian matrix; Λ is the diagonal matrix of eigenvalues, where the elements on the diagonal represent the eigenvalues. Based on this feature decomposition, we can obtain the positional encoding for each node and obtain the final structure-aware representation through linear transformation:

$$\lambda_i = \mathbf{W}\lambda_i + \mathbf{b} \tag{6}$$

In this transformation, λ_i denotes the original position encoding of node i, $\mathbf{W}$ and $\mathbf{b}$ represent the learnable parameter matrix and bias vector. This encoding enables the model to clearly distinguish the essential differences between core and peripheral elements within the graph topology, providing crucial structural prior information for subsequent graph neural networks. Specifically, core elements such as fault liability – due to their highly central position in the graph and complex logical connections with multiple other elements – will receive coding representations significantly distinct from peripheral elements. The core element occupying a central position in the graph and maintaining complex logical connections with numerous other elements exhibits a distinct positional encoding derived from Laplace feature decomposition, which differs significantly from that of sparsely connected auxiliary elements. This structural awareness is crucial for accurate legal reasoning, as elements of differing structural significance carry varying weight in legal argumentation.

3.4 Quality assessment

To comprehensively validate the construction quality of the knowledge graph, a systematic quality evaluation framework is established. The evaluation framework encompasses three dimensions – completeness, accuracy, and consistency – and employs a combined quantitative and qualitative approach for comprehensive assessment. The selection of this framework is grounded in classical theories of knowledge graph quality assessment. These three dimensions respectively correspond to the fundamental requirements for knowledge graphs in terms of content coverage, factual accuracy, and logical consistency, comprehensively reflecting their reliability across various judicial application scenarios. Completeness evaluation is achieved by comparing automatically extracted entities against manually annotated gold standards, ensuring the knowledge graph covers core concepts within the field of tort liability. The evaluation formula is as follows:

$$\text{Completeness} = \frac{\left|E_{\text{extracted}} \cap E_{\text{manual}}\right|}{\left|E_{\text{manual}}\right|} \tag{7}$$

In this formula, the numerator represents the number of entities where automatically extracted entities intersect with manually annotated entities, while the denominator denotes the total number of manually annotated entities. This metric reflects the knowledge graph's coverage of domain-specific knowledge. During actual evaluation, we invited three legal experts to independently complete manual annotation, using the union of their annotations as the gold standard to ensure comprehensive and authoritative assessment. Accuracy assessment employs a sampling verification method, where domain experts evaluate the correctness of randomly sampled entities and relationships to calculate metrics such as accuracy, precision, and recall. Consistency checks validate the logical coherence of entity relationships within the knowledge graph, particularly scrutinising the logical consistency of 'requires' and 'excludes' relationships to ensure the graph supports reliable logical reasoning.

Through the organic integration of ontology design, knowledge extraction, graph enhancement, and quality assessment, the tort liability elements knowledge graph not only accurately captures the semantic information of legal elements but also deeply depicts the logical relationships and structural characteristics among these elements. This provides high-quality data support for legal prediction research based on graph neural networks.

4 Predictive model for tort liability

Building upon the completion of the tort liability knowledge graph construction, this chapter will elaborate on a tort liability prediction model based on graph convolutional networks. Through an innovative case-element alignment mechanism and integrated graph construction technology, this model organically combines unstructured case descriptions with structured legal knowledge graphs to achieve end-to-end tort liability prediction. The fusion of Laplace coding and attention mechanisms yields an alignment mechanism. Its theoretical foundation lies in enhancing information from two complementary perspectives: structural perception and semantic perception. Laplacian coding captures nodes' global structural roles within the entire graph spectrum from a graph-spectral perspective, enabling the model to distinguish core elements from auxiliary ones. Meanwhile, the attention mechanism dynamically calculates the strength of associations between case facts and legal elements from a semantic perspective. The entire model architecture is shown in Figure 2.

The primary component of the model is feature extraction, which is responsible for learning deep representations of legal elements from the knowledge graph. In traditional graph neural networks, node features often contain only basic semantic information, lacking consideration of the node's structural role within the overall graph. To address this, we introduce a representation learning method that integrates semantic and structural features. This method first utilises Laplacian position encoding to obtain the structural features of nodes. It then employs a multi-head attention mechanism to learn dependencies between nodes. Layer normalisation techniques are applied to stabilise the training process. The complementary advantages of combining GCNs with attention mechanisms. GCNs enable message passing across graph structures, aggregating information from neighbouring nodes to learn feature representations aligned with legal logic networks, yet their receptive fields are constrained by the number of layers. Attention mechanisms dynamically assign weights to different semantic units, capturing long-range dependencies. In this model, GCN handles structural information propagation and inference, while attention mechanisms filter and align key semantics. Working synergistically, they jointly achieve precise reasoning. The mathematical expression for this process is as follows:

$$\mathbf{h}_i^{\text{fuse}} = \text{LayerNorm}\left(\mathbf{h}_i^{\text{lap}} + \text{MultiAttn}\left(\mathbf{h}_i^{\text{lap}}, \mathbf{X}^{\text{lap}}\right)\right) \tag{8}$$

where $\mathbf{h}_i^{\text{lap}}$ denotes the Laplace position encoding for the i^{th} feature, $\mathbf{X}^{\text{lap}}$ represents the feature matrix for all features, MultiAttn denotes the multi-head attention mechanism, and LayerNorm refers to the layer normalisation operation. Under the feature fusion mechanism, the representation of each element integrates its global structural information within the graph and its associative features with other elements, providing a rich representational foundation for downstream tasks.

The core innovation of the model lies in its case element alignment mechanism, designed to address the limitations of traditional legal text analysis that treats case details as a whole while neglecting precise mapping to legal elements. Based on attention methods, we automatically calculate the semantic relevance between case semantic units and legal elements using the following formula:

$$\alpha_{ij} = \text{softmax}\left(\frac{\mathbf{W}_1\mathbf{h}_i^{\text{case}}\left(\mathbf{W}_2\mathbf{h}_j^{\text{element}}\right)^{\top}}{\sqrt{d}}\right) \tag{9}$$

where $\mathbf{h}_i^{\text{case}}$ denotes the vector representation of the i^{th} semantic unit in the case description, obtained through a pre-trained language model; $\mathbf{h}_j^{\text{element}}$ denotes the feature representation of the j^{th} legal element; $\mathbf{W}_1$ and $\mathbf{W}_2$ denote learnable parameter matrices; d represents the vector dimension; α_{ij} denotes the association weight between the i^{th}

semantic unit and the j^{th} legal requirement. This attention mechanism automatically focuses on key information within case descriptions that pertains to specific legal elements, effectively bridging the semantic gap between natural language descriptions and legal concepts.

After obtaining the alignment relationship between case elements and knowledge graph features, the next critical step is to construct a fusion graph and perform graph convolutional inference. The fusion graph construction aims to incorporate the current case being processed as a special node into the existing knowledge graph, forming a new dynamic inference graph. Specifically, we fuse the aligned features of case descriptions with the element features in the knowledge graph through a multi-layer perceptron to generate the initial representation of the case node:

$$\mathbf{h}_i^{\text{fusion}} = \text{MLP}_f\left(\mathbf{h}_i^{\text{element}} \| \mathbf{h}_i^{\text{case}}\right) \tag{10}$$

where $\|$ denotes vector concatenation, and MLP_f represents a multilayer perceptron network. This operation enables case nodes to incorporate both prior knowledge of legal requirements and the specific characteristics of the current case. Building upon this foundation, we construct an extended adjacency matrix containing case nodes and utilise a multi-layer graph convolutional network for information propagation and feature learning.

The feature aggregation process for image convolutions is as follows:

$$\mathbf{X}^{(l+1)} = \sigma\left(\tilde{\mathbf{D}}^{-\frac{1}{2}}\tilde{\mathbf{A}}\tilde{\mathbf{D}}^{-\frac{1}{2}}\mathbf{X}^{(l)}\mathbf{W}^{(l)} + \mathbf{b}^{(l)}\right) \tag{11}$$

where $\tilde{\mathbf{A}}$ is the adjacency matrix for adding self-connections, $\tilde{\mathbf{D}}$ is the corresponding degree matrix, $\mathbf{X}^{(l)}$ is the node feature matrix for layer l, $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ are learnable weight and bias parameters, with σ representing the nonlinear activation function. Through multi-layer graph convolutional operations, the model enables information exchange between case nodes and legal requirement nodes, allowing case representations to progressively integrate legal knowledge while simultaneously adapting legal requirement representations to specific case contexts.

The model simultaneously accomplishes two related tasks – predicting liability for infringement and determining the fulfilment of essential elements – through a multi-task prediction head. This multi-task learning design not only enhances the performance of the primary task but also improves the model's interpretability. For the primary task of predicting whether liability for infringement is established, we perform pooling on the final representations of all nodes and then classify them using a softmax function:

$$\mathcal{L}_{\text{total}} = \lambda_1\mathcal{L}_{\text{pred}} + \lambda_2\mathcal{L}_{\text{element}} + \lambda_3\mathcal{L}_{\text{relation}} \tag{12}$$

where $\mathcal{L}_{\text{pred}}$ $\mathcal{L}_{\text{element}}$ and $\mathcal{L}_{\text{relation}}$ represent the cross-entropy loss for the main task, the loss for the feature recognition task, and the loss for the relationship prediction task, respectively. λ_1 λ_2 and λ_3 represent the weighting coefficients for each task, with optimal values determined through grid search. This multi-task learning mechanism enables the model to learn shared feature representations from multiple supervisory signals. During training, the model is forced to learn the causal relationships between the states of input factors and the final responsibility determination. During inference, decision-makers can view the final conclusion and trace back to the specific states of individual factors that collectively led to this conclusion, transforming the decision into a transparent chain of reasoning. This approach not only enhances the model's generalisation capabilities but also allows predictions to be traced back to the specific conditions under which they were established, significantly improving the model's legal interpretability.

5 Experiments and results analysis

5.1 Dataset construction and selection

Regarding dataset selection, this study employed three representative public legal datasets to ensure the reliability and reproducibility of experimental results. The CAIL2018-Small dataset, sourced from the China AI for justice challenge, encompasses ten categories of criminal charges with a total of 10,000 legal cases. With an average text length of 512 characters, it features authentic cases and diverse criminal charges. The JEC-QA dataset, released by Tsinghua University, comprises 6,000 training samples covering five types of legal questions. Its question-answering format closely mirrors real-world legal consultation scenarios. For legal reasoning benchmarks, this study uses the legal-bench dataset. It comprises 5,000 training samples and offers cross-language testing capabilities. Multiple internal sub-datasets each focus on distinct aspects, enabling comprehensive evaluation of model performance across various legal tasks.

In knowledge graph construction, we systematically built a specialised legal knowledge graph based on the tort liability section of the Civil Code. This graph comprises 1,045 entity nodes and 2,387 relationship edges, with entity types covering multiple core legal concepts. The construction process combined automated extraction with manual verification: entities and relationships were extracted from legal texts, then reviewed and refined by legal experts. This approach not only ensures knowledge quality but also provides a clear, reproducible technical pathway for subsequent research.

5.2 Experimental setup and baseline model selection

The experiments were conducted on a hardware environment comprising Tesla V100s GPUs and Intel Xeon Silver 4216 CPUs, utilising the PyTorch 1.9.0 and PyTorch Geometric 2.0.4 deep learning frameworks. A fixed random seed was set to ensure result consistency. All relevant code and scripts have been open-sourced to guarantee the reproducibility of the experiments.

Figure 3 Performance comparison of primary experimental categories (see online version for colours)

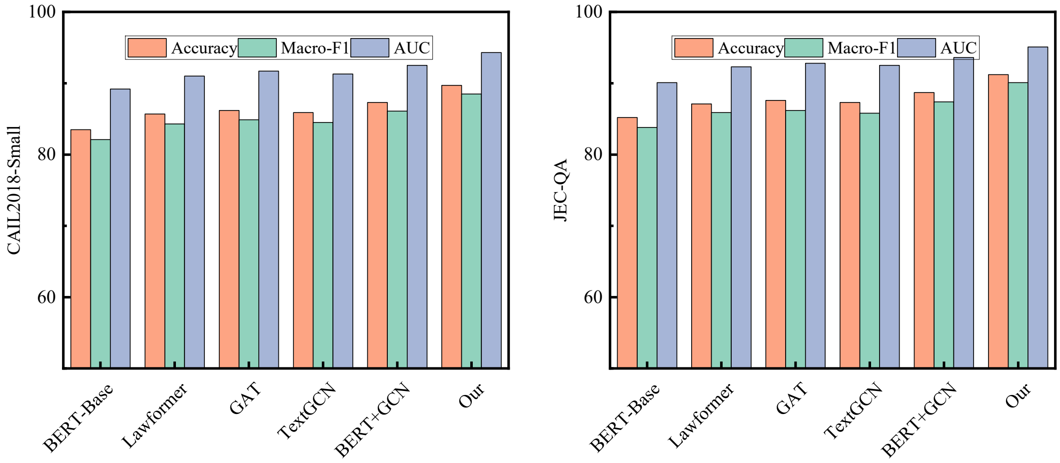

Figure 4 Loss on the CAIL2018-Small dataset (see online version for colours)

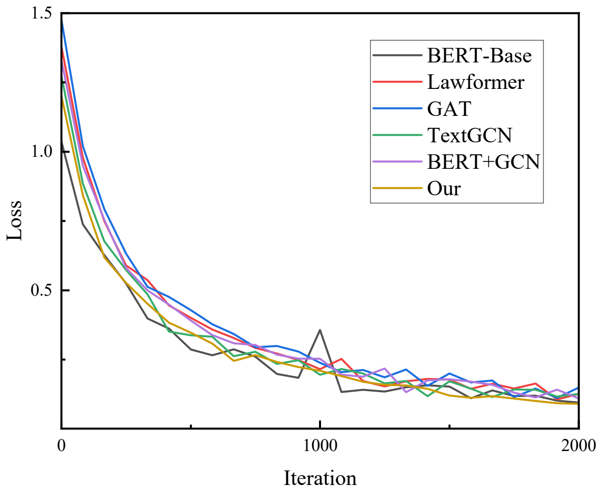

Table 1 Key metrics in the training process

Training stage	*Iteration count*	*Training loss*	*Validation loss*	*Validation accuracy (%)*	*Learning rate*
Initial stage	0–150	1.200→0.672	1.232→0.734	72.3→83.5	0.001
Rapid convergence	150–400	0.672→0.421	0.734→0.523	83.5→87.2	0.001
Stability optimisation	400–1,100	0.421→0.243	0.523→0.367	87.2→89.1	0.0005
Final fine-tuning	1,100–2,000	0.243→0.090	0.367→0.101	89.1→89.7	0.0002

This study incorporates five representative models as a baseline. BERT-Base, as a general-purpose pre-trained model, provides a fundamental performance reference (Liao et al., 2024). Lawformer is a pre-trained model specifically optimised for long-text legal documents, better handling the peculiarities of legal texts (Xiao et al., 2021). GAT and TextGCN represent attention-based graph neural networks and word-document graph-based text classification methods, respectively (Lv et al., 2024; Aras et al., 2024). BERT+GCN validated the effectiveness of a simple fusion strategy (She et al., 2022). The selection of these baseline models represents several mainstream technical approaches in the current field of legal AI: BERT, as a general-purpose pre-trained language model, serves as a performance benchmark for numerous NLP tasks; Lawformer is a domain-specific pre-trained model optimised for lengthy legal texts, embodying the domain adaptation approach; GAT and TextGCN are graph-based models that respectively explore the role of attention mechanisms and document-word graphs in text classification; BERT+GCN demonstrates the effectiveness of simple hybrid strategies. The selection of these baseline models balances both technical representativeness and applicability within the legal domain, providing comprehensive comparative benchmarks for evaluating the model.

Model parameters were meticulously tuned using the AdamW optimiser with an initial learning rate of 0.001, combined with a weight decay of 0.01 to prevent overfitting. The batch size was set to 32, with a maximum training cycle of 100 rounds. Training was terminated via early stopping when validation set performance failed to improve for ten consecutive rounds. Specific parameters of the model include a two-layer GCN network, 256-dimensional hidden layer representations, a dropout rate of 0.5, and an eight-head attention mechanism. All these parameters were determined through grid search to achieve optimal values.

5.3 Experimental results and analysis

To comprehensively evaluate the effectiveness of the proposed model in predicting legal judgements, this chapter designed systematic experiments to validate it from multiple perspectives.

The main experimental results demonstrate that our model achieves optimal performance across two datasets. On the CAIL2018-Small dataset, the model achieves an accuracy of 89.7%, a Macro-F1 score of 88.5%, and an AUC of 94.3%, representing improvements of 6.2%, 6.4%, and 5.1% respectively over the BERT-Base model. From a judicial practice perspective, the 89.7% prediction accuracy demonstrates a high level of reliability. In the early stages of case adjudication, the system provides highly consistent preliminary assessments, effectively assisting judges in identifying logical inconsistencies. As an efficient decision-making support tool, it enhances judicial efficiency and maintains consistency in sentencing standards. This significant improvement is primarily attributed to the introduction of knowledge graphs. Structured legal knowledge provides the model with rich semantic constraints, effectively reducing ambiguity issues in text comprehension. Compared to the specialised legal pre-training model Lawformer, our model still maintains a 4.0% accuracy advantage, indicating that purpose-built legal knowledge graphs capture the specificity of the legal domain more effectively than general-purpose legal pre-training models.

On the JEC-QA dataset, our model achieved an accuracy of 91.2%, representing a significant improvement over baseline models. This outcome demonstrates the effectiveness of the case-element alignment mechanism, which precisely connects factual case descriptions with the legal conceptual framework, thereby enhancing the model's legal reasoning capabilities. Particularly when addressing complex legal issues, the model employs multi-task learning to simultaneously predict both the establishment of liability and the fulfilment of specific elements. This design not only improves performance on the primary task but also enhances the model's interpretability.

Analysis of the training process reveals the model's convergence characteristics on the CAIL2018-Small dataset. The training loss value changes with the number of iterations as shown in Figure 4.

During the first training epochs, the model converged rapidly. This accelerated convergence was driven by the strong semantic guidance provided by the knowledge graph, enabling the model to learn effective feature representations more efficiently. During subsequent training, the model entered a stable optimisation phase. Fine-tuning was achieved through gradual learning rate reduction, ultimately reaching a stable performance of 89.7%. The gap between training and validation losses remained within a reasonable range throughout, indicating strong generalisation capabilities and no significant overfitting. Key metrics in the training process are shown in Table 1.

The ablation experiment conducted an in-depth analysis of each component's contribution to the model's performance, with the results shown in Table 2. On the CAIL2018-Small dataset, removing the knowledge graph component resulted in the most significant decline in model performance, with accuracy dropping by 7.6%. This fully demonstrates the pivotal role of structured legal knowledge in legal judgement prediction. The knowledge graph not only provides rich prior knowledge but also captures complex logical relationships among legal concepts through its entity-relationship network. Removing the case-element alignment mechanism resulted in a 4.5% performance drop. This mechanism achieves semantic alignment between case facts and legal elements through attention calculations, serving as a crucial bridge connecting textual descriptions to structured knowledge. The multi-task learning framework contributed a 2.4% performance boost by simultaneously optimising the primary task and the auxiliary legal element prediction task, providing additional supervisory signals that enhanced the model's generalisation capability. Although Laplace position encoding contributed relatively less at 1.6%, it played a crucial role in distinguishing the structural roles of core elements versus peripheral elements. Knowledge graphs make the most significant contribution because they infuse models with rich structural prior knowledge. Through entity-relationship networks, knowledge graphs explicitly define logical constraints between elements, providing models with inductive biases that transcend textual statistical patterns and align with legal reasoning.

Table 2 Melting experiment results (CAIL2018-Small accuracy/%)

Model	*Accuracy*	*Macro-F1*	*Decrease relative to full model*
Our	89.7	88.5	-
W/o knowledge graph	82.1	80.0	7.6
W/o case – essential alignment	85.2	83.9	4.5
W/o multitask learning	87.3	86.2	2.4
W/o Laplace coding	88.1	87.0	1.6

Cross-lingual experiments validated the model's generalisation capability. On the English Legal-Bench dataset, the model achieved an accuracy of 86.7%, maintaining 96.7% of its performance relative to the Chinese dataset. This result demonstrates that despite linguistic differences, knowledge graph-based legal representations exhibit strong cross-lingual transferability.

Structured legal knowledge reduces reliance on specific linguistic expressions, enabling the model to better adapt to judgement prediction tasks across different legal systems.

Through systematic experimental validation, the model demonstrates significant advantages in legal judgement prediction tasks. The integration of a knowledge graph endows the model with rich legal prior knowledge. The case-element alignment mechanism achieves effective fusion of textual data and knowledge, while the multi-task learning framework enhances the model's generalisation capabilities. The synergistic interaction of these technical components not only renders the model highly valuable for practical applications but also offers novel insights and methodologies for legal AI research. All experimental code and dataset information has been made publicly available, ensuring complete reproducibility of the research findings and laying a solid foundation for subsequent studies.

6 Conclusions

This paper systematically explores a technical approach integrating knowledge graphs and graph neural networks to address the intelligent analysis requirements for tort liability elements under the Civil Code. The constructed knowledge graph of tort liability elements establishes a machine-understandable foundation for legal knowledge. The proposed case element alignment method bridges the semantic gap between textual descriptions and legal concepts. The multi-task prediction model not only ensures accuracy but also enhances the transparency and interpretability of the decision-making process. Experiments demonstrate the method's outstanding performance in prediction tasks, validating the technical approach's effectiveness while providing a practical solution for developing intelligent judicial assistance systems.

Future research can be deepened in multiple dimensions: first, expanding the scope of knowledge graphs to encompass other legal domains to validate the method's generalisation capabilities; second, exploring the integration of multimodal data such as evidentiary documents to enhance comprehensive analysis of complex cases; third, focusing on model efficiency optimisation and lightweight deployment – a critical step for moving beyond the laboratory. The continuous evolution of legal artificial intelligence will undoubtedly contribute more core technological strength to the modernisation of judicial practice.

Declarations

All authors declare that they have no conflicts of interest.

References

Abu-Salih, B. and Alotaibi, S. (2024) 'A systematic literature review of knowledge graph construction and application in education', *Heliyon*, Vol. 10, No. 3, pp.27178–27187.

Aras, A.C., Alikasifoglu, T. and Koç, A. (2024) 'Text-RGNNs: relational modeling for heterogeneous text graphs', *IEEE Signal Processing Letters*, Vol. 31, pp.1955–1959.

Benedetto, I., Koudounas, A., Vaiani, L., Pastor, E., Cagliero, L., Tarasconi, F. and Baralis, E. (2025) 'Boosting court judgment prediction and explanation using legal entities', *Artificial Intelligence and Law*, Vol. 33, No. 3, pp.605–640.

Bi, S., Ali, Z., Wu, T.X. and Qi, G.L. (2024) 'Knowledge-enhanced model with dual-graph interaction for confusing legal charge prediction', *Expert Systems with Applications*, Vol. 249, pp.116–122.

Bin, L. (2022) 'GCN-BERT and memory network based multi-label classification for event text of the Chinese Government hotline', *IEEE Access*, Vol. 10, pp.109267–109276.

Branting, L.K., Pfeifer, C., Brown, B., Ferro, L., Aberdeen, J., Weiss, B., Pfaff, M. and Liao, B. (2021) 'Scalable and explainable legal prediction', *Artificial Intelligence and Law*, Vol. 29, No. 2, pp.213–238.

Cui, J.Y., Shen, X.Y. and Wen, S.C. (2023) 'A survey on legal judgment prediction: datasets, metrics, models and challenges', *IEEE Access*, Vol. 11, pp.102050–102071.

Gao, Y., Xiong, G.Y., Li, H.J. and Richards, J. (2024) 'Exploring bridge maintenance knowledge graph by leveraging GrapshSAGE and text encoding', *Automation in Construction*, Vol. 166, pp.10–17.

He, C.Q., Tan, T.P., Zhang, X.B. and Xue, S. (2023) 'Knowledge-enriched multi-cross attention network for legal judgment prediction', *IEEE Access*, Vol. 11, pp.87571–87582.

Li, Z.F., Jian, Y., Xue, Z.C., Zheng, Y.M., Zhang, M., Zhang, Y., Hou, X.J. and Wang, X.G. (2024) 'Text-enhanced knowledge graph representation learning with local structure', *Information Processing & Management*, Vol. 61, No. 5, pp.2151–2166.

Liao, W.X., Liu, Z.L., Dai, H.X., Wu, Z.H., Zhang, Y.Y., Huang, X.K., Chen, Y.Z., Jiang, X., Liu, D.V., Zhu, D.J., Li, S., Liu, W., Liu, T.M., Li, Q.Z., Cai, H.M. and Li, X. (2024) 'Mask-guided BERT for few-shot text classification', *Neurocomputing*, Vol. 610, p.1.

Liu, Z.G., Li, F.S., Hao, G.Y., He, X.Q. and Zhang, Y.H. (2024) 'GCN-LSTM: multi-label educational emotion prediction based on graph convolutional network and long and short term memory network fusion label correlation in online social networks', *Computer Science and Information Systems*, Vol. 21, No. 4, pp.587–592.

Lv, S.Q., Dong, J.A., Wang, C.C., Wang, X.H. and Bao, Z.Q. (2024) 'RB-GAT: a text classification model based on RoBERTa-BiGRU with graph attention network', *Sensors*, Vol. 24, No. 11, pp.15–22.

Shang, X.R. (2022) 'A computational intelligence model for legal prediction and decision support', *Computational Intelligence and Neuroscience*, Vol. 2022, pp.17531–17540.

She, X.R., Chen, J.P. and Chen, G. (2022) 'Joint learning with BERT-GCN and multi-attention for event text classification and event assignment', *IEEE Access*, Vol. 10, pp.27031–27040.

Tang, M.W., Zhao, S.S., Zhu, Y., Chen, J. and Zhang, D. (2024) 'A semi-automatic annotation system for building a psychological anxiety relief-oriented legal knowledge graph', *Medicine*, Vol. 103, No. 14, pp.4690–4699.

Tong, S.X., Yuan, J.L., Zhang, P.L. and Li, L. (2024) 'Legal judgment prediction via graph boosting with constraints', *Information Processing & Management*, Vol. 61, No. 3, pp.1–18.

Wu, Y.L., Pan, X., Li, J.H., Dou, S.M., Dong, J.H. and Wei, D. (2024) 'Knowledge graph-based hierarchical text semantic representation', *International Journal of Intelligent Systems*, Vol. 2024, p.1.

Xiao, C.J., Hu, X.Y., Liu, Z.Y., Tu, C.C. and Sun, M.S. (2021) 'Lawformer: a pre-trained language model for Chinese legal long documents', *AI Open*, Vol. 2, pp.79–84.

Yang, R., Zhu, J.H., Man, J.P., Fang, L. and Zhou, Y. (2024) 'Enhancing text-based knowledge graph completion with zero-shot large language models: a focus on semantic enhancement', *Knowledge-Based Systems*, Vol. 300, pp.2817–2822.

Yang, Z.P., Ye, X. and Xu, H.T. (2023) 'TGCN-Bert emoji prediction in information systems using TCN and GCN fusing features based on BERT', *International Journal on Semantic Web and Information Systems*, Vol. 19, No. 1, pp.35–48.

Zhang, J.C., Zain, A.M., Zhou, K.Q., Chen, X. and Zhang, R.M. (2024) 'A review of recommender systems based on knowledge graph embedding', *Expert Systems with Applications*, Vol. 250, pp.91–102.

Zhong, L.F., Wu, J., Li, Q., Peng, H. and Wu, X.D. (2024) 'A comprehensive survey on automatic knowledge graph construction', *ACM Computing Surveys*, Vol. 56, No. 4, pp.16240–16244.

Zhu, G.C., Hao, M.H., Zheng, C.L. and Wang, L.L. (2022) 'Design of knowledge graph retrieval system for legal and regulatory framework of multilevel latent semantic indexing', *Computational Intelligence and Neuroscience*, Vol. 2022, pp.116–122.