

# LexRel: Benchmarking Legal Relation Extraction for Chinese Civil Cases

Anonymous ACL submission

## Abstract

Legal relations form a highly consequential analytical framework of civil law system, serving as a crucial foundation for resolving disputes and realizing values of the rule of law in judicial practice. However, legal relations in Chinese civil cases remain underexplored in the field of legal artificial intelligence (legal AI), largely due to the absence of comprehensive schemas. In this work, we firstly introduce a comprehensive schema, which contains a hierarchical taxonomy and definitions of arguments, for AI systems to capture legal relations in Chinese civil cases. Based on this schema, we then formulate legal relation extraction task and present **LexRel**, an expert-annotated benchmark for legal relation extraction in Chinese civil law. We use **LexRel** to evaluate state-of-the-art large language models (LLMs) on legal relation extractions, showing that current LLMs exhibit significant limitations in accurately identifying civil legal relations. Furthermore, we demonstrate that incorporating legal relations information leads to promising performance gains on other downstream legal AI tasks.

## 1 Introduction

Legal relations constitute the foundational analytical unit of the Chinese civil law system. Defined as the relationships between individuals as regulated by legal norms, legal relations serve as the crucial interface between the normative world of law and the empirical world of social facts (Yuan, 2024). In judicial practice, legal professionals frequently rely on legal relations as a key reference for several tasks, such as retrieving legal information, statute prediction and case outcome analysis (Shen, 2004). However, while legal relations are central to both legal theory and practice, they remain underexplored in the field of legal AI, particularly in the context of Chinese civil law.

Legal information extraction in legal AI primarily targets factual entities (e.g., persons, goods, contracts) or general social relations (e.g., employment, ownership) (Cui et al., 2024; Huang et al., 2021; Li et al., 2024b, 2021; Xu, 2021). Such approaches overlook that legal relations—a concept grounded in statutory rules and judicial practice—are qualitatively distinct from ordinary semantic associations in natural language. Moreover, in real-world practice, legal relations are seldom explicitly stated in judgments or other legal documents, hence models trained solely on large-scale legal texts struggle to capture the ability to understand and use legal relations. Furthermore, the current legal relation schema is coarse-grained, broadly classified according to civil rights and obligations (Wan et al., 2022). To enable AI systems to better distinguish and master legal relations, developing a schema of legal relations is therefore essential.

In this paper, we present the first comprehensive schema to model legal relations in Chinese civil law. Building upon established legal theory, real-world judicial practice, and the guidance of legal experts, our schema consists of a hierarchical taxonomy spanning 9 major domains of civil law and 265 relation types, together with precise definitions of their core arguments (subject, object and content).

Based on this schema, we define the task of *legal relation extraction*, which involves extracting legal relations from factual text. Building on this task, we develop **LexRel**, a benchmark consisting of 1,140 expert-annotated legal relations extracted from real-world cases.

We use **LexRel** to evaluate capabilities of state-of-the-art large language models (LLMs) on the task of *legal relation extraction*. Our evaluation consists of two approaches: a zero-shot baseline and a relation-enhanced baseline. In addition, we apply this task’s approach to other legal AI tasks, demonstrating that integrating structured legal relation knowledge enhances downstream task performance.

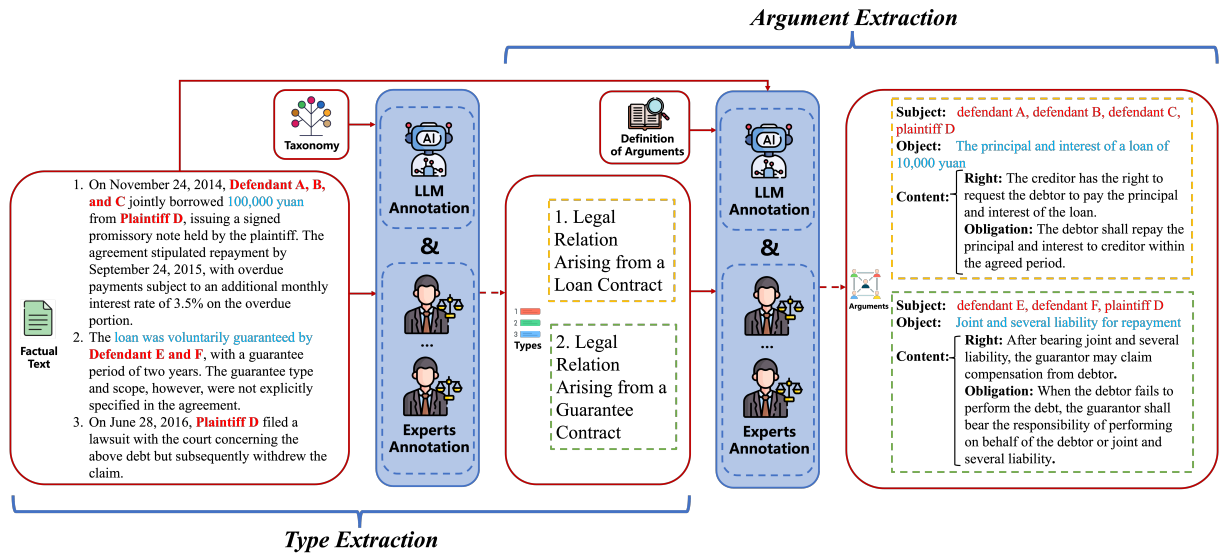


Figure 1: The construction workflow of **LexRel** is shown (English translation of the original Chinese text). The *type extraction* task involves extracting types from factual text by referencing taxonomy, while the *argument extraction* task involves extracting arguments from factual text and types by referencing definition of arguments. The red and blue annotations denote subject and object, as well as information that may represent them within the factual text.

Our contributions are summarized as follows:

- We propose the first structured schema of legal relations in Chinese civil law, with a taxonomy covering 265 relation types across 9 major domains, and clear definitions of subjects, objects, and content.
- We propose the task of *legal relation extraction*, and construct **LexRel**, the first expert-annotated benchmark for *legal relation extraction* in Chinese civil cases.
- We evaluate existing LLMs on **LexRel**, revealing significant model limitations in extracting legal relations from case facts.
- We demonstrate that integrating legal relations improves performance on downstream legal AI tasks.

## 2 Related Work

### 2.1 Legal Relation

The concept of legal relation refers to a social relation regulated by law, defined by the rights and obligations that arise between parties by operation of law (Liang, 2021). As used in this article, the term primarily denotes civil legal relations under the Civil Code of the People’s Republic of China (Wan et al., 2022). Within the new Civil Code, legal relations structure the organization of all chapters

and form the basis for the logical framework of civil law (Wang, 2018). In judicial practice, legal relations also serve as a highly consequential methodological approach: judges situate the parties within a specific relation, analyze the subjects, objects, and their respective rights and obligations, and determine the creation, alteration, or extinction of rights (Wang, 2017). While prior scholarship and doctrine have outlined general definitions of legal relations and their arguments (subject, object and content), a systematic taxonomy and precise operationalization of these elements has not previously been established.

### 2.2 Legal Knowledge Graph

Legal knowledge graphs (KGs) are widely used in legal AI, powering applications such as legal reasoning (Luo et al., 2024; Li et al., 2021), question answering (Liang et al., 2024), and decision support (Cui et al., 2024). These graphs typically organize entities (e.g., persons, organizations, statutes) and their relationships into structured formats to enable downstream machine reasoning. However, most existing legal KGs in Chinese civil law focus on general social relations (e.g., employment, kinship, transaction) (Cui et al., 2024; Li et al., 2021) or entity co-occurrence patterns (Li et al., 2024b). That is, while legal KGs often capture who interacts with whom, they rarely characterize how they are related under law, such as the presence

of a creditor-debtor relationship or a contractual obligation. Some works (Xu, 2021) adopt stricter definitions of civil legal relations to construct relation types, but their coverage remains limited and lacks a comprehensive taxonomy.

## 2.3 Legal Large Language Models

Large language models (LLMs) show remarkable promise in legal applications, including legal consultation, trial assistance, and document analysis (Lai et al., 2023). Several works explore adapting general LLMs to legal domain through continued pretraining and fine-tuning (Huang et al., 2023; Xiao et al., 2021; Yue et al., 2023). In addition, some works focus on evaluating capabilities of LLMs in the legal field (Dai et al., 2023; Li et al., 2024a). These advancements demonstrate the growing capability of legal LLMs and highlight their strong potential to support a wide range of legal applications.

## 3 Schema for Civil Legal Relations

We present the first comprehensive schema for legal relations in Chinese Civil Law, which comprises a hierarchical taxonomy of relation types and precise definitions of their arguments.

### 3.1 Taxonomy for Civil Legal Relations

We construct the taxonomy of civil legal relations through a combination of draft taxonomy construction and expert-guided refinement:

- **Draft taxonomy construction.** We begin by applying a keyword-matching approach to extract relational expressions from Chinese civil court judgments (China Judgments Online, 2013), producing an initial list of candidate legal relation terms. These candidates are then refined by referencing specialized legal terminology references (Sun and Luo, 1996), retaining only those whose prefixes correspond to legally recognized concepts. These works result in a draft taxonomy comprising 123 candidate relation types, categorized into 6 major domains based on the rights-obligations framework of Chinese civil law (Wan et al., 2022): *Legal Relations of Personality Rights*, *Legal Relations of Identity Rights*, *Legal Relations of Intellectual Property Rights*, *Legal Relations of Succession Rights*, *Legal Relations of Property Rights*, and *Legal Relations of Obligations*.

- **Expert-guided refinement.** Two senior legal scholars (professors or lecturers in law) then review and refine the draft taxonomy. Drawing on legal textbooks, academic literature, and their professional expertise, they add 3 new domains—*Bill Relations*, *Letter of Credit Relations*, and *Independent Guarantee Relations*—and expand the number of legal relation types to 265.

These two stages, combining empirical induction from litigated cases with normative grounding in established legal doctrine, ensure that the resulting taxonomy is both data-informed and conceptually rigorous. The finalized taxonomy covers 9 broad domains in Chinese civil law, consisting of 265 fine-grained legal relation types.

However, the distribution of relation types across these domains is not uniform. Among these domains, *Legal Relations of Obligations* encompasses the largest number of relation types with 200. This predominance reflects the inclusion of two categories that correspond to the most common types of civil disputes in practice—*Contractual Legal Relations* and *Legal Relations of Tort Liability*—which account for the majority of cases in judicial proceedings (Supreme People’s Court of the People’s Republic of China, 2024).

### 3.2 Arguments of Civil Legal Relations

Having established the taxonomy of relation types, we next define their core arguments: *subject*, *object*, and *content* (Wang, 2018). In Chinese civil law, there are no precise definitions for different arguments in different types of legal relation. With the help of the specific provisions of the Civil Code of the People’s Republic of China regarding different categories of rights (Wan et al., 2022) and senior legal scholars, we design precise definitions of the legal relations contained in main domains of our taxonomy.

Specifically, we select 8 of the 9 domains in the taxonomy, as well as 4 major subdomains under the domain *Legal Relations of Obligations* to standardize 12 distinct definitions. The 4 subdomains are: *Contractual Legal Relations*, *Legal Relations of Tort Liability*, *Legal Relations of Negotiorum Gestio* and *Legal Relations of Unjust Enrichment*. All definitions follow the following coarse-grained definitions (Wang, 2018):

- *Subject* refers to the party initiating or holding the right or obligation, typically a natural

| Definitions of Arguments in Legal Relations of Property Rights |   |
|--|---|
| <b>Subject</b>   | The right holder, together with any other subjects involved in the rights-and-obligations relationship, are all considered obligors. In practice, this includes the right holder himself/herself and any other parties who directly bear legal obligations under the relevant facts.  |
| <b>Object</b>  | Corporeal things, including real property, movable property and special items.  |
| <b>Content</b>   | In terms of rights, the right-holder legally enjoys ownership, usufructuary right, security interest, etc. The right-holder exercises direct control over the object through powers such as possess, use, seek profits from and dispose of the real property or movable property according to law. Real property rights become effective upon registration, whereas rights in movable property become effective upon delivery. In terms of obligations, it includes completing the registration or delivery procedures in accordance with the law; not infringing upon the legitimate property rights of others; and assuming the protection obligations such as removing obstacles, restoring the original state, and returning the original object when the property rights are unlawfully interfered with. |

Table 1: An example of the definitions of the subject, object, and content for *Legal Relations of Property Rights* (English translation for the original Chinese text).

person, legal entity, or organization.

- *Object* denotes the counterparty or legal target of the relation, such as another individual, a piece of property, or an intellectual asset.
- *Content* denotes a paragraph that captures substantive nature of the legal relation, such as ownership, liability, or contractual duty.

Table 1 shows an example of definitions of arguments in a legal relation (for the complete set of definitions, refer to Appendix C). As a result, each relation type in our taxonomy is associated with a specific configuration of these triplet elements, forming a unified schema that serves as the foundation for dataset construction and model evaluation in this work.

## 4 LexRel

As shown in Fig. 1, with the structured legal relation schema in place, we first define the task of extracting explicit legal relations from legal text that may contain implicit legal relation information as *legal relation extraction*. Based on this task definition, we construct **LexRel** benchmark. Each entry in **LexRel** consists of three components: factual text containing implicit legal relation information, a list of legal relation types, and the corresponding arguments (subject, object, and content) for each legal relation. **LexRel** is carefully curated from real-world Chinese civil court decisions through a two-step process: first, use LLMs to annotate candidate relations and their corresponding arguments, and second, legal experts review and refine the annotations to ensure both coverage and correctness.

### 4.1 Task Definition

Specifically, *legal relation extraction* comprises two subtasks: *type extraction*, which aims to identify the relevant legal relation types, and *argument extraction*, which seeks to determine the specific knowledge in each relation.

#### 4.1.1 Step 1: Type Extraction

Given input case facts  $x$ , which consists of the fact statements in a court judgment, the type extraction stage aims to extract the legal relation type  $\hat{r}$ :

$$\hat{r} = f_{\text{type}}(x), \quad \hat{r} \in \mathcal{R} \quad (1)$$

Here,  $f_{\text{type}}$  denotes relation type classifier, and  $\mathcal{R}$  is the legal relation set predefined in our taxonomy.

#### 4.1.2 Step 2: Argument Extraction

Conditioned on the predicted relation type  $\hat{r}$ , the model then extracts a set of associated subjects  $\hat{S} = \{\hat{s}_1, \hat{s}_2, \dots\}$ , a set of associated objects  $\hat{O} = \{\hat{o}_1, \hat{o}_2, \dots\}$ , and a detailed content (i.e., rights and obligations of the parties)  $\hat{c}$ . The extractor  $f_{\text{arg}}$  produces argument sets based on both input facts  $x$  and relation type  $\hat{r}$ :

$$(\hat{S}, \hat{O}, \hat{c}) = f_{\text{arg}}(x, \hat{r}) \quad (2)$$

### 4.2 Benchmark Construction

In order to ensure the quality of extraction and cost effectiveness, we adopt the method of extracting first with LLM and then correcting by experts.

**Data preprocessing.** We draw on court judgments for real-world adjudicated civil trial cases to construct the **LexRel**. Firstly, We utilize DeepSeek-V3 (DeepSeek-AI, 2025b) to extract legal relation

| Type Extraction   | Argument Extraction   |
|---|---|
| <p>Please extract the existing legal relations from the following judgment text. Only output the names, no need for explanation or description.</p> <p><b>Factual text:</b> {}</p> <p><b>Legal relation candidate set:</b> {}</p> <p>Output requirements: 1. Based on the content of the judgment, find the names of the legal relationships existing in the judgment text. 2. Only output the name of the most important legal relationship existing in the judgment text, and the output relationship name should be within 10 characters. 3. Each legal relationship name output must be in the legal relation candidate set. 4. Output strictly in accordance with the output format, and the output format is: [Legal relation 1, Legal relation 2.....]</p> | <p>Please extract the subject, object, and content of {<b>relation type</b>} from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> {}.</p> <p><b>Object definition:</b> {}.</p> <p><b>Content definition:</b> {}.</p> <p>Output requirements: 1. Each set of results is presented as a standard JSON object. 2. Required fields: {'subject': '', 'object': '', 'content': ''}. 3. Each set is in a separate paragraph without numbering or sorting. 4. Only one set needs to be output for the same subject and object.</p> |

Table 2: Prompt templates used in *type extraction* and *argument extraction* are shown (English translation for the original Chinese text). In *type extraction*, the additional inputs required are factual text and the collection of all relation types from the taxonomy. In *argument extraction*, supplementary inputs include factual text, one relation type from the set of previously extracted relation types, together with the definitions of three arguments.

types and their arguments from the full-text judgment as draft annotations to guide human experts during their annotations. The extraction follows two-step task structure defined above, with prompting instructions provided in Table 2.

And then, we extract the factual text authenticated by the court with DeepSeek-V3 as input for our dataset from each judgment as the final input for our dataset. The LLM is prompted to identify factual paragraphs that reflect adjudicated evidence, excluding legal analysis and procedural history.

**Expert annotation.** We engage six legal experts for annotation, all of whom hold law degrees and have prior experience with data annotation. They are provided with a detailed guideline including our schema and are briefed to ensure consistent understanding. Each annotator is assigned 200 items, each contains its facts and court reasoning extracted from judgments, together with draft relation types and arguments. They are tasked with verifying and refining input facts, types and arguments. The annotation process is overseen by a senior expert in legal AI research with extensive knowledge of legal relations. Analysis of annotators’ agreement is provided in Appendix D. After annotation, 60 items without identifiable legal relations are removed. The final **LexRel** benchmark comprises 1,140 annotated samples, each consisting of factual text paired with legal relations and their corresponding arguments, forming the **LexRel** evaluation dataset for assessing models’ ability to extract legal relations from case facts.

## 5 Experiment

### 5.1 Evaluation on LLMs

We begin by using **LexRel** to evaluate the performance of state-of-the-art large language models (LLMs) on the task of *legal relation extraction*, under two settings: a zero-shot baseline and a relation-enhanced baseline. The zero-shot baseline assesses the models’ ability to extract legal relations without any task-specific fine-tuning. In contrast, the relation-enhanced baseline involves providing large models with the full judgment text, rather than just the factual text. This allows large models to generate relation types and arguments, which are then used to fine-tune smaller models through supervised fine-tuning (SFT).

#### 5.1.1 Settings

**Models.** We select widely-used state-of-the-art LLMs for evaluation, including closed-source models: GPT-4o (OpenAI, 2024), o3-mini (OpenAI, 2025) and Claude-Sonnet-4 (Anthropic, 2025), as well as open-source models: DeepSeek-V3 (DeepSeek-AI, 2025b), DeepSeek-R1 (DeepSeek-AI, 2025a)<sup>1</sup>, Llama3.1-8B-Instruct, Llama3.1-70B-Instruct (Grattafiori et al., 2024), InternLM3-8B-Instruct (Cai et al., 2024), MiniCPM4-8B (MiniCPM-Team, 2025), Qwen3-8B, Qwen3-14B and Qwen3-32B (Yang et al., 2025). Specifically, for relation-enhanced baseline, we select DeepSeek-R1 and GPT-4o as generators to produce the training data.

<sup>1</sup>Although DeepSeek-R1 and DeepSeek-V3 are open-source models, we experimented via API calls due to computational resource considerations.

| Models                       | Methods                  | Type Extraction |              |              |              | Argument Extraction |              |              |              |
|------------------------------|--------------------------|-----------------|--------------|--------------|--------------|---------------------|--------------|--------------|--------------|
|                              |                          | P               | R            | micro-F1     | macro-F1     | P                   | R            | micro-F1     | macro-F1     |
| <b>GPT-4o</b>                | zero-shot                | 0.582           | <b>0.789</b> | 0.670        | 0.314        | 0.188               | 0.278        | 0.224        | 0.068        |
| <b>o3-mini</b>               | zero-shot                | <b>0.768</b>    | 0.756        | <b>0.762</b> | <b>0.441</b> | <b>0.413</b>        | 0.355        | <b>0.382</b> | <b>0.129</b> |
| <b>Claude-Sonnet-4</b>       | zero-shot                | 0.619           | 0.564        | 0.590        | 0.330        | 0.206               | 0.345        | 0.258        | 0.088        |
| <b>DeepSeek-R1</b>           | zero-shot                | 0.647           | 0.746        | 0.693        | 0.376        | 0.299               | 0.242        | 0.268        | 0.065        |
| <b>DeepSeek-V3</b>           | zero-shot                | 0.440           | 0.781        | 0.563        | 0.308        | 0.164               | 0.430        | 0.237        | 0.083        |
| <b>Llama3.1-70B</b>          | zero-shot                | 0.192           | 0.384        | 0.256        | 0.074        | 0.022               | 0.203        | 0.039        | 0.008        |
| <b>Qwen3-32B</b>             | zero-shot                | 0.500           | 0.700        | 0.583        | 0.158        | 0.030               | <b>0.551</b> | 0.057        | 0.015        |
| <b>Llama3.1-8B-Instruct</b>  | zero-shot                | 0.170           | 0.472        | 0.250        | 0.052        | 0.014               | 0.337        | 0.027        | 0.006        |
|                              | RE w/ <i>GPT-4o</i>      | 0.488           | 0.741        | 0.588        | 0.270        | 0.131               | 0.287        | 0.180        | 0.054        |
|                              | RE w/ <i>DeepSeek-R1</i> | 0.586           | 0.771        | 0.666        | 0.338        | 0.236               | 0.339        | 0.279        | 0.093        |
| <b>InternLM3-8B-Instruct</b> | zero-shot                | 0.310           | 0.644        | 0.418        | 0.139        | 0.055               | 0.042        | 0.048        | 0.002        |
|                              | RE w/ <i>GPT-4o</i>      | 0.479           | 0.773        | 0.592        | 0.288        | 0.150               | 0.349        | 0.209        | 0.061        |
|                              | RE w/ <i>DeepSeek-R1</i> | 0.599           | <u>0.817</u> | 0.691        | 0.324        | 0.277               | 0.387        | 0.323        | 0.100        |
| <b>MiniCPM4-8B</b>           | zero-shot                | 0.353           | 0.417        | 0.382        | 0.218        | 0.095               | 0.128        | 0.109        | 0.053        |
|                              | RE w/ <i>GPT-4o</i>      | 0.519           | 0.760        | 0.617        | 0.314        | 0.162               | 0.314        | 0.213        | 0.054        |
|                              | RE w/ <i>DeepSeek-R1</i> | 0.599           | 0.769        | 0.674        | 0.343        | 0.243               | 0.310        | 0.272        | 0.086        |
| <b>Qwen3-8B</b>              | zero-shot                | 0.338           | 0.745        | 0.464        | 0.149        | 0.061               | 0.275        | 0.099        | 0.014        |
|                              | RE w/ <i>GPT-4o</i>      | 0.536           | 0.769        | 0.632        | 0.311        | 0.175               | 0.346        | 0.233        | 0.078        |
|                              | RE w/ <i>DeepSeek-R1</i> | 0.576           | 0.816        | 0.675        | 0.337        | 0.255               | 0.375        | 0.304        | 0.098        |
| <b>Qwen3-14B</b>             | zero-shot                | 0.564           | 0.603        | 0.583        | 0.313        | 0.072               | 0.173        | 0.102        | 0.035        |
|                              | RE w/ <i>GPT-4o</i>      | 0.614           | 0.749        | 0.675        | 0.357        | 0.212               | 0.351        | 0.264        | 0.096        |
|                              | RE w/ <i>DeepSeek-R1</i> | <u>0.719</u>    | 0.748        | <u>0.733</u> | <u>0.430</u> | <u>0.373</u>        | <u>0.389</u> | <u>0.381</u> | <u>0.146</u> |

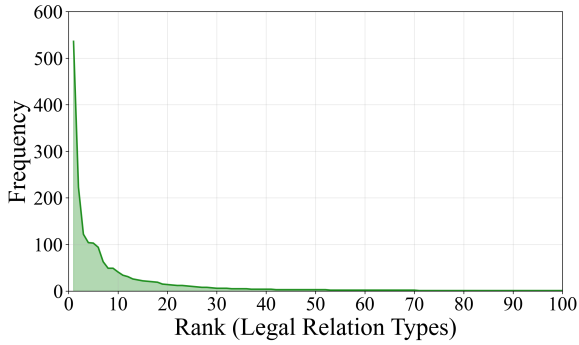
Table 3: Evaluation results of state-of-art LLMs on **LexRel** under precision (P), recall (R), micro-F1 and macro-F1 to *type extraction* and *argument extraction* are shown. The first 7 models are evaluated only on the zero-shot baseline due to the consideration of training costs, while the remaining models with 8B and 14B parameters are evaluated under both zero-shot and relation-enhanced (RE) baselines. The notation “w/” specifies which model is used to generate the training data. Across all results, the best zero-shot performance for each metric is shown in **bold**, and the best performance under relation-enhanced baseline is underlined.

**Configuration.** For the zero-shot setting, we directly evaluate the models on the **LexRel** benchmark. For the relation-enhanced setting, we first generate 5,500 synthetic training samples, and then finetune and evaluate the resulting models. Note that, due to training cost constraints, only the 8B and 14B variants of the open-source models are finetuned. Prompts used for both settings are presented in Table 2. All experiments are conducted on 4×A800 GPUs (each with 40GB memory).


**Metrics.** We use precision, recall, micro-F1 score and macro-F1 score (computing by averaging F1 scores across different relation types) to assess model performance on both *type extraction* and *argument extraction* tasks. For *type extraction*, evaluation is conducted by directly matching each predicted legal relation type against the gold-standard label. Considering the inherent ambiguity in legal

language and the cost of large-scale human evaluation, we adopt a LLM-as-a-Judge strategy for evaluating *argument extraction* correctness. Specifically, we pair each model-predicted argument with its corresponding gold argument and use DeepSeek-V3 to assess whether the two convey the same meaning.

To validate this choice, we randomly select 60 items to conduct 193 pairs for manual inspection. A human expert independently judges whether each predicted-gold pair conveys matching information, and these expert decisions serve as the reference gold standard for evaluating the LLM-as-a-Judge strategy. Based on this analysis, we obtain a subject accuracy of 0.954, an object accuracy of 0.969, and a content accuracy of 0.810. Taken together, these results show that DeepSeek-V3 is sufficiently reliable as an automatic evaluator in our setting.



(A) Distribution of top 100 most frequent legal relation types in **LexRel**.



(B) Distribution of top 100 most frequent causes of action in 26.6 million real-world civil court judgments.

Figure 2: (A) and (B) show the distributions of top 100 most frequent legal relation types in **LexRel** and causes of action in real-world civil judgments. (A) shows a clear downward trend, which may reflect the long-tail pattern. (B) shows the frequency of the top cause types compared to those that appear less frequently, emphasizing the imbalance in occurrence and illustrating a clear long-tail pattern.

### 5.1.2 Results and Discussion

Results of our evaluation are shown in Table 3. In zero-shot *type extraction* task, the best overall performer is o3-mini (micro-F1 = 0.762), followed by DeepSeek-R1 (micro-F1 = 0.693). Among open-source models, Qwen3 series demonstrates the strongest performance, with Qwen3-8B, Qwen3-14B and Qwen3-32B achieving micro-F1 scores of 0.464, 0.583 and 0.583 respectively. These results suggest that reasoning LLMs are better suited for zero-shot *type extraction*. In contrast, *argument extraction* proves significantly more challenging for the models. The highest zero-shot micro-F1 is again achieved by o3-mini (0.382), followed by DeepSeek-R1 (0.268) and Claude-Sonnet-4 (0.258). Most open-source models perform poorly (micro-F1 < 0.2), indicating the difficulty of identifying arguments spans in unstructured legal text under zero-shot conditions.

In the SFT experiments, we distill responses from GPT-4o and DeepSeek-R1 into smaller open-source models, and their performance improves significantly<sup>2</sup> across both tasks. For *type extraction*, the best distillation result is achieved by Qwen3-14B (micro-F1 = 0.733), which closely matches the performance of o3-mini in zero-shot baseline. And *argument extraction* also benefits much from SFT, for example, InternLM3-8B improves from 0.048 to 0.323 in micro-F1. This large gain suggests that SFT can effectively equip smaller models with the ability to learn structured legal relations.

<sup>2</sup>In this paper, significance is tested through t-test ( $\alpha = 0.05$ ).




Figure 3: Pareto distribution of causes of action, demonstrating that the top 25 case types account for 80% of all civil judgments. This graph highlights the dominance of a small number of case types in the overall dataset and identifies the 80% threshold.

Furthermore, we observe that the macro-F1 scores of all LLMs are lower than their corresponding micro-F1 scores across all settings. This disparity suggests that LLMs exhibit varying capabilities across different types of legal relations.

### 5.2 Long-Tail Analysis

**LexRel** contains 1,140 civil cases and 1,863 fully annotated legal relations spanning all major civil law domains. Fig 2(A) shows top 100 most frequent legal relation types in **LexRel**, revealing a clear downward trend. We further examine whether this downward pattern faithfully mirrors real-world distributions. Because real-world cases do not explicitly specify legal relations, we use **civil cause of action** as a proxy. We first validate reliability of this proxy by measuring association between legal relations and causes of action in **LexRel**, which

shows a strong correlation ( $\chi^2$  test,  $p < 0.05$ ; Cramér’s  $V = 0.528$ )<sup>3</sup>. We then examine a comprehensive set of 26.6 million real-world civil court judgments from China Judgments Online (China Judgments Online, 2013) and extract cause of action information for each case. Fig 2(B) shows that causes of action for real-world civil cases follow a long-tail distribution. Consistently, Pareto analysis (Fig 3) shows that a small number of frequent cause of action types (25 types) account for the majority of instances, while a large number of types occur sparsely. Taken together, these results indicate that the long-tail distribution of legal relations in LexRel echoes the intrinsic pattern observed in real-world litigation data.

### 5.3 Effects on Downstream Legal AI Tasks

To demonstrate the effects of explicitly incorporating legal relations in legal AI tasks, we select a set of related tasks from LawBench (Fei et al., 2023) for further analysis.

#### 5.3.1 Settings

**Tasks.** We select three tasks—Case Analysis, Consultation and Criminal Damages Calculation—from the LawBench dataset (Fei et al., 2023), based on their inherent dependency on legal relation identification. This dependency primarily manifests as a prerequisite step; for example, identifying the existence and type of legal relations prior to structured case analysis or legal consultation, or determining ownership or contractual relations as a basis for computing criminal damages.

**Models.** We experiment with existing models including GPT-4o, DeepSeek-V3, MiniCPM4-8B, and Qwen3-8B.

**Hyper-parameter.** Considering both efficiency and effectiveness, we use legal relation distilled from DeepSeek-V3. Specifically, we extract relation types and arguments of the input and insert them directly into the task instructions during inference. All other evaluation parameters remain consistent with the previous section.

**Metrics.** We follow the official evaluation guidelines of LawBench, where each task is scored on a scale from 0 to 100 based on correctness and completeness of the model’s output. We compare

<sup>3</sup>Statistical significance and effect size are evaluated using a chi-square test and Cramér’s  $V$ . A result is conventionally regarded as strong correlation when  $p < 0.05$  and  $V > 0.5$ .

| Models             | Case        | Consultation | Calculation |
|--------------------|-------------|--------------|-------------|
| <b>GPT-4o</b>      | 55.8        | 18.2         | 84.4        |
| w/ LR              | <b>56.6</b> | <b>19.2</b>  | <b>85.8</b> |
| <b>DeepSeek-V3</b> | 66.2        | 16.4         | 85.0        |
| w/ LR              | <b>68.2</b> | <b>16.5</b>  | <b>86.4</b> |
| <b>MiniCPM4-8B</b> | 32.0        | 7.6          | 65.0        |
| w/ LR              | <b>45.0</b> | <b>8.4</b>   | <b>76.8</b> |
| <b>Qwen3-8B</b>    | 57.2        | 15.0         | 91.8        |
| w/ LR              | <b>57.6</b> | 14.2         | <b>93.2</b> |

Table 4: Results on downstream legal AI tasks. The three tasks presented are: Case Analysis (Case), Consultation and Criminal Damages Calculation (Calculation). The notation ‘w/ LR’ refers to the setting where the model is evaluated using relation-aware inputs. Results which are improved by legal relation information are highlighted in **bold**.

results under two settings: (1) vanilla zero-shot, where the model receives only the task instruction and input; and (2) relation-aware zero-shot, where the input is enhanced with distilled legal relation types and arguments.

#### 5.3.2 Results and Analysis

As shown in Table 4, incorporating legal relations (“w/ LR”) consistently enhances performance across nearly all tasks and models. For Case Analysis and Consultation, legal relations provide essential cues for identifying and reasoning over relevant statutory knowledge. In Criminal Damages Calculation, certain types of legal relations (e.g., *Legal Relation of Ownership*) enable the model to more accurately identify property that has been subject to infringement, which can then be used to calculate damages. Although current LLMs show limited performance in extracting legal relations, incorporating these relations into downstream tasks still yields performance benefits, suggesting that legal relation modeling remains a promising direction for further exploration.

## 6 Conclusion

In this work, we explicitly model legal relations in the context of Chinese civil law. By constructing a structured schema of civil legal relations and curating an evaluation dataset, LexRel, for the legal relation extraction task, we evaluate large language models on their ability to effectively extract legal relations and associated arguments, and further demonstrate the benefits of incorporating such structured knowledge across multiple tasks.

## 7 Limitations

Our schema and dataset focus specifically on civil legal relations under Chinese law, which may limit direct transferability to other jurisdictions. Although the conceptual structure of legal relations is broadly consistent across civil law systems, each jurisdiction embodies distinct doctrines and terminology, requiring localized adaptation.

## 8 Ethical Considerations

All data used in this study are sourced from publicly available Chinese civil court decisions released through official judicial channels ([China Judgments Online, 2013](#)). We strictly adhere to data privacy and research ethics standards, ensuring that no confidential information is disclosed. All human annotators participated voluntarily and are compensated fairly according to institutional guidelines.

## References

Anthropic. 2025. Introducing Claude 4. <https://www.anthropic.com/news/claude-4>.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, and 81 others. 2024. InternLM2 technical report. *arXiv preprint arXiv:2403.17297*.

China Judgments Online. 2013. China Judgments Online. <https://wenshu.court.gov.cn>.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.

Jiayi Cui, Munan Ning, Zongjian Li, Bohua Chen, Yang Yan, Hao Li, Bin Ling, Yonghong Tian, and Li Yuan. 2024. ChatLaw: A multi-agent collaborative legal assistant with knowledge graph enhanced mixture-of-experts large language model. *arXiv preprint arXiv:2306.16092*.

Yongfu Dai, Duanyu Feng, Jimin Huang, Haochen Jia, Qianqian Xie, Yifang Zhang, Weiguang Han, Wei Tian, and Hao Wang. 2023. LAiW: A Chinese legal large language models benchmark. *arXiv preprint arXiv:2310.05620*.

DeepSeek-AI. 2025a. DeepSeek-R1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

DeepSeek-AI. 2025b. DeepSeek-V3: Technical report. *arXiv preprint arXiv:2412.19437*.

Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Songyang Zhang, Kai Chen, Zongwen Shen, and Jidong Ge. 2023. LawBench: Benchmarking legal knowledge of large language models. *arXiv preprint arXiv:2309.16289*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Quzhe Huang, Mingxu Tao, Chen Zhang, Zhenwei An, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. 2023. Lawyer LLaMA: Technical report. *arXiv preprint arXiv:2305.15062*.

Zhigang Huang, Xinqiang Xie, Dong Ge, Chenqiu Cai, Lili Dou, and Tianyi Wang. 2021. Case recommendation based on knowledge graph of judicial cases. *Journal of Nanjing University Natural Science*, 57(6):1053–1063.

Jinqi Lai, Wensheng Gan, Jiayang Wu, Zhenlian Qi, and Philip S. Yu. 2023. Large language models in law: A survey. *arXiv preprint arXiv:2312.03718*.

J Richard Landis and Gary G Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33(1):159–174.

Haitao Li, You Chen, Qingyao Ai, Yueyue Wu, Ruizhe Zhang, and Yiqun Liu. 2024a. Lexeval: A comprehensive chinese legal benchmark for evaluating large language models. *arXiv preprint arXiv:2409.20288*.

Jun Li, Lu Qian, Peifeng Liu, and Taoxiong Liu. 2024b. Construction of legal knowledge graph based on knowledge-enhanced large language models. *Information*, 15(11):666.

Luoqiu Li, Zhen Bi, Hongbin Ye, Shumin Deng, Hui Chen, and Huaixiao Tou. 2021. Text-guided legal knowledge graph reasoning. *arXiv preprint arXiv:2104.02284*.

Huixing Liang. 2021. *General Theory to Civil Law*. Law Press China.

Lei Liang, Mengshu Sun, Zhengke Gui, Zhongshu Zhu, Zhouyu Jiang, Ling Zhong, Yuan Qu, Peilong Zhao, Zhongpu Bo, Jin Yang, Huaidong Xiong, Lin Yuan, Jun Xu, Zaoyang Wang, Zhiqiang Zhang, Wen Zhang, Huajun Chen, Wenguang Chen, and Jun Zhou. 2024. Kag: Boosting llms in professional domains via knowledge augmented generation. *arXiv preprint arXiv:2409.13731*.

Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2024. Reasoning on graphs: Faithful and interpretable large language model reasoning. *arXiv preprint arXiv:2310.01061*.

|     |  |
|-----|--|
| 620 | MiniCPM-Team. 2025. MiniCPM4: Ultra-efficient LLMs on end devices. <i>arXiv preprint arXiv:2506.07900</i> .  |
| 621 |  |
| 622 |  |
| 623 | OpenAI. 2024. GPT-4o system card. <i>arXiv preprint arXiv:2410.21276</i> .   |
| 624 |  |
| 625 | OpenAI. 2025. OpenAI o3-mini. <a href="https://openai.com/index/openai-o3-mini/">https://openai.com/index/openai-o3-mini/</a> .  |
| 626 |  |
| 627 | Weixing Shen. 2004. Reconsideration on the constitution of civil legal relationship. <i>Journal of Comparative Law</i> , (1):42–54.  |
| 628 |  |
| 629 |  |
| 630 | Wanzhong Sun and Haocai Luo. 1996. <i>The Comprehensive Dictionary of Chinese Law</i> . China Procuratorial Press.   |
| 631 |  |
| 632 |  |
| 633 | Supreme People’s Court of the People’s Republic of China. 2024. Gazette of the supreme people’s court (2024). <a href="http://gongbao.court.gov.cn/Details/3ece7439305fa8bf1a7aae143d1598.html">http://gongbao.court.gov.cn/Details/3ece7439305fa8bf1a7aae143d1598.html</a> .                                      |
| 634 |  |
| 635 |  |
| 636 |  |
| 637 |  |
| 638 | Meng Wan, Feng Zhu, Benedict Amour, and Hailong Tang. 2022. <i>The Civil Code of the People’s Republic of China</i> . Springer.  |
| 639 |  |
| 640 |  |
| 641 | Liming Wang. 2017. <i>General Principles of Civil Law</i> . Renmin University Press.   |
| 642 |  |
| 643 | Liming Wang. 2018. <i>Civil Law</i> . Renmin University Press.   |
| 644 |  |
| 645 | Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. Lawformer: A pre-trained language model for chinese legal long documents. <i>arXiv preprint arXiv:2105.03887</i> .   |
| 646 |  |
| 647 |  |
| 648 |  |
| 649 | Jiayu Xu. 2021. Legal text knowledge mining for civil law. Master’s thesis, Harbin Institute of Technology.  |
| 650 |  |
| 651 | An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. <i>arXiv preprint arXiv:2505.09388</i> . |
| 652 |  |
| 653 |  |
| 654 |  |
| 655 |  |
| 656 |  |
| 657 |  |
| 658 | Zhonghua Yuan. 2024. Reflections on the paradigm of legal relations. <i>Chinese Journal of Law</i> , 46(1):129–147.  |
| 659 |  |
| 660 |  |
| 661 | Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, Chenchen Shen, Shujun Liu, Yuxuan Zhou, Yao Xiao, Song Yun, Xuanjing Huang, and Zhongyu Wei. 2023. Disc-LawLLM: Fine-tuning large language models for intelligent legal services. <i>arXiv preprint arXiv:2309.11325</i> .                                       |
| 662 |  |
| 663 |  |
| 664 |  |
| 665 |  |
| 666 |  |
| 667 | Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. <i>arXiv preprint arXiv:2403.13372</i> .  |
| 668 |  |
| 669 |  |
| 670 |  |
| 671 |  |

|   |  |
|---|--|
| <b>A Taxonomy</b>   | 672  |
| The complete taxonomy can be found in the <b>Taxonomy</b> folder in attachment.   | 673<br>674   |
| <b>B LexRel and SFT data</b>  | 675  |
| The LexRel we constructed and SFT data generated by GPT-4o and DeepSeek-R1 are included in the <b>LexRel</b> folder in attachment.  | 676<br>677<br>678  |
| <b>C Prompts</b>  | 679  |
| The complete definition of arguments and the prompts used in <i>argument extraction</i> are as follows in Tables 5-8. The original Chinese version is in the <b>Prompts</b> folder in attachment.   | 680<br>681<br>682<br>683   |
| <b>D Inter-annotator Agreement</b>  | 684  |
| Because evaluation of <i>argument extraction</i> in <i>legal relation extraction</i> involves semantic alignment between generated textual arguments rather than exact label matching, inter-annotator agreement (IAA) is assessed based on expert judgment. Accordingly, a senior legal expert is designated as the reference annotator for agreement assessment.  | 685<br>686<br>687<br>688<br>689<br>690<br>691  |
| Owing to the substantial cost of fully annotating and validating the entire dataset, we randomly sample 300 instances for agreement assessment. These instances are evenly distributed among other annotators for independent annotation. The senior expert subsequently reviews each annotation by comparing it with a reference interpretation derived from the same annotation guidelines. When discrepancies arise, the instance is adjudicated by the senior expert to establish a new reference annotation. IAA is then computed between the senior expert and the other annotators, yielding a Cohen’s Kappa (Cohen, 1960) of 0.706 <sup>4</sup> for <i>legal relation extraction</i> in <b>LexRel</b> . | 692<br>693<br>694<br>695<br>696<br>697<br>698<br>699<br>700<br>701<br>702<br>703<br>704<br>705 |
| <b>E Code</b>   | 706  |
| Code is located in the <b>Code</b> folder <sup>5</sup> . It includes code for generating via API and vllm to perform type and argument extraction, as well as code for conducting two types of evaluations. Additionally, it contains code for evaluating downstream tasks. The training is implemented using the LlamaFactory (Zheng et al., 2024) framework, and the SFT configuration files are also included.   | 707<br>708<br>709<br>710<br>711<br>712<br>713<br>714   |
| <sup>4</sup> According to Landis and Koch (1977), a Cohen’s Kappa score between 0.61 and 0.80 signifies “substantial agreement”   |  |
| <sup>5</sup> Our schema, dataset and code will be released upon acceptance.   |  |

| Types                                 | Prompt   |
|---------------------------------------|--|
| Legal Relations of Personality Rights | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The right holder himself/herself and all other subjects are the obligors (generally only the right holder is output).</p> <p><b>Object definition:</b> The corresponding personal interests (output xx rights interests only).</p> <p><b>Content definition:</b> In terms of rights, the subject enjoys the right to have their life, body, health, name, portrait, reputation, privacy and other personal interests not infringed upon by law. In terms of obligations, other civil subjects bear negative obligations such as not infringing upon, not insulting or defaming, and not disclosing relevant information. When necessary, they also have the obligation to actively protect and provide relief.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>   |
| Legal Relations of Status Rights      | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The right holder themselves and all other subjects are the obligors (generally, multiple right holders and all obligors related to the facts should be output simultaneously).</p> <p><b>Object definition:</b> Corresponding identity interests (output xx relationship interests only).</p> <p><b>Content definition:</b> In terms of rights, the subject enjoys the rights of spousal, parental, guardianship, support and inheritance, etc. in accordance with the law due to the identity relationship. In terms of obligations, the corresponding subjects undertake identity obligations such as loyalty, assistance, support, education, and guardianship.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>   |
| Tortious Legal Relations              | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The infringer is the subject that implements the infringement act, and the infringed party is the subject that suffers damage due to the infringement act (generally, both should be output simultaneously).</p> <p><b>Object definition:</b> The legitimate rights and interests of the infringed party (just output the name of the infringed right).</p> <p><b>Content definition:</b> The rights enjoyed by the creditor: The right to request the debtor to perform or refrain from acting in accordance with the agreement or law. When a debtor fails to exercise his rights or infringes upon his creditor's rights, the right of subrogation or revocation may be exercised in accordance with the law. The obligations borne by the debtor: They should fulfill the obligation of payment as agreed, including delivering money, property or providing services, and may also bear the obligation of inaction.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol> |

Table 5: Definitions and prompt templates used in *argument extraction* for *Legal Relations of Personality Rights*, *Legal Relations of Status Rights* and *Tortious Legal Relations* are shown (English translation for the original Chinese text).

| Types                                    | Prompt   |
|--|--|
| Legal Relations of Intellectual Property | <p>Please extract the subject, and object of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The right holder himself/herself and all other subjects are the obligors (generally only the right holder is output).</p> <p><b>Object definition:</b> Intellectual achievements, including works, inventions and creations, trademarks, trade secrets, etc. (Output the name of the intellectual achievement and the corresponding rights to this achievement).</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>   |
| Legal Relations of Real Rights           | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The right holder himself/herself and all subjects related to the rights and obligations relationship are the obligors (generally the output right holder and other obligors who bear direct legal obligations in this fact).</p> <p><b>Object definition:</b> Tangible objects, including immovable property, movable property and special objects. (Name of the output).</p> <p><b>Content definition:</b> In terms of rights, the right holder enjoys ownership, usufructuary rights, security interests, etc. in accordance with the law, exercises direct control over the object through rights such as possession, use, income, and disposal, and can legally exclude interference from others. Immovable property rights come into effect through registration, while movable property rights take delivery as a prerequisite for their effectiveness. In terms of obligations, it includes completing the registration or delivery procedures in accordance with the law; It is prohibited to infringe upon the legitimate property rights of others. When property rights are illegally interfered with, it undertakes the obligation of protection such as removing obstacles, restoring the original state and returning the original property.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol> |
| Legal Relations of Success               | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> Heir and decedent. (Generally, both need to be output simultaneously).</p> <p><b>Object definition:</b> Specific heritage content.</p> <p><b>Content definition:</b> In terms of rights, heirs enjoy rights such as inheriting the estate, accepting bequests, and subrogation inheritance in accordance with the law or will. In terms of obligations, the heirs should divide the property in the spirit of friendly consultation, fulfill the obligations stipulated in the conditional will, notify the relevant parties and properly keep the estate, and also settle debts and taxes in accordance with the law.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>   |

Table 6: Definitions and prompt templates used in *argument extraction* for *Legal Relations of Intellectual Property*, *Legal Relations of Real Rights* and *Legal Relations of Success* are shown (English translation for the original Chinese text).

| Types                               | Prompt   |
|-------------------------------------|--|
| Contractual Legal Relations         | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The parties to the contract (generally all parties need to be output simultaneously).</p> <p><b>Object definition:</b> The subject matter pointed to by the contractual obligations, including goods, acts, rights, etc. (output the performance of the contract).</p> <p><b>Content definition:</b> The rights enjoyed by the creditor: The right to request the debtor to perform or refrain from acting in accordance with the agreement or law. When a debtor fails to exercise his rights or infringes upon his creditor's rights, the right of subrogation or revocation may be exercised in accordance with the law. The obligations borne by the debtor: They should fulfill the obligation of payment as agreed, including delivering money, property or providing services, and may also bear the obligation of inaction.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>  |
| Legal Relation of Unjust Enrichment | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> A beneficiary is the subject that gains benefits without legal basis, while a victim is the subject that suffers losses due to the beneficiary's acquisition of benefits (generally, both should be output simultaneously).</p> <p><b>Object definition:</b> Interest, which refers to the property or benefit obtained by the beneficiary without legal basis. (Content of Output Benefits).</p> <p><b>Content definition:</b> The rights enjoyed by the creditor: The right to request the debtor to perform or refrain from acting in accordance with the agreement or law; When a debtor fails to exercise his rights or infringes upon his creditor's rights, the right of subrogation or revocation may be exercised in accordance with the law. The obligations borne by the debtor: They should fulfill the obligation of payment as agreed, including delivering money, property or providing services, and may also bear the obligation of inaction.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>   |
| Legal Relation of Negotiorum Gestio | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> A manager (creditor) refers to a person who, without being entrusted or bound by legal obligations, takes the initiative to manage the affairs of others to prevent their interests from being damaged. The beneficiary (debtor) refers to the party whose interests are attributed to the management of the affairs (usually both of them should be output simultaneously).</p> <p><b>Object definition:</b> The actions of the manager (such as repairing houses on behalf of others, providing assistance to the sick and injured, etc.) and the necessary expenses or compensation for losses arising therefrom. (Content of the Output act).</p> <p><b>Content definition:</b> The rights enjoyed by the creditor: The right to request the debtor to perform the act or refrain from acting in accordance with the agreement or law; When a debtor fails to exercise his rights or infringes upon his creditor's rights, the right of subrogation or revocation may be exercised in accordance with the law. The obligations borne by the debtor: They should fulfill the obligation of payment as agreed, including delivering money, property or providing services, and may also bear the obligation of inaction.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol> |

Table 7: Definitions and prompt templates used in *argument extraction* for *Contractual Legal Relations*, *Legal Relation of Unjust Enrichment* and *Legal Relation of Negotiorum Gestio* are shown (English translation for the original Chinese text).

| Types                                    | Prompt   |
|--|--|
| Legal Relation of Letters of Credit      | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> The primary parties are the Applicant, the Issuing Bank, and the Beneficiary, but there are eight parties in total, including: the Advising Bank, the Confirming Bank, the Negotiating Bank, the Paying Bank, and the Reimbursing Bank.</p> <p><b>Object definition:</b> The delivery of documents and the promise of payment.</p> <p><b>Content definition:</b> 1. The contractual sales relationship between the Applicant and the Beneficiary. 2. The principal-agent contractual relationship between the Issuing Bank and the Applicant. 3. The relationship between the Issuing Bank and the Beneficiary (a subject of theoretical and practical debate). Generally, when a revocable letter of credit is issued, there is no contractual relationship between the Issuing Bank and the Beneficiary. However, when an irrevocable letter of credit is issued and received by the Beneficiary, a binding independent contractual relationship is formed between them. 4. The principal-agent relationship between the Advising Bank and the Issuing Bank. 5. The contractual relationship between the Issuing Bank and the Paying/Negotiating Bank arises when the Issuing Bank designates or authorizes another bank to pay or negotiate, and that bank accepts the role. Otherwise, no legal relationship exists. 6. Typically, there is no legal relationship between the Beneficiary and the Paying/Negotiating Bank, unless a bank (other than the Issuing Bank) pays or negotiates the Beneficiary's draft under the Issuing Bank's authorization. In such cases, a creditor-debtor legal relationship is established under negotiable instruments law.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol> |
| Legal Relation of Independent Guarantees | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> Issuer (bank/financial institution), beneficiary, applicant (can be an instructor).</p> <p><b>Object definition:</b> The "document" on which the claim request is based.</p> <p><b>Content definition:</b> It refers to the bill rights that the subject enjoys and the bill obligations it should undertake in accordance with the law. The right to a bill is the right enjoyed by the right holder to request the obligor to pay the amount of the bill, including the right to request payment and the right of recourse. Bill obligations refer to the responsibilities that the obligor must fulfill to satisfy the rights claims enjoyed by the right holder based on the bill, such as payment obligations, acceptance obligations, and guarantee payment obligations, etc.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>  |
| Legal Relation of Bills                  | <p>Please extract the subject, object, and content of <b>{relation type}</b> from the given judgment documents without explanation or description.</p> <p><b>Factual text:</b> {}.</p> <p><b>Subject definition:</b> A subject is specific and mainly includes: the drawer, the payee, the payer, the holder of the bill, the acceptor, the endorser, the guarantor, the participant, etc.</p> <p><b>Object definition:</b> A certain amount of currency.</p> <p><b>Content definition:</b> The "underlying transaction contract" relationship between the applicant and the beneficiary; The "independent guarantee contract" relationship between the issuer and the beneficiary; The "application/recovery contract" relationship between the issuer and the applicant. considerations correspond to: the issuer's unconditional payment obligation, the principle of document review and "apparent conformity", the right of recovery after payment, the right of defense against fraud exceptions, and the conditions for termination or invalidation of the guarantee.</p> <p>Output requirements:</p> <ol style="list-style-type: none"> <li>1. Each set of results is presented as a standard JSON object.</li> <li>2. Required fields: {'subject': '', 'object': '', 'content': ''}.</li> <li>3. Each set is in a separate paragraph without numbering or sorting.</li> <li>4. Only one set needs to be output for the same subject and object.</li> </ol>  |

Table 8: Definitions and prompt templates used in *argument extraction* for *Legal Relation of Letters of Credit*, *Legal Relation of Independent Guarantees* and *Legal Relation of Bills* are shown (English translation for the original Chinese text).