



PDF Download
3731443.3771354.pdf
25 January 2026
Total Citations: 0
Total Downloads: 171

 Latest updates: <https://dl.acm.org/doi/10.1145/3731443.3771354>

RESEARCH-ARTICLE

Capturing Legal Reasoning Paths from Facts to Law in Court Judgments using Knowledge Graphs

RYOMA KONDO, The University of Tokyo, Tokyo, Japan

RIONA MATSUOKA, The University of Tokyo, Tokyo, Japan

TAKAHIRO YOSHIDA, Canon Inc., Ohta-ku, Tokyo, Japan

KAZUYUKI YAMASAWA

RYOHEI HISANO, The University of Tokyo, Tokyo, Japan

Open Access Support provided by:

The University of Tokyo

Canon Inc.

Published: 10 December 2025

[Citation in BibTeX format](#)

K-CAP '25: Knowledge Capture Conference 2025

December 10 - 12, 2025

Ohio, Dayton, USA

Conference Sponsors:

SIGAI

Capturing Legal Reasoning Paths from Facts to Law in Court Judgments using Knowledge Graphs

Ryoma Kondo*
Graduate School of Information
Science and Technology
The University of Tokyo
Tokyo, Japan
kondor@g.ecc.u-tokyo.ac.jp

Riona Matsuoka
Graduate Schools for Law and Politics
The University of Tokyo
Tokyo, Japan
m-riona248050@g.ecc.u-tokyo.ac.jp

Takahiro Yoshida
The Canon Institute for Global
Studies
Tokyo, Japan
amearashi08241995@gmail.com

Kazuyuki Yamasawa
TKC Corporation
Tokyo, Japan
yamasawa-kazuyuki@tkc.co.jp

Ryohei Hisano
Graduate School of Information
Science and Technology
The University of Tokyo
Tokyo, Japan
hisanor@g.ecc.u-tokyo.ac.jp

Abstract

Court judgments reveal how legal rules have been interpreted and applied to facts, providing a foundation for understanding structured legal reasoning. However, existing automated approaches for capturing legal reasoning, including large language models, often fail to identify the relevant legal context, do not accurately trace how facts relate to legal norms, and may misrepresent the layered structure of judicial reasoning. These limitations hinder the ability to capture how courts apply the law to facts in practice. In this paper, we address these challenges by constructing a legal knowledge graph from 648 Japanese administrative court decisions. Our method extracts components of legal reasoning using prompt-based large language models, normalizes references to legal provisions, and links facts, norms, and legal applications through an ontology of legal inference. The resulting graph captures the full structure of legal reasoning as it appears in real court decisions, making implicit reasoning explicit and machine-readable. We evaluate our system using expert annotated data, and find that it achieves more accurate retrieval of relevant legal provisions from facts than large language model baselines and retrieval-augmented methods.

CCS Concepts

• **Information systems** → **Knowledge representation and reasoning; Semantic web description languages**; • **Computing methodologies** → *Artificial intelligence*; • **Applied computing** → Law.

*Corresponding author



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

K-CAP '25, Dayton, OH, USA

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1867-0/25/12

<https://doi.org/10.1145/3731443.3771354>

Keywords

Legal Knowledge Graph, Legal Epistemology, Legal Search Engine, Retrieval-Augmented Generation

ACM Reference Format:

Ryoma Kondo, Riona Matsuoka, Takahiro Yoshida, Kazuyuki Yamasawa, and Ryohei Hisano. 2025. Capturing Legal Reasoning Paths from Facts to Law in Court Judgments using Knowledge Graphs. In *Knowledge Capture Conference 2025 (K-CAP '25)*, December 10–12, 2025, Dayton, OH, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3731443.3771354>

1 Introduction

Court judgments provide a uniquely rich source of structured knowledge about how legal reasoning unfolds in practice. Rather than merely stating legal rules, they show how those rules are interpreted and applied to specific factual situations. This reasoning process is shaped by precedent and expert judgment, and reflects patterns that are essential for understanding how legal knowledge is constructed and used. Within each decision lies an implicit, but identifiable, structure that includes evidence, circumstantial facts, primary facts, legal provisions, interpretations, legal applications, and final conclusions [7].

Specifically, legal reasoning emerges from the interaction of two structured axes: one that builds facts from evidence and one that derives legal meaning from statutes [7, 8, 13, 15, 16]. Along the factual axis, courts identify pieces of evidence, evaluate circumstantial facts, and infer primary facts that support legal evaluation. On the legal axis, legal provisions are interpreted to define the applicable norms and conditions for judgment. These two axes meet at the point of application, where the court determines whether the legal norms apply to the established facts. This intersection defines the core structure of legal inference.

Legal reasoning structures are rarely made explicit. Hence, previous law-focused research in the semantic web community has often overlooked this crucial aspect of legal knowledge. For example, the Legal Knowledge Interchange Format [6], which is a structured framework for representing legal rules and concepts, focuses on static rules, and thus cannot capture the nuanced, context-specific

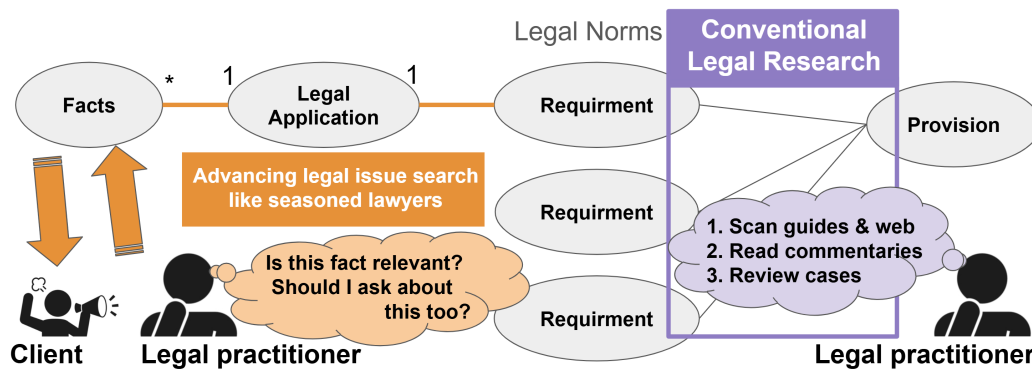


Figure 1: Comparison between conventional legal research and our fact-centered retrieval approach. Traditional workflows begin with legal provisions and commentaries, relying heavily on expert experience to infer relevant facts. Our method reverses this pipeline by starting from observed facts and following structured reasoning paths that connect them to legal norms and provisions. This enables issue-spotting behavior, even by junior practitioners, by making the reasoning steps explicit.

interpretations commonly found in court judgments. Sakhaee and Wilson [10] focused on dynamic legislation networks, but did not account for the context-dependent reasoning processes that are central to court judgments. Garcia-Godinez et al. [4] emphasized that court judgments are not merely declarations of legal rules, but also illustrate how those rules are interpreted and applied to specific factual situations. However, their work does not provide a formal ontological definition or develop a knowledge graph to represent this structure.

Given the recent rise of large language models (LLMs), it is tempting to ask whether such structured representations are still necessary for the construction of legal search engines. Could a general-purpose LLM be sufficient to answer queries such as “Given the facts brought before the court, what laws and past precedents should a lawyer or judge consider?”

Although LLMs can generate plausible legal text, they face three fundamental obstacles in supporting reliable legal reasoning. First, they lack the ability to establish legal framing. That is, they cannot determine which jurisdiction, time period, or legal system applies. For example, in defamation cases, whether truth constitutes a defense depends on the jurisdiction. Under US law, truth typically negates unlawfulness, whereas under Japanese law, additional conditions such as public interest and purpose are required [11]. LLMs often fail to recognize such legal contexts and may produce arguments that conflate incompatible legal standards. Second, LLMs lack the structural grounding needed to connect facts, norms, and provisions through explicit reasoning paths. Although fluent, their outputs are often shallow, mimicking surface-level textual patterns without reconstructing the logic of legal inference. Third, even within a single jurisdiction, courts sometimes adopt multi-layered or divergent perspectives. For instance, in US abortion cases, some lines of judicial opinion support the right, while others oppose it [14]. Simply querying an LLM will not yield a detailed, quantitative summary of how judges construct and weigh these competing legal viewpoints.

Recent studies have addressed these limitations using retrieval-augmented generation (RAG) techniques, which blend LLMs with external document retrieval to anchor the generated answers in

authoritative sources [17]. In the legal domain, [3] developed a RAG-based system that retrieves supporting fragments from statutes, enabling more accurate and transparent statutory reasoning. However, while these methods improve factual grounding, they do not reconstruct the underlying logic of judicial decision-making or model the structured inferential processes that connect facts, norms, and legal provisions. Systematic evaluation tools such as LegalBench [5] reveal that, even when assisted with additional context or citation retrieval, LLMs often fail on legal tasks requiring precise, multi-step reasoning and structured context-awareness. These findings suggest that explicit structured representations, such as legal knowledge graphs (LKGs), remain essential for supporting robust and transparent legal search and reasoning systems.

Hence, in this paper, we build on [7] by defining a legal ontology that captures the core reasoning structure found in court decisions. We adopt the resource description framework to encode this graph structure, enabling an explicit and machine-readable representation of legal inference paths. We populate an LKG through a three-step process. First, we extract the key elements of evidence, circumstantial facts, main facts, laws, and legal interpretations from court decisions, refining the prompts used in [7]. Second, we normalize the references to legal provisions so they can be treated as unified nodes. Finally, we create new prompts that link the extracted elements according to the legal ontology, forming structured connections from facts to legal rules.

Using the constructed LKG, we evaluate performance on a novel class of legal search queries that require factual circumstances to be connected with relevant legal provisions. As shown in Figure 1, conventional legal search engines typically begin with legal provisions and commentaries, and rely heavily on expert intuition to infer the relevant facts. In contrast, our method starts from the observed facts and traces the structured reasoning paths that link them to the applicable legal norms and provisions.

We compare our approach against several baselines, including LLMs and RAG systems that do not incorporate structured knowledge graphs. Our results show that, based on evaluation by expert human judges, the proposed method significantly outperforms these alternatives in both precision and recall.

To summarize, our contributions are as follows:

- We define a schema for representing structured legal reasoning, capturing how facts, provisions, legal norms, and legal applications are connected within court judgments.
- We construct an LKG from district-level court decisions in Japanese administrative litigation. The LKG explicitly encodes the factual elements, legal norms, and reasoning paths found in real-world judgments.
- We develop a novel legal search task and evaluation framework to assess how effectively our system replicates and explains the reasoning patterns observed in expert judicial decisions.
- All code, data, the legal ontology, and the LKG are available as open-source resources to support future research in legal knowledge modeling and reasoning¹.

2 Method: LKG Construction and Legal Search

Our main methodological contribution is a structured pipeline for constructing an LKG from judgments in Japan. This pipeline captures the structure of legal reasoning. Unlike prior approaches that focus primarily on surface-level legal text or static ontologies, our method explicitly models the inferential relationships among facts, legal norms, legal applications, and provisions. This structured representation enables a fine-grained, interpretable legal search and supports tasks requiring multi-step legal inference. The methodology consists of three main components: (i) schema design, (ii) node and edge extraction for constructing the LKG, and (iii) provision retrieval using the resulting graph.

2.1 Schema Design

The LKG schema is designed to capture legal entities and the inferential structure underlying legal reasoning. Many existing ontologies encode legal entities such as statutes and court decisions, but few attempt to model the reasoning process itself, especially the intermediate steps involved in applying norms to facts. Our schema introduces a minimal, but highly expressive, set of classes and properties as a means of representing this structure in a faithful and interpretable way.

The core classes in our RDF-based schema include “Fact”, “Legal-Norm”, “LegalApplication”, and “Provision”. LegalApplication represents an explicit reasoning step that connects a legal norm to a fact. Each LegalApplication node has two outbound RDF properties: “appliesNorm”, which link to one or more LegalNorms, and “toFact”, which point to one or more Facts.

A key feature of our schema is its ability to represent the hierarchical structure of legal interpretation. Courts often derive higher-level conclusions by combining multiple factual elements. For instance, a person’s “ownership” of land and their “physical proximity” to a loud construction site may together support the conclusion that the person has “standing” to file a lawsuit. Similarly, abstract legal norms can be inferred from more specific statutory obligations, such as deriving a “duty to prevent specific risks” from a general “duty to ensure safety”. To support these forms of reasoning, our graph allows connections not only between different categories (such as from LegalNorm to Fact), but also within the

same category (such as from Fact to Fact or from LegalNorm to LegalNorm), enabling the modeling of layered legal reasoning.

This explicit modeling of reasoning steps distinguishes our schema from existing legal ontologies, which have traditionally focused on vocabularies and legal references while omitting the logical structure that connects them [6, 10]. Our approach treats these inferential steps as first-class elements in the graph, enabling downstream tasks such as a traceable legal search and fine-grained legal reasoning. An illustration of our schema, written in JSON-LD format, is shown in Figure 2. The complete schema is available online.

```

...{
  "@id": "LKG:LegalApplication",
  "@type": "owl:Class",
  "rdfs:label": "Legal Application",
  "rdfs:comment": "A reasoning step that applies a legal norm to a fact",
  "rdfs:subClassOf": [
    { "@id": "LKG:LegalNode" },
    { "@id": "schema:Action" }
  ],
},
{
  "@id": "LKG:appliesNorm",
  "@type": "rdf:Property",
  "rdfs:label": "applies norm",
  "rdfs:domain": "LKG:LegalApplication",
  "rdfs:range": "LKG:LegalNorm"
},
{
  "@id": "LKG:toFact",
  "@type": "rdf:Property",
  "rdfs:label": "to fact",
  "rdfs:domain": "LKG:LegalApplication",
  "rdfs:range": "LKG:Fact"
},
...

```

Figure 2: JSON-LD representation of a legal application and its properties, compatible with the schema.org [12] vocabulary.

2.2 Node and Edge Extraction for LKG

In constructing the LKG, our method for node and edge extraction consists of three steps: (i) preprocessing, (ii) node extraction, and (iii) edge construction. Each step incrementally transforms the raw court judgment text into a structured representation that is aligned with our LKG schema.

(i) Preprocessing. We focus on district court judgments from Japan, using HTML versions provided in the NII legal dataset [9] (one of the most widely used legal resources in the country). Although the HTML structure appears machine-readable at first glance, it is often inconsistent because of artifacts introduced by automatic generation tools.

The problem of malformed or inconsistent tags and nesting means that we had to manually traverse the HTML node tree and

¹GitHub repository URL will be added upon publication.

... extract the portions corresponding to the labels below ...

- **Judicial Fact:** Case-specific facts derived from submitted evidence or undisputed matters. Do not include legal judgments or interpretations.
Example: “Zalk Noxis was born on Saturn, but currently resides on Mars.”
- **Citation or Reference to Law, Ordinance, Regulation, or Contract:** Only the article/paragraph number or cited source itself. Do not include surrounding text.
Example: “Articles 1 and 2 of the Law on Coexistence with Martians.”
- **Legal Norm:** The rule presented by law itself. In particular, when the purpose or objective of a law leads to a generalized standard (norm), that portion should be labeled as “Legal Norm”, even if it explicitly discusses the law’s purpose.
Example: “The Law on Coexistence with Martians aims to promote harmony between humans and extraterrestrial life in light of frequent clashes. Therefore, a ‘Martian’ (Article 2) refers to any non-human lifeform related to Mars.”
- **Legal Application:** Application of a law or legal norm to specific facts. When a judgment is made by applying the norm or provision to a concrete situation.
Example: “Because Zalk Noxis currently lives on Mars and is non-human, he qualifies as a ‘Martian’.”

...

Output must be in JSON format ...
Case Overview: [CASE OVERVIEW TEXT]
Excerpt from Judgment: [JUDGMENT TEXT]

Figure 3: Prompt for node extraction (Japanese original available on GitHub).

apply regex-based rules to identify and extract meaningful content blocks². Headings are detected using a combination of HTML tags, textual features (such as the absence of final punctuation or the presence of initial parentheses), and length-based heuristics. These signals allow the document to be segmented into semantically coherent parts.

In our approach, the section-level structure and source provenance are preserved, as they enable section-aware prompting and maintain traceability to the original document. In particular, we treat the case overview as an independent, structurally significant section, since it is reused in downstream tasks to improve contextual coherence and justification in GPT-based inferences later on.

(ii) Node Extraction. We experimented by extracting graph nodes from segmented court judgments using GPT-4o and Claude Sonnet 3.5. However, Claude did not perform reliably, so we opted to

²The original PDFs in the dataset are generated via OCR, and the HTML often contains noisy text, such as misaligned punctuation and character misrecognitions; for example, “四” (four) may be misread as “厶” (animal counter). The quality of the NII conversion is not always superior to our own reprocessed PDFs, but we rely on the NII HTML for consistency.

We now provide a list of laws and a list of legal norms. For each law, output the corresponding legal norm(s) in JSON format. Use laws as keys and legal norms as values. Do not output anything other than JSON.

Laws: Legal Norms: Original excerpt:

Example:
“The Law on Coexistence with Martians...” “‘Martians’ (Article 2) refers...”

Output:
{ "...Article 1": "The Law on Coexistence with Martians...", "...Article 2": "'Martians' refers to ..." }

Figure 4: Prompt for linking laws to legal norms.

use only GPT-4o. The prompt followed a structured format consisting of role instructions, node category definitions (e.g., **Fact**, **Legal Norm**, **Legal Application**), and prototypical examples. These were carefully aligned with our LKG schema.

To improve consistency, we included a *case overview* at the beginning of the prompt. This helped the model resolve cross-references and interpret the role of each sentence more effectively. We observed a significant improvement in accuracy, particularly in distinguishing between **Fact** and **Legal Application** nodes.

Providing few-shot examples typically improves performance [2]. However, instead of using real examples, we employed a one-shot fictional scenario based on a fabricated “Martian Law.” Because the scenario was fictional, it was easier to detect surface copying, such as the model using the word “Martian,” rather than genuine generalization, which is more difficult to assess with real law names. This design also provided two additional benefits: (1) it enabled a controlled demonstration of the desired reasoning process, and (2) it reduced the risk of overfitting to artifacts present in actual legal data. The Martian Law was crafted to reflect edge cases commonly encountered in real court judgments, with guidance from two legal experts.

Finally, statutory reference nodes were re-submitted to the model for normalization. Phrases such as “the Act” or “Article 2” were standardized into canonical legal titles to enhance downstream linking. For illustrative purposes, a snippet of the prompt is shown in Figure 3. The full prompt is available online.

(iii) Edge Construction. To capture the legal reasoning paths in court judgments, we again used GPT-based prompts. We focused on three types of directed relations that correspond to stages in legal interpretation:

- **Provision** → **Legal Norm:** capturing interpretive derivation from statutory text
- **Legal Norm** → **Legal Application:** linking abstract norms to legal conclusions
- **Fact** → **Legal Application:** grounding legal conclusions in concrete case facts

Language models are subject to input length constraints. Therefore, we partitioned the input contexts and used tailored prompting strategies for each edge type. For **Provision** → **Legal Norm** edges, we assumed local proximity within sections and paired “law” and

Below is a list of legal norms and a sentence of legal application. Determine which norm the application is based on and return the most appropriate one. Only establish a correspondence if the application clearly and logically relies on a specific norm. If no such relationship exists, return an empty list []. Do not link procedural descriptions or mere factual summaries. Respond strictly in JSON format; do not use any other formats.

Legal Norms:

...

(28) Even considering that the noise guidelines ... the train noise in such sections is reasonably assumed to be minimal.

Legal Application:

(1) The plaintiffs qualify as persons who may suffer ... due to the operation of the Nishi–Osaka extension line ...

Output:

```
{ "Application 1": ["Norm 28"] }
```

Figure 5: Prompt for linking legal norms to applications.

We provide a list of facts and a list of legal application sentences.

A sentence qualifies as a legal application **only if it satisfies both** of the following conditions:

- It is based on explicit legal norms (e.g., statutes, ordinances, ministerial regulations)
- It legally evaluates concrete facts (e.g., lawful, unlawful)

Do **not** link the following types of sentences to any facts. Return an empty list [] instead:

- Mere references to technical guidelines or standards
- Evaluations of party arguments (e.g., “The claim is groundless”)
- Summaries or organization of facts without legal judgment
- Explanations of legal frameworks or prerequisites without legal conclusions
- Descriptions of administrative procedures (e.g., approval by a committee) lacking legal evaluation

Facts:

...

(25) The plaintiffs reside near the Nishi–Osaka extension line and have filed complaints . . .

Applications:

(1) The plaintiffs qualify as persons who may suffer . . . due to the operation of the Nishi–Osaka extension line.

(2) Therefore, Plaintiff 51 is not eligible to bring the lawsuit.

Output:

```
{ "Application 1": ["Fact 25"], "Application 2": [] }
```

Figure 6: Prompt for linking facts to legal applications.

“norm” nodes from the same block. The model, prompted as a legal assistant, returned matching pairs in structured JSON format.

In contrast, “Legal Norm → Legal Application” and “Fact → Legal Application” links often span distant sections (sometimes more than eight segments apart). To address this, we adopted a scoped-history prompting strategy: each Legal Application node was paired with all preceding Legal Norm or Fact nodes, allowing the model to reason over the accumulated context. This setup emulates how legal arguments unfold in court judgments, where conclusions frequently depend on reasoning steps distributed across the document.

We validated this design choice by analyzing the section-wise distances between related nodes and found that, in some cases, the distances can be quite large. This finding reinforces the need for a history-aware approach over fixed-size local windows. Although the full analysis is omitted here for brevity, the result supports the modular structure of our edge construction strategy. This modular design enables LLMs to capture multi-step reasoning processes, providing a structured foundation for interpretable and reusable legal inference.

2.3 Legal Search with LKG

The final component of our methodology defines legal provision retrieval as the task of identifying relevant statutory articles for a given fact extracted from a court judgment. Our method performs this retrieval by leveraging the LKG described in the previous sections. Each Fact node is embedded using the text-embedding-3-large model and indexed with Annoy [1] for an approximate nearest-neighbor search. Embeddings are generated at the level of individual fact sentences, enabling fine-grained retrieval of legally similar cases. For each query fact, the system retrieves the top-*k* most similar facts and collects the provisions linked to them. To avoid trivial matches, the query fact itself is excluded from retrieval by masking its outgoing edges, referred to as the **Fact-Masked** setting.

For comparison, we implemented three baselines using GPT 4o. In the **GPT Simple** setting, the model receives only a single fact sentence and is prompted to return relevant legal provisions. This setup tests whether a pretrained model can identify applicable statutes based on minimal input.

The **GPT With Context** baseline augments the input with a case overview excerpted from the same judgment. These overviews, which summarize the involved parties and legal issues, are a common feature of Japanese court decisions. By including this context, the model is given a broader legal frame in which to interpret the fact.

The **GPT With RAG** configuration incorporates RAG. The input fact is embedded using text-embedding-3-large, and semantically similar segments are retrieved from the judgment corpus. Unlike our LKG-based method, which embeds individual fact sentences, this baseline uses **larger textual units**, such as full judgment sections containing multiple sentences. The retrieved segments are concatenated with the input fact and passed to GPT 4o, allowing the model to generate provision predictions based on retrieved evidence.

To ensure a fair comparison, all baselines were prompted to return relevant statutory provisions in a standardized JSON format.

Table 1: Evaluation of GPT with designed prompts for node extraction in LKG construction, as assessed by legal experts

Node Type	TP	FP	FN	Precision	Recall	F1
Provision	119	0	0	1.0000	1.0000	1.0000
Norm	64	8	1	0.8889	0.9846	0.9347
Application	60	0	9	1.0000	0.8696	0.9302
Fact	193	1	1	0.9948	0.9948	0.9948

This consistent output format enabled systematic evaluation across all configurations.

2.4 Dataset

We focus on judgments issued after 2008, following the introduction of a new document formatting standard in Japan. These judgment documents typically include structured sections such as the main ruling, an overview of the case, and the reasoning of the court. Among these components, we exclusively target the “reasoning” section for our analysis because it provides rich textual material connecting facts and legal norms.

We further limit our scope to District Court civil cases, focusing on administrative matters. These judgments offer more detailed factual and legal reasoning than those from higher courts. This choice is driven by practical considerations, because covering all areas of civil law would incur prohibitive computational costs. More importantly, administrative litigation presents favorable conditions for analyzing legal reasoning, typically revolving around the interpretation and application of law rather than fact-intensive disputes.

3 Evaluation

This section evaluates whether the proposed LKG effectively supports legal search and retrieval. We focus on three aspects: (1) the accuracy of the extracted graph components, (2) whether the resulting knowledge graph exhibits a coherent structure, and (3) how well the LKG enables provision retrieval from factual inputs.

3.1 Extraction Accuracy

We evaluated the accuracy of GPT’s node and edge extractions, prompted using our designed instructions, through a manual annotation study conducted with two legal experts. The experts reviewed 57 randomly selected sections of court judgments and annotated nodes corresponding to the four categories in our LKG: **Fact**, **Provision**, **Legal Norm**, and **Legal Application**. These categories follow the same definitions used in our prompt design: Fact refers to case-specific factual findings without legal interpretation; Provision refers to explicit references to statutes, regulations, or contracts; Legal Norm refers to normative propositions derived from the law; and Legal Application refers to the application of a legal norm or law to a specific fact.

Each GPT output was compared with the human annotations, and we recorded the numbers of true positives (TPs), false positives (FPs), and false negatives (FNs) per category. Table 1 indicates that, overall, the model achieved high precision and recall, with perfect scores for Provision and near-perfect performance for Fact. Recall

Table 2: Evaluation of GPT with designed prompts for edge extraction in LKG construction, as assessed by legal experts

Edge Type		TP	FP	FN	Precision	Recall	F1
Provision	→	135	8	7	0.9434	0.9507	0.9470
Norm							
Norm	→ Application	78	18	5	0.8125	0.9398	0.8710
Application	→	105	5	8	0.9545	0.9292	0.9417
Fact							

was slightly lower for Legal Application, primarily because of frequent misclassification as Legal Norm. In these cases, segments that should have been identified as Legal Application were instead labeled as Legal Norm by the model. This error pattern aligns with our prompt design, which reflects the domain knowledge that Legal Application generally presupposes the presence of a corresponding Legal Norm. Therefore, we do not regard this behavior as an extraction failure, but consider it consistent with the structure the prompt was designed to elicit.

After evaluating node extraction, we proceeded to assess the quality of edge extraction. The legal experts manually reviewed the predicted edges across 92 randomly selected judgment sections. For each edge, they assigned one of three labels: TP, FP, or FN. An edge was labeled as TP if it correctly connected the intended pair of nodes with the appropriate relation type. Edges that linked incorrect node pairs or used the wrong relation type were labeled as FP, while necessary edges that had been omitted were labeled as FN.

According to Table 2, GPT demonstrated strong overall performance in edge extraction, achieving precision above 0.94 for most edge types. In the case of *Legal Norm* → *Legal Application*, a higher number of FP cases occurred. These were mainly the result of two factors: overextension of *Legal Norm* spans, and missing links to supporting norms in layered legal reasoning structures. For example, GPT often captured the primary norm applied to a fact, but failed to link to additional norms that support or constrain the main one. This limitation reflects a structural constraint in our current schema, which does not support edges between norms (e.g., *Legal Norm* → *Legal Norm*).

From the perspective of a legal retrieval system, minimizing omissions and ensuring dense graph connectivity are important. In this respect, the model’s output is largely desirable. In particular, the capture of hierarchical reasoning structures in which legal norms and facts are layered across multiple levels will be enhanced by expanding the edge schema to support more nuanced connections.

3.2 Structural Analysis of the LKG

Subsequently, we assess whether the overall structure of the constructed LKG reflects desirable properties for downstream legal reasoning. Specifically, we analyze whether the graph exhibits hierarchical composition, element reuse, and localized modularity. Tables 3, 4, and 5 provide the relevant statistics.

At the node level, *Legal Application* nodes tend to have high in-degree and low out-degree, consistent with their role as endpoints of reasoning. *Provision* and *Fact* nodes more often serve as sources.

Table 3: Graph statistics by node class

Node Type	# Nodes	Average In-Degree	Average Out-Degree	# Self-Loops
Provision	10,552	0.00	1.70	0
Norm	15,445	1.17	0.95	25
Application	8,957	3.80	0.13	746
Fact	11,199	0.18	1.70	472

Table 4: Graph statistics by edge type

Edge Type (Source → Target)	# Edges	Average Multiplicity
Provision → Norm	18,814	1.69
Norm → Application	15,356	2.20
Fact → Application	19,242	1.63

Table 5: Graph statistics at the network level. WCC stands for weakly connected component.

Metric	Value
Number of Nodes	44,447
Number of Edges	51,296
Number of WCCs with diameter ≥ 2	1,478
Avg. Diameter of WCCs (≥ 2)	3.10
Std. Dev. of Diameter	2.39
Std. Dev. of Degree	5.28
Graph Density	2.59×10^{-5}

Legal Norm nodes act as intermediaries. Self-loops occur primarily in *Fact* and *Application* nodes, suggesting that certain segments fulfill multiple roles.

At the edge level, we observe frequent reuse of *Fact* and *Norm* nodes across applications. This pattern indicates that conclusions are often supported by multiple inputs and that key legal elements participate in multiple reasoning chains.

At the global level, the graph is sparse and fragmented, consisting of many weakly connected components. This structure reflects the document-level nature of construction: each judgment forms an isolated subgraph. The graph’s lack of strong connectivity is consistent with the forward-only, acyclic flow of legal logic.

3.3 Evaluation as Legal Search Engine

Finally, we evaluate the effectiveness of LKG-based retrieval using both quantitative and qualitative methods. First, we compare retrieval accuracy against GPT-based baselines. Then, we analyze representative reasoning paths to assess whether the system captures judicial logic beyond surface-level matching. We evaluated LKG-based retrieval in comparison with three GPT-4o-based baselines introduced in Section 2.3. The benchmark dataset was derived from Japanese administrative litigation judgments, where each fact sentence was linked to gold-standard provisions via the LKG.

Table 6 compares performance across methods. We report predicted provisions (Pred), true positives (TPs), and the resulting

precision, recall, and F1 scores under both macro and micro evaluation metrics. Macro scores treat each case equally by averaging over judgments, whereas micro scores aggregate globally across predictions. **LKG Retrieval under the Fact-Masked setting significantly outperforms all GPT-based methods in both macro and micro recall.** For example, at $k = 3$, the macro recall is 0.311 and the micro recall is 0.667, far exceeding GPT With Context (0.137 and 0.271, respectively). Although $k = 2$ yields slightly higher micro recall (0.671), $k = 3$ provides better macro recall and more stable performance. Full results for $k = 1-7$ are provided in the appendix. Among GPT baselines, GPT With Context outperforms GPT Simple, highlighting the value of case overviews. However, **GPT With RAG fails to bridge the gap with LKG Retrieval**, despite access to retrieved context. This suggests that unstructured retrieval lacks the structural alignment required for multi-step legal reasoning.

While these quantitative results validate the overall performance of LKG-based retrieval, they do not reveal how well the system captures the logic of legal reasoning. To address this, we turn to qualitative evaluation. We examine complete reasoning paths extracted by the system and assess whether they faithfully mirror judicial interpretive structure. The following examples demonstrate how the LKG supports both straightforward and nuanced inference patterns.

We begin with a simple case that the system correctly extracted without manual correction. The court addressed a resident audit request governed by *Article 242 of the Local Autonomy Act*. It examined a **Fact** where *the request was filed on June 1, 2007*, and concluded in its **Application** that *the request was inadmissible due to expiration of the statutory deadline*. This conclusion rested on a **Norm** specifying that *the deadline must be calculated from the date of the boundary confirmation consultation or its completion*. This example shows that the LKG correctly linked the fact, norm, and statutory provision to form a legally valid path.

More complex examples reveal how factual variation can yield different legal outcomes even under the same norm and provision. In urban planning disputes, one case presented a **Fact** where *the plaintiffs lived, worked, or owned property near the project site and claimed harm to their environment and property use*. The court held that they had standing to seek cancellation, applying a **Norm** that *only those with a legally protected interest may file such a suit*, grounded in *Article 9 of the Administrative Case Litigation Act*. We also identified a contrasting case in which the plaintiff resided *556 meters from the site*, a distance the court found insufficiently direct, leading to an **Application** denying standing. These cases illustrate how minor factual differences can materially affect legal reasoning.

Finally, we highlight a different type of variation, where legal norms interact with constitutional principles. In one case, the **Fact** was that *the plaintiff acquired the nationality of Country A through a document submitted by their parents, which met the legal requirements under the former nationality law of Country A*. The court deemed this acquisition voluntary and applied the **Norm** that *a person who voluntarily acquires a foreign nationality loses Japanese nationality*, pursuant to *Article 11 of the Nationality Act*. The LKG-based retrieval system also surfaced a structurally corresponding case in which the same norm and provision were challenged on constitutional grounds. In that case, the **Fact** was that *the plaintiffs expressed their intent to acquire the nationalities of Country B and*

Table 6: Evaluation results for GPT baselines and LKG retrieval ($k = 1-7$). The total number of gold-standard labels is 1242. Bold indicates the best value in each column. Underline indicates GPT outputs and k settings in LKG retrieval that yield a similar number of outputs.

Method	Pred	TP	Macro Recall	Micro Recall	Macro Precision	Micro Precision	Macro F1	Micro F1
GPT Simple	4922	123	0.046	0.099	0.037	0.025	0.041	0.040
GPT With Context (baseline)	<u>7377</u>	336	0.137	0.271	0.054	0.046	0.077	0.078
GPT With RAG	<u>7792</u>	436	0.176	0.351	0.062	0.056	0.092	0.097
LKG Retrieval (Fact-Masked, $k = 1$)	1253	497	0.203	0.400	0.199	0.397	0.201	0.398
LKG Retrieval (Fact-Masked, $k = 2$)	2403	712	0.278	0.573	0.235	0.296	0.255	0.391
LKG Retrieval (Fact-Masked, $k = 3$)	3470	829	0.311	0.667	0.227	0.239	0.263	0.352
LKG Retrieval (Fact-Masked, $k = 4$)	4415	887	0.330	0.714	0.213	0.201	0.259	0.314
LKG Retrieval (Fact-Masked, $k = 5$)	5423	920	0.340	0.741	0.194	0.170	0.247	0.276
LKG Retrieval (Fact-Masked, $k = 6$)	6464	953	0.354	0.767	0.179	0.147	0.238	0.247
LKG Retrieval (Fact-Masked, $k = 7$)	<u>7482</u>	965	0.359	0.777	0.161	0.129	0.222	0.221

Country C, and filed lawsuits claiming that Article 11 violated their freedom to change nationality. The court rejected the claim, applying the **Application** that the provision did not exceed legislative discretion and was consistent with the constitutional guarantee of nationality renunciation, affirming the statute’s aim of preventing dual nationality.

4 Conclusion

We have introduced a structured LKG that captures the multi-step reasoning paths embedded in Japanese administrative court judgments. By modeling the connections between facts, legal norms, and legal applications, our LKG makes the legal reasoning explicit and machine-interpretable. To evaluate the utility of this LKG, we designed a legal provision retrieval task and compared the performance of GPT- and graph-based methods. The results show that structured legal knowledge enables more accurate and consistent reasoning than language models alone, even when they have access to relevant source content. These findings demonstrate that legal AI systems benefit from explicitly grounded knowledge structures, reaffirming the essential role of LKGs in interpretable and generalizable legal reasoning.

Acknowledgments

We thank TKC Corporation for helpful discussions. This work uses judicial precedent data from the NII Informatics Research Data Repository. R.K. is funded by JST, ACT-X JPMJAX23CA. R.H. is supported by the JST FOREST Program (JPMJFR216Q), JST PRESTO Program (JPMJPR2469), a Grant-in-Aid for Scientific Research (KAKENHI) (JP24K03043), and the UTEC-UTokyo FSI Research Grant Program.

References

- [1] Erik Bernhardsson and Spotify. 2015. Annoy: Approximate Nearest Neighbors in C++/Python. <https://github.com/spotify/annoy>. Accessed: 2025-07-29.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, Vol. 33. 1877–1901.
- [3] Ilias Chalkidis, Antonis Mylonas, Peter Fergus, and Nikolaos Aletras. 2024. LEXFRAG: Legal Evidence Fragment Retrieval to Support Statutory Reasoning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*.
- [4] Miguel Garcia-Godinez. 2024. A Deflationary Approach to Legal Ontology. *Synthese* 203, 2 (Feb. 2024), 64.
- [5] Neel Guha, Julian Nyarko, Daniel Ho, Christopher Ré, Adam Chilton, Aditya K. Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory Dickinson, Haggai Porat, Jason Hegland, Jessica Wu, Joe Nudell, Joel Niklaus, John Nay, Jonathan Choi, Kevin Tobia, Margaret Hagan, Megan Ma, Michael Livermore, Nikon Rasumov-Rahe, Nils Holzenberger, Noam Kolt, Peter Henderson, Sean Rehaag, Sharad Goel, Shang Gao, Spencer Williams, Sunny Gandhi, Tom Zur, Varun Iyer, and Zehua Li. 2023. LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 44123–44279.
- [6] Rinke Hoekstra, Joost Breuker, Marcello Di Bello, and Alexander Boer. 2007. The LKIF Core Ontology of Basic Legal Concepts. In *Proceedings of the 2nd Workshop on Legal Ontologies and Artificial Intelligence Techniques (LOAIT 2007) (CEUR Workshop Proceedings, Vol. 321)*. 43–63.
- [7] Ryoma Kondo, Tomohiro Watanabe, Takahiro Yoshida, Kazuyuki Yamasawa, and Ryohei Hisano. 2024. Collaborative System Synergizing Human Expertise and Large-Scale Language Models for Legal Knowledge Graph Construction. In *Proceedings of the ISWC 2024 Posters, Demos and Industry Tracks*, Vol. 3828. CEUR-WS.org, 53–56.
- [8] John Leubsdorf. 2001. The Structure of Judicial Opinions. *Minnesota Law Review* 86 (2001), 447–500.
- [9] Research Organization of Information and Systems, Center for Juris-Informatics. 2024. HTML Data of Japanese Judicial Precedents. Informatics Research Data Repository, National Institute of Informatics. doi:10.32130/rdata.7.1 Dataset: <https://doi.org/10.32130/rdata.7.1>.
- [10] Neda Sakhaee and Mark C. Wilson. 2021. Information Extraction Framework to Build Legislation Network. *Artificial Intelligence and Law* 29, 3 (2021), 361–385.
- [11] Ruben E. Rodriguez Samudio. 2023. Japanese Law and Social Media.
- [12] Schema.org. 2011–. Schema.org. <https://schema.org>. Accessed: 2025-07-29.
- [13] Jan-Reinard Sieckmann. 2020. The Dual-Nature Thesis: Which Dualism? *Ratio Juris* 33, 3 (2020), 271–282.
- [14] Joseph William Singer. 2023. Conflict of Abortion Laws. *SSRN Electronic Journal* (2023).
- [15] S.I. Strong. 2015. Writing Reasoned Decisions and Opinions: A Guide for Novice, Experienced, and Foreign Judges. *Journal of Dispute Resolution* (2015), 93–109.
- [16] Rujing Yao, Yang Wu, Chenghao Wang, Jingwei Xiong, Fang Wang, and Xiaozhong Liu. 2025. Elevating Legal LLM Responses: Harnessing Trainable Logical Structures and Semantic Knowledge with Legal Reasoning. In *Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Mexico City, Mexico, 5630–5642.
- [17] Xinyu Zheng, Md Tahmidul Islam, Kelvin Guu, and Jason Wei. 2024. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks: A Survey. *Transactions of the Association for Computational Linguistics* (2024).