adorsys
Technology Innovation

Case Study

THEMIS: AI-Assisted Legal Knowledge Graph for International Law Firms

Tool for Handling, Extracting, Mapping & Intelligent Structure-analysis

Augmented AI | 2025

Case Study Summary

What was the challenge?

Manual legal document analysis takes 5-8 hours per case[7], provides no structured querying capability, and existing AI tools offer no audit trail linking answers back to source text, making them unusable in professional legal environments.

What did we build?

THEMIS: an AI-powered legal knowledge graph platform that converts unstructured lawsuit documents into structured, queryable intelligence with full provenance back to the original document. Built on a purpose-designed knowledge graph stack, deployed to Azure in 4 months.

Key Results

↓ 85–94 % reduction in analysis time per document
↓ Cost per document reduced from €905–€2,208 to ~€90–€140

- 448 entities and 1,630 relationships extracted automatically from a 53-page GDPR lawsuit
- Average semantic similarity: 0.846 (above the 0.8 quality threshold)
- Estimated annual savings: €75,000–€205,000 per firm (100 docs/year)

Why it Matters

AI becomes usable in legal workflows only when every result is traceable to its source. THEMIS builds that trust chain by design where every answer links to the exact passage in the original document, satisfying the auditability standard and professional legal practice demands.

Executive Summary

Large international law firms manage hundreds of complex lawsuit documents annually — dense, multi-party instruments that may span 150 to 300 pages and require 5 to 8 hours of manual analysis per case at billing rates between €181 and €276 per hour[1]. A leading international law firm with operations across more than 30 countries partnered with adorsys to explore whether AI could fundamentally transform this process. Over four months, we built THEMIS (Tool for Handling, Extracting, Mapping & Intelligent Structure-analysis): an AI-powered knowledge graph platform that converts unstructured legal documents into structured, queryable intelligence with full source traceability. Processing an anonymised 53-page lawsuit document through a fully automated pipeline, THEMIS extracted 448 entity nodes and 1,630 relationships — validated against 10 competency questions across three difficulty levels, achieving an average semantic similarity of 0.846. The client described the outcome as having “enormous potential to disrupt the market.”

At a Glance: Manual Process vs. THEMIS

The following table provides a direct comparison across the key dimensions that matter to a law firm evaluating AI-assisted legal document intelligence.

Dimension	Dimension	Dimension
Analysis time per document	5–8 hours (associate + senior review)	~30 minutes (structured output review)
Cost per document	€905–€2,208 (€181–€276/hr)	~€90–€140
Annual cost — 100 documents	€90,500–€220,800	~€15,000 (incl. infrastructure)
Structured querying	Full-text search only	Graph-structured Q&A with semantic retrieval
Evidence traceability	Manual cross-referencing	Automatic; every answer links to source passage
Audit trail for AI insights	None	Full provenance chain (Art. 5(2) GDPR-ready)
Scalability	Limited by headcount	Parallel, document-volume-independent
Cross-case knowledge reuse	None; per-case notes only	Queryable graph across all ingested cases
Processing speed (automated pipeline)	N/A	~39 min / 53-page document (fully automated)
Estimated efficiency gain	—	85–94 % reduction in analysis time

1. Project Overview

A leading international law firm with a presence across more than 30 countries and over 1,600 legal professionals — serving clients across financial services, technology, life sciences, and digital industries — recognised a growing tension in its practice. Despite its global scale and strategic focus on innovation, its legal teams faced a fundamentally manual workflow when processing lawsuit documents: reading, extracting parties, mapping claims to evidence, and identifying relevant legal foundations, all carried out page by page, without structured tooling.

As caseloads grew — Germany's civil and labour courts alone processed over 1.45 million new cases in 2024[3][4] — the firm sought to understand whether modern AI technologies could meaningfully accelerate this process without sacrificing the rigour and auditability[6] that legal practice demands. They approached adorsys with a focused goal: validate the feasibility of an AI-powered legal document intelligence system, purpose-built for professional legal workflows.

Figure 1. Themis, ancient Greek goddess of law, impartial justice and a symbol used in courthouses throughout the world, namesake of the THEMIS platform

2. Challenge / Problem Statement

The core challenge was a lack of uniform structure across different lawsuit documents from different sources. Lawsuit documents (statements of claim) are dense, multi-party legal instruments that may run to 40–80 pages, and depending on the type of document, spanning to a maximum of 200-300 pages, each containing layered claims, factual assertions, evidentiary references, legal foundations, and case timelines all interwoven in free-form prose with document structure differing between different parties and a law firm's own personal practices. Lawyers had no tool to navigate this structure other than manual reading and full-text search.

Three pain points defined the challenge:

1 Time intensity.

Initial analysis of a single lawsuit document required 5 to 8 hours of associate time, at billing rates of €181–€276 per hour[1], a cost that compounded significantly across a firm's entire caseload.

2 No structured querying.

Without a structured representation of a document's contents, finding specific claims, parties, or evidence relationships meant reading the entire document from start to finish. Without machine-readable, formalised legal structure, Law as Code approaches — which encode statutory obligations as directly executable, computable logic — cannot be applied to the specific facts of individual cases.

3 No audit trail for AI-generated insights.

Existing general-purpose AI tools provided answers without linking back to source text, making them unsuitable for legal practice where verifiability is non-negotiable.

> The scale of the addressable problem is significant. Germany's labour courts alone received approximately 330,000 new cases in 2024[4][5], with employment termination protection claims accounting for around 23.7 % of all labour court proceedings — approximately 78,000 cases per year[5] . Civil courts added a further 1.12 million new cases across district and regional courts in 2024[3]. For firms specialising in labour or civil litigation, this translates to a high-volume, document-heavy workflow that existing tools could not adequately support.

3. Solution Approach

adorsys approached the challenge with a lean, research-first methodology. Rather than immediately building a system, the team began by working closely with the client's domain expert to understand exactly what questions a lawyer asks when reading a lawsuit document and to formalise these as competency questions: a structured set of queries that any useful system must be able to answer correctly. Because real-world documents could not be annotated with ground truth labels within the confidentiality constraints of the engagement, a representative set of synthetic test documents was purpose-built for this validation. This framework served as both the design anchor and the evaluation benchmark throughout the project.

Why a graph? Legal documents have hidden graph structures. Every statute follows an internal schema — paragraphs, sub-sections, and clauses — in which legal rules are expressed through conditional references to other elements: a right under Art. 15 GDPR triggers only when the conditions of Art. 22 are satisfied; a damages claim under Art. 82 depends on the chain of proof linking a specific processing act to a documented harm. These dependencies are not decorative — they are the legal reasoning itself, and they are invisible to any system that treats a lawsuit as a flat sequence of words.

Making these relationships first-class citizens — named, typed edges between named nodes — transforms a legal document into something fundamentally more powerful: a structured database that can be queried. Instead of asking a language model to "read" the document and hope it remembers the right facts, THEMIS encodes the legal content as a knowledge graph in which every claim, every statutory reference, every piece of evidence, and every party is a queryable node. A lawyer can then ask precise questions about structure ("which claims cite Art. 82?"), about evidence chains ("what proof supports the damages claim?"), or about temporal logic ("which events preceded the termination notice?") — and receive answers grounded in the exact passages of the original document.

From this foundation, the team iteratively designed a domain-specific legal ontology, a formal schema of 13 entity types (including lawsuits, claims, factual assertions, evidence, persons, events, and legal statutes) and 23 relationship types that capture the semantic structure of legal documents as a knowledge graph. The ontology was revised iteratively, tested against anonymised lawsuit documents, and refined until it accurately represented real-world legal reasoning patterns.

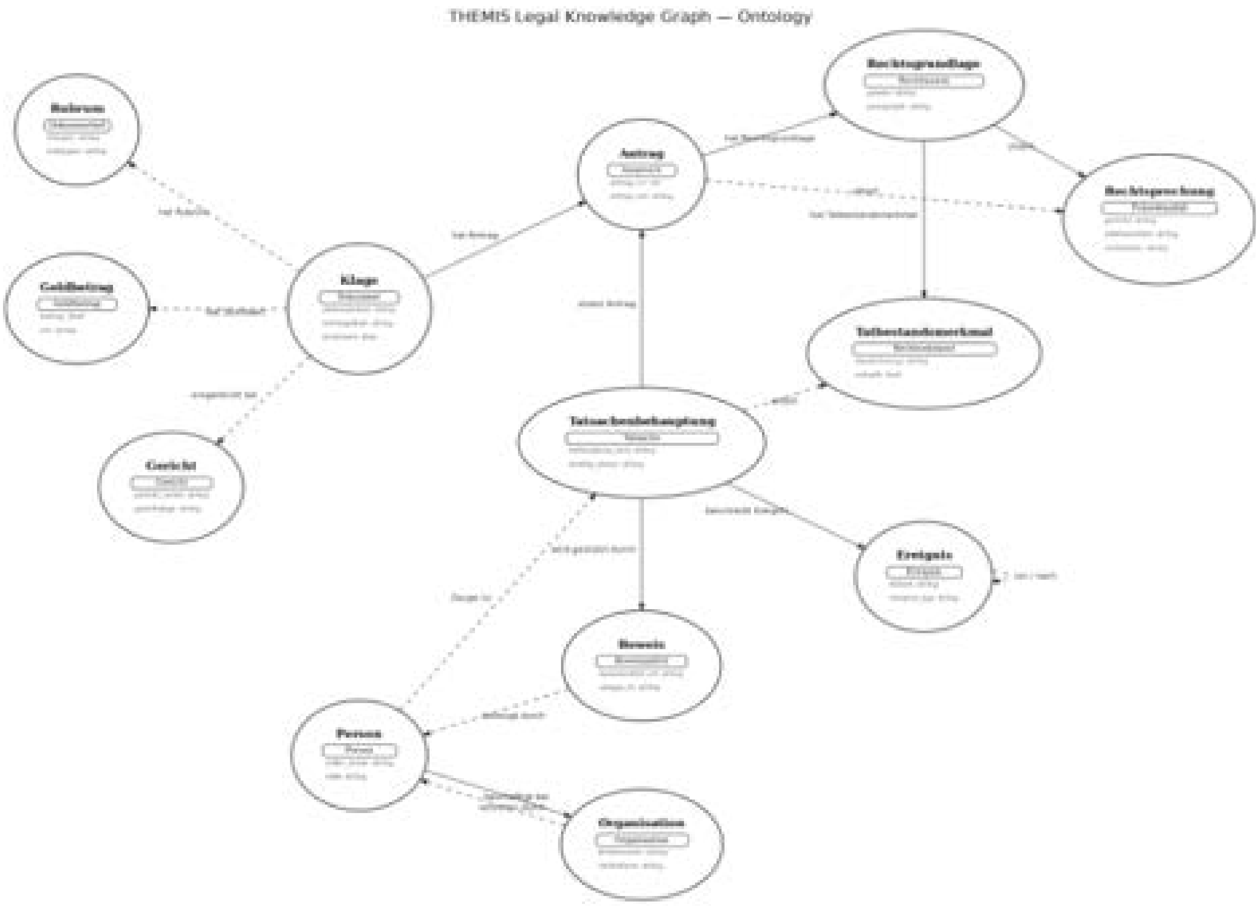

Figure 2. THEMIS legal knowledge graph ontology: 13 entity types and 23 relationship types. Main attributes and selected relationships shown.

The ontology THEMIS built represents a critical step towards Law as Code (the emerging practice of encoding legal rules as machine-executable logic). By formalizing legal entities such as GDPR obligations (Art. 22, Art. 15, Art. 82), the relationships between them, and the 387 type constraints governing valid legal arguments, THEMIS produces the structured legal knowledge layer that frameworks such as LegalRuleML, Catala, and the SALI Legal Matter Specification Standard require as input. In the GDPR context, where statutory obligations are explicitly enumerated, bounded, and directly linked to sanctions, this foundation makes automated compliance reasoning tractable and positions THEMIS as a bridge between document intelligence and executable law.

With the ontology validated, the team built THEMIS: a full-stack knowledge graph platform that parses legal PDFs, extracts entities and relationships using large language models, stores them in a property graph database with full provenance back to source text, and exposes the result through a professional web interface. Lawyers can ask natural-language questions, view cited answers with highlighted evidence passages in the original PDF, explore interactive analytics dashboards, and navigate the case structure through an ontology mind-map. At every step, THEMIS maintains a complete trust chain: every answer links to a specific passage in the original document.

Tech Stack / Tools Used

Knowledge graph construction layer

LLM-powered entity extraction, relationship inference, and episodic memory

Property graph database

Structured storage of entities, relationships, and full provenance back to source text

Large language model APIs

Entity extraction, relationship inference, and answer synthesis

Web interface

Lawyer-facing query, evidence highlighting, analytics dashboards, and ontology mind-map

PDF parsing pipeline

Document ingestion with page-level metadata for evidence highlighting

Azure (App Service, Container Registry, Container Instances)

Cloud deployment and infrastructure

Containerised deployment

Portable, reproducible environment for local development and cloud delivery

Data processing and validation layer

Structured pipeline from raw PDF to queryable knowledge graph

4. Implementation

THEMIS was delivered over four months (July–November 2025) by a focused team of two adorsys engineers — one AI Engineer and one Product Owner — working alongside one technical domain expert from the client. The small team size was deliberate: tight collaboration and short feedback loops enabled rapid iteration, with the client's domain expert embedded throughout ontology design and validation.

5. Results & Impact

(This case study has been anonymised due to confidentiality requirements.)

5.1. Knowledge Graph Construction

THEMIS was applied to an anonymised 53-page GDPR lawsuit document (statement of claim against a credit-scoring agency, filed at a German regional court). The automated extraction pipeline, running without manual intervention, produced the following results:

Artefact	Count	Description
Entity Nodes	448	All nodes in the knowledge graph
Domain-typed Entities	270	Entities with an ontology-specific type label
Relationships	1,630	Semantic links between entities (with fact annotations)
Document Chunks	80	Source-linked episodes with full provenance
Entity Mentions	2,158	Episode-to-entity links preserving source context
Entity Types	13	Custom legal ontology types
Type Constraints	387	Valid source-target-relationship type triples

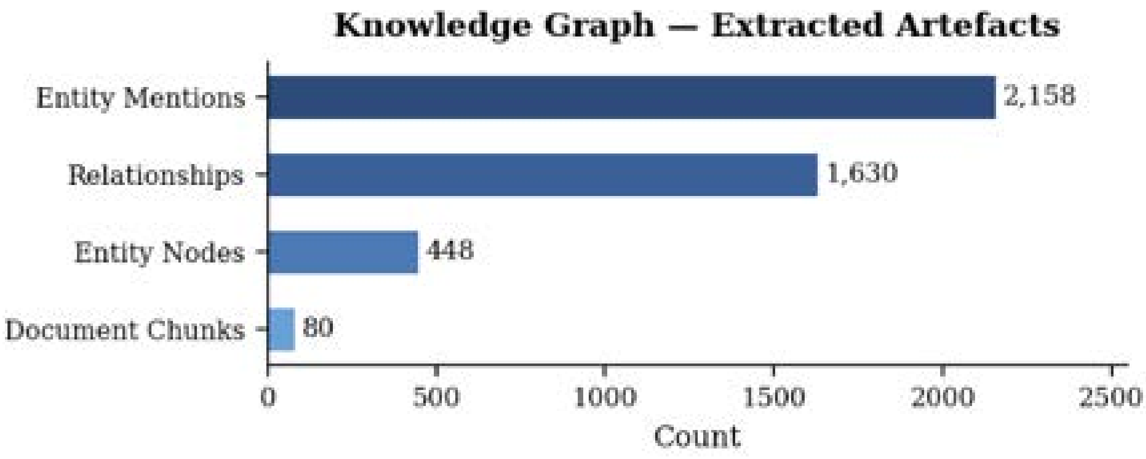

Figure 3. Knowledge graph artefacts extracted from the anonymised 53-page GDPR lawsuit document. The pipeline produced 448 entity nodes, 1,630 semantic relationships, and 2,158 entity mentions across 80 document chunks without any manual intervention, demonstrating end-to-end automated extraction from unstructured legal prose.

Entity extraction covered all 13 ontology types. Legal statutes (Rechtsgrundlage, 73 entities) dominated, reflecting the dense GDPR citation structure of the document. Organisations (50 entities) and case law references (Rechtsprechung, 39 entities) were also prominent, consistent with the regulatory and comparative-law argumentation typical of data protection litigation.

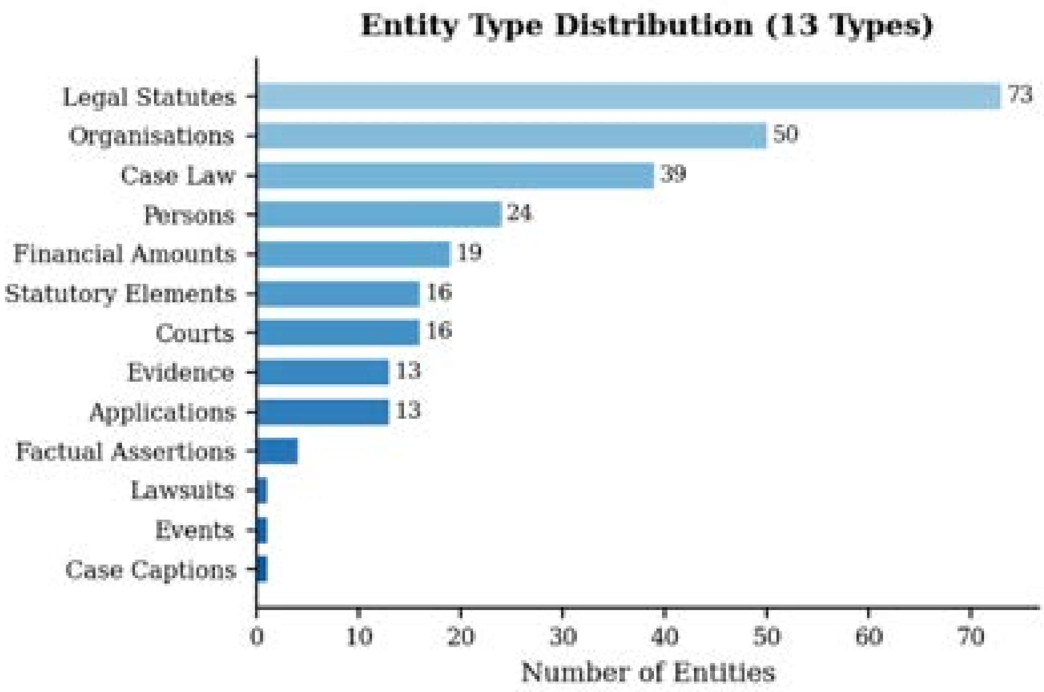

Figure 4. Distribution of extracted entities across the 13 legal ontology types (GDPR / credit-scoring lawsuit document). Legal Statutes dominate (73 entities), reflecting the dense GDPR citation structure of data protection litigation; Organizations (50) and Case Law references (39) follow, consistent with the regulatory and comparative-law argumentation of the claim.

5.2 Document Processing Performance

The full ingestion pipeline was measured end-to-end on the anonymised lawsuit document. Processing is fully automated — no manual steps between PDF upload and queryable knowledge graph.

Stage	Measured Time
Document	GDPR / Credit Scoring lawsuit (53 pages)
PDF Parsing	29.5s (OCR)
Text Chunking	1.02s (80 chunks)
Graph Ingestion	2308s (80 episodes)
Total Pipeline	2338s (~39 minutes)
Per-chunk (mean)	28.9s (min 3.3s / max 50.8s / p95 43.0s)

PDF parsing completed in 29.5 s. Graph ingestion — the dominant stage — required 2308 s (38 minutes) to process 80 text chunks through LLM-powered entity extraction pipeline, at a mean of 28.9 s per chunk. Per-chunk time varies with entity density: legal statute citations and multi-party relationships are extracted in a single LLM call per chunk, with latency driven by LLM API response time. A human analyst would typically spend 5 to 8 hours on initial analysis of a 53-page document of equivalent complexity.

5.3 Competency Question Validation

10 competency questions were defined collaboratively with the client's domain expert against the ingested lawsuit document, covering 39 expected answer triples (subject–predicate–object tuples representing a verified fact in the knowledge graph) grounded in [MK1] actual Neo4j graph nodes and relationships. Every triple was verified against the live graph before evaluation. Questions were classified by reasoning depth (single-hop = Easy: direct entity lookup; two-hop = Medium: requires one intermediate inference step; multi-hop = Hard: requires chaining two or more reasoning steps across the graph):

> Methodology note: Ground truth triples were derived directly from the knowledge graph built from the real (anonymised) lawsuit document. Each expected triple references named graph nodes and their RELATES_TO edges as they exist in Neo4j, ensuring that the evaluation reflects genuine extraction quality on production legal content rather than synthetic approximations.

Difficulty	Depth	Questions	Expected Triples	Examples
Easy	Single-hop	3 (30)	6	Who are the parties and at which court was the case filed?
Medium	Two-hop	5 (50%)	24	What GDPR rights does the plaintiff assert for the damages claims?
Hard	Multi-hop	2 (20&)	11	What is the full legal basis for challenging the automated credit scoring?
Total	—	10	39	—

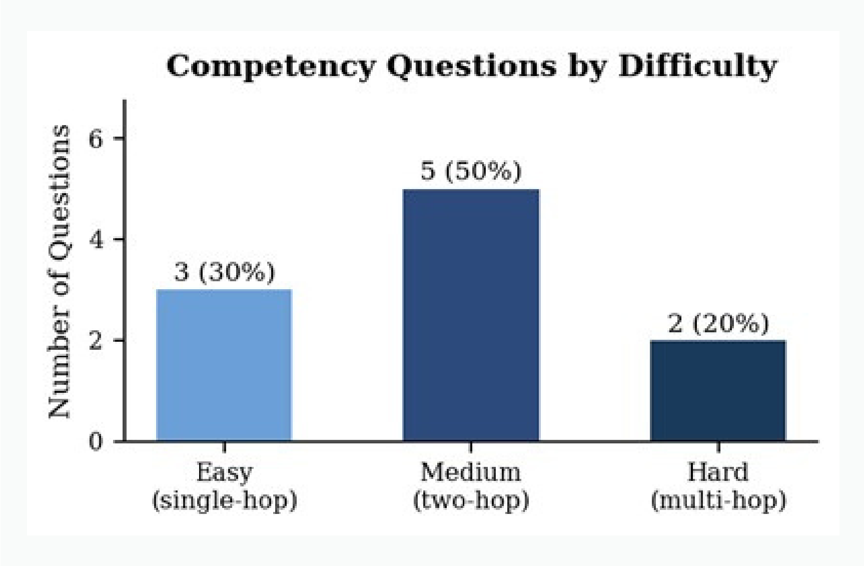

Figure 5. Distribution of the 10 competency questions by reasoning difficulty: single-hop (Easy), two-hop (Medium), and multi-hop (Hard) graph traversal. Difficulty mirrors the inferential depth a human lawyer must apply when reading the same document, from straightforward party identification to multi-step reasoning over injunction grounds and legal challenge frameworks.

The questions span the full legal structure of the GDPR / credit-scoring complaint: party identification and court filing (Easy), GDPR rights, scoring products, discriminatory factors, and pre-litigation steps (Medium), and multi-hop reasoning over injunction grounds and the complete legal challenge framework (Hard).

5.4 Evaluation Results

All evaluation metrics were computed against the 10 competency questions grounded in the real (anonymised) lawsuit document. The evaluator queried the knowledge graph, retrieved the semantically closest nodes to each question, and measured cosine semantic similarity between retrieved and expected nodes (threshold 0.8) and graph distance (shortest-path hop count between matched nodes).

Evaluation completed across 10 questions in 39.3 seconds. Average semantic similarity: 0.846. Average graph distance: 0.03 hops. Easy (single-hop) questions achieved perfect semantic similarity of 1.000.

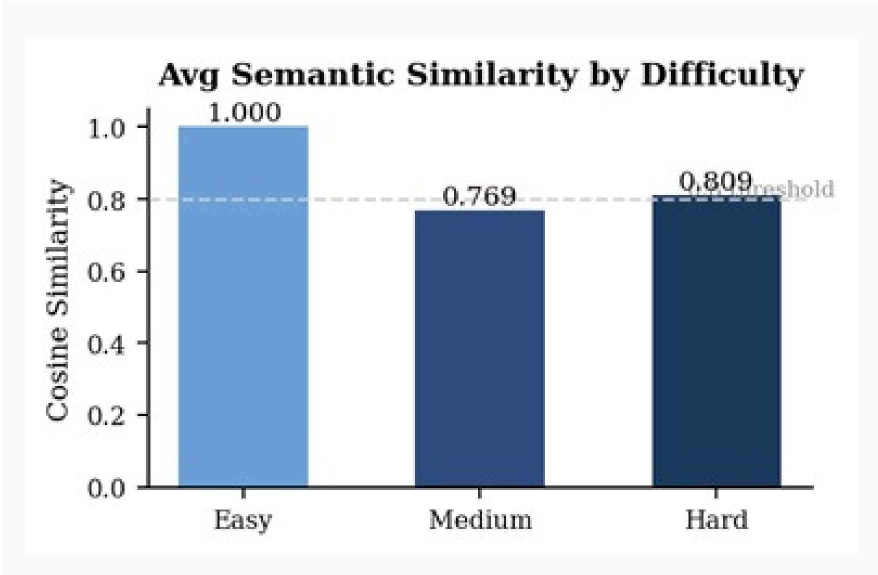

Figure 6. Average cosine semantic similarity between retrieved and expected knowledge graph nodes, by question difficulty, reflecting precise single-hop entity retrieval; the overall average of 0.846 confirms that THEMIS's embeddings capture the semantic content of legal entities reliably across all question types. Measured values.

5.5 Estimated Business Impact

Thefollowing impactestimates arebased onofficial German Federal Bar Association (BRAK) billing rate data, Federal Statistical Office civil court statistics for 2024, and Federal Ministry of Labour court statistics.

Metric	Manual Process	THEMIS
Time per lawsuit document	5–8 hours (associate + senior review)	~30 minutes (structured output review)
Automated pipeline time per document	N/A	~29s PDF parsing + graph ingestion
Cost per document (€181–€276/hr)	€905–€2,208	~€90–€140
Annual cost (100 documents)	€90,500–€220,800	~€15,000 (incl. infrastructure)
Estimated efficiency gain	—	~85–94% reduction in analysis time

Market context: Germany's labour courts received approximately 330,000 new cases in 2024[4][5], including ~78,000 employment termination protection claims, while civil courts recorded over 1.12 million new cases (Federal Statistical Office, 2024)[3]. With 165,776 licensed lawyers in Germany and 46,035 specialist lawyers (German Federal Bar Association, January 2024)[2], a firm handling 100 lawsuit documents per year could realise estimated annual savings of €75,000–€205,000 by deploying THEMIS.

> Note: Business impact figures are projections based on published market data. PDF parsing times are measured from the PoC. Graph ingestion times and lawyer-facing efficiency gains are pending full empirical measurement.

6. Key Learnings / Success Factors

Several factors drove the success of THEMIS as a proof of concept.

01 Competency questions as a validation framework.

Defining concrete, answerable questions before building the system ensured every design decision was grounded in real lawyer needs — and gave the team a clear, objective benchmark throughout development.

02 Iterative ontology design is non-negotiable.

The first ontology draft was never the right one. Multiple revision cycles — driven by testing against real documents and direct feedback from the domain expert — were essential to capturing the semantic structure of legal documents across five distinct legal domains.

03 Trust through provenance.

In legal AI, an answer without a source is not an answer. Building the complete evidence chain — from query result through knowledge graph node back to a highlighted passage in the original PDF — was the single most important design decision for user acceptance.

04 Cross-domain generalisation.

A single 13-type ontology successfully covered employment, construction, tenancy, traffic, and family law — demonstrating that a domain-specific but domain-agnostic schema design is achievable in legal AI.

05 Small team, high impact.

A focused team of three delivered a production-quality proof of concept in four months. Tight collaboration and short feedback loops were more valuable than scale.

06 Law as Code as a natural foundation.

The domain-specific ontology and knowledge graph THEMIS[MK1] produced are not solely retrieval tools — they constitute the formalised, machine-readable legal knowledge layer that Law as Code frameworks (such as Catala and LegalRuleML) require.

In effect, THEMIS makes a legal document queryable as a structured database: statutory obligations, party relationships, and factual assertions become addressable nodes, enabling rule logic to evaluate compliance automatically rather than relying on manual review. Regulatory domains like GDPR, where rights and obligations are explicitly enumerated and bounded, are prime candidates for this transition: from knowledge-assisted legal reasoning to fully computable rule execution that can automatically evaluate whether a controller's actions comply with statutory obligations.

7. Future Directions: From Knowledge Graph to Executable Law

THEMIS demonstrates that AI-powered legal knowledge graphs are not just retrieval systems — they are the data foundation for the next generation of legal technology: Law as Code. Law as Code encodes legislation, regulations, and contractual obligations as machine-executable logic, enabling automated compliance checking, regulatory sandbox testing, and self-executing legal rules at scale. Implementations such as OpenFisca (tax and benefit policy simulation), the Catala programming language (designed specifically for statutory logic), and LegalRuleML (the OASIS standard for rule interchange) are actively deployed by governments and regulated industries across Europe.

> THEMIS's 13-type legal ontology and 448-node knowledge graph, built from a 53-page GDPR lawsuit document, encode the factual and normative structure of data protection litigation in a form that Law as Code frameworks can directly consume. The graph already captures GDPR statutory articles as named, queryable nodes — Art. 22 GDPR on automated decision-making, Art. 15 on access rights, Art. 82 on damages — that are one engineering step away from executable legal predicates capable of automatically assessing whether a data controller's scoring practices violate those articles.

The competency questions framework THEMIS used for formally specifying the legal queries a system must answer correctly is itself analogous to Law as Code's formal specification of what legal rules must produce. The next engineering step is to layer computable rule logic over the knowledge graph: encoding GDPR obligations as executable predicates using tools such as Catala or LegalRuleML, enabling THEMIS to move from answering questions about case facts to automatically evaluating whether a controller's actions comply with statutory obligations — generating audit trails that satisfy Art. 5(2) GDPR's accountability principle by design.

8. Client Feedback

The client's stakeholder responded with strong enthusiasm upon reviewing the final THEMIS prototype. adorsys received feedback that the team had gone above and beyond the originally agreed scope, delivering not just a functional proof of concept but a thoughtfully designed, practice-ready system. The client specifically highlighted THEMIS's potential to fundamentally change how law firms approach document-intensive casework, describing the platform as having enormous potential to disrupt the legal market.

Frequently Asked Questions

What is a legal knowledge graph?

A legal knowledge graph is a structured database that represents the entities, relationships, and facts extracted from legal documents — such as parties, claims, evidence, legal statutes, and court filings — as interconnected nodes and edges. Unlike a traditional database, it captures the semantic meaning of legal content and allows lawyers to query complex relationships (e.g. 'which evidence supports which claim?') rather than searching by keyword alone.

How does AI reduce legal document analysis time?

THEMIS uses large language models to automatically extract entities and relationships from PDF documents and store them in a property graph database. A 53-page lawsuit that would take a lawyer 5–8 hours to analyse manually is processed in approximately 39 minutes by the automated pipeline — with every extracted fact linked back to its source passage, ready for structured querying. This reduces per-document analysis effort by an estimated 85–94 %.

Is this usable in real legal environments?

Yes. The central design requirement for professional legal practice is verifiability: every AI-generated insight must be traceable to its source. THEMIS satisfies this by maintaining a complete provenance chain — from query result, through knowledge graph node, back to the highlighted passage in the original PDF. Lawyers can inspect the evidence behind every answer without leaving the interface.

What types of legal documents does THEMIS support?

THEMIS was validated on German civil lawsuit documents (statements of claim), including GDPR / data protection litigation, employment law, construction, tenancy, traffic, and family law cases. The 13-type legal ontology is designed to be domain-agnostic within civil and labour law, and can be extended to cover additional document types and jurisdictions with targeted ontology refinement.

What is Law as Code, and how does THEMIS relate to it?

Law as Code is the practice of encoding legal rules — statutes, regulations, contractual obligations — as machine-executable logic that software can evaluate automatically. THEMIS

How long does it take to deploy THEMIS for a new document type?

The core pipeline — PDF parsing, chunking, graph ingestion, and Q&A interface — is document-type-agnostic. Adapting THEMIS to a new legal domain primarily involves ontology refinement (adding or modifying entity types and relationship constraints), which the adorsys team iterated across multiple legal domains within the four-month proof-of-concept timeline.

How accurate is THEMIS?

Evaluated against 10 competency questions grounded in a real anonymised GDPR lawsuit, THEMIS achieved an average cosine semantic similarity of 0.846 (above the 0.8 quality threshold). Easy (single-hop) questions — which cover most common lawyer queries such as party identification and court details — achieved perfect semantic similarity of 1.000. Accuracy scales with the density and quality of the ingested document.

What does this cost to build and run?

The THEMIS proof of concept was delivered by a team of two adorsys engineers and one client domain expert over four months. Running costs include Azure cloud infrastructure, OpenAI/Anthropic API usage for graph ingestion (approximately 39 minutes of LLM processing per 53-page document), and a Neo4j graph database instance. Full cost modelling is available on request at adorsys.com/contact.

9. About adorsys

We are a technology company! We are passionate about developing innovative, customised software and architecture solutions for our customers, supporting them in their innovation and transformation projects. Our services portfolio includes cybersecurity solutions and AI-enhanced software engineering, as well as cloud services, digital finance & identity and AI & data solutions.

Contact Us

Are you interested in achieving similar results? Let’s talk.

www.adorsys.com

Glossary of Key Terms

Knowledge Graph.

A structured representation of entities (people, organisations, concepts, legal statutes) and the typed relationships between them, stored as nodes and edges in a graph database. Unlike a relational database, a knowledge graph captures semantic meaning and supports multi-hop reasoning across connected facts.

THEMIS.

Tool for Handling, Extracting, Mapping & Intelligent Structure-analysis. An AI-powered legal knowledge graph platform built by adorsys.

Ontology.

A formal specification of the types of entities and relationships that exist within a domain. THEMIS's legal ontology defines 13 entity types (e.g. Lawsuit, Legal Statute, Person, Evidence, Financial Amount) and 387 valid source–target–relationship type triples that constrain which connections are semantically meaningful.

Competency Questions.

A structured set of questions that a knowledge system must be able to answer correctly — used as both a design anchor and an evaluation benchmark. In THEMIS, 10 competency questions were defined collaboratively with the client's legal domain expert and grounded in actual graph triples, enabling objective measurement of extraction quality.

Law as Code.

The practice of encoding legal rules — statutes, regulations, and obligations — as machine-executable logic that software can evaluate automatically. Frameworks include Catala (a programming language for statutory logic), LegalRuleML (the OASIS rule interchange standard), and OpenFisca (open-source policy simulation). THEMIS produces the structured legal knowledge layer these frameworks require as input.

Semantic Similarity.

A measure of how closely two pieces of text match in meaning, computed as the cosine distance between their vector embeddings. THEMIS uses a threshold of 0.8: retrieved nodes with similarity ≥ 0.8 to an expected answer are counted as correct matches. The overall average across all 10 competency questions was 0.846.

Sources

[1] German Federal Bar Association (BRAK), STAR 2020 — Billing Rate Statistics. https://www.brak.de/presse/zahlen-und-statistiken/star/star-2020/abrechnung-ueber-zeithonorare/

[2] German Federal Bar Association (BRAK), Membership Statistics January 2024. https://www.brak.de/presse/zahlen-und-statistiken/statistiken/

[3] Federal Statistical Office (Destatis), Civil Courts Report 2024. https://www.destatis.de/DE/Themen/Staat/Justiz-Rechtspflege/Publikationen/Downloads-Gerichte/statistischer-bericht-zivilgerichte-2100210247005.html

[**4]** Federal Statistical Office (Destatis), Labour Courts Report 2024. https://www.destatis.de/DE/Themen/Staat/Justiz-Rechtspflege/Publikationen/Downloads-Gerichte/statistischer-bericht-arbeitsgerichte-2100280247005.html

[5] Federal Ministry of Labour and Social Affairs (BMAS), Statistics on Labour and Social Courts. https://www.bmas.de/DE/Service/Statistiken-Open-Data/Statistik-zur-Sozial-und-Arbeitsgerichtsbarkeit/statistik-zur-sozial-und-arbeitsgerichtsbarkeit.html

[6] Wolters Kluwer, Legal AI Adoption Report 2024: Time Savings and Revenue Growth. https://www.wolterskluwer.com/en/expert-insights/legal-ai-adoption-time-savings-revenue-growth

[7] Everlaw, Lawyers Report Saving up to 32.5 Working Days per Year with Generative AI (2025). https://www.everlaw.com/blog/ai-and-law/lawyers-report-saving-up-to-32-5-working-days-per-year-with-generative-ai/

Contact us.

www.adorsys.com

info@adorsys.com

@ADORSYS

adorsys
Technology Innovation