FourCorners: Grounded Thai Legal Research over a Temporal Knowledge Graph

Pawitsapak Akarajaradwong[1], Thitiwat Nopparatbundit[1], Treephop Saeteng[1], Kasidit Phoncharoen[1], Sarana Nutanong[2], Chompakorn Chaksangchaichot[1]
[1]VISAI AI, Thailand [2]Vidyasirimedhi Institute of Science and Technology, Thailand
{pawitsapaka,thitiwatn}_visai@vistec.ac.th
{treephops,kasiditp,chompakornc}_pro@vistec.ac.th
{sarana.n}@vistec.ac.th

Abstract

Legal practitioners in Thailand must navigate fragmented government websites to research over 3,800 active laws and 87,000 Supreme Court decisions, with no unified tool for cross-referencing, version tracking, or structural navigation. We present FourCorners, a deployed platform that addresses five practitioner pain points through three modules built on a temporal legal knowledge graph covering 552K nodes and 6.3M edges: (1) an AI legal assistant that performs grounded generation via structured graph retrieval, streaming verified source content inline with responses; (2) an interactive law reader that translates graph structure into navigation and comparison features; and (3) a court decision explorer with version-aware citations produced by temporal entity resolution across 87,394 rulings. The system discovers implicit cross-corpus relationships through co-citation analysis of Supreme Court decisions. Interviews with 20 legal professionals reveal that inline source verification fundamentally changes how practitioners interact with AI-generated legal content, and that cross-corpus enrichment surfaces legal relationships that existing tools leave invisible.

[1]All screenshots are browser-translated from Thai. Thai legal dates use the Buddhist Era (BE = CE + 543); decision codes retain their original BE years (e.g., 1263/2567 = 2024 CE). Some translated labels may differ slightly from the original Thai.

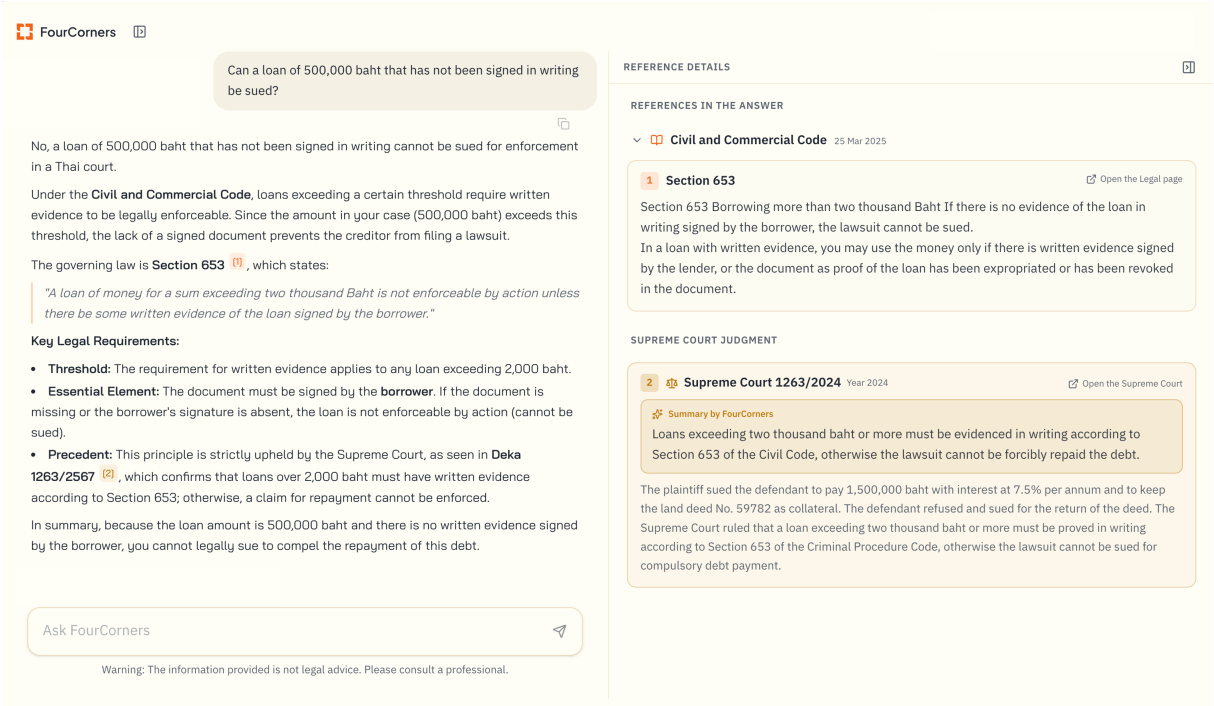

Figure 1: FourCorners AI Legal Assistant.[1] (a) The practitioner asks whether an unsigned 500,000-baht loan is enforceable. Expanding the source card for Section 653 reveals verified legal text. (b) Clicking Decision 1263/2567 (2024) opens a detail drawer showing the court's holding that evidence in writing is mandatory for loans exceeding two thousand baht.

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations), pages 463–472

1 Introduction

Thailand's legal corpus spans over 3,840 active laws and 87,394 Supreme Court decisions. Legal practitioners must regularly navigate this corpus to advise clients, draft documents, and resolve disputes, yet the tools available to them remain fragmented. Through interviews with legal practitioners and feedback from 3,500 users during a 2025 beta deployment, we identified five recurring challenges in Thai legal research:

P1. **No unified source.** Laws are published on the OCS (Office of the Council of State) website,[2] Supreme Court decisions on the DEKA website,[3] and regulations across dozens of ministry sites. Practitioners must search each source independently. From an NLP perspective, this requires cross-corpus entity linking across heterogeneous legal document formats.

P2. **Hard to read.** Official sources render laws as flat HTML or PDF with no structural navigation. The Thai Civil and Commercial Code, for example, contains over 1,700 sections that must be read by scrolling linearly.

P3. **No cross-corpus insight.** When Section A references Section B of another law, practitioners must manually locate it. Implicit relationships are entirely invisible: provisions frequently applied together in court decisions, or how courts interpret a given section, can only be discovered by reading thousands of rulings. Surfacing these relationships requires implicit relation discovery via citation network analysis across the full decision corpus.

P4. **No temporal awareness.** Thai laws are amended frequently. When a 2018 court decision cites a section, practitioners cannot determine which version was in force at the time. Resolving this requires temporal entity resolution linking each citation to a specific law version.

P5. **Unreliable data.** Official sources contain errors that practitioners cannot easily detect. The DEKA website lists cited law sections but citations are sometimes wrong and never version-specific. The OCS website contains structural indexing errors. Neither source cross-validates against the other or against its own records. Thai law has no canonical naming convention; the same law may be referenced by its formal title, a common abbreviation, or a colloquial name, making entity normalization across sources a non-trivial extraction task.

We present **FourCorners**, a deployed platform that addresses all five pain points through three integrated modules:

- **AI Legal Assistant** (§2.1) that answers legal questions in natural language while streaming verified source content alongside its responses, enabling practitioners to verify every claim without leaving the conversation.
- **Interactive Law Reader** (§2.2) that translates four structural properties of legal data into interactive navigation and comparison features.
- **Court Decision Explorer** (§2.3) with corrected, version-aware law citations that link each cited section to the version in force at the time of the ruling.

These capabilities are enabled by aggregating fragmented legal sources into a structured knowledge graph, constructed through a pipeline of NLP methods including regular expression pattern matching over Thai legal text (which lacks whitespace-delimited tokens), heuristic correction of noisy official data, and LLM-assisted reference extraction.

A core design principle is *groundedness*: every feature is grounded in official source data derived from the knowledge graph. Recent work has shown that even RAG-augmented legal AI tools hallucinate 17-33% of the time (Magesh et al., 2025). FourCorners addresses this by allowing the language model to traverse the knowledge graph via tool calling, retrieving legal definitions and provisions from the graph rather than relying on parametric knowledge, which may become outdated as laws are amended. Responses include inline citations that guarantee every cited provision exists in the graph and is not hallucinated (§3). The underlying knowledge graph covers 3,840 active laws with 6,273 versions and 87,394 Supreme Court decisions; its design and construction are detailed in Akarajaradwong et al. (2026).

Unlike medical or financial domains where knowledge transfers across borders, legal data is inherently domestic, requiring each jurisdiction to build its own infrastructure. While FourCorners is built for Thai law, its design principles are jurisdiction-agnostic: decoupling verification from generation, deriving implicit knowledge through cross-corpus citation analysis, and translating graph structure into domain-specific interfaces.

[2] https://www.ocs.go.th

[3] https://deka.supremecourt.or.th

2 System Overview

FourCorners consists of three modules, each addressing specific practitioner pain points: the AI Assistant addresses P1 (unified source), P3 (cross-corpus insight), and P5 (data quality); the Law Reader addresses P2 (readability), P3, and P4 (temporal awareness); and the DEKA Explorer addresses P4 and P5. We describe each module with a walkthrough on real legal research tasks.

2.1 AI Legal Assistant

A practitioner types a legal question in natural language. A real-time activity feed shows the system's research steps (searching laws, retrieving decisions, cross-referencing provisions). The response streams incrementally, and *inline citations* appear alongside the text where specific laws or decisions are referenced. Each source card expands to reveal: (a) verified legal text, (b) outgoing cross-references, (c) Supreme Court decisions citing it with an interpretation summary, and (d) resolved legal definitions for specialized terms used in the provision.

Walkthrough. A practitioner asks: "Can a loan of 500,000 baht that has not been signed in writing be sued?" The assistant streams its response, referencing provisions across the Civil and Commercial Code (Figure 1). Expanding the inline citation [1] (Section 653) reveals the verified text, and Decision 1263/2024 which applied this section. Clicking the decision opens a detail drawer showing that the court voided the lawsuit under Section 653.

2.2 Interactive Law Reader

FourCorners translates four structural properties of the knowledge graph into domain-specific UX features (Table 1). The reader uses a three-panel layout (Figure 2): a collapsible hierarchy tree (left), section content with clickable cross-references (center), and an enrichment panel (right) showing co-cited sections, court interpretations, and resolved legal definitions for specialized terms.

Walkthrough: Cross-references. A practitioner reading Section 67 of the Revenue Code (corporate income tax rates) sees that it references Sections 40, 65 bis, and 66. Each reference is a clickable link. The enrichment panel shows that Sections 69 and 71 reference Section 67 back, and that Sections 22, 27, 65, and 65 bis are frequently co-cited with it across court decisions. These bidirectional and co-citation relationships are invisible in statute text alone.

Graph Property	UX Feature	Practitioner Benefit
Hierarchy	Collapsible tree	Navigate 1,700+ section laws by structure
Temporal	Timeline slider + word-level diff	See what changed in each amendment
Reference	Clickable cross-references + legal definitions	Jump to referenced sections; resolve specialized terms in context
Sequential	Document-order navigation	Read law as intended, find by label

Table 1: Four knowledge graph properties expressed as UX features in the Interactive Law Reader.

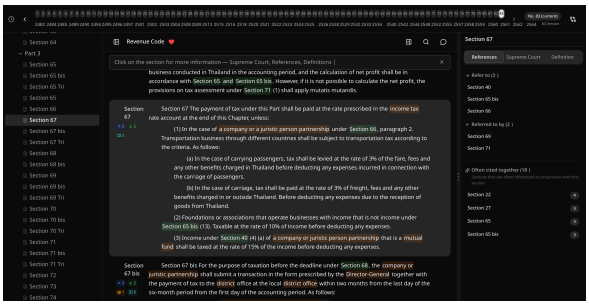

Figure 2: Interactive Law Reader. Three-panel layout: (left) collapsible hierarchy tree, (center) Section 67 with clickable cross-references highlighted, (right) enrichment panel showing outgoing references, incoming references, and co-cited sections.

Walkthrough: Temporal comparison. The practitioner compares the two most recent versions of the Penal Code (v.34 vs. v.35). The comparison view operates at two levels (Figure 3): the *component level* shows that Section 284/4 was added in v.35; the *word level* shows inline annotations on Section 397, with deleted text in red strikethrough.

2.3 Court Decision Explorer

The Court Decision Explorer provides faceted search across Supreme Court decisions by year, keyword, and cited law section. Each decision displays an AI-extracted "golden paragraph": the core legal holding, highlighted for quick reading.

Walkthrough: Version-aware citations. A practitioner finds Decision T.99/2568, which concerns an erroneous land deed number in an enforcement petition (Figure 4). The golden paragraph summarizes the holding: the error cannot be remedied under Section 143 of the Code of Civil Procedure.

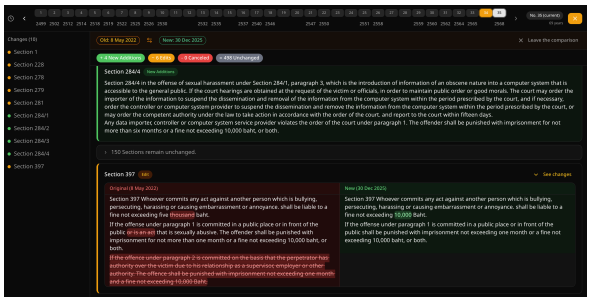

Figure 3: Temporal version comparison of the Penal Code (v.34 vs. v.35). The component-level view shows Section 284/4 was added; the word-level view shows inline edits to Section 397 with red strikethrough for deleted text.

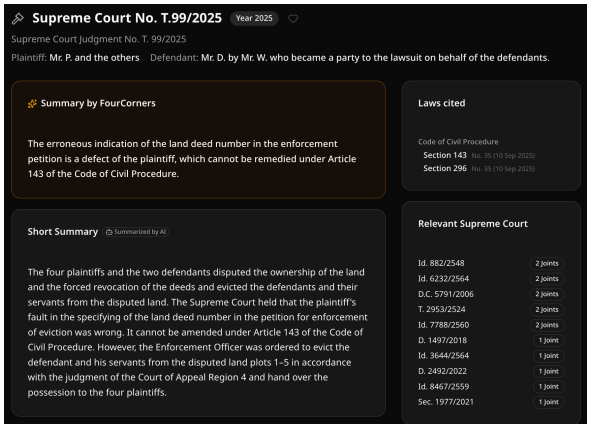

Figure 4: Court Decision Explorer. Decision T.99/2568 displays an AI-extracted golden paragraph, version-aware citations linking Section 143 to the version in force at the time, and related decisions ranked by co-citation overlap.

The citation panel shows that Section 143 links to version No. 35 (10 Sep 2025), the version in force at the time of the ruling, and lists related decisions ranked by co-citation overlap.

3 System Components

Legal information extraction. The knowledge graph (552K nodes, 6.3M edges) encodes the four structural properties in Table 1 that directly power the UX features in §2. It contains 231K cross-reference edges extracted from statute text, 237K court-to-statute citation edges, 24K co-citation edges, 3,427 court interpretation profiles, and 9,468 legal term definitions. Constructing this graph from official Thai sources poses language-specific challenges: Thai text lacks whitespace word boundaries, requiring dictionary-based tokenization for reference extraction; section labels mix Thai and Arabic numerals (e.g., "Section 77/1" may appear as *maatraa 77/1* in Thai script); and legal terminology diverges substantially from general Thai corpora, necessitating domain-specific extraction patterns. The graph design and construction are detailed in a Akarajaradwong et al. (2026).

Grounded generation with structured retrieval. The assistant uses a large language model with structured tool access to the knowledge graph, operating as a multi-step research agent: given a query, it executes structured graph queries (fetch provisions, trace cross-references, resolve temporal versions). Individual graph queries return in under 50ms; end-to-end time to first token is typically 1–2 seconds. As the language model generates its response, each provision reference triggers an independent graph lookup producing a source card with verified legal text, cross-references, court interpretations, and definitions. This ensures *decoupled verification* (source cards are produced by graph queries, not by the language model; a card cannot exist for a nonexistent provision) and *inline delivery* (verified content appears at the moment of relevance, not as end-of-response lists).

Implicit relation discovery via co-citation. By analyzing citation patterns across 87,394 Supreme Court decisions, FourCorners identifies pairs of law sections cited together in three or more rulings, yielding 24,010 co-citation edges. For example, Revenue Code Section 118 (stamp duty as evidence requirement) and Civil & Commercial Code Section 653 (written evidence for loans) share no cross-reference in statute text, yet 33 Supreme Court decisions cite both, because loan agreements require proper stamp duty to be admissible. These relationships require analysis across the full decision corpus; they are invisible in any single source. Co-citation edges are precomputed and surfaced in the enrichment panel.

Temporal entity resolution for citations. DEKA records only the year a decision was published, and its law citations are plain text with no version information. FourCorners uses the publication year to identify which version of a cited law was in force at the time of the ruling. When a law was amended multiple times within the same year, the system selects the latest version in that year as the most likely applicable text. Of all court-to-statute citations in the DEKA corpus, 30% reference laws absent from the digital repository (typically pre-digitization statutes);

Capability	OCS	DEKA	AI†	FC
Unified law + court search			○	✓
Hierarchical navigation	○			✓
Temporal version comparison				✓
Clickable cross-references				✓
Version-aware court citations				✓
Citation error correction				✓
Co-citation analysis				✓
AI legal Q&A			✓	✓
Inline source verification			○	✓
Court interpretation summaries				✓
Daily data updates	✓	✓	○	✓

Table 2: Feature comparison. OCS: statutory law website. DEKA: court decision website. AI†: general-purpose AI (e.g., ChatGPT) and Thai commercial legal AI tools (e.g., Khorn AI, Thanoy). FC: FourCorners. ✓ = full support; ○ = partial (e.g., OCS provides flat navigation without hierarchy labels; AI tools provide inline citations without structured verification).

within covered laws, multi-stage validation discards approximately 7% of constructed links.

4 Evaluation

We evaluate FourCorners through semi-structured interviews with legal professionals and retrieval benchmark evidence from Akarajaradwong et al. (2026).

4.1 Practitioner Interviews

We conducted semi-structured interviews with 20 legal professionals (3 to 30 years of experience, including practicing lawyers, in-house counsel, and legal researchers) who used FourCorners on self-chosen research tasks. After each session, participants rated five statements on a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree); Table 3 reports the results. Three consistent qualitative findings emerged.

Groundedness drives adoption. All participants identified data reliability as the primary value: every AI response traces to official provisions retrievable through the knowledge graph. Inline citations shifted the verification model from post-hoc (checking each AI citation on a separate government website) to concurrent (reading AI synthesis and verified legal text side by side). The unified platform eliminated the multi-source workflow of searching OCS, DEKA, and external search engines independently.

Enrichment reveals hidden relationships. The most unexpected finding was participants' response to the enrichment panel. Co-citation analysis surfaced practical relationships between provisions that neither reading statute text nor querying other AI tools can reveal, giving participants context about how laws operate in practice. No existing Thai legal platform provides these data-driven relationships.

Statement	Mean	SD
Inline citations increase my trust in AI-generated legal responses	4.65	0.49
I can verify cited provisions faster than with my current workflow	4.50	0.61
Temporal comparison helps me identify risks I would otherwise miss	4.20	0.77
The enrichment panel surfaces useful relationships I did not expect	4.55	0.51
I would adopt this tool for my daily legal research	4.30	0.73

Table 3: Post-session Likert ratings (1 = strongly disagree, 5 = strongly agree; $N = 20$ legal professionals, 3–30 years experience).

Temporal awareness. Temporal comparison proved most valuable for compliance tasks, where participants discovered they had been citing subsequently amended provisions. Participants also identified limitations discussed in §6.

4.2 Retrieval Quality

We evaluate on NitiBench-Tax (Akarajaradwong et al., 2025), a 50-question Thai legal QA benchmark with section-level citation ground truth (Table 4). The first four rows reproduce a controlled evaluation from Akarajaradwong et al. (2026), which isolates each retrieval backend in a single-agent setup to enable fair comparison. The final row evaluates the deployed production system as-is: the full chat pipeline with web search, cross-reference traversal, and co-citation expansion, with domains hosting the benchmark blocked to prevent data leakage. All systems search 269,996 sections (53× larger than baselines).

We report F2 alongside F1 because practitioners value recall over precision: missing a relevant provision risks incorrect advice, while an extra provision costs only seconds of reading. The production system surfaces related provisions via cross-reference traversal and co-citation expansion, achieving near-ceiling recall (.990) and the highest F2 (.901). The knowledge graph outperforms vector retrieval because structured edge traversal discovers related provisions that embedding simi-

Retrieval Method	Prec.	Rec.	F1	F2
Golden Context (ceiling)	.962	.993	.977	.987
Web Search	.616	.836	.666	.739
Agentic RAG (vector)	.591	.952	.685	.799
Microsoft GraphRAG	.720	.868	.761	.811
FourCorners (single-agent)	**.751**	.983	**.812**	.886
FourCorners (production)	.738	**.990**	.810	**.901**

Table 4: Citation accuracy on NitiBench-Tax (50 questions). Rows above the line use a controlled single-agent setup from Akarajaradwong et al. (2026). The production row evaluates the deployed chat system with benchmark-hosting domains blocked to prevent data leakage. All systems search a 53× larger corpus than the vector and GraphRAG baselines.

larity misses, particularly when relevant sections use different terminology from the query.

5 Related Work

Legal AI systems. Commercial platforms (Westlaw, LexisNexis, CoCounsel, Harvey) serve common law jurisdictions with established digital infrastructure. In the Thai market, Khorn AI[4] and Thanoy[5] provide AI legal Q&A via agentic RAG over government web content, but do not advertise temporal versioning, citation correction, or cross-corpus enrichment features. In legal NLP, Indian Kanoon (Indian Kanoon), ChatLaw (Cui et al., 2024), and HyPA-RAG (Kalra et al., 2024) address legal search and QA. FourCorners differs in combining a temporal knowledge graph, interactive law reading, and AI QA with inline source verification.

Legal information visualization and knowledge graphs. LegalViz (Onami et al., 2025) generates structured diagrams from legal documents. ELLA (Hu et al., 2024) presents similarity scores between retrieved legal articles and LLM responses for interactive grounding. Colombo et al. (2025) build an Italian legislative KG from laws published in XML format. The European Legislation Identifier (ELI) (Council of the European Union, 2012) and Akoma Ntoso (Palmirani and Vitali, 2011) provide standards for legislative data interchange. These systems address individual aspects of legal data; FourCorners is, to our knowledge, the first to translate four structural graph properties into a unified practitioner workflow.

[4] https://khorn.ai

[5] https://iapp.co.th/th/products/thanoy

Temporal legal document systems. The UK's Legislation.gov.uk (The National Archives) and EUR-Lex (Publications Office of the European Union) provide point-in-time access to historical statute versions. However, these systems support version browsing only; they do not resolve court citations to the law version in force at the time of the ruling or surface how amendments affect cross-law relationships. FourCorners extends temporal awareness to citation correction and cross-corpus temporal comparison.

Grounded AI and source verification. Retrieval-augmented generation (Lewis et al., 2021) grounds language model responses in retrieved documents. Microsoft GraphRAG (Edge et al., 2025) extends this to graph-structured knowledge. OLMoTrace (Liu et al., 2025) traces AI outputs to training data for transparency. FourCorners takes a complementary approach: it delivers verified source content directly alongside AI responses, allowing practitioners to inspect actual legal text supporting each claim. To our knowledge, FourCorners is the first deployed legal AI system that architecturally decouples source verification from text generation via structured graph queries.

6 Conclusion

We present FourCorners, a deployed platform for Thai legal research that addresses five practitioner pain points through three modules built on a temporal knowledge graph. Practitioner interviews reveal that inline source verification changes how legal professionals interact with AI-generated content, and that cross-corpus enrichment surfaces relationships that current tools leave invisible. Three design principles may generalize to other jurisdictions: (1) *decouple verification from generation* so hallucinated citations are architecturally impossible; (2) *link corpora to derive enrichment*, since cross-corpus analysis reveals patterns invisible in any single source; and (3) *translate graph structure into domain-specific UX*. Because legal data is inherently domestic, each jurisdiction must solve its own information extraction, entity normalization, and temporal resolution challenges. We hope FourCorners demonstrates that a knowledge graph-grounded approach can make these tasks tractable even for under-resourced legal systems.

Limitations

Coverage. FourCorners covers OCS laws and DEKA Supreme Court decisions. Ministerial regulations, departmental orders, and local ordinances fall outside the current scope. From the 87,394 court decisions, 30% cite laws absent from the digital repository (typically pre-digitization statutes), resulting in citation links that cannot be resolved.

Temporal resolution. DEKA records only the publication year of decisions, not the exact date. When a law was amended multiple times within the same calendar year, the system selects the latest version in that year as the most likely applicable text. This heuristic is correct in most cases but may resolve to the wrong version for decisions published shortly after a mid-year amendment.

AI reasoning. Source cards guarantee that every cited provision exists in the knowledge graph, but do not guarantee that the model selects the most relevant provisions for a given question. The model may cite a real but irrelevant section, and a source card will still appear for it. Practitioners retain final judgment on legal relevance.

Language and accessibility. The system is Thai-only; all legal text, court decisions, and the user interface are in Thai, as faithful translation of legal text remains challenging. Non-Thai speakers can use built-in browser translation for basic access. The platform requires internet access and is usable on mobile browsers but optimized for desktop.

Ethical Considerations

Not a substitute for legal judgment. FourCorners is a research tool, not a legal advisor. Users may over-rely on AI-generated analysis; the system mitigates this by presenting verified legal text alongside AI synthesis rather than AI output alone.

Data quality. Official sources themselves contain errors (§3); the system corrects what it can detect but cannot guarantee all source errors are caught.

Data licensing and privacy. All legal texts and court decisions used are explicitly excluded from copyright under Thailand's Copyright Act B.E. 2537, Section 7(2)–(5), which places legislation, court decisions, and government orders in the public domain. The FourCorners platform is proprietary software developed by VISAI AI; the knowledge graph and application code are not open-sourced. The platform is publicly accessible; user conversations are stored to enable session continuity and are not shared with third parties.

References

Pawitsapak Akarajaradwong, Sarana Nutanong, and Chompakorn Chaksangchaichot. 2026. Fourcorners: A production knowledge graph unifying thailand's legal system. In *The 64th Annual Meeting of the Association for Computational Linguistics – Industry Track*.

Pawitsapak Akarajaradwong, Pirat Pothavorn, Chompakorn Chaksangchaichot, Panuthep Tasawong, Thitiwat Nopparatbundit, Keerakiat Pratai, and Sarana Nutanong. 2025. NitiBench: Benchmarking LLM frameworks on Thai legal question answering capabilities. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 34304–34327, Suzhou, China. Association for Computational Linguistics.

Andrea Colombo, Anna Bernasconi, and Stefano Ceri. 2025. An llm-assisted etl pipeline to build a high-quality knowledge graph of the italian legislation. *Information Processing & Management*, 62(4):104082.

Council of the European Union. 2012. Council conclusions inviting the introduction of the European Legislation Identifier (ELI). 2012/C 325/02, Official Journal of the European Union.

Jiaxi 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. *Preprint*, arXiv:2306.16092.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2025. From local to global: A graph rag approach to query-focused summarization. *Preprint*, arXiv:2404.16130.

Yutong Hu, Kangcheng Luo, and Yansong Feng. 2024. ELLA: Empowering LLMs for interpretable, accurate and informative legal advice. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 374–387, Bangkok, Thailand. Association for Computational Linguistics.

Indian Kanoon. Indian kanoon: Search engine for indian law.

Rishi Kalra, Zekun Wu, Ayesha Gulley, Airlie Hilliard, Xin Guan, Adriano Koshiyama, and Philip Colin Treleaven. 2024. HyPA-RAG: A hybrid parameter adaptive retrieval-augmented generation system for AI legal and policy applications. In *Proceedings of*

the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U), pages 237–256, Miami, Florida, USA. Association for Computational Linguistics.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Preprint*, arXiv:2005.11401.

Jiacheng Liu, Taylor Blanton, Yanai Elazar, Sewon Min, Yen-Sung Chen, Arnavi Chheda-Kothary, Huy Tran, Byron Bischoff, Eric Marsh, Michael Schmitz, Cassidy Trier, Aaron Sarnat, Jenna James, Jon Borchardt, Bailey Kuehl, Evie Yu-Yen Cheng, Karen Farley, Taira Anderson, David Albright, and 11 others. 2025. OLMoTrace: Tracing language model outputs back to trillions of training tokens. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 178–188, Vienna, Austria. Association for Computational Linguistics.

Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, and Daniel E. Ho. 2025. Hallucination-free? assessing the reliability of leading ai legal research tools. *Journal of Empirical Legal Studies*, 22(2):216–242.

Eri Onami, Taiki Miyanishi, Koki Maeda, and Shuhei Kurita. 2025. LegalViz: Legal text visualization by text to diagram generation. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6657–6676, Albuquerque, New Mexico. Association for Computational Linguistics.

Monica Palmirani and Fabio Vitali. 2011. Akoma-ntoso for legal documents. In Giovanni Sartor, Monica Palmirani, Enrico Francesconi, and Maria Angela Biasiotti, editors, *Legislative XML for the Semantic Web: Principles, Models, Standards for Document Management*, pages 75–100. Springer Netherlands, Dordrecht.

Publications Office of the European Union. EUR-Lex: Access to European Union law.

The National Archives. Legislation.gov.uk.

A Reviewer Guide

The live demo is available at https://fourcorners-preview.visai.ai. Each reviewer has a dedicated account according to the single blind:

Reviewer	Username	Password
Reviewer 1	acl26-demo-r1	Fc-Rev1@Acl26
Reviewer 2	acl26-demo-r2	Fc-Rev2@Acl26
Reviewer 3	acl26-demo-r3	Fc-Rev3@Acl26

Suggested walkthrough:

1. **AI Assistant.** Ask: "Can a loan of 500,000 baht that has not been signed in writing be sued?" Observe the activity feed, streamed response, and inline citations. Click a source card to see verified legal text, cross-references, court decisions, and definitions.
2. **Law Reader.** Navigate to the Revenue Code via the law repository. Expand the hierarchy tree (left panel). Click Section 67 to read it. Click a highlighted cross-reference to jump across laws. Open the enrichment panel (right) to see co-cited sections and court interpretations.
3. **Temporal comparison.** Navigate to the Penal Code. Click the timeline icon and compare v.34 vs. v.35. Observe that Section 284/4 was added and Section 397 has word-level edits.
4. **Court Decision Explorer.** Search for Decision T.99/2568 (the translated interface may display 2025). Note the golden paragraph and version-aware citation links. Click Section 143 to open it in the Law Reader at the historical version.

The interface is in Thai. We recommend using the built-in browser translation for the best experience.

B Existing Tools

Figures 5–7 show the three baseline tools that Thai legal practitioners currently use (browser-translated from Thai; see footnote 1). Comparing these with FourCorners (Figures 1–4) illustrates the gaps addressed in §1.

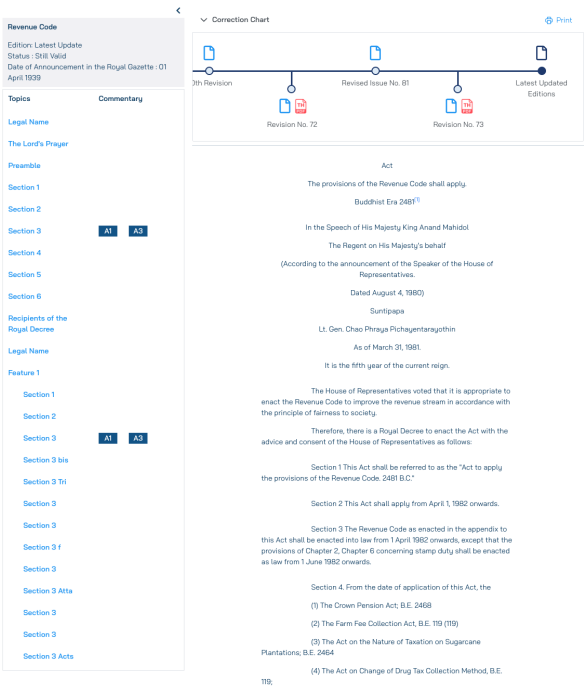

Figure 5: OCS website. A navigation sidebar provides anchor links to jump to components, and a timeline bar allows switching to older versions, but there is no hierarchical tree, no clickable cross-references, and no version comparison.

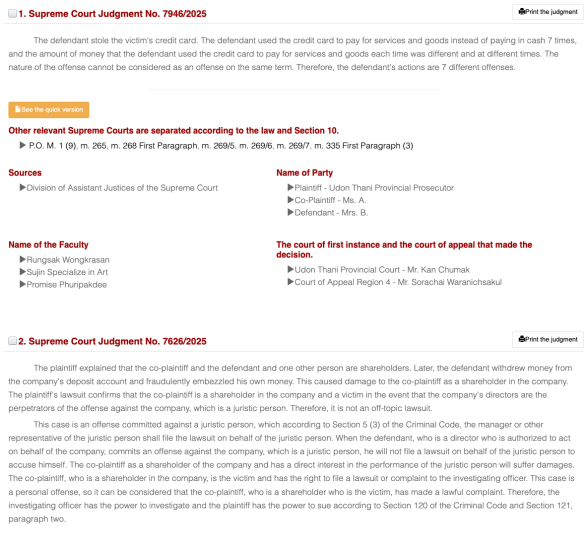

Figure 6: DEKA website. Court decisions display cited law sections as unlinked plain text with no version information. Practitioners cannot determine which version of a cited law was in force at the time of the ruling.

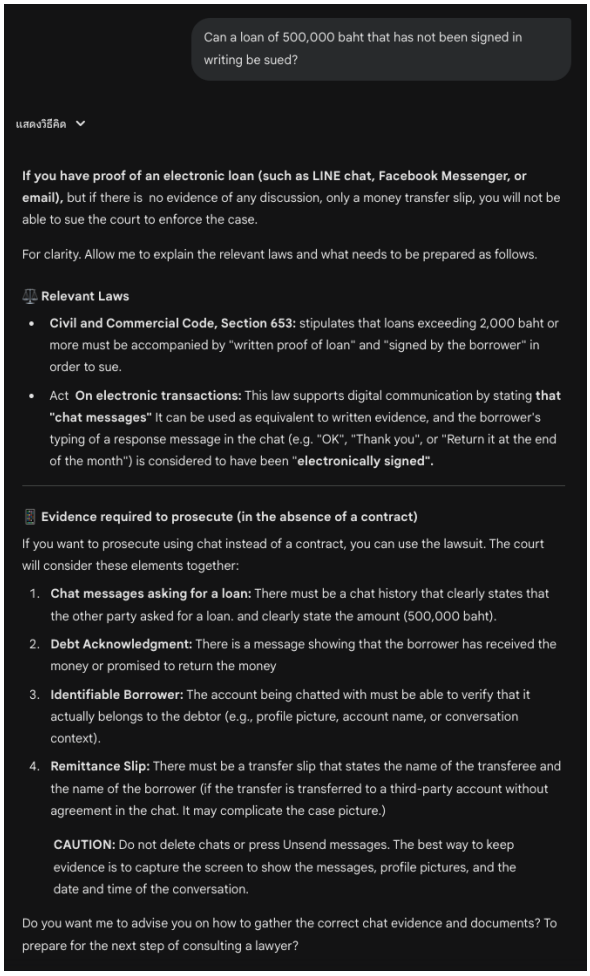

Figure 7: General-purpose AI (Gemini). The response cites law sections but provides no structured source verification. Practitioners must manually look up each citation on a separate government website to verify it exists and is correctly quoted.