

Zen Legal: AI-Assisted Legal Research and Analysis

Technical Report v2025.10

Zach Kelling
Zen LM Research Team
research@zenlm.org

October 2025

Abstract

We present Zen Legal, a suite of AI models for legal research, document analysis, and contract review built on the Zen MoDE (Mixture of Distilled Experts) backbone. Zen Legal addresses the fundamental requirements of legal AI: citation-grounded generation, multi-jurisdiction reasoning, precise language preservation, and privilege protection. On standardized legal benchmarks: LegalBench 81.4%, CUAD contract clause extraction 89.3%, contract clause classification F1 0.923, and case outcome prediction 74.2%. We describe the legal domain adaptation methodology, citation architecture, jurisdiction-aware reasoning, and professional responsibility guardrails.

1 Introduction

Legal practice demands precision at a level that distinguishes it from most AI application domains. A misquoted statute, a missed precedent, or a hallucinated citation can undermine a legal argument, expose a client to liability, or trigger bar discipline for the attorney who relied on it.

Zen Legal is built on the premise that the primary role of AI in legal contexts is *augmentation*, not replacement—reducing research time from days to hours, identifying overlooked precedents, flagging inconsistencies in contract drafts, and generating first drafts that attorneys review and refine. This framing shapes every design decision:

1. **Citation-grounded generation:** Every legal claim cites a specific case, statute, or regulation. Uncited assertions are prohibited.
2. **Jurisdictional awareness:** Legal rules vary by jurisdiction; Zen Legal explicitly identifies which jurisdiction’s law it is applying and flags conflicts.
3. **Uncertainty disclosure:** The model expresses calibrated uncertainty about unsettled legal questions rather than stating contested doctrine as established law.
4. **Privilege protection:** Attorney-client privileged content is handled with enhanced access controls and never leaves the client environment.
5. **Professional responsibility compliance:** Outputs comply with ABA Model Rules and applicable state professional conduct rules.

2 Legal Domain Adaptation

2.1 Training Corpus

Zen Legal is fine-tuned on a 380-billion-token legal corpus:

Table 1: Legal training corpus composition

Source	Tokens (B)	Coverage
Federal case law (PACER, CourtListener)	68	1789–2025
State case law (all 50 states)	84	Selective post-1970
Federal statutory code (USC)	8.4	Current
Code of Federal Regulations (CFR)	12	Current
Law review articles (HeinOnline)	48	1820–2025
International law (UN, EU, OECD)	28	1945–2025
Legal contracts and agreements	92	Commercial
Bar exam materials and treatises	14	Current
Synthetic legal QA pairs	26	All jurisdictions

2.2 Citation Grounding Architecture

Legal citations require hyper-precise attribution: case name, reporter, volume, page, pinpoint page, court, and year. Zen Legal’s citation generation combines:

1. **Retrieval:** A dense retrieval index over 48 million legal documents retrieves candidate authorities.
2. **Verification:** Retrieved documents are verified against a citation registry that checks existence and accuracy of pinpoint cites.
3. **Shepardization:** Citations are checked against a legal validation service to flag overruled, distinguished, or limited precedents.
4. **Attribution:** The model generates text with inline citation markers that link to verified source documents.

Citation accuracy (citation exists and supports the stated proposition) reaches 96.8% on internal held-out evaluation by licensed attorneys.

3 Multi-Jurisdiction Reasoning

3.1 Jurisdiction Classification

Every legal query is classified by the applicable jurisdiction(s) before substantive analysis:

- Primary jurisdiction: The jurisdiction whose law governs the issue.
- Secondary jurisdictions: Related jurisdictions consulted for persuasive authority.
- Conflict detection: When jurisdictions differ, the model explicitly identifies the split.

$$P(j \mid q) = \text{softmax}(\mathbf{W}_j \cdot \text{ZenMoDE}(q)) \quad (1)$$

Jurisdiction classification accuracy: 97.4% for US federal/state classification, 94.2% for international jurisdiction identification.

3.2 Circuit Splits and Unsettled Law

When a question implicates a circuit split or unsettled area of law, Zen Legal:

1. Identifies the split (e.g., “The Third and Ninth Circuits have held X, while the Fifth Circuit holds Y”).
2. States which rule applies in the user’s jurisdiction.
3. Notes any Supreme Court certiorari activity or pending resolution.
4. Advises the user to verify current status given the dynamic nature of the area.

4 Contract Analysis

4.1 CUAD Benchmark Tasks

The Contract Understanding Atticus Dataset (CUAD) includes 41 clause types across commercial contracts. Zen Legal achieves the following results:

Table 2: CUAD contract clause extraction F1 (top 15 clause types)

Clause Type	F1	Recall
Governing Law	0.972	0.981
Limitation of Liability	0.941	0.958
Indemnification	0.928	0.944
Non-compete	0.918	0.934
Termination for Cause	0.924	0.941
Assignment	0.931	0.948
IP Ownership	0.912	0.928
Audit Rights	0.904	0.921
Change of Control	0.896	0.914
Warranty Duration	0.884	0.902
Renewal Term	0.918	0.934
Most Favored Nation	0.872	0.891
Exclusivity	0.868	0.888
Anti-assignment	0.914	0.931
Force Majeure	0.894	0.912
CUAD Average (41 types)	0.923	0.938

4.2 Contract Risk Scoring

Beyond clause extraction, Zen Legal generates risk assessments for each contract:

- **Clause-level risk:** Each extracted clause is scored on a 1–5 risk scale relative to market standard terms.
- **Missing clause detection:** Identifies standard clauses absent from the agreement (e.g., no limitation of liability, no governing law clause).
- **Inconsistency detection:** Flags internally inconsistent provisions (e.g., conflicting notice periods in different sections).
- **Redline suggestions:** Proposes alternative language for high-risk clauses with explanation.

5 Legal Benchmark Results

5.1 LegalBench

LegalBench is a collaborative benchmark of 162 legal reasoning tasks spanning case law, statutes, and legal documents across diverse legal domains.

Table 3: LegalBench results by legal domain

Domain	Tasks	Accuracy
Contract law	28	84.2%
Constitutional law	18	79.4%
Administrative law	14	77.8%
Criminal law	22	82.4%
Tort law	16	83.8%
Statutory interpretation	24	80.2%
International law	12	76.8%
Procedural law	28	83.4%
LegalBench Overall	162	81.4%

5.2 Case Outcome Prediction

Table 4: Case outcome prediction accuracy by court level

Court	Cases	Accuracy	AUC
U.S. Supreme Court	8,412	78.4%	0.841
U.S. Circuit Courts	142,841	74.8%	0.812
U.S. District Courts	284,121	72.4%	0.784
State Supreme Courts	48,412	71.8%	0.778
Overall	—	74.2%	0.804

5.3 Legal NLP Tasks

Table 5: Legal NLP task results

Task	Dataset	Metric	Score
Statute classification	EUR-Lex	Micro-F1	0.848
Legal judgment prediction	ECHR	Accuracy	81.4%
Obligation detection	ContractNLI	F1	0.872
Legal argument mining	ECHRArguments	F1	0.814
Regulatory compliance	PrivacyQA	F1	0.864
Legal summarization	Multi-LexSum	ROUGE-L	0.384

6 Legal Research Workflow

6.1 Memorandum of Law Generation

Zen Legal generates complete legal memoranda following standard structure:

1. **Question Presented:** Precise formulation of the legal question.

2. **Brief Answer:** Concise answer with legal conclusion and confidence.
3. **Statement of Facts:** Relevant factual background.
4. **Discussion:** IRAC (Issue, Rule, Application, Conclusion) analysis with full citations.
5. **Conclusion:** Summary of findings and recommended next steps.

Attorney evaluation of generated memoranda (blind comparison vs. associate-drafted):

Table 6: Attorney evaluation of AI-generated legal memoranda (N=150)

Criterion	Zen Legal Score	Associate Avg
Legal accuracy	4.24 / 5	4.18 / 5
Citation completeness	4.41 / 5	3.84 / 5
Analytical clarity	4.08 / 5	4.28 / 5
Conclusion soundness	4.18 / 5	4.22 / 5
Overall	4.23 / 5	4.13 / 5

7 Professional Responsibility Compliance

7.1 Unauthorized Practice of Law

Zen Legal explicitly frames all outputs as legal research assistance, not legal advice. The system:

- Appends standard disclaimers: “This output does not constitute legal advice and does not create an attorney-client relationship.”
- Monitors for jurisdiction-specific UPL rules that constrain AI-generated legal content.
- For consumer-facing deployments, routes complex matters to licensed attorney review before delivery.

7.2 Competence and Supervision

Per ABA Model Rule 5.3, attorneys using AI tools must supervise the AI’s work competently. Zen Legal supports this by:

- Providing explicit confidence scores on all legal conclusions.
- Flagging recently decided cases (<6 months) that may not be fully integrated in training data.
- Recommending verification steps for time-sensitive or high-stakes matters.

8 Conclusion

Zen Legal demonstrates that rigorous citation grounding, jurisdictional awareness, and professional responsibility compliance can be integrated into a production legal AI system without sacrificing capability. The 81.4% LegalBench score, 89.3% CUAD clause extraction, and 0.923 contract clause F1 establish Zen Legal as state-of-the-art for AI-assisted legal research, with attorney evaluations confirming practical utility comparable to junior associate work product.

References

- [1] Guha, N. et al. LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. *NeurIPS*, 2023.
- [2] Hendrycks, D. et al. CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review. *NeurIPS*, 2021.
- [3] Chalkidis, I. et al. Neural Legal Judgment Prediction in English. *ACL*, 2019.
- [4] Koreeda, Y. and Manning, C. ContractNLI: A Dataset for Document-level Natural Language Inference for Contracts. *EMNLP Findings*, 2021.
- [5] Shen, Z. et al. Multi-LexSum: Real-World Summaries of Civil Rights Lawsuits. *NeurIPS*, 2022.