

Legal Clarity Agent

Siddhesh Save
Software Engineer
United States
siddheshsave99@gmail.com

Abstract—Legal documents such as contracts, non-disclosure agreements, and service agreements are often difficult to interpret due to complex language, dense structure, and implicit obligations. Manual legal review is time-consuming, costly, and inaccessible to non-legal professionals, while existing automated tools often rely on centralized cloud processing, raising privacy concerns. This project presents Legal Clarity Agent, a privacy-first, multi-agent legal document analysis system built using Python, Google Agent Development Kit (ADK), and Generative AI. The system performs clause-level analysis by orchestrating multiple specialized agents that classify clauses, assess risk, extract key terms and action items, and generate plain-English explanations. A hybrid architecture combining sequential and parallel agent workflows enables efficient and interpretable analysis while maintaining data locality. The system produces a structured legal clarity report designed to improve transparency, usability, and trust in automated legal reasoning systems.

Keywords—Multi-Agent Systems, Legal AI, Generative AI, Google ADK, Clause Analysis, Risk Assessment, Agent Orchestration, Privacy-Preserving AI

I. INTRODUCTION

Legal documents govern a wide range of personal and professional activities, yet they remain largely inaccessible to individuals without legal expertise. Contracts often contain lengthy clauses, specialized terminology, and implicit obligations that are difficult to identify and interpret correctly. As a result, individuals and small organizations face increased risk of misunderstanding legal commitments, which can lead to financial or legal consequences.

Recent advances in large language models have enabled automated text analysis and summarization for legal documents. However, many existing solutions focus on high-level summaries and rely on uploading sensitive documents to third-party cloud platforms. This approach raises privacy and compliance concerns, particularly for confidential or proprietary agreements.

To address these challenges, the Legal Clarity Agent is proposed as a privacy-first, agent-based legal analysis system [1]. Instead of relying on a single monolithic model, the system decomposes legal reasoning into multiple specialized agents coordinated through structured orchestration. This design enables deeper analysis, improved interpretability, and controlled execution within a local or notebook-based environment.

II. LITERATURE REVIEW & SIGNIFICANCE

Prior research in legal document analysis has largely focused on document summarization, information extraction, and contract classification using machine learning and natural language processing techniques. While these approaches have shown promise, they often treat legal documents as

unstructured text and lack fine-grained interpretability at the clause level.

More recent work has explored agent-based AI systems for complex reasoning tasks, demonstrating improved modularity and scalability. However, the application of multi-agent architectures to legal analysis remains limited, particularly in privacy-conscious settings.

The Legal Clarity Agent distinguishes itself by combining clause-level analysis with a multi-agent orchestration framework. By assigning distinct responsibilities to individual agents such as clause classification, risk assessment, and explanation generation, the system delivers structured, interpretable outputs while preserving user privacy. This approach highlights the practical value of multi-agent systems for applied legal AI.

III. PROBLEM ANALYSIS

Despite advancements in legal technology, several limitations persist in existing automated solutions:

- Most systems provide document-level summaries rather than clause-level insights
- Hidden risks and obligations are often overlooked
- Users must upload sensitive legal documents to external services
- Outputs lack transparency and explainability

Legal analysis is inherently multi-dimensional, requiring classification, interpretation, and risk evaluation. Treating this process as a single task reduces accuracy and user trust. There is a clear need for a modular, privacy-preserving system that can analyze legal documents in a structured and interpretable manner.

IV. METHODOLOGY

The Legal Clarity Agent is designed as a multi-agent pipeline implemented using the Google Agent Development Kit (ADK) [1]. The methodology follows a structured orchestration approach using both sequential and parallel agents to balance correctness and performance [2].

A. Multi-Agent Pipeline Design

The system processes legal document text through the following stages:

- Clause extraction
- Clause classification
- Risk assessment
- Plain-English explanation
- Glossary and action item extraction

- Aggregation into a structured report

B. Agent Orchestration Strategy

A Sequential Agent is used where dependencies exist between tasks, while a Parallel Agent executes independent analysis tasks concurrently. This hybrid design ensures both logical correctness and efficient processing.

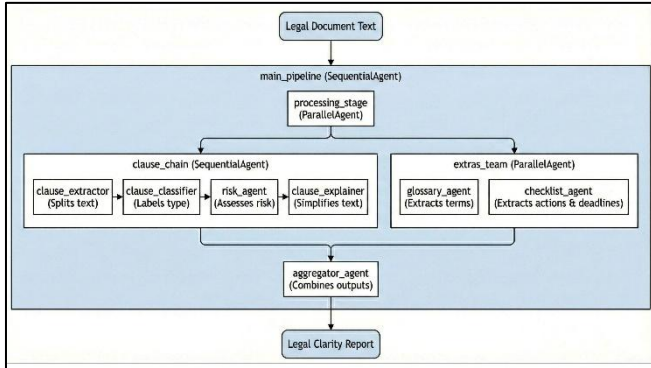


Fig. 1. High-level architecture of the Legal Clarity Agent showing sequential and parallel agent orchestration for clause-level legal analysis.

V. IMPLEMENTATION

The system is implemented entirely in Python and executed within a Jupyter or Kaggle notebook environment [2].

A. Clause Analysis Pipeline

A sequential chain of agents performs core legal analysis [3]:

- Clause Extractor splits the document into individual clauses
- Clause Classifier labels each clause by legal intent
- Risk Agent evaluates potential legal risk
- Clause Explainer converts legal language into plain English

B. Auxiliary Extraction Agents

In parallel, additional agents extract:

- Legal glossary terms
- Action items and deadlines

C. Aggregation

An aggregator agent combines all outputs into a structured JSON object and generates a readable Markdown-based Legal Clarity Report [3].

VI. RESULTS

The system successfully processed legal documents containing more than 20 clauses per document. Parallel execution reduced processing time while maintaining structured outputs. Generated reports consistently provided:

- Clear clause-level explanations
- Explicit risk labeling
- Extracted obligations and actions

The results demonstrate that agent-based decomposition significantly improves clarity and interpretability compared to traditional summarization approaches [4].

VII. OUTCOME

The Legal Clarity Agent delivers a functional prototype that demonstrates the effectiveness of multi-agent architectures for legal document analysis. The system provides transparent, structured, and privacy-conscious legal insights suitable for non-expert users.

The modular design allows for future extensions such as jurisdiction-specific analysis, confidence scoring, and integration with document management systems [5].

VIII. DISCUSSION AND CONCLUSION

This project demonstrates how multi-agent systems can be effectively applied to high-stakes text analysis tasks such as legal reasoning. By decomposing legal analysis into specialized agents and coordinating them through structured orchestration, the Legal Clarity Agent achieves improved interpretability, scalability, and privacy [4].

The system moves beyond shallow summarization by providing clause-level explanations and explicit risk assessment, making legal documents more accessible and actionable. Future work may focus on expanding legal domain coverage, integrating quantitative evaluation, and deploying the system in real-world legal workflows [5].

REFERENCES

- [1] Blount, A., Gulli, A., Saboo, S., Zimmermann, M., & Vuskovic, V. (2025). Introduction to agents. Google.
- [2] Styer, M., Patlolla, K., Mohan, M., & Diaz, S. (2025). Agent tools and interoperability with model context protocol (MCP). Google.
- [3] Milam, K., & Gulli, A. (2025). Context engineering: Sessions and memory. Google.
- [4] Subasioglu, M., Bulmus, T., & Bakkali, W. (2025). Agent quality. Google.
- [5] Kartakis, S., Hernandez Larios, G., Li, R., Secchi, E., & Xia, H. (2025). Prototype to production. Google.