Argonne National Laboratory, a federal research center, relies on secure AI systems to support operations and scientific research. Its generative AI tool, Argo, provides contextual responses from internal documents while maintaining data confidentiality. Currently, Argo needs to be extended to better meet user needs, especially for streamlining common task types and ensuring that compliance standards are met.

The team integrated lightweight AI agents to interpret user intent, query internal databases, and generate accurate responses. The team first developed a literature review generator to summarize internal documents. After uploading files and entering a research question, users receive structured summaries in four sections: introduction, findings, gaps, and conclusion (Figure 1).

Our team also developed a requirements extraction pipeline to streamline Argonne's compliance process. The system retrieves relevant internal contract information and external legal references from a vector database. The retrieved content is then passed into a large language model (LLaMA 3), extracting explicit obligations from a structured prompt (Figure 2). Expert reviewers from Argonne found that the outputs aligned well with the contract's legal expectations. Across the evaluation set, the model achieved an average similarity score of ~0.65 against human-annotated ground truth, demonstrating strong potential for enabling traceable, AI-assisted compliance workflows.

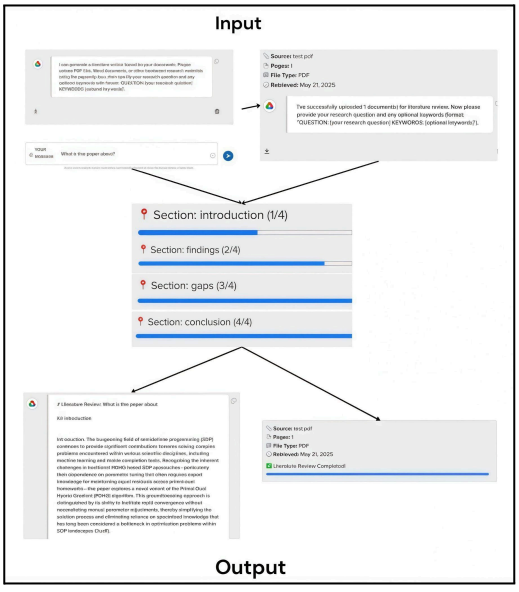

Figure 1: Example literature review workflow in Argo: users upload research documents and receive structured summaries

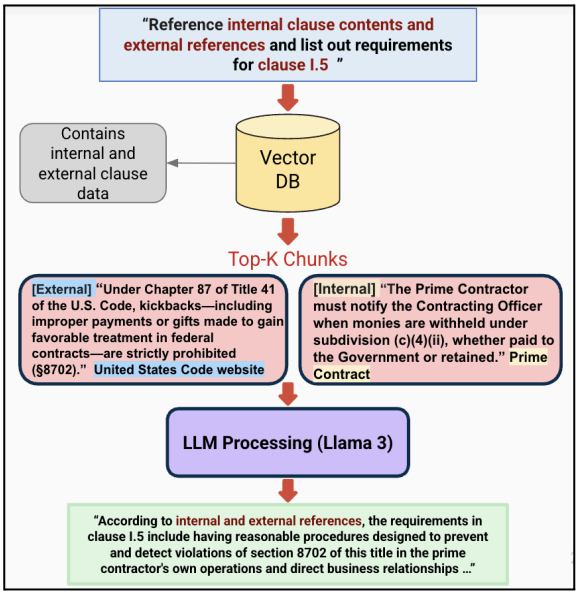

Figure 2: Schema of requirement management process for extracting internal and external content from Clause I.5 of Argonne's legal document