AI legal due diligence: how an LLM workflow handles M&A data rooms under lawyer review

26 November 2026 · 4 min read · blogs.iictp.org/ai-legal-due-diligence-data-room-llm-workflow-explained

Short answer: AI legal due diligence can classify data room documents, extract facts with page-level evidence and prepare report entries, while lawyers keep control of review and judgment, as described in the LEXI system paper posted on 25 November 2026. The authors, working with a corporate law firm, call it work in progress and report early results on two small data rooms. It matters because the paper shows where human checkpoints belong in an LLM review workflow.

TL;DR The short version

- LEXI, described in a paper posted on 25 November 2026, uses an LLM to classify data room documents, extract facts with page-level evidence and draft report entries.
- Lawyers keep control of review and judgment, and the authors call it work in progress, with early results on two small data rooms.
- The useful lesson is where to put human checkpoints in an LLM review workflow, not the accuracy figures.

Read online, listen to it, or change the text size: https://blogs.iictp.org/ai-legal-due-diligence-data-room-llm-workflow-explained

What is LEXI?

LEXI stands for Legal Extraction and Intelligence. Johannes R. Hübers and colleagues describe a system that takes virtual data room documents from ingestion to a draft report. The paper was accepted as a workshop paper at IEEE BigData 2026.

How the workflow works

- **Ingestion:** optical character recognition and layout analysis keep text with page and position.
- **Classification:** documents receive candidate labels from a taxonomy of 72 document types in 14 due diligence areas, plus a label for unreadable or restricted files.
- **Extraction:** each answer is marked exact, inferred or not found, with verbatim snippets and page numbers.
- **Verification of quotes:** a snippet is highlighted on the source page only if it can be located in the parsed document.
- **Report:** only reviewed findings feed the draft; lawyers' comments and risk ratings are fixed inputs, and the output is an editable Word document.

What is new compared with contract-review benchmarks

The authors note that well-known legal datasets study individual English contracts. A real data room mixes shareholder agreements, filings, employment contracts and more, often with unreliable folders. LEXI joins classification, extraction, review and drafting with two lawyer gates: routing and factual review.

Results reported by the authors

- Classification was tested on two real German transaction data rooms with 76 and 73 documents.
- With reasoning enabled, the largest model tested ranked the primary label first for 98.1% and 95.2% of documents, each figure a mean of two runs.
- Reasoning raised first-choice accuracy in all six model and data room comparisons.
- A human expert checked 41 German extraction answers from two models and found all matched, comparing answers only, not the cited snippets.
- The authors say that sample is too small for a robust accuracy estimate.

Limitations and open questions

- Two data rooms from one legal system cannot show performance elsewhere.
- Report preparation has not been formally evaluated, and no review-time saving is measured.
- The taxonomy reflects German practice.
- Missed documents and missed clauses are the costly errors, and these results do not measure them across a full workflow. This is an arXiv preprint or workshop paper. Treat its numbers as the authors' own results, not as independently confirmed findings. Sources were read on 26 November 2026. Products, rules and standards change, so check the linked originals before you rely on any detail.

How to apply the design in your firm

1. Fix your document taxonomy and the questions per document type before choosing a model.
2. Require every extracted fact to carry a quotation and page, and verify the quotation exists in the source.
3. Allow "not found" and forbid guessing in the instructions, then test that behaviour.
4. Block unreviewed findings from reaching the report.
5. Measure on closed matters: classification accuracy, missed items, edit rate and reviewer time.
6. Confirm confidentiality, privilege and data location before any client file is processed.

What it means for transactional lawyers

Ask vendors to show the two gates: who confirms classification, and who approves each fact. A tool that drafts conclusions directly from files makes supervision hard. Keep legal assessment, recommendations and risk ratings with the lawyer, which is also the division the paper uses.

Frequently asked questions

- **What is AI legal due diligence?** Using AI to organise and extract information from transaction documents under lawyer supervision.
- **Can it replace lawyer review?** The paper keeps lawyers responsible for assessment and final wording.
- **How accurate is it?** Early classification results are high on two small data rooms; extraction testing covers only 41 answers.
- **Is the research peer reviewed?** The authors say it was accepted as a workshop paper; it remains early work.

Related reading

- What is RAG?
- Chatbots and unauthorised practice of law
- Private AI for regulated industries
- AI Practitioner certification

> Related credential: AI Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Hübers et al. From Data Room to Report: End-to-End AI Support for M&A Legal Due Diligence. arXiv 2611.34120. 25 November 2026 (IEEE BigData 2026 workshop paper). arxiv.org/abs/2611.34120

Published by IICTP, an independent technical credential body, on 26 November 2026. The current version of this post is at https://blogs.iictp.org/ai-legal-due-diligence-data-room-llm-workflow-explained. More guides: blogs.iictp.org.