Избасар Асем Бахытжановна
Казахстанский институт менеджмента, экономики и прогнозирования

Измеримые эффекты использования документальной аналитики на основе LLM в процессе M&A due diligence

Аннотация. В обзоре количественно оценивается влияние аналитики документов на основе больших языковых моделей (LLM), включая retrieval augmented generation (RAG) и агентный искусственный интеллект, на процесс due diligence со стороны покупателя при сделках слияний и поглощений среднего рынка. Проведена быстрая оценка доказательств по регуляторным и профессиональным вопросам, отраслевым кейсам, рецензируемым исследованиям и вендорским отчетам. Анализ сосредоточен на измеримых показателях, таких как время проверки документов, качество выявления рисков и механизмах участия человека. Сопоставленные данные показывают сокращение времени проверки на 70-75% по сравнению с ручным анализом. Оптимизированные RAG-конфигурации достигают точности анализа договоров до 95%, что сопоставимо или превосходит предыдущие решения машинного обучения. Банки, фонды прямых инвестиций и юридические фирмы отмечают ускорение сделочных циклов и повышение эффективности выявления рисков при сохранении качества. Участие человека остается ключевым элементом, обеспечивающим надежность и управляемость технологии.
Ключевые слова: due diligence в сфере M&A, большие языковые модели, генеративный искусственный интеллект, анализ документов, retrieval-augmented generation, эффективность, legaltech, сделки среднего рынка.

Izbassar Assem Bakhytzhankyzy
KIMEP University

Measurable impacts of LLM-powered document intelligence on M&A due diligence

Abstract. This review quantifies how large language model (LLM) powered document intelligence, including retrieval augmented generation (RAG) and agentic AI, is reshaping buy side M&A due diligence in mid market deals. We conduct a rapid evidence assessment across regulatory and professional surveys, industry case studies, peer reviewed studies, and vendor whitepapers. We focus on measurable indicators: time to review, issue detection quality, and human in the loop controls. Triangulated findings show 70–75% reductions in document review time versus manual baselines. Optimized RAG configurations reach contract analysis accuracy up to 95%, matching or surpassing earlier machine learning tools. Banks, private equity firms, and law firms report faster deal cycles, stronger risk flagging, and maintained or improved quality with appropriate oversight. Humans in the loop remain essential to manage hallucinations and privacy constraints without erasing efficiency gains. The study consolidates metrics from real deployments and highlights governance practices required for adoption in the mid market.
Keywords: M&A due diligence, large language models, generative AI, document review, retrieval-augmented generation, efficiency, legaltech, mid-market transactions.

1. Introduction
Due diligence requires intensive review of contracts, financials, and compliance materials to surface pre-close risks [1]. In mid-market buy-side contexts, constrained resources and deadlines heighten the risk of oversight [2]. LLMs (e.g., GPT-3.5/4) now enable automated contract analysis and corpus-level Q&A, promising faster evaluations without sacrificing rigor.

Surveys suggest normalization of AI in diligence: 97% of U.S. M&A lawyers (82% Canada) agree AI is becoming standard [3], while Bain reports generative-AI use rising from 16% (2023) to 21% (2024) with >50% expected by 2027 [4]. Compared with earlier ML clause-extraction, LLMs are more adaptable, handling novel provisions and nuanced context [4][5]. Risks persist—hallucinations, opacity, data-privacy constraints, and emerging regulation (e.g., EU AI Act)—necessitating HITL and secure deployments [2][6]. This review aggregates 2023–2025 outcomes to answer: How much faster is AI-assisted diligence? Does quality improve? Which controls work in practice? We focus on mid-market transactions where AI's leverage is greatest.

2. Methodology

2.1 Rapid Evidence Assessment Approach

We conducted a rapid evidence assessment of 2023–2025 sources spanning (i) regulator/central-bank and professional surveys; (ii) consulting and industry reports; (iii) peer-reviewed papers and preprints; (iv) documented bank/PE use cases; and (v) vendor technical studies. Inclusion required quantitative metrics on (1) time efficiency; (2) quality (error/recall/issue-detection); (3) HITL controls and productivity changes. Searches covered IEEE Xplore, SSRN, arXiv, and the open web using terms such as "AI due diligence M&A metrics," "LLM document review case study," and "due diligence efficiency AI 2024," plus backward citation tracking. Twenty-eight sources met screening criteria. We treat results as indicative given pilot-stage maturity, and we report times as relative percentages or deltas, following source conventions .

3. Results

3.1 Adoption of LLM-Powered Due Diligence and Use Cases (2023–2025)

Adoption is accelerating. Bain (2025) reports 21% of companies using gen-AI in M&A (36% among frequent acquirers), with >50% expected by 2027 [8][9]. Over 60% of large PE firms used at least one AI tool by late-2024 for sourcing, screening, or diligence [10]. Prime diligence applications are (a) document review/analysis (contracts, disclosures) via LLM+RAG; and (b) targeted Q&A over data-room corpora. Law-firm pilots (e.g., Addleshaw Goddard, 2024) applied GPT-4 with optimized retrieval to hundreds of contracts; industry surveys indicate parallel use in financial diligence (e.g., parsing semi-structured ledgers for anomalies) [11]. The hybrid model—AI for volume, humans for judgment—dominates early deployments.

Table 1. Selected deployments of LLM-powered diligence

Organization / Context	Use Case	AI Approach	Reported Outcome
Addleshaw Goddard (UK, 2024)	Legal DD—contracts in data room	GPT-4 + RAG (custom chunking/prompts)	Accuracy ↑ from 74% to ~95%; +16% recall via targeted prompts.
Deloitte / Mid-market PE (2024)	Buy-side contracts & disclosures	GenAI vs. manual & prior ML	~75% faster than manual; ~25% faster than older AI; comparable risk findings.
PE fund example (2024)	~10–15k-page data room	LLM-based VDR summarizer	DD in ~2 weeks vs. ~6; earlier close; no critical issues missed.

These cases show repeatable gains in mixed human-AI workflows, with mid-market teams benefiting disproportionately due to smaller headcounts and compressed timelines.

3.2 Efficiency Gains: Time and Cost Reduction

Across sources, AI-assisted review reduces effort by roughly one-half to three-quarters. Savings concentrate on document-heavy tasks (contracts, financial statements, compliance checks). Deal calendars compress from weeks to days, directly affecting competitiveness.

Table 2. Reported efficiency gains

Source (Year)	Context	Efficiency Gain	Note
Deloitte (2024)	Legal DD (contracts)	~75% faster vs. manual; ~25% faster vs. prior AI	Maintained quality with fewer reviewers.
Open-source synthesis (2025)	Mid-market financial DD	~70% faster; 20–30% cost savings	Example: 6–8 weeks → 2 weeks.
Bain (2024)	Diligence summarization	~85% faster for summarization	"One day vs. one week" for deep-dive summaries.

Optimization matters: robust chunking, hybrid retrieval, and prompt engineering are prerequisites for speed and accuracy [11]. Some studies note increased deal throughput post-adoption (more simultaneous deals with same headcount) [9].

3.3 Quality, Accuracy, and Human-in-the-Loop Controls

Quality remains high when workflows are engineered and supervised. In Addleshaw Goddard's experiments, optimized RAG prompts raised accuracy to ~95%, matching or slightly exceeding prior ML tools; recall rose with keyword-guided prompts and "follow-up" instructions [12]. Deloitte cases cite multilingual analysis, prioritization via risk scoring, and clearer summaries as qualitative boosts [13]. Hallucination risk is mitigated through HITL checks, auditability (source-linked evidence), and conservative deployment. Emerging evidence suggests improved breadth of coverage and occasional detection of obscure issues (e.g., scattered environmental liabilities) [15]. Yet qualitative "soft signals" (culture, management probity) remain human-led [6].

Table 3. HITL controls and best practices

Control	Description	Reported Benefit
Task-specific prompts & follow-ups	Explicit instructions; "double-check" behaviors	↑ Accuracy/recall; fewer careless errors.
Chunked RAG with hybrid search	Right-sized chunks; semantic + keyword retrieval	Faster, more relevant context; ↑ precision.
Human validation & sampling	Review of high-risk flags; spot-checks of "no issue" docs	Catches residual errors; supports legal sign-off.
Audit trails & traceability	Source links to pages/paragraphs; logs	Trust and regulatory comfort.
Secure deployment	On-prem/private models; VDR-integrated AI	Data-privacy compliance; client assurance.

4. Discussion

LLM-powered document intelligence delivers large efficiency gains in buy-side M&A due diligence without proportional accuracy loss; in several respects, quality improves. Echoing broader knowledge-work trends, properly applied GenAI raises productivity while maintaining output quality. Field studies outside M&A report 20–50% gains; in due diligence we find 50–75% time savings, as AI excels at standardized text analysis. Mid-market deals benefit most: compressed timelines and lean teams make overnight first-pass reviews transformative. Measured impacts (e.g., 70% faster reviews) create competitive advantages—quicker bids, parallel diligence on more targets, lower fees, and reduced drudgery that aids retention. Over time, routine tasks may enable "always-on" diligence.

Limits remain. Confidentiality is paramount: firms weigh cloud LLMs versus private/on-prem solutions; VDRs now embed AI to mitigate risk, and compliance with NDAs and professional rules is essential. Tools alone don't create value—workflow integration, training, and change management matter; efficiency may dip during a learning curve. Liability stays with

supervising professionals (e.g., ABA guidance), reinforcing human-in-the-loop oversight; bias risks must be monitored.

LLMs differ from earlier, narrowly trained clause-finders (2015-2020) by understanding context with less setup. Evidence suggests GenAI adds 25% efficiency over "traditional AI-enabled diligence," lowering barriers for mid-market adopters. Vendor claims broadly align with independent findings: high time savings and 90-95% accuracy are achievable, with the caveat that humans must review AI outputs.

LLMs are leveling the field: mid-market acquirers can do in days what once took weeks, accelerating deal cycles and increasing scrutiny; sell-side teams will also adopt AI to pre-empt issues. Adoption momentum is strong; late adopters may face disadvantages as benefits compound.

Also, AI will permeate target identification, diligence, execution, and integration. Expect deeper links to deal platforms, regulatory checks, and multimodal analysis (text, images, tables). Ethical and workforce effects include shifting junior roles and pressure on billable-hour models, potentially toward value-based fees. Net effect-human+AI should raise diligence quality and reduce post-merger surprises. Triangulated sources consistently indicate order-of-magnitude time reductions (weeks to days), even if not universal across all deals.

5. Conclusion

Between 2023 and 2025, LLM-enabled document intelligence has materially improved mid-market M&A due diligence. Time-to-review commonly falls by 50–75%, while tuned RAG approaches achieve ~95% contract-analysis accuracy. Crucially, HITL controls and traceability preserve legal standards and mitigate hallucination risk. Rather than trading speed for thoroughness, teams increasingly achieve both, accelerating deal cycles and broadening coverage. This consolidated, metric-focused review provides a baseline for practitioners benchmarking adoption and for researchers mapping where quality and governance require further work. With responsible deployment, LLMs function as force multipliers—expanding what lean deal teams can accomplish within tight timelines and budgets—while leaving judgment and accountability squarely with human experts.

References
1.	D. Acemoglu, L. Bursztyn, M. Demirer, L. Fahey, R. Galbiati, I. Méjean, S. Noy, S. Peng, J. Regier, and W. Zhang, "The simple macroeconomics of AI," *NBER Working Paper* No. 32487, 2024, doi: 10.3386/w32487.
2.	J. Villasenor, "Generative artificial intelligence and the practice of law: Impact, opportunities, and risks," *Minnesota Journal of Law, Science & Technology*, vol. 25, p. 25, 2023.
3.	Litera, "Technology in M&A report: AI, tech adoption, and talent management in US and Canada," 2023. [Online]. Available: https://info.litera.com/north-america-tech-in-mergers-and-acquisitions-survey-report. Accessed: Sep. 25, 2025.
4.	Bain & Company, "Generative AI in M&A: You're not behind—yet," *2025 M&A Report*, Feb. 4, 2025. [Online]. Available: https://www.bain.com/insights/generative-ai-m-and-a-report-2025/. Accessed: Sep. 28, 2025.
5.	M. Jang and G. Stikkel, "Leveraging natural language processing and large language models for assisting due diligence in the legal domain," in *Proc. NAACL-HLT 2024, Industry Track*, June 2024, pp. 155–164. Association for Computational Linguistics.
6.	EY, "How AI will impact due diligence in M&A transactions," 2024. [Online]. Available: https://www.ey.com/en_ch/insights/strategy-transactions/how-ai-will-impact-due-diligence-in-m-and-a-transactions. Accessed: Sep. 30, 2025.
7.	M. Deshpande, "Accelerating M&A: AI-driven data acquisition guide for integration program managers," *Journal of Artificial Intelligence, Machine Learning & Data Science*, vol. 1, no. 4, pp. 642–647, 2023.
8.	Addleshaw Goddard, "Harnessing the power of large language models for legal document review—The RAG Report: Can LLMs be good enough for legal due diligence?,"

Addleshaw Goddard LLP Innovation Research, 2024. [Online]. Available: https://www.addleshawgoddard.com. Accessed: Sep. 20, 2025.

9.	Artificial Lawyer, "AG proves LLMs can do well on due diligence," Oct. 8, 2024. [Online]. Available: https://www.artificiallawyer.com/2024/10/08/ag-proves-llms-can-do-well-on-due-diligence/. Accessed: Sep. 22, 2025.

10.	Bain & Company, "Generative AI in M&A: Where hope meets hype," 2024. (Referenced in: I. Käyhkö, *The Disruption of Due Diligence*, Aalto University, 2025).

11.	Deloitte Legal, "The future of legal work? The use of generative AI by legal departments," Deloitte (UK) Survey Report, June 2024. [Online]. (Excerpt via Lexology). Accessed: Oct. 8, 2025.

12.	Z. Han, "Does AI adoption in M&A teams improve deal performance?," *SSRN Working Paper* No. 5128530, Dec. 27, 2024.

13.	M. A. Bedekar *et al.*, "AI in mergers and acquisitions: Analyzing the effectiveness of artificial intelligence in due diligence," in *Proc. 2024 Int. Conf. on Knowledge Engineering and Communication Systems (ICKECS 2024)*, Apr. 2024, pp. 1–5, doi: 10.1109/ICKECS.2024.10616599.

14.	Y. Zhao *et al.*, "Utilizing large language models for information extraction from real estate transactions," *arXiv*:2404.18043 [cs.CL], Apr. 2024.

15.	OpenLedger, "AI in M&A accounting: Transforming financial due diligence in 2025," 2025. [Online]. Available: https://www.openledger.com/future-of-ai-in-accounting/ai-in-m-a-accounting-transforming-financial-due-diligence-in-2025. Accessed: Sep. 28, 2025.

Сведения об авторе
Избасар Асем Бахытжановна, студентка, Казахстанский институт менеджмента, экономики и прогнозирования, Алматы, Казахстан.

Научный руководитель
Хамзина Аяго, кандидат технических наук, Школа цифровой государственной политики, Департамент науки и инноваций, Астанинский университет информационных технологий, Астана, Казахстан

Information about the author
Izbasar Asem Bakhytzhanovna, student, Kazakhstan Institute of Management, Economics and Forecasting, Almaty, Kazakhstan.
Scientific supervisor
Khamzina Ayago, PhD, School of Digital Public Policy, Department of Science and Innovation, Astana University of Information Technology, Astana, Kazakhstan