Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 00 (2019) 000–000

Procedia
Computer Science

www.elsevier.com/locate/procedia

The 17th International Conference on Ambient Systems, Networks and Technologies (ANT)
April 14-16, 2026, Istanbul, Türkiye

TermSheetPilot: An LLM-Powered Agent for Filling Venture Term Sheets from Contracts for Founders and Venture Capitalists

Amir Sartipi[a,*], Muriel-Larissa Frank[a], Afshin Khadangi[a], Igor Tchappi[a], Gilbert Fridgen[a]

[a] *SnT - Interdisciplinary Center for Security, Reliability and Trust, University of Luxembourg, Luxembourg*

Abstract

Venture capital term sheets are critical instruments for structuring early-stage investment agreements, yet their manual preparation remains time-consuming and dependent on specialized legal expertise. Recent advances in natural language processing (NLP), particularly large language models (LLMs), have opened new opportunities for automating the interpretation of complex financial and legal documents. This paper frames the challenge of term sheet drafting as an information extraction problem, aiming to generate standardized term sheets automatically by extracting and structuring data from unstructured contract texts. To support this process, we constructed a domain-specific term dictionary derived from real investment agreements obtained through the U.S. Securities and Exchange Commission's (SEC) Electronic Data Gathering, Analysis, and Retrieval (EDGAR) system. The dictionary organizes terms under a set of umbrella categories and provides clear operational definitions for each element, enabling consistent and interpretable extraction.

Building on this foundation, we developed an intelligent LLM empowered agent specialized in financial contracts that can extract key deal terms with contextual awareness. The system was evaluated on a corpus of 59 contracts sourced from EDGAR, benchmarking it against expert annotations across multiple contractual categories. Though the model had difficulty identifying context-dependent clauses in certain categories, it achieved an overall correspondence rate of 86.92% with human evaluations. These findings underscore both the promise and the limitations of scalable automation in venture-financing document analysis, suggesting that further domain calibration and fine-tuning are necessary for optimal accuracy. Overall, the results highlight the potential of LLM-driven systems to structure and surface relevant information, streamline due diligence, and support the development of AI-assisted tools for contract intelligence and venture-finance decision-making.

Keywords: Large Language Models (LLMs); Information Extraction; Entrepreneurship; Venture Capital

* Corresponding author.
E-mail address: amir.sartipi@uni.lu

1877-0509 © 2020 The Authors. Published by Elsevier B.V.
This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/)
Peer-review under responsibility of the Conference Program Chairs.

1. Introduction

Entrepreneurship is widely recognized as a fundamental driver of innovation, economic growth, and job creation, particularly through the establishment of startups, which are new ventures often distinguished by high growth potential and innovative business models [25]. In their formative stages, startups typically lack sufficient internal capital and thus turn to external financing sources to scale operations and commercialize innovations. The entrepreneur–investor relationship is reciprocal: entrepreneurs contribute innovative opportunities and managerial effort, while venture capitalists provide capital, expertise, and governance support. One prominent mechanism for securing such financing is through venture capital (VC), where venture capitalists provide capital in exchange for equity and a stake in the company's future success [8, 15]. Yet, venture capital is more than a financial transaction: it is a contractual relationship structured to address uncertainty, asymmetric information, and the need for control in dynamic entrepreneurial settings. Today's so-called term sheet originated in the late 1950s as an informal letter outlining investment terms, with one of the earliest examples found in the funding of Digital Equipment Corporation (DEC) by American Research and Development Corporation (AR&D). Since then, it has evolved and expanded into a widely used document in entrepreneurial finance, incorporating a range of complex and negotiated provisions [7]. In early-stage venture financing, the term sheet, while non-binding in most cases, functions as a key contractual scaffold through which venture capitalists and founders define the economic and governance framework of a financing round. Although it only outlines the basic terms of the investment agreement, it plays an essential role in shaping trust and control by clarifying the perceived risks for both parties. Thus, serving as a tool for risk allocation and the distribution of control rights, the term sheet also influences value creation and governance outcomes in entrepreneurial ventures [9].
Prior empirical analysis and research shows that various provisions embedded in term sheets, such as liquidation preferences, anti-dilution protections and investor control rights, can materially affect startup outcomes and the split of value between parties [14, 4, 23, 18]. In addition, research shows that increasingly complex and opaque clauses, for example stacked liquidation preferences, participating preferred structures and IPO ratchets, can render venture term sheets "toxic" by shifting value in non-obvious ways across financing rounds [6]. In such cases, the purported transparency and a shared understanding of risk and payoff provided by the term sheet may be undermined, with founders, employees and early venture capitalists underestimating their true exposure to downside risk [9].
These challenges highlight the need for better tools to decode and demystify complex legal-financial documents in venture financing. The past decade has witnessed significant progress in NLP technologies, particularly through transformer-based architectures and Large Language Models (LLMs). These models have been applied in various applications [29, 17, 13] and have expanded the capability of machines to interpret, generate, and reason over unstructured textual data at scale [12, 26]. Applied to the legal domain, this shift has enabled tasks like contract clause extraction [24], document summarization [2], and legal text classification [10] with unprecedented precision and speed [27].
In this paper, we bring these advances to venture capital contracting by processing full contract texts and assigning concrete values to predefined deal terms. As a foundation, we build a dictionary of terms grouped into 16 overarching topics and use it to guide expert annotation on 59 contracts. An intelligent agent equipped with an LLM then uses this dictionary to search for and extract term values, and the authors evaluate its behavior across 59 contracts in total, all sourced as PDF filings from the SEC's EDGAR system.

2. Related Work

Across legal and financial domains, LLM-based extraction methods show that unstructured documents can be converted into structured data with relatively lightweight supervision [5, 19, 28, 22]. In legal term extraction, LLMs can be prompted and weakly supervised to identify fine-grained contractual clauses with limited annotated data, outperforming traditional sequence-labeling baselines in low-resource settings [5]. Complementary work in accounting and information systems proposes a design-science LLM framework that combines text mining with prompt engineering to automatically extract key financial indicators from PDF-formatted governmental annual reports and corporate ESG disclosures, demonstrating that LLM-based extraction can reliably support both academic analysis and regulatory or industrial workflows [19]. Closely related research on hybrid long documents introduces frameworks where long, text–table financial reports are chunked, routed, and summarized before LLM-based extraction, and where simple

table serialization plus targeted prompting suffices to recover complex numerical facts at scale [28]. At the document level, general-purpose architectures for information extraction and localization integrate textual content with layout features, enabling clause-level field filling and bounding-box prediction across diverse forms and contracts [22].
Within the legal domain more broadly, surveys synthesize tasks, datasets, and modeling strategies for legal NLP, highlighting challenges such as extremely long documents, domain-specific terminology, limited open corpora, and privacy constraints [3]. Against this backdrop, case studies on attribute extraction from legal documents evaluate LLMs as schema-driven extractors, systematically analyzing prompt design, calibration, and robustness when mapping contracts to structured attributes [1]. Beyond prompting general models, domain-adapted transformers for legal named entity recognition and anonymization have been proposed, showing that pretraining and fine-tuning on legal corpora yields substantial gains in both entity coverage and GDPR-compliant redaction, with throughput suitable for production workflows [16]. In parallel, benchmark efforts for legal reasoning assemble large suites of tasks covering statutory interpretation, contract reasoning, and procedural judgments, establishing standardized evaluations for both open-source and commercial LLMs and underscoring the gap between surface extraction and deeper legal reasoning capabilities [11].
For example, Narendra et al. [21] examines how LLMs can support end-to-end contract workflows rather than only single-document parsing. The authors introduce an LLM-based comparison system that aligns contracts with templates, identifies deviations, categorizes differences by risk or negotiability, and generates natural-language explanations tailored to negotiation scenarios [21]. Such a system illustrates how structured representations of contract terms, once extracted and normalized, can be leveraged for tasks such as redlining, playbook enforcement, and review across multiple drafts or counterparties. Our focus on schema-guided extraction of venture financing terms from real-world contracts is situated in this landscape: it adopts LLM-based extraction and hybrid-document processing ideas from general legal and financial settings [5, 19, 28, 22, 1], builds on insights from legal NLP surveys, domain-adapted models, and benchmarks [3, 16, 11], and targets representations that can directly feed into comparison and negotiation tools for venture term sheets [21].

3. Data Collection

This study required access to a corpus of real-world venture-financing contracts to enable the development and validation of an intelligent agent capable of extracting investment terms from unstructured legal documents. The dataset was designed to capture the diversity of contractual language and structural variations typical of early-stage financing agreements.
Contracts were gathered from the *Electronic Data Gathering, Analysis, and Retrieval (EDGAR)* system[1], which provides open access to official filings submitted to the U.S. Securities and Exchange Commission (SEC). Within this database, *Form C* filings[2] were targeted because they disclose comprehensive investment documentation for crowdfunding rounds below five million USD. Each filing contains multiple attached documents; only those explicitly outlining investment terms were extracted for inclusion. From these, a total of 59 contracts were selected after screening for completeness and relevance. The collected contracts formed the foundation for constructing a domain-specific dictionary of recurring venture-finance terms observed across various contractual types. This dictionary subsequently guided the identification of corresponding clauses and value expressions within the contracts. Wixdom domain xperts[3] annotated all documents, and their expertise ensured precise term mapping. The annotated corpus established the empirical basis for testing the proposed intelligent extraction agent presented in Section 4.

4. Intelligent Agent

The following subsections describe each step in the pipeline, starting from document preprocessing and continuing to evaluation and performance metrics, before presenting the experimental results and evaluation.

[1] `https://www.sec.gov/edgar/search/`
[2] **Form C** refers to filings under *Regulation Crowdfunding* for fundraising rounds capped at $5 million USD, which include the disclosure of key investment terms and governance details.
[3] `https://next.wixdom.io/`

4.1. Document Preprocessing

Contracts from the EDGAR system are provided as heterogeneous PDFs (born-digital and scanned), complicating direct processing and accurate text extraction. We therefore built a preprocessing pipeline that converts all PDFs to machine-readable text using a custom parser with OCR for image-based files. Structural markers such as section headers, numbering, and page identifiers are retained, as they provide essential context for identifying relevant provisions. Text normalization procedures remove artifacts, unify whitespace, and standardize encoding to ensure that subsequent extraction operates on clean and coherent text.

4.2. Schema Construction

Following the analysis of collected contracts, a structured term dictionary was developed to support schema-based extraction. Each entry in this dictionary represents a distinct contractual concept categorized under one of several overarching categories. Within each category, multiple sub-terms were defined to capture recurring legal and financial expressions identified across the dataset. These categories collectively form the semantic foundation of the extraction schema and serve as the primary reference for identifying corresponding term values in the contracts.
For each term, its allowable data type (e.g., binary, string, or list of specified values) was considered to ensure consistent representation during extraction. Additionally, for every coded term, a concise definition was generated automatically by GPT-4 to provide supplementary context to the intelligent agent, thereby narrowing its interpretive space and improving the precision of semantic reasoning. Table 1 summarizes the investment term categories and their respective counts.

Table 1. Overview of dictionary categories and corresponding term counts.

Category	No. of Terms	Category	No. of Terms
ANTI-DILUTION RIGHTS	3	CAPITALIZATION RATIO	12
LIQUIDATION PREFERENCE	24	TEMPLATE	11
VOTING PROXYS	11	REVENUE SHARES	38
TYPE OF SECURITY	35	OTHER SECURITIES	11
PURPOSE OF INVESTMENT	8	INFORMATION RIGHTS	6
CAPITAL STRUCTURE	17	DIVIDENDS	5
CONVERSION MECHANISM IDENTIFICATION	53	IPO LOCKUP	6
OPTION POOL / EMPLOYEE STOCK OWNERSHIP PLANS (ESOPs)	6	PRO-RATA INVESTMENT RIGHTS	7
PRE-EMPTIVES	3	IPO RATCHET	2
PARTICIPATION PREFERRED	3		
		Total Terms	**261**

4.3. Guided Reasoning

The intelligent agent was designed to execute schema-based information extraction through a structured reasoning process, extending beyond surface-level pattern recognition. Conceptually, the agent operationalizes a hybrid paradigm that integrates symbolic schema alignment with contextual inference. Its architecture was developed to achieve three core objectives: (i) the accurate extraction of investment terms in compliance with a formally defined schema, (ii) interpretability of outputs through explicit reasoning traces, and (iii) mitigation of model hallucination by enforcing evidence-based grounding within the source text. In this formulation, the agent functions as a domain-aware analytical system that both identifies contractual attributes and justifies their extraction through verifiable linguistic cues.
The reasoning behavior of the agent is guided by data-type constraints, and term descriptions formulated in Subsection 4.2. These predefined semantic constructs delimit the inference space of the model, establishing an explicit mapping between the abstract legal-financial lexicon and machine-interpretable categories. Each definition operates as an inductive prior that steers the model's contextual understanding, thereby constraining interpretive ambiguity and promoting semantic consistency across documents. In effect, the schema serves as a controlled reasoning scaffold that transforms free-text legal language into structured, operational data representations. During the extraction process,

Table 2. Category-level extraction performance of the schema-guided agent across 59 venture-financing documents. The table summarizes per-category weighted and average accuracies, coverage metrics, and dispersion indicators.

Category	Total Contracts	Weighted Acc. (%)	Avg. Acc. (%)	Terms Eval.	Matches	Min (%)	Max (%)	Median (%)	SD (%)
ANTI-DILUTION RIGHTS	1.00	100.00	100.00	2.00	2.00	100.00	100.00	100.00	0.00
CAPITALIZATION RATIO	1.00	100.00	100.00	1.00	1.00	100.00	100.00	100.00	0.00
LIQUIDATION PREFERENCE	28.00	86.92	85.33	214.00	186.00	0.00	100.00	100.00	27.27
TEMPLATE	17.00	81.25	79.41	96.00	78.00	0.00	100.00	83.33	20.24
VOTING PROXYS	23.00	68.00	72.10	50.00	34.00	0.00	100.00	100.00	41.67
REVENUE SHARES	7.00	66.67	62.96	48.00	32.00	25.00	100.00	57.14	24.52
TYPE OF SECURITY	59.00	62.80	61.89	293.00	184.00	0.00	100.00	60.00	22.82
CONVERSION MECHANISM IDENTIFICATION	33.00	61.88	52.69	425.00	263.00	0.00	100.00	60.00	28.38
OTHER SECURITIES	11.00	53.33	68.18	15.00	8.00	0.00	100.00	100.00	44.07
PURPOSE OF INVESTMENT	56.00	52.50	45.09	80.00	42.00	0.00	100.00	25.00	47.11
IPO LOCKUP	4.00	50.00	60.42	12.00	6.00	0.00	100.00	70.83	36.98
INFORMATION RIGHTS	1.00	50.00	50.00	2.00	1.00	50.00	50.00	50.00	0.00
CAPITAL STRUCTURE	56.00	47.95	47.02	415.00	199.00	0.00	100.00	30.95	32.92
DIVIDENDS	7.00	45.45	50.00	11.00	5.00	0.00	100.00	50.00	37.80
OPTION POOL / ESOPs	23.00	20.00	26.09	60.00	12.00	0.00	100.00	0.00	35.04
PRO-RATA INVESTMENT RIGHTS	5.00	16.67	15.00	12.00	2.00	0.00	50.00	0.00	20.00
TOTAL	59.00	60.77	61.40	1736.00	1055.00	17.65	100.00	64.00	19.70

the agent produces for each field both a predicted value and an accompanying justification, which explicitly cites the contractual evidence supporting the output. When a field cannot be located, the model generates a null value and an explanatory rationale describing the absence of relevant text. This dual-output formulation enforces interpretability at the instance level, preventing spurious inference and ensuring traceability of model decisions. By embedding justification generation into the reasoning workflow, the system preserves a transparent audit trail that enables human evaluators to examine the textual grounding and logical coherence of every extracted element.
The prompt design governing the agent encapsulates these guided reasoning principles. It instructs the model to extract only the fields specified within the schema, while mandating that each predicted value be accompanied by a succinct justification grounded in the contractual text. This paradigm redefines extraction as an evidential reasoning task—each output must be supported by linguistic evidence that explicitly substantiates the model's interpretive claim. Such evidence-centric reasoning not only improves reliability but also reduces epistemic opacity, converting the traditionally black-box process of language model inference into a verifiable analytical procedure. For empirical implementation, the agent employed the *GPT-5* model configured for moderate reasoning capability.

4.4. Evaluation and Performance Metrics

We evaluated the intelligent agent's performance using field-level accuracy and dispersion statistics across all venture-financing agreements. For each field, we compared the extracted value to its annotated ground truth after normalization and type harmonization. We then computed the overall extraction accuracy as:

$$\text{Accuracy} = \frac{m}{n} = \frac{\text{Number of Correctly Extracted Terms}}{\text{Total Number of Evaluated Terms}} \tag{1}$$

Weighted accuracy across contracts accounted for the number of annotated terms per document, while dispersion metrics (minimum, maximum, median, and standard deviation) quantified performance variability across categories.

5. Results

The schema-guided intelligent agent was evaluated across sixteen categories of contractual clauses derived from 59 venture-financing agreements. Table 2 reports the detailed performance metrics, including weighted and average accuracies, total evaluated terms, and dispersion indicators.
Across all documents, the model achieved a weighted accuracy of 60.77% and an average accuracy of 61.40%, based on 1,736 evaluated terms and 1,055 correct matches. Accuracy varied across categories, reflecting the linguistic diversity and structural variability of contractual language. The *Liquidation Preference* category demonstrated one of

the strongest results, achieving a weighted accuracy of 86.92% with 186 correct matches out of 214 evaluated terms, indicating consistent recognition of this standardized clause type.
Categories including *Type of Security* (62.80%), *Conversion Mechanism Identification* (61.88%), and *Revenue Shares* (66.67%) also maintained relatively stable accuracy levels, showing that the agent can handle moderately variable but lexically structured provisions with reasonable consistency.
Figure 1 provides a granular view of the extraction results in individual contracts, examining how exact-match accuracy interacts with the scope of the schema, the density of the distribution, and the length of the document.

Exact-match performance across documents. Exact-match scores span the full 0–1 range for all schema sizes, with no strong visual dependence on the number of evaluated terms. The fitted regression line has a shallow negative slope (correlation $r \approx -0.07$), consistent with at most a very weak tendency for scores to decrease as schema size increases. High-scoring contracts (Exact Match ≥ 0.8) occur across a wide range of schema sizes, whereas the lowest scores are somewhat more concentrated among contracts with fewer than about 25 terms. Marker size, which reflects document length, is similarly dispersed across the vertical axis: both short and long contracts appear at higher and lower accuracy levels, and longer documents cluster around intermediate scores rather than at one extreme.

Number of terms vs. exact match. The relationship between schema scope and extraction accuracy shows little apparent dependence, with a small correlation of $r \approx -0.07$. Performance remains broadly distributed across contracts regardless of whether the schema includes around 15 or close to 50 evaluated fields. Larger dots in this plot represent longer documents, and their wide dispersion across accuracy levels suggests that document length does not systematically influence extraction precision when schema size is held constant. This pattern is consistent with the idea that performance variation may arise more from clause-level ambiguity or drafting differences than from either document scale or schema breadth, though this cannot be established causally from the figure alone.

Distributional density of extraction outcomes. The kernel density of document-level exact-match scores is unimodal, with its maximum around 0.6–0.7 and most mass between 0.5 and 0.8. The lower and upper tails below 0.3 and above 0.9 are comparatively thin, so both very low and near-perfect extraction outcomes are rare relative to mid-range performance.

Contract length vs. exact match. Exact-match scores exhibit a weak positive association with contract length ($r \approx 0.27$). Shorter contracts (below roughly 10k words) span almost the full accuracy range, from below 0.2 to above 0.9, while longer agreements cluster more tightly between about 0.6 and 0.9. The fitted regression line rises only gradually, so length accounts for limited variation in performance. Marker size reflects the number of evaluated terms; larger points are distributed across both higher and lower scores, indicating that expanded schema coverage is not systematically linked to either.

6. Implications and Limitations

Building on LLM-based approaches to contract clause extraction, financial-report parsing, and legal document comparison [5, 19, 28, 22], the intelligent agent developed in this work introduces a new application of LLM-powered automation in venture financing. It offers practical benefits to entrepreneurs, founders, venture capitalists, and legal advisors by automating the extraction and organization of contractual terms. In doing so, the tool can serve both as a validation mechanism for drafted term sheets and as an assistive drafting aid, enabling entrepreneurs to work more efficiently without having to master specialized legal terminology. This capability has the potential to reduce negotiation friction, enhance transparency, and accelerate deal-making processes. Moreover, the agent's justification-based outputs support traceable decision-making, strengthening accountability and compliance in financial documentation workflows.
However, several limitations must be acknowledged. The definitions of terms and their corresponding descriptions require rigorous peer review to ensure conceptual accuracy and alignment with legal practice. The annotation process should be evaluated by multiple domain experts to assess inter-annotator consistency (for instance, using Cohen's κ as a measure of annotator cohesion) [20] and minimize subjective bias in label generation. Moreover, the current dataset is limited in both scale and representativeness; expanding it to include contracts from different industries, and

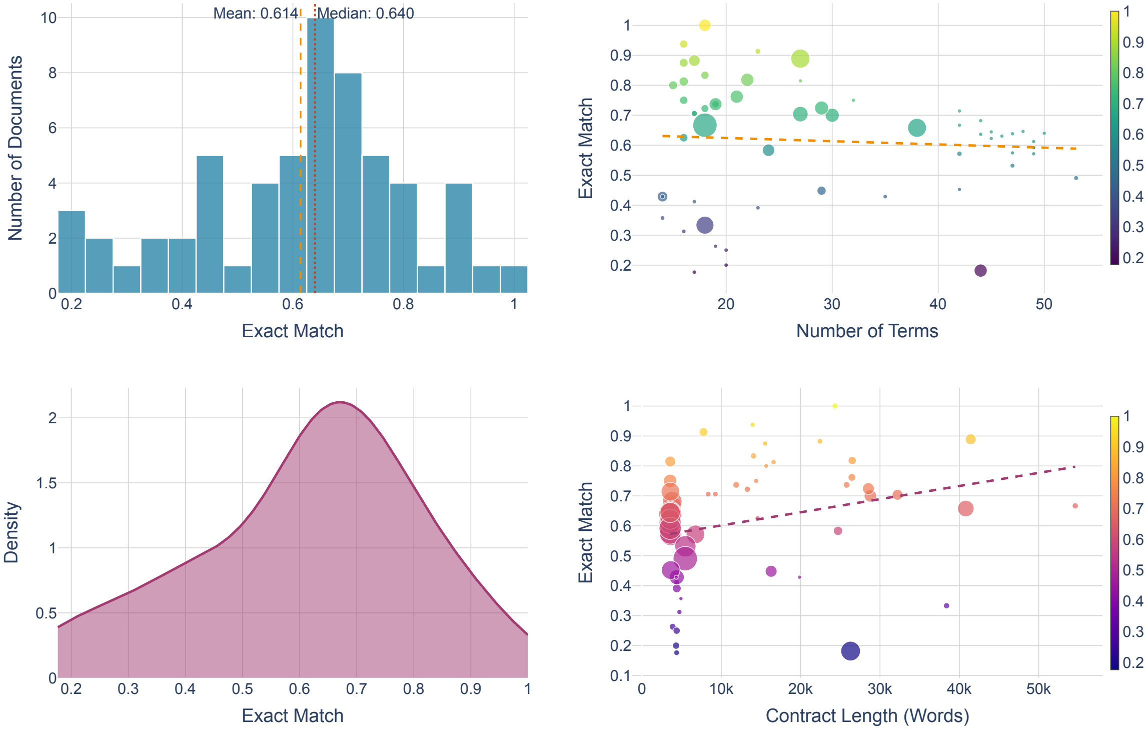

Fig. 1. Exact match analysis across 59 documents showing (top-left) the distribution of exact match scores, (top-right) the relationship between number of terms and exact match where dot size represents document length (larger dots indicate longer documents), (bottom-left) the kernel density estimate of exact match scores, and (bottom-right) the relationship between contract length and exact match where dot size represents the number of terms (larger dots indicate more terms).

investment stages would enhance generalizability. Future research should also experiment with alternative model architectures, reasoning configurations, and prompt-engineering strategies to test robustness across document types and improve the precision of context-sensitive extractions.

7. Conclusion

This study established a baseline for extracting information from venture-financing term sheets using an intelligent agent capable of identifying key contractual values and providing text-grounded justifications. The agent demonstrated consistent performance across diverse clauses, indicating that LLM can support interpretable and reproducible analyses of financial agreements when guided by domain-specific definitions. While the overall accuracy remains moderate, the findings confirm the feasibility of automated extraction within complex legal-financial texts. Further investigation is needed to refine the definitions, improve contextual reasoning, and enhance the robustness of extraction across heterogeneous document formats.

Acknowledgements

This research was conducted as part of the “LLM Cognition Toolbox for SMEs” project in partnership with the crowdfunding analytics platform Wixdom (https://www.wixdom.io/). It was also partially funded by the Luxembourg National Research Fund (FNR) under grant reference NCER22/IS/16570468/NCER-FT and by the PayPal-FNR

PEARL grant under reference number 13342933/Gilbert Fridgen.

References

[1] Adhikary, S., Sen, P., Roy, D., Ghosh, K., 2024. A case study for automated attribute extraction from legal documents using large language models. Artificial Intelligence and Law .

[2] Akter, M., Çano, E., Weber, E., Dobler, D., Habernal, I., 2025. A comprehensive survey on legal summarization: Challenges and future directions. ACM Comput. Surv. Just Accepted.

[3] Ariai, F., Mackenzie, J., Demartini, G., 2024. Natural language processing for the legal domain: A survey of tasks, datasets, models, and challenges. ACM Computing Surveys .

[4] Bratton, W.W., 2002. Venture capital on the downside: preferred stock and corporate control. Michigan Law Review 100, 891–945.

[5] Breton, J., Billami, M.M., Chevalier, M., Nguyen, H.T., Satoh, K., Trojahn, C., Zin, M.M., 2025. Leveraging LLMs for legal terms extraction with limited annotated data. Artificial Intelligence and Law .

[6] Brown, K.C., Wiles, K.W., 2016. Opaque financial contracting and toxic term sheets in venture capital. Journal of Applied Corporate Finance 28, 72–85.

[7] Feld, B., Mendelson, J., 2019. Venture deals: Be smarter than your lawyer and venture capitalist. John Wiley & Sons.

[8] Gompers, P., Lerner, J., 2001. The venture capital revolution. Journal of economic perspectives 15, 145–168.

[9] Gordon, J.M., Orozco, D., 2014. Trust and control: the value effect of venture capital term sheet provisions as risk allocation tools. Mich. Bus. & Entrepreneurial L. Rev. 4, 195.

[10] Graham, S.G., Soltani, H., Isiaq, O., 2025. Natural language processing for legal document review: categorising deontic modalities in contracts. Artificial Intelligence and Law 33, 79–100.

[11] Guha, N., Nyarko, J., Ho, D., Ré, C., Chilton, A., Chohlas-Wood, A., Peters, A., Waldon, B., Rockmore, D., Zambrano, D., et al., 2023. Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models. Advances in neural information processing systems 36, 44123–44279.

[12] Hagos, D.H., Battle, R., Rawat, D.B., 2024. Recent advances in generative ai and large language models: Current status, challenges, and perspectives. IEEE transactions on artificial intelligence .

[13] Hornek, T., Sartipi, A., Tchappi, I., Fridgen, G., 2025. Benchmarking pre-trained time series models for electricity price forecasting, in: 2025 21st International Conference on the European Energy Market (EEM), IEEE. pp. 1–7.

[14] Kaplan, S.N., Strömberg, P., 2003. Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. The review of economic studies 70, 281–315.

[15] Kaplan, S.N., Strömberg, P.E., 2004. Characteristics, contracts, and actions: Evidence from venture capitalist analyses. The journal of finance 59, 2177–2210.

[16] Karamitsos, I., Roufas, N., Al-Hussaeni, K., Kanavos, A., 2025. Legner: a domain-adapted transformer for legal named entity recognition and text anonymization. Frontiers in Artificial Intelligence 8, 1638971.

[17] Khadangi, A., Sartipi, A., Tchappi, I., Fridgen, G., 2025. Cognartive: Large language models for automating art analysis and decoding aesthetic elements. arXiv preprint arXiv:2502.04353 .

[18] Korsmo, C., 2013. Venture capital and preferred stock. Brooklyn Law Review 78.

[19] Li, H., Gao, H.H., Wu, C., Vasarhelyi, M.A., 2025. Extracting financial data from unstructured sources: Leveraging large language models. Journal of Information Systems 39, 135–156.

[20] McHugh, M.L., 2012. Interrater reliability: the kappa statistic. Biochemia medica 22, 276–282.

[21] Narendra, S., Shetty, K., Ratnaparkhi, A., 2024. Enhancing contract negotiations with llm-based legal document comparison, in: Proceedings of the Natural Legal Language Processing Workshop 2024, pp. 143–153.

[22] Perot, V., Kang, K., Luisier, F., Su, G., Sun, X., Boppana, R.S., Wang, Z., Wang, Z., Mu, J., Zhang, H., Lee, C.Y., Hua, N., 2024. LMDX: Language model-based document information extraction and localization, in: Ku, L.W., Martins, A., Srikumar, V. (Eds.), Findings of the Association for Computational Linguistics: ACL 2024, Association for Computational Linguistics, Bangkok, Thailand. pp. 15140–15168.

[23] Sanga, S., Talley, E., 2024. "don't go chasing waterfalls": Fiduciary duties in venture-capital-backed start-ups. The Journal of Legal Studies 53, 21–65.

[24] Savelka, J., Ashley, K.D., 2023. The unreasonable effectiveness of large language models in zero-shot semantic annotation of legal texts. Frontiers in Artificial Intelligence 6, 1279794.

[25] Shane, S., 2003. A General Theory of Entrepreneurship: The Individual-Opportunity Nexus. Edward Elgar Publishing, Cheltenham, UK.

[26] Shao, M., Basit, A., Karri, R., Shafique, M., 2024. Survey of different large language model architectures: Trends, benchmarks, and challenges. IEEE Access .

[27] Siino, M., Falco, M., Croce, D., Rosso, P., 2025. Exploring llms applications in law: A literature review on current legal nlp approaches. IEEE Access 13, 18253–18276.

[28] Yue, C., Xu, X., Ma, X., Du, L., Ding, Z., Han, S., Zhang, D., Zhang, Q., 2025. Extract information from hybrid long documents leveraging llms: A framework and dataset, in: ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE. pp. 1–5.

[29] Zang, Q., Zgrzendek, C., Tchappi, I., Khadangi, A., Sedlmeir, J., 2025. Kg-htc: Integrating knowledge graphs into llms for effective zero-shot hierarchical text classification. arXiv preprint arXiv:2505.05583 .