

Contractsphere AI: A Smart Contract Management System Using Artificial Intelligence And Blockchain

Harshada Magar¹, Yash Bhalekar², Sarthak Belvalkar³, Om Jadhav⁴
^{1,2,3,4} Department of Information Technology, AISSMS Institute of Information Technology,
Pune, India

Abstract- This paper presents ContractSphere AI, a system designed to help organizations manage their contracts more easily and securely. Managing contracts in companies involves many steps such as writing, reviewing, checking legal rules, and storing the final signed document. Doing all these steps manually takes a lot of time and often leads to mistakes. ContractSphere AI uses artificial intelligence to automate these steps and uses blockchain technology to make sure that signed contracts cannot be changed or faked. The system can understand contract language in multiple languages and can handle contracts from different countries with different legal rules. It uses a language model trained on legal documents together with a search system that finds relevant rules and contract examples. The final signed contract is stored securely by saving its unique hash on the blockchain, which proves the contract is genuine. This paper describes how the system works, explains the main processing steps, and discusses how well the system performs in terms of speed, security, and cost.

Index Terms- contract management, legal AI, large language models, retrieval-augmented generation, blockchain, compliance checking

I. INTRODUCTION

Every company, big or small, has to deal with legal contracts. These include agreements with suppliers, employment contracts, software licenses, and many other types of documents. In large companies, thousands of contracts may be handled every year. Writing and reviewing these contracts by hand is slow, expensive, and often leads to errors or missed legal requirements. When contracts are poorly written or contain incorrect clauses, companies can face legal trouble or financial losses.

In many organizations, contracts are still managed using basic tools like spreadsheets, email, and physical files. There is no proper system to track deadlines, remind teams about renewals, or check if a contract follows the law in a particular country. This becomes even more difficult when a company works across multiple countries, each with its own legal rules.

Another big problem is that contracts can be faked or changed after signing. If a signed contract is stored only in a regular file system or database, it can be edited by someone with access, which makes the document untrustworthy. There is a need for a system that can prove that a contract has not been changed after it was signed.

To solve these problems, we built ContractSphere AI. This system uses AI to help write, review, and check contracts

automatically. It also uses blockchain technology to store a unique fingerprint, called a hash, of every final contract. This way, anyone can verify that the contract is real and unchanged. The system works in multiple languages and supports different legal environments.

The main contributions of this work are as follows. First, we provide a complete system that handles all stages of a contract from writing to storage and verification. Second, we use a language model fine-tuned on legal texts combined with a search system to find relevant legal rules. Third, we describe the processing steps for writing contracts, checking risks in contract clauses, and verifying compliance. Fourth, we build a security system using role-based access, encryption, and blockchain to protect contracts. Finally, we discuss how the system performs and where it can be improved in the future. The rest of the paper is organized as follows. Section II reviews related work on contract management and blockchain verification. Section III explains how the system is designed and built. Section IV discusses the system performance. Section V talks about challenges and limitations. Section VI gives the conclusion and future work.

II. RELATED WORK

Problems With Traditional Contract Management

Contract management is not a single task. It involves many steps: receiving a request, writing a draft, reviewing the draft, getting approvals, signing, storing, and then checking

that all agreed terms are followed over time. In traditional organizations, each of these steps is done manually, often by different teams using different tools. This leads to delays, inconsistencies, and mistakes.

When a company has contracts with parties in different countries, the problem becomes even harder. Legal rules differ from country to country, and keeping track of all these rules manually is nearly impossible. This is why many researchers and companies have tried to build AI-based tools to help with contract management.

Existing Legal AI Tools And Their Limitations

Several tools have been developed to help with contracts. Some tools help with writing contract templates. Others focus on reviewing contracts to find risky clauses. Some use large language models to answer legal questions or summarize documents. However, most of these tools handle only one part of the contract process and do not provide a complete solution from writing to verification.

Template-based drafting tools are useful for creating simple contracts but they cannot handle complex or customized agreements. Review-focused tools help find risky clauses but do not help with writing or approval workflows. GPT-based legal assistants are good at explaining contract language but often do not have access to up-to-date legal information and lack proper security mechanisms [9]. Common problems with existing tools include the absence of support for all stages of the contract lifecycle in one system, limited support for multiple languages, no automatic checking of whether a contract follows the laws of a specific country, weak security features, and no use of blockchain to verify that contracts are genuine.

Recent research has shown that language models trained specifically on legal texts perform better on legal tasks than general-purpose models [11], [13]. Adding a retrieval system that finds relevant legal documents makes these models even more accurate and reliable [12]. However, combining such models with a full lifecycle management system and blockchain verification remains an open area of work.

Blockchain For Document Verification

Blockchain is a technology where records are stored in a chain of blocks. Once a record is added, it cannot be changed. This makes blockchain very useful for verifying documents. Several systems have been built to verify academic certificates using blockchain [1], [2]. These systems store the hash of a document on the blockchain. When someone wants to verify the document, they calculate the hash again and compare it with the one stored on the blockchain. If both hashes match, the document is genuine

Early systems used the Ethereum blockchain, but Ethereum has some limitations. It can handle only around 15 to 45 transactions per second and transaction fees can be very high during busy periods [4]. Newer blockchain platforms can handle much more traffic at a much lower cost [5]. For contract verification, it is important to choose a blockchain that is fast, affordable, and secure. The idea of using blockchain for contract verification is not new, but combining it with an AI-based contract management system is still an area that needs more exploration. ContractSphere AI attempts to bridge this gap.

III. SYSTEM DESIGN AND IMPLEMENTATION

ContractSphere AI is a web-based application that manages the full lifecycle of a contract. The system is divided into several main parts: a user-facing web interface, an AI processing module, a compliance and risk checking module, and a blockchain verification module. Contracts and their documents are stored in a database, while only the hash of final signed contracts is stored on the blockchain.

How The System Is Organized

The system has four main layers working together. The first is the user interface layer, which is the web page that users interact with. Legal professionals can submit contract requests, view drafts, review clause analysis results, and verify contracts from this layer. The second is the AI processing layer, which handles contract writing, clause analysis, compliance checking, and translation using a language model and a legal document search engine. The third is the storage layer, where full contract documents are stored in a regular database. For optional distributed access, they can also be linked to an external storage system. The fourth is the blockchain verification layer, where when a contract is finalized and signed, its hash along with a timestamp and issuer information is stored on the blockchain, creating a permanent and tamper-proof record.

The backend of the system is built using the Django web framework. It handles all the logic for generating contracts, running checks, managing user roles, and communicating with the blockchain.

Contract Lifecycle In The System

A contract goes through several steps in the system. A user first submits a request with the type of contract needed, the country it applies to, and any specific requirements. The system then generates a draft contract using the AI language model. The draft is reviewed by the AI to identify any risky or problematic clauses. The compliance module checks whether the contract follows the applicable legal rules. Authorized

users then review and approve the contract. The final signed contract is stored in the database and its hash is recorded on the blockchain. After that, the system continues to track obligations and renewal deadlines over time.

Ai-Based Contract Drafting

The system uses a large language model that has been trained on legal texts including contracts, laws, and court judgments from different countries. This training helps the model understand legal language and generate proper contract clauses.

When writing a contract, the system also uses a Retrieval-Augmented Generation approach, often called RAG. This

means that before generating text, the system searches a database of legal documents and past contracts to find relevant information. This information is then given to the language model as context, which helps it write more accurate and legally sound content [12].

The contract is written section by section. For each section, the model considers the type of contract, the country it applies to, and the specific requirements of the organization. The system first retrieves templates related to the jurisdiction and contract type, converts organizational rules into a format the model can use, and then generates each section while checking that it follows the relevant legal rules.

Clause Analysis And Risk Scoring

After a contract is drafted, the system checks each clause for potential risks. This is done by converting each clause into a numerical representation called an embedding and comparing it to a database of known risky clauses and past contract examples [14], [15].

A risk score is calculated for each clause based on how similar it is to previously flagged clauses, how ambiguous its language is, and whether it is sensitive in the applicable jurisdiction. If a clause has a risk score above a set limit, the system automatically suggests an alternative version of that clause. This approach is more powerful than simple keyword matching because it can find problematic clauses even when the exact wording is different from known examples.

Multilingual And Compliance Support

ContractSphere AI can process contracts in multiple languages. It uses a multilingual language model trained on legal texts in different languages [10], [16]. This helps the system handle contracts for international businesses without losing the meaning of legal terms during translation.

The compliance module contains a structured database of legal rules organized by country and contract type. When checking a contract, the system identifies which rules apply

and automatically compares the contract content against those rules. It checks whether required clauses are present and whether any forbidden terms have been included. The system then produces a compliance report showing which requirements are satisfied and which are not.

Security And Workflow Automation

The system uses several security features to protect sensitive contract data and ensure proper access control. Role-based access control means that different users have different permissions. For example, a legal officer can approve contracts but a regular employee can only submit requests. All contract data is encrypted using AES-256 encryption before storage.

When a contract is finalized, the system calculates a SHA-256 hash and stores it on the blockchain along with a time-stamp and the identity of the issuer. When someone wants to verify the contract later, the system recalculates the hash and compares it with the stored value on the blockchain. If the values match, the contract is confirmed as genuine and

unchanged. Every action taken on a contract such as viewing, editing, approving, or signing is logged so that there is a complete history of what happened.

A rule-based workflow engine automatically routes contracts to the right people for review and approval. It also sends reminders and escalates tasks if deadlines are missed. This combination of role-based access, encryption, blockchain hashing, and automated routing provides a strong security and process management foundation for enterprise use.

IV. PERFORMANCE AND EVALUATION

This section discusses how ContractSphere AI performs in terms of speed, cost, security, and usefulness. Since the system has not yet been deployed in a large organization, the evaluation is based on a design-level analysis using published benchmarks and results from similar systems in the literature [6], [7].

Blockchain Performance

One important design choice was which blockchain platform to use for storing contract hashes. Ethereum-based systems are widely known but have significant limitations for large-scale use. They can handle only about 15 to 45 transactions per second and fees can be quite high during busy periods [4]. For a company that processes hundreds or thousands of contracts per day, this is too slow and too expensive.

Newer high-throughput blockchain platforms can handle thousands of transactions per second and charge only a very small fee per transaction [5]. The design of ContractSphere AI

is compatible with such platforms, which makes it more practical for real-world enterprise use.

By storing only the hash on the blockchain and not the full document, the system keeps transaction sizes small and avoids exposing private contract content on a public ledger. The full document is kept securely in an off-chain database. This design gives a good balance between transparency and privacy while keeping costs low.

Contract Writing Quality

The AI model can generate contract sections that are appropriate for the given country and type of contract. Compared to simple template-based tools, the AI approach produces more accurate and flexible contracts because it understands the context of the request. Research on fine-tuned legal language models shows that domain-specific training significantly improves the quality of generated legal text compared to using a general-purpose model [11], [13].

Clause Risk Detection

The clause analysis module can identify clauses that are similar to known risky patterns. By using embedding-based search, the system finds semantically similar clauses even when the exact wording is different. This makes the risk detection more reliable than approaches that rely only on fixed rules or keyword lists [14].

Compliance Checking

The system can check contracts against a library of legal rules and flag missing or prohibited clauses. This helps companies make sure their contracts are legally valid before signing. The structured regulatory knowledge base enables systematic checking across different legal environments and jurisdictions.

Security Strength

The SHA-256 hashing algorithm used in this system is considered highly secure. The chance of two different documents producing the same hash is extremely small, which means that any change to a contract after signing will produce a completely different hash and will be detected immediately. Combined with encryption and access control, the system provides a strong security foundation for enterprise contract management [6].

Limitations of the Evaluation

It is important to note that this evaluation is based on system design and published literature rather than a live deployment with real users. A complete evaluation would require testing the system with real contracts across multiple organizations and measuring performance through actual user studies. This is planned as part of future work.

V. DISCUSSION

What the System Achieves

ContractSphere AI brings together several technologies to solve a real-world problem. Managing legal contracts in large organizations is time-consuming, error-prone, and insecure. By combining a legal AI model with an automated workflow and blockchain verification, the system provides a single platform that covers the entire contract process.

The design addresses several gaps found in existing tools. Unlike tools that only support one phase of the contract lifecycle, ContractSphere AI handles everything from drafting to storage and verification. The blockchain component adds a level of trust that traditional database-only systems cannot provide. Even if the database is compromised, the blockchain hash can still be used to prove whether a contract is genuine. The on-chain and off-chain hybrid design is an important architectural decision. Storing full documents on a blockchain would be expensive and slow. By storing only the hash on the blockchain and keeping the document off-chain, the system achieves the best of both worlds, which is a tamper-proof verification record together with fast and affordable document storage.

Known Challenges and Limitations

Despite its advantages, the system has several limitations that should be acknowledged. The system depends on a database of legal rules for compliance checking, and these rules must be regularly updated as laws change. This is a significant ongoing effort and may be difficult to scale across all countries and legal systems.

The current system supports standard contract workflows but does not handle complex multi-party negotiations where clauses may depend on each other or where multiple rounds of editing are needed between several parties. Handling such negotiations automatically remains a difficult problem.

When the AI flags a clause as risky or generates a new clause, it should explain why in a clear and understandable way. This is important for legal professionals who need to trust and verify AI decisions. The current system provides basic explanations but this area could be improved significantly.

Deploying a blockchain component also requires some technical setup and ongoing maintenance. Organizations must manage access keys carefully and ensure that the blockchain network is properly configured. Additionally, connecting the system to existing enterprise tools like ERP or CRM platforms may require custom development, which adds to the adoption effort.

VI. CONCLUSION

This paper presented ContractSphere AI, a system for managing legal contracts using artificial intelligence and blockchain technology. The system automates the main steps of the contract lifecycle including drafting, clause review, compliance checking, and secure verification. It uses a language model trained on legal texts, a retrieval system for finding relevant legal rules, and blockchain to ensure that signed contracts are genuine and unaltered.

The system addresses important limitations of existing tools by providing end-to-end lifecycle support, multilingual processing, automated compliance checking, and enterprise-grade security. The hybrid on-chain and off-chain design keeps costs and processing time low while still providing strong verification guarantees.

The design presented in this paper is suitable for local and academic deployment as an initial version. In the future, we plan to conduct real-world testing with actual contracts and users, expand the legal rule database to cover more countries, add support for complex negotiation workflows, improve the explainability of AI decisions, and explore advanced privacy techniques such as zero-knowledge proofs. These improvements will help make the system more practical and trustworthy for widespread use in real organizations.

ACKNOWLEDGMENTS

The authors would like to thank their institution, AISSMS Institute of Information Technology, Pune, for providing the academic environment needed to carry out this research. The authors also acknowledge the open-source tools and prior research in legal AI and blockchain-based document verification that made this work possible.

REFERENCES

1. N. Vikhankar, A. Andhare, I. Barne, P. A. Dhawale, and S. Kauchali, "E-Certificate Verification Using Blockchain," *International Journal of Engineering Research & Technology*, vol. 13, no. 5, pp. 1–6, May 2024.
2. S. Gangwar and A. Chaurasia, "Blockchain-Based Authentication and Verification System for Academic Certificate Using QR Code and Decentralized Applications," *International Journal of Computer Applications*, vol. 186, no. 32, pp. 1–7, 2024, doi: 10.5120/ijca2024923722.
3. M. Hasan, A. Rahman, and M. Islam, "DistB-CVS: A Distributed Secure Blockchain-Based Online Certificate Verification System from Bangladesh Perspective," in *Proc. Int. Conf. on Computer and Information Technology (ICCIT)*, Dhaka, Bangladesh, 2020, pp. 1–6, doi: 10.13140/RG.2.2.14973.59369.
4. A. Rustemi, F. Dalipi, V. Atanasovski, and A. Risteski, "A Systematic Literature Review on Blockchain-Based Systems for Academic Certificate Verification," *IEEE Access*, vol. 11, pp. 64679–64696, 2023, doi: 10.1109/ACCESS.2023.3289598.
5. R. Priyadarshini, S. K. Panda, A. K. Jena, and P. K. Dash, "A Faster, Integrated, and Trusted Certificate Authentication and Issuer Validation System Based on Blockchain," *IEEE Access*, vol. 13, pp. 27037–27049, 2025, doi: 10.1109/ACCESS.2025.3539180.
6. M. N. M. Bhutta, A. A. Khwaja, A. Nadeem, H. F. Ahmad, M. K. Khan, and C. Han, "A Survey on Blockchain Technology: Evolution, Architecture, and Security," *IEEE Access*, vol. 9, pp. 61048–61073, 2021, doi: 10.1109/ACCESS.2021.3072849.
7. S. Islam, M. J. Islam, M. Hossain, S. Noor, K.-S. Kwak, and S. M. R. Islam, "A Survey on Consensus Algorithms in Blockchain-Based Applications: Architecture, Taxonomy, and Operational Issues," *IEEE Access*, vol. 11, pp. 39066–39082, 2023, doi: 10.1109/ACCESS.2023.3267047.
8. L. Chen, Y. Wang, Z. Liu, and Q. Zhang, "Deep Learning for Contract Analysis: Methods and Applications," *Artificial Intelligence and Law*, vol. 31, no. 2, pp. 189–215, 2023.
9. R. Williams, "Retrieval-Augmented Generation for Legal Applications: A Comprehensive Survey," *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–38, 2024.
10. M. Garcia and H. Schmidt, "Cross-Lingual Legal NLP: Challenges and Solutions," *Computational Linguistics*, vol. 49, no. 3, pp. 567–592, 2023.
11. Y. Zhang, X. Liu, J. Wang, and H. Chen, "Fine-Tuning Large Language Models for Legal Domain Applications," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, Bangkok, Thailand, 2024, pp. 3421–3436.
12. P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, and S. Riedel, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 9459–9474.
13. I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, and I. Androutopoulos, "LEGAL-BERT: The Muppets Straight Out of Law School," in *Findings of the Association for Computational Linguistics: EMNLP*, 2020, pp. 2898–2904.
14. M. Hendrycks, D. Burns, S. Basart, A. Zou, M. Mazeika, and D. Song, "CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review,"

- in Proceedings of the NeurIPS 2021 Datasets and Benchmarks Track, 2021.
15. J. Johnson, M. Douze, and H. Je'gou, "Billion-Scale Similarity Search with GPUs," *IEEE Transactions on Big Data*, vol. 7, no. 3, pp. 535–547, 2021.
 16. A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, and E. Grave, "Unsupervised Cross-Lingual Representation Learning at Scale," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 8440–8451.