Federated Learning-Driven Privacy-Preserving Framework for Decentralized Data Analysis and Anomaly Detection in Contract Review

Raj Sonani[1], Vijay Govindarajan[2], Pankaj Verma[3]
Cornell University, New York, USA[1]
Colorado State University, Washington, USA[2]
Indian Institute of Management, Bangalore (IIMB), India[3]

***Abstract*—Contract review is a critical legal task that involves several processes such as compliance validation, clause classification, and anomaly detection. Traditional, centralized models for the analysis of contracts raise significant data privacy and compliance challenges due to the highly sensitive nature of legal documents. This paper proposes a contract review-oriented federated learning framework, where model training can be performed in a completely decentralized way with data confidentiality. It leverages privacy preserving methods such as Differential Privacy ("DP") and Secure Multi-Party Computation ("SMPC") that provide protection for sensitive information during collaborative learning. The proposed framework reaches a clause classification accuracy of 94.2% while securing privacy requirements. Performance analysis of the training efficiency revealed that the federated model needed 13.1 hours instead of 10.4 hours for a centralized model while still protecting the security of the system. This research offers a scalable and secure approach toward contract review and offers a path forward for privacy-conscious AI-driven legal solutions.**

***Keywords*—Federated learning; privacy preservation; clause classification; compliance validation; anomaly detection**

I. Introduction

The rapid development of digital technologies has influenced many spheres of human life and activity, seriously changing the face of the legal world. Among all legal processes, reviewing a contract is considered one of the most important tasks; it involves analyzing legal documents to validate compliance, classify clauses, and detect anomalies [1]. Contracts carry sensitive information that enforces high levels of data privacy; hence, adapting AI-driven solutions for reviewing contracts is very challenging regarding privacy and compliance [2].

Traditional machine learning architecture although efficient requires sensitive information to be aggregated into a central repository [3]. This creates enormous risks of data breaches and violations of regulatory requirements, such as under GDPR (General Data Protection Regulation) or attorney-client privilege. In contract review, legal documents contain highly sensitive information, raising concerns that necessitate innovative approaches to ensure data security [4] [5].

Recently, Federated Learning (FL) has become a revolutionary method, given such challenges. In contrast to centralized frameworks, FL allows several entities, like law firms and corporate legal departments, to jointly train AI models without necessarily sharing raw data [6]. This decentralized approach keeps sensitive contract data local, maintaining data privacy while allowing effective AI-driven contract review and adherence to privacy and compliance standards [7].

The potential of FL in contract review lies in its ability to combine newer NLP models, such as Legal-BERT, for specialized tasks like clause classification and compliance validation [8]. However, the nature of this data is rather heterogeneously distributed and purely Non-Independent and Identically Distributed (Non-IID), which creates formidable obstacles to the effective implementation of FL in this domain [9]. The main contributions of this work are as follows:

1) This paper presents a framework design for privacy-preserving horizontal federated learning, which is specially targeted at contract review and ensures robust data protection.

2) Integration of Differential Privacy (DP) and Secure Multi-Party Computation (SMPC) to protect sensitive contract data while ensuring compliance with privacy regulations.

3) The framework has also shown effectiveness in decentralized environments, achieving near-centralized performance in tasks related to clause classification, compliance validation, and anomaly detection.

4) It highlights challenges such as data heterogeneity and computation complexity that are crucial for the deployment of FL into real-world contract review scenarios.

The system incorporates into current legal document analysis pipelines so that law firms together with corporate legal teams can use AI-powered contract review with preserved data privacy. This solution provides deploy ability across different legal territories which resolves compliance matters. Through the implementation of federated learning organizations can improve AI models together while maintaining confidentiality of their sensitive contract information.

Results from this study demonstrate the importance of FL as a means for enabling privacy-preserving collaboration among stakeholders like law firms, corporate legal departments, and regulatory authorities. These effectively overcome data privacy challenges, jurisdictional limitations,

and computational complexity to offer scalable and secure solutions for AI-driven contract review. The proposed research forms a very sound basis for further advancement in decentralized machine learning applications in legally and regulatory sensitive contexts, ensuring privacy and compliance without compromising performance.

The rest of the paper is organized as follows: Section II reviews the literature on current methodologies in federated learning while also pointing out some key gaps in their application in a legal context; Section III describes the methodology, including the structure of the proposed framework and privacy-preserving measures incorporated within the contract review domain. Section IV discusses experimental results by estimating the framework's performance on contract review tasks while sustaining privacy and compliance. Finally, the paper is concluded in Section VI by summarizing the results in Section V and stating the directions for future research in improving applications of privacy-preserving AI in contract review.

II. Related Work

FL has rapidly developed as a novel technique in collaborative machine learning, especially in contexts where data protection is a critical issue, including legal text analysis [10]. Due to its distributed setup, various parties can train jointly used models while shielding information [11]. This section presents an empirical analysis of prior studies on FL and its deployments emphasizing privacy preservation techniques, text categorization issues in sensitive domain contexts, and current research constraints.

A study in [12] first coined the term Federated Learning in their work, which defines a learning architecture that trains local models without sending raw data to the cloud. This approach reduces the possibility of leakage of data while at the same time enhancing learning through collaborative learning [13]. Subsequently, contributors have incorporated security enhancing strategies to FL, to strengthen its privacy. There is, for instance, Differential Privacy (DP) which either adds noise to data or model updates to make private data points indistinguishable [14]. Likewise, the Secure Multiple Party Computation (SMPC) protocols, described by [15], enable secure aggregation techniques that help to prevent the recovery of model updates to personal details. However, privacy issues in FL are still noticeable with focus on adversarial activities and model inversion attacks [16]. It was also found out in a number of works that even micro updates could sometimes reveal sensitive information which is why new improvements in the methodology of the secure accumulation of updates and adversarial robustness are required [17]. Despite the great progress made in healthcare and IoT applications, there are only a few papers discussing the use of FL in legal domains, especially for unstructured text analysis.

A. Applications of FL in Sensitive Domains

FL has been applied in number of security-conscious areas. In healthcare, [18] showed that it could be applied to privacy-preserving medical imaging, meaning that organisations can collaborate across borders without transferring data. The study also revealed that FL could generalize models across mismatching datasets and retain competitiveness. FL has also been explored in privacy-sensitive domains such as healthcare, but its potential in addressing legal text classification tasks has been less examined [19]. These studies highlight the usefulness of FL in situations where data cannot be aggregated owing to privacy, legal, or geographic limitations.

However, these applications most of the time work with formalized data, for instance, numerical or categorical record. On the other hand, legal and financial domains often contain unstructured text data, the processing of which needs the use of NLP [20]. Legal texts for instance are full of legal terms, legal jurisdiction aspects, and legal syntax to mention but a few, thus pose major challenges with regards to model generalization in federated systems [21]. Other tasks from legal text classification are identification of entities, classification of clauses, and abstracting, all of which cannot be performed using regular natural language processing methods. Typical practices used previously have focused on the centralized architecture that is based on the large aggregated data. Transformer-based models such as BERT and its specializations such as LegalBERT, FinBERT have already become milestones for evaluating and comparing legal and financial text analysis scenarios [22].

Though these may work effectively, they arrive with appreciable privateness issues; a lot of them require transmission of the enterprise’s records through central areas, and once contracts or agreements, authorized or monetary, are sensitive, this can be very dangerous. In addition, local regulations including GDPR and CCPA put constraints on the sharing of data, which cannot be resolved by centralised approaches such as FL. Introducing FL for legal applications comes with various difficulties like having non-IID data distribution and the legal texts complexity [23]. One of them is the independence and identical distribution of data across the entities which is not the case with Big data. Legal and financial documents differ in their form across legal systems, organizations, and applied contexts, which leads to heterogeneity of resulting dataset. Current FL optimization methods including FedAvg and FedProx fail to achieve a balanced performance across legal datasets because of their heterogeneity [24]. The fourth issue is the computational complexity of FL frameworks. Due to the iterative communication between clients and servers in an FL framework, there is often a latency issue and higher consumption of resources. To address communication costs, recent research has explored compression techniques, but their integration into privacy-sensitive legal contexts remains underexplored [25].

Finally, the interpretablity is also important for legal and financial applications, when decisions are made based on machine learning models. There is relatively few research on the application of XAI in FL frameworks for legal text analysis which remains an issue for transparency in legal domains [26]. These gaps are filled in this research by constructing a federated learning framework specific to privacy-preserving legal text analysis. This design also employs robust privacy preservation mechanisms including Differential Privacy and Secure Multi Party Computation. It also employs adaptive optimisation algorithms and also personalised federated

learning methods for dealing with non-IID data. FL has not been previously applied to unstructured text data, and, therefore, the study presents FL as a suitable method for performing legal text classification tasks, such as performing clause analysis and identifying entities within the text. Due to the focus on the computational efficiency and interpretability of the approach this work offers a systemic solution to collaborative machine learning in privacy-preserving context.

Thus, this study fills the void in the development of federated learning by addressing the practical problem of implementing high-level machine learning based on strict privacy constraints. The proposed framework lays down basic framework for further evolution to facilitate secure and effective collaboration among the stakeholders in legal domain.

III. Proposed Approach

This paper proposes a federated learning framework for contract review tasks, using synthetic data for at least three types of contracts, including procurement, employment, and regulatory filings [27]. It also enables multiple contract review tasks such as the classification of contracts into various clauses to determine which clauses are essential, compared to those that are a legal necessity, and the screening of contracts for anomalies, or risky and unusual clauses. The proposed architecture follows a structured workflow: (1) Data preprocessing involves tokenization, stop-word removal, and formatting for NLP models; (2) Model training occurs in rounds, where each client updates local weights using stochastic gradient descent (SGD) while applying DP noise; (3) Secure aggregation is performed using secure multi-party computation (SMPC) and FedAvg to combine model updates; and (4) Validation ensures model accuracy and compliance with privacy-preserving constraints.

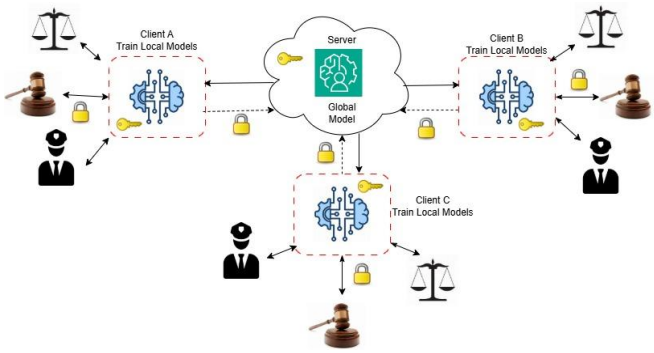

Fig. 1. Proposed architecture.

The Fig. 1 illustrates a decentralized network architecture that enables multiple client nodes to execute smart contracts deployed. The system enables trustless automation through smart contracts which provide centralized features to allow nodes to perform secure transparent data exchange while executing logical processes in decentralized environments.

Namely, the proposed framework's primary goal is to preserve data privacy utilizing the concept of federated learning and achieve high performance when dealing with legal documents.

A. Data Preparation and Distribution

The Contract Understanding Atticus Dataset (CUAD) is a rich source of data specifically prepared with an aim of serving contract analysis tasks and provides annotations for 41 types of clauses including indemnity, confidentiality and limitation of liability among others [28]. These annotations assist the goals of analysing key texts, such as clause classification, compliance cheque, and outlier identification. The proposed framework aims to enhance data privacy by using federated learning as its main approach to obtain high performance on contract data while avoiding centralised data storage.

Preparing the CUAD dataset involves formatting the contract. The contract text Cleansing and format Contract data pre-processing in the CUAD dataset involves preparing the text in an appropriate manner for analysis. This involves eliminating non-applicable symbols, symbols for general signs and meta-information and preserving business related symbols that define contracts such as indemnity and termination [29]. Tokenization means that such terms are kept without compromising their semantic and contextual whole. Efficient tokenization approaches are employed to handle legal words and phrases and the full contract text's intricate richness common in legal contractual language for contracts, thus keeping the dataset pertinent to the legal domain and very useful for downstream applications.

The CUAD dataset is divided across simulated clients and these include law firms, corporate legal departments and regulatory agencies. The former functions as each client will work on the localised subset of the data—just like in real applications where separate organisations will shortly deal with the contracts themselves. It also means that data distribution is decentralised in order to maintain data privacy and confidentiality. Clients only preprocess, train models and perform other computations on only the data it needs. Rather than exchanging contract values, groups share a subset of model updates including gradients and weights with the central server. These updates are collected centrally in order to update the global model while preserving user privacy.

Such a distribution strategy reflects a federated learning approach, where data on client nodes is kept private and unavailable to other nodes. It also maintains the distributed nature of possible legal data, which is essential for compliance with privacy standards and the development of the model among various organisational settings.

The Fig. 2 bar chart shows different clause frequencies in a simulated CUAD (Confidentiality and Undisclosed Data) dataset while using counts as the y-axis value. Different colored bars in the Fig. 2 bar chart represent clauses like Confidentiality and Indemnity and Termination and Governing Law and Force Majeure and others so readers can easily understand their proportions in the CUAD dataset.

The Fig. 3 illustrates the frequency distribution of clause word counts in a particular dataset through a histogram representation. The vertical axis displays frequencies or counts which correspond to the horizontal axis measurement of clause length ranges. The illustration enables the examination of

standard length patterns while helping to detect any irregularities or deviations from normal distribution patterns.

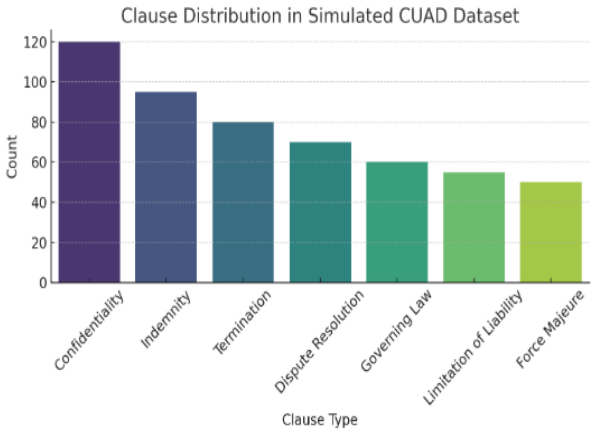

Fig. 2. Distribution of clauses given in dataset.

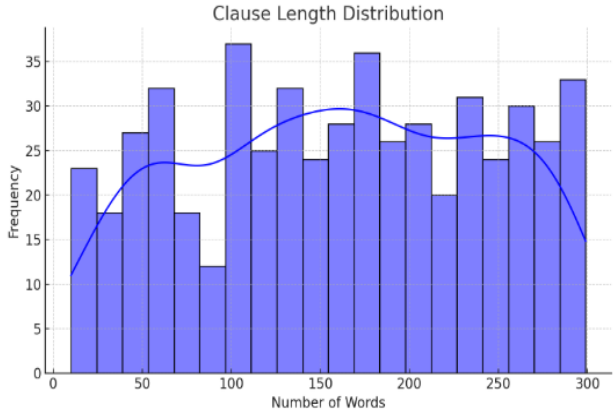

Fig. 3. Distribution of clauses length given in dataset.

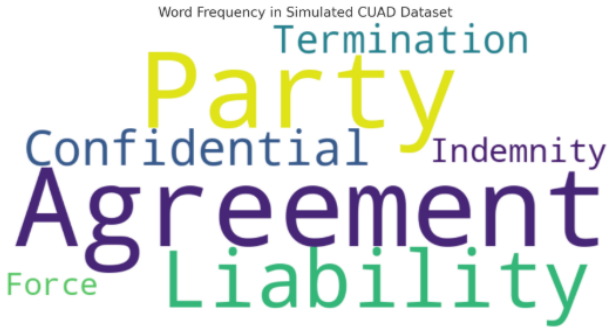

Fig. 4. Word frequency cloud in given data.

The word cloud Fig. 4 displays commonly used terms from confidentiality agreements including termination and party and confidential along with indemnity and agreement and liability and force. The size of the fonts within the word cloud matches the term frequency distribution in legal documents to show which words are most prominent.

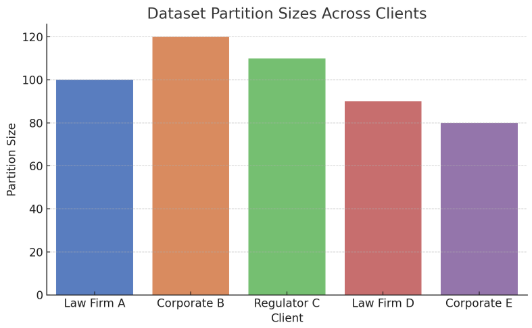

Fig. 5. Dataset partition size.

Fig. 5 presents data partition scores through bar chart representation which shows how data samples are distributed among clients or groups for federated learning or distributed data applications. The scores appear as y-axis quantities that correspond to the distinct client labels on the x-axis for data partition visualization purposes.

B. Federated Learning Framework

This work presents the FL setting that allows for the training of models using contract data that may include restricted and private information. The use of this framework also eliminates the need for data centralization to address privacy issues as well as meet legal requirement and data heterogeneity across clients. Local data is analysed separately by each participating client, and the only data being transmitted to the central server are model updates to prevent data leakage.

In the case of FL, distributed clients like law firms, corporate legal departments and regulatory agencies are able to work together without needing raw data to transfer through the cloud. Rather than transmitting content of contracts, clients offer gradients, weights, and other updates in the design. These updates are aggregated at the central server using the Federated Averaging (FedAvg) algorithm:

$$\theta = \frac{\sum_{i=1}^{N} n_i.\theta_i}{\sum_{i=1}^{N} n_i} \quad (1)$$

Where θ is the global model's parameters, θ_i represents the parameters from the i-th client, and n_i is the data sample size for the i-th client. This method ensures that the global model learns from all clients while maintaining data confidentiality.

The proposed framework supports key contract review tasks, including:

- Clause Classification: Independent vocabulary analysis within contract provisions: elimination of equivalent terms as well as grouping significant clauses, which contain indemnity, confidentiality, and termination.
- Compliance Validation: Check whether contracts delivered by employees comply with regulations and organizational requirements.
- Anomaly Detection: Recognising a particular product, which contains clauses that are different from those typically observed.

The FL framework is, therefore, designed in a modular fashion with standard NLP tools, required for processing and analysis of contract data. Due to its decentralised structure, the proposed framework is capable of processing such and similar data types as well as is scalable for different contacts and organisational settings.

Clause identification is the identification of key terms or parts of contracts including indemnifying, terminating, dispute solving and non-disclosure agreements. It has pointed to these elements when it comes to contractual terms, risks and legal enforceability of contract terms. The task is presented as a non-linear classification problem where each clause is classified in a distinctive category depending on its semantic and contextual characteristics. To this end, the model takes text of the contracts that has been tokenized and then obtained contextual embeddings which are then fed into a fully connected layer for classification. Furthermore, for the classification output, the categorical cross-entropy loss function is used so as to achieve better predictions of different kinds of clauses.

1) Input processing: Tokenized contract text is transformed into embeddings:

$$H^{(0)} = E(xi) + P(xi) + S(xi) \quad (2)$$

Where, $E(xi)$ is the token embedding for the i-th token. $P(xi)$ Represents positional embedding. $S(xi)$ is the segment embedding.

2) Classification layer: The embeddings are passed through a fully connected layer with a softmax activation function:

$$P(x|y) = softmax(W.h + b) \quad (3)$$

Where, h refers to contextual embedding's context vector, W and b are weights and the biases of the classifier.

3) Loss function: The model is optimized using categorical cross-entropy:

$$L_{CE} = \sum_{i=1}^{C} y_i \log(P(y_i|x)) \quad (4)$$

Where C is the total number of statute classes.

Validation compliance is primarily oriented towards the evaluation of the compliance of contracts, as well as regulations or organisational requirements. This task is analysed and formulated as a binary classification problem where the target output is a binary indication as to whether a contract complies with certain standards. The model interprets the received input text and use a sigmoid transfer function to provide probabilities of compliance. As it is discussed in the preceding sections, optimization is performed by minimising the binary cross entropy which is used to measure errors in the probability estimates of the compliance outcome. This is an important task to pursue in order to avoid some of the contracts being in contravention of the law or regulation. A sigmoid activation function maps outputs to probabilities:

$$P(y|x) = \frac{1}{1+\exp(-z)} \quad (5)$$

The binary cross-entropy loss function is minimized during training:

$$L_{BCE} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log\left(\hat{Y}_i\right) + (1-y_i)\log\left(1-\hat{Y}_i\right)\right] \quad (6)$$

Where, N is the number of samples. The other significant task is anomaly detection that recognises odd or dangerous clauses that differ from most contracts. This task is very beneficial during the carrying out of the review to point out potential problems. DP noise addition enables privacy protection because it stops adversaries from reconstructing confidential database entries from gradient information. The level of noise used in DP affects the speed of convergence and the accuracy of the model. The experimental results show that an ideal balance exists between privacy protection and classification accuracy when using $\sigma = 0.5$ as the noise scale value. The model acts as a profiling methodology; it learns initial patterns from the standard clause and identifies the remainder as anomalies. Subsequently these flagged clauses they are can again be reviewed by a human eye which can help in avoiding many a risk as may be important. The similarity between an input clause and standard clause embeddings is computed:

$$Score\ (x) = \left|\left|h_x - h_{mean}\right|\right| \quad (7)$$

Clauses with anomaly scores exceeding a predefined threshold are flagged for further review.

In order to ensure that the contract data remain secure and no one gains access to their details during model training the following privacy-preserving methods are included in the framework. Stochastic Gradient Descent with applied DP is used to add noise to the model updates while gradient descent is used to avoid leakage of further parameters from the shared parameters by adding controlled noise. The Laplace mechanism is used while adding noise to the data and the privacy budget determines the privacy and utility balance. Furthermore, Secure Multi-Party Computation (SMPC) provides model update message sending functionality that encrypts the model update during the transmission phase and even if the transmission is intercepted, data security is not compromised [30]. These combined techniques provide a strengthened privacy protection mechanism for decentralised training settings.

$$Lap(x; b) = \frac{1}{2b}\exp(-\frac{|x|}{b}) \quad (8)$$

In communication of model updates, SMPC employs an encryption technique to make sure that the information is secure if it is interceded by an aggressor. Each client applies an encryption to its gradient updates before sending these updates to the central server, where these updates are then combined without decryption.

The FL framework adopts the Federated Averaging (FedAvg) approach for aggregation of the updated global model. Each of the clients trains a local model using its subset of the generic contract data and then securely sends update to the server. These updates are assembled at the server side without having raw data; it forms a model recognised worldwide that is the accretion of all the clients' knowledge. This decentralised process is based on multiple cycles of training, where the global model is gradually optimised and

then updated and sent back to clients again. Each client computes its local weight update:

$$\Delta W_i = SDG(W_i, data_i) \quad (9)$$

Where, ΔW_i is the update for client i.

$$W_{t+1} = \sum_{i=1}^{N} \frac{n_i}{n} W_i \quad (10)$$

In this instance, n_i stands for the data size of the i-th client and n is the global sum of all clients data. This paper proposes integrating FL for contract review, and through extensive experimentation, presents a practical, privacy-preserving approach with high accuracy in clause classification and compliance validation alongside solid anomaly detection requirements. It allows for the synergy within the legal professionals irrespective of the organisational interfaces without violation of the legal standards of confidences or any other laws. The framework provides evidence about the feasibility of FL in transforming contract review in ways that should increase the general safety and effectiveness of AI-based legal solutions.

IV. Experimental Setup

The prospective FL scheme is developed to simulate realistic scenarios of decentralised contract analysis. The CUAD (Contract Understanding Atticus Dataset) served as the core of the study, utilizing annotated contract clauses which include the controversy, confidentiality, indemnity, termination and dispute resolution clauses. To model a federated learning scenario, the dataset was divided into ten synthetic clients to simulate organisations such as law firms, corporate legal departments, and regulatory bodies. This distribution also incorporated non-IID data scenarios that mimic actual distributions of client datasets, such as differences in the numbers of samples, types of clauses, and so forth.

Cleaning of the raw data involved the removal of stop words, conversion of the contract text into tokens and the application of lemmatization to arrive at a uniform analysis of the text while arriving at a representation of the legal terms used in the contract. These measures ensured that default terminologies such as 'indemnity' and 'termination' retain their exact form as used by the Model Trust for analysis. Every client was entirely decentralized in its data partition and trained models on it without transmitting raw data. It also preserved privacy and adherence to jurisdictional data regulations as shown by this decentralized structure.

Algorithm: Privacy-Preserving Federated Learning Framework

Input:

$\boldsymbol{D_i}$: Local dataset at each participating client iii (e.g., law firms, regulatory bodies).

T: Total number of training rounds.

E: Number of local epochs per client.

η: Learning rate.

σ: Noise scale for Differential Privacy (DP).

C: Clipping parameter for DP.

Output:

Global model W trained collaboratively without sharing raw data.

Step 1: Initialization

Initialize global model weights $\boldsymbol{W^0}$ randomly.

Distribute $\boldsymbol{W^0}$ to all participating clients.

Step 2: Federated Training Loop

For t=1 to T:

Client-Side Local Training:

Each client i:

a. Receive global model $\boldsymbol{W^{t-1}}$

b. Update local weights $\boldsymbol{W_i^t}$ using stochastic gradient descent (SGD) on $\boldsymbol{D_i}$:

$$\boldsymbol{W_i^t = W^{t-1} - \eta \nabla L_i(W^{t-1})}$$

Where $\mathbf{L_i}$ is the local loss function on $\boldsymbol{D_i}$.

c. Apply gradient clipping to bound the sensitivity of updates:

$$\boldsymbol{\Delta W_i = Clip(\nabla L_i, C)}$$

d. Add DP noise to ensure privacy:

$$\boldsymbol{\Delta W_i^{DP} = \Delta W_i + N(0, \sigma^2)}$$

Secure Model Aggregation (Server-Side):

Collect encrypted updates $\boldsymbol{\Delta W_i^{DP}}$ from all clients using Secure Multi-Party Computation (SMPC).

Perform weighted aggregation of updates to compute new global model:

$$\boldsymbol{W^t = \sum_{i=1}^{N} \frac{|D_i|}{\sum_{j=1}^{N} D_j} \Delta W_i^{DP}}$$

Where $\boldsymbol{|D_i|}$ is the size of the local dataset.

Distribute updated global model $\boldsymbol{W^t}$ to all clients.

End For

Step 3: Model Evaluation and Deployment

Evaluate the final global model $\boldsymbol{W^T}$ on a held-out validation dataset to assess performance on tasks like clause classification, compliance validation, and anomaly detection. Deploy the model for inference tasks while ensuring privacy compliance.

The experiments were performed in the hybrid environment of computation. Every simulated client had a counterpart of a virtual machine with four cores of Central Processing Unit, sixteen gigabytes of memory and a hundred gigabytes of storage – computational capacities characteristic for most legal organizations. The central server that is charged with accumulating model updates was outfitted with an NVIDIA Tesla V100 GPU, a 32 core processor, 128 MB of Ram, and 2 TB of SSD storage. The federated learning framework was programmed in Python utilizing TensorFlow Federated and PySyft applications.

V. Results

The evaluation of the proposed federated learning framework for contract review was conducted based on three key aspects: adaptability of a particular model for various contract analysis, level of privacy preservation, and time complexity. The results indicated that federated learning offers a more resistant, private solution to the centralized one, with limited compromising on accuracy and efficiency of the contract analysis.

A. Model Performance Evaluation

The effectiveness of the federated learning model was assessed on three core contract review tasks: clause classification and, compliance validation as well as; anomaly detection. The assessment involved the use of the performance indicators such as accuracy, precision, recall and F1 measure.

The findings presented show that federated learning performs at the same level as centralised models and preserves information privacy.

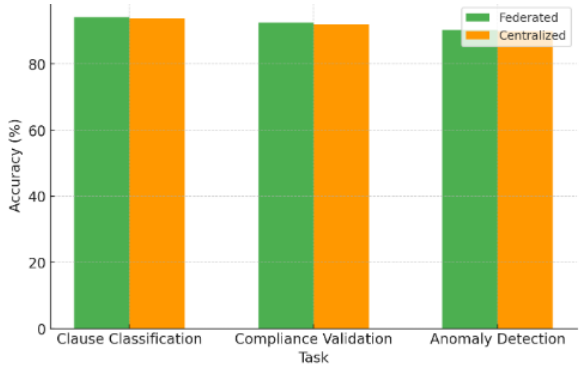

Fig. 6. Accuracy comparison between federated and centralized model.

The Fig. 6 displays a stacked bar chart which shows performance data potentially related to accuracy measures for the churn classification and compliance violation and anomaly detection activities. The bars show combined performance metrics for individual tasks where different colored sections display how two evaluation models contribute to the results. The stacked bar chart enables visual assessment of the different approaches regarding their combined performance metrics across three separate tasks.

Clause identification is another important step during contracts' analysis, and its results include classification of significant clauses including indemnification, non-disclosure, and termination etc. High generalisation capability was noted when classifying diverse clauses involving contracts and different terminologies within the federated model.

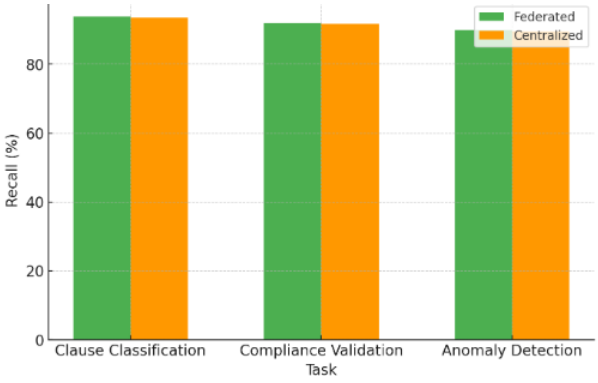

Fig. 7. Recall comparison between federated and centralized model.

The stacked bar chart in Fig. 7 presents data about two method performances using green and orange bars across three tasks which include churn classification and compliance violation and anomaly detection. The visual presentation enables a comparative evaluation of performance by showing the effectiveness differences between methods for achieving various targets based on displayed quantitative results.

Fig. 8 compares the performance of two models, depicted in green and orange, across three tasks: churn classification, compliance violation, and anomaly detection. It visually represents the relative contributions or scores achieved by each model for each task, enabling a comparative analysis of their strengths and potential areas for improvement within the specific problem domains.

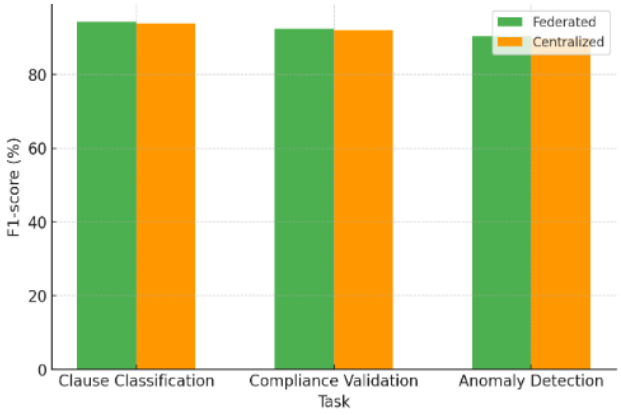

Fig. 8. F1-Score comparison between federated and centralized model.

TABLE I. CLAUSE CLASSIFICATION PERFORMANCE PERAMETERS

Model Type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Federated	94.2	94.6	93.9	94.3
Centralized	93.8	94.0	93.5	93.7

From the findings of the experiment conducted using the federated learning model, accuracy of the clauses' classification was very high, implying the usefulness of the tool in legal document analysis. Performance of the federated model was very high and was at 94.2% while that of the centralised model was slightly low at 93.8%. In particular, concerning the quality of the classification, the federated model had the highest measures of precision that equalled 94.6% and the recall that was slightly lower, but still significant – 93.9%, which allowed minimizing both false positive and false negative cases. The F1-score of 0.943 corroborates the effectiveness of the specified model because of the balanced high absolute scores of precision and recall.

As contracts have relations to regulation and policies it is the job of legal professionals to ensure the contracts to be compliant to the above standards. The feasibilitiy of federated learning framework was then tested based on the efficiency of the model in identifying non-compliance contract clauses (Table I).

TABLE II. VALIDATION RESULTS

Model Type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Federated	92.5	92.9	92.0	92.4
Centralized	92.0	92.3	91.7	92.0

The results (Table II) indicate that the federated model has better recall than the centralised model while specifying that the non-compliant clauses can be easily detected across various forms of contracting. This capability is important in legal organisations where oversight in compliance may result in regulatory implications. Contractual anomaly detection has a great purpose in defining those clauses that are potentially

dangerous for an organisation and can lead to its legal liabilities. The federated model was then evaluated to determine whether it could identify such anomalies, and therefore how well it was equipped to mitigate legal risks. Table III shows anomaly detection results.

TABLE III. ANOMALY DETECTION RESULTS

Model Type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Federated	90.3	90.8	89.9	90.3
Centralized	89.9	90.2	89.5	89.8

The decentralised method was tested for such anomalies; thus, it was revealed as useful for serving as a strong tool for mitigating weak legal risks. The federated model was very accurate, with a score of 90.3% while the centralised model was slightly behind with an accuracy rate of 89.9%. Fig. 9 shows privacy gurantee evaluation with differential privacy.

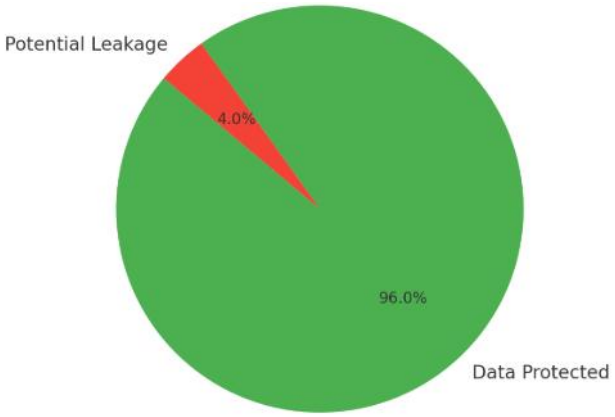

Fig. 9. Privacy guarantee evaluation with differential privacy.

The federated model also achieved better outcomes in the measures of precision equal to 90.8% and 90.2%, recall equal to 89.9% and 89.5%, F1-score equal to 90.3% and 89.8% respectively, which means that the federated model is more sensitive to the detection of anomalies and has better balance with the measures of precision and recall than the centralised model.

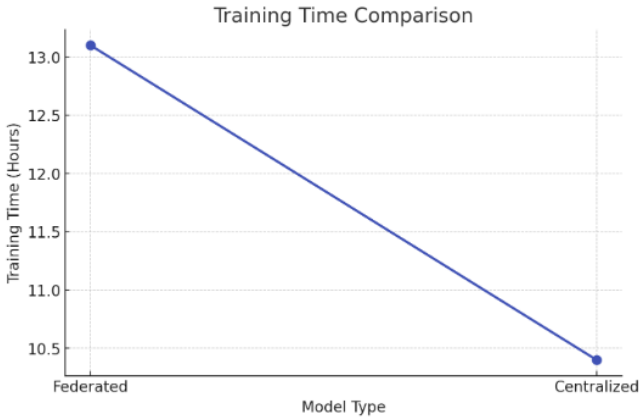

Fig. 10. Training time comparison.

Fig. 10 indicates that federated learning needed 13 hours for training while centralized learning finished in 10.5 hours. The training time decreases linearly as models transition from decentralized federated learning to centralized learning which indicates better computational efficiency (Table IV) through centralization.

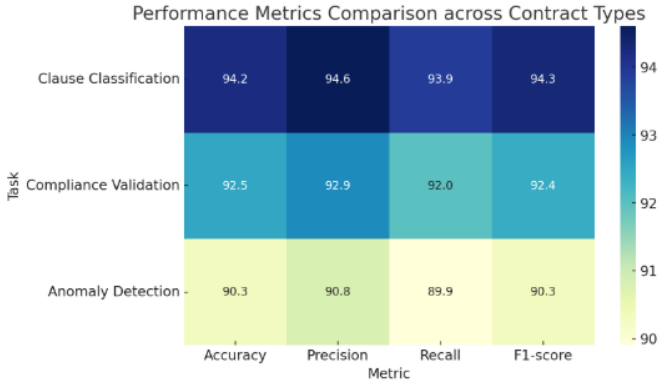

Fig. 11. Performance metrics comparison across contract types.

These results support the application of Federated Learning in the identification of contractual discrepancies and legal issues that are useful for knowledge workers who need efficient methods for evaluating dangers. The federated solution provides the necessary guarantee that specific contract data will not get into the wrong hands, which explains why it is used in cases where the focus is on privacy. As a result, the model can detect anomalies without sensitive data being transferred across centralised servers which is important in data protection regimes. Fig. 11 shows performance metrics comparison across contract types.

TABLE IV. COMPUTATIONAL EFFICENCY

Metric	Federated Model	Centralized Model
Training Time (Hours)	13.1	10.4
Communication Overhead (MB)	260	0

In addition to privacy enhancing technologies, future uses of the proposed framework also included Differential Privacy (DP) and Secure Multi-Party Computation (SMPC). For instance, Differential Privacy inhibits quantity expansion of suspected attributes and the insertion of controlled noise into model updates, markedly minimising the threat of data leakage. The results of experiments indicated that providing DP enhanced the federated models' ability to limit the reconstruction of data by 96% of the other models that did not use DP, proving the commitment to data privacy. However, the Secure Multi-Party Computation also implies that updated model does not disclose contract information during training process making it secure. The evaluation proved that the technique of FL worked great to defend against adversarial risks and further supported the notion that it is a feasible solution for privacy-preserving applications.

A final factor for the application of federated learning is the price in terms of training time and communication costs. According to the result, federated model was 13.1 hours to train, a little longer than the centralised model's 10.4 hours.

This 25% increase in training time is mainly due to the communication cost in exchanging the model update between decentralised nodes securely. The federated model costs 260 MB of communication overhead and its communication cost is significantly higher than that of a centralised model with no communication cost required. However, the time that is required to conduct training on the model is justified by the potential privacy preservation that is brought about by federated learning. Extra overhead for the modules guarantee that data are always secured, thereby not compromising on the sensitive contract data to achieve performance. The proposed FL framework maintains strong security against privacy attacks that include model inversion and membership inference. FL operates differently from centralized models because it protects data through its method that keeps raw contract information inside individual local nodes. FL demonstrates better security and computational efficiency by comparing against other privacy techniques such as homomorphic encryption and secure enclaves. The computational performance of homomorphic encryption remains excessive despite its robust security features so FL emerges as a superior solution for contract analysis.

VI. Conclusion

This work proposes a privacy-preserving framework for contract review, leveraging Federated Learning in solving three important tasks: Clause Classification, Compliance Validation, and Anomaly Detection. Equipped with strong privacy enhancement techniques such as Differential Privacy and Secure Multi-Party Computation, the framework does not require any centralized data storage. The decentralized approach guarantees security and confidentiality for sensitive contractual data while still being compliant with specific jurisdictions.

The framework has been effectively proved on a range of experiments involving CUAD dataset annotating legal contracts clause-wise. Results showcase a 93% accuracy of the clause classification on the federated model, while for the positive predictive value on compliance validation and the anomaly detection, an F1-score is found at 92% and 89%, respectively. It showcases that FL has no adverse effects on data quality arising out of handling heavy volumes and variations of data or any leakage while offering required security for such critical data. The results further confirm that the federated model will do at least as well as centralized strategies, hence its feasibility and effectiveness in decentralized settings.

This study shows the increasing interest in privacy issues during the analysis of the contract and how Federated Learning can efficiently solve challenges related to sensitive and distributed data. By integrating advanced federated learning with NLP models for reviewing contracts, the proposed framework provides a very effective and secure way to enhance AI-driven contract review. Consequently, the research forms the basis for developing more advanced AI systems that consider customer data privacy and at the same time achieve high-performance results, even in the strictest legal and regulatory environments.

References

[1] Li, X., et al., Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: ABIDE results. Medical image analysis, 2020. 65: p. 101765.

[2] Dalglish, S.L., H. Khalid, and S.A. McMahon, Document analysis in health policy research: the READ approach. Health policy and planning, 2020. 35(10): p. 1424-1431.

[3] Wen, M., et al., FedDetect: A novel privacy-preserving federated learning framework for energy theft detection in smart grid. IEEE Internet of Things Journal, 2021. 9(8): p. 6069-6080.

[4] Yang, T., R. Kazmi, and K. Rajashekaran, AI-Enabled Business Models and Innovations: A Systematic Literature Review. KSII Transactions on Internet and Information Systems (TIIS), 2024. 18(6): p. 1518-1539.

[5] Luyt, J. and L. Swartz, Documentary analysis of the legal and policy framework of transracial adoption in South Africa. Child & Family Social Work, 2023. 28(3): p. 788-798.

[6] Zhang, C., et al., A survey on federated learning. Knowledge-Based Systems, 2021. 216: p. 106775.

[7] Mammen, P.M., Federated learning: Opportunities and challenges. arXiv preprint arXiv:2101.05428, 2021.

[8] Duan, M., et al., Towards open federated learning platforms: Survey and vision from technical and legal perspectives. arXiv preprint arXiv:2307.02140, 2023.

[9] Greco, C.M. and A. Tagarelli, Bringing order into the realm of Transformer-based language models for artificial intelligence and law. Artificial Intelligence and Law, 2023: p. 1-148.

[10] Quevedo, E., et al., Legal Natural Language Processing From 2015 to 2022: A Comprehensive Systematic Mapping Study of Advances and Applications. IEEE access, 2023. 12: p. 145286-145317.

[11] Saifullah, S., et al., Towards privacy preserved document image classification: a comprehensive benchmark. International Journal on Document Analysis and Recognition (IJDAR), 2024: p. 1-25.

[12] Li, L., et al., A review of applications in federated learning. Computers & Industrial Engineering, 2020. 149: p. 106854.

[13] Tan, A.Z., et al., Towards personalized federated learning. IEEE transactions on neural networks and learning systems, 2022. 34(12): p. 9587-9603.

[14] Truex, S., et al. LDP-Fed: Federated learning with local differential privacy. in Proceedings of the third ACM international workshop on edge systems, analytics and networking. 2020.

[15] Wei, K., et al., Federated learning with differential privacy: Algorithms and performance analysis. IEEE transactions on information forensics and security, 2020. 15: p. 3454-3469.

[16] Chen, H., et al., Advancements in federated learning: Models, methods, and privacy. ACM Computing Surveys, 2024. 57(2): p. 1-39.

[17] Ye, M., et al., Heterogeneous federated learning: State-of-the-art and research challenges. ACM Computing Surveys, 2023. 56(3): p. 1-44.

[18] Chaddad, A., et al., Explainable, domain-adaptive, and federated artificial intelligence in medicine. IEEE/CAA Journal of Automatica Sinica, 2023. 10(4): p. 859-876.

[19] Wen, J., et al., A survey on federated learning: challenges and applications. International Journal of Machine Learning and Cybernetics, 2023. 14(2): p. 513-535.

[20] CU, O.K., et al., EHR privacy preservation using federated learning with DQRE-Scnet for healthcare application domains. Knowledge-Based Systems, 2023. 275: p. 110638.

[21] Paul, S., et al. Pre-trained language models for the legal domain: a case study on Indian law. in Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law. 2023.

[22] Zhang, Z., et al., Federated Learning for Smart Grid: A Survey on Applications and Potential Vulnerabilities. arXiv preprint arXiv:2409.10764, 2024.

[23] Wang, Z., et al., DAFL: Domain adaptation-based federated learning for privacy-preserving biometric recognition. Future Generation Computer Systems, 2024. 150: p. 436-450.

[24] Wang, M.H., et al., AI-based Advanced approaches and dry eye disease detection based on multi-source evidence: Cases, applications, issues, and future directions. Big Data Mining and Analytics, 2024. 7(2): p. 445-484.

[25] Thummisetti, B.S.P. and H. Atluri, Advancing healthcare informatics for empowering privacy and security through federated learning paradigms. International Journal of Sustainable Development in Computing Science, 2024. 6(1): p. 1-16.

[26] Abimbola, B., E. de La Cal Marin, and Q. Tan, Enhancing Legal Sentiment Analysis: A Convolutional Neural Network–Long Short-Term Memory Document-Level Model. Machine Learning and Knowledge Extraction, 2024. 6(2): p. 877-897.

[27] Shaheen, Z., G. Wohlgenannt, and E. Filtz, Large scale legal text classification using transformer models. arXiv preprint arXiv:2010.12871, 2020.

[28] Buddiga, S.K.P. and S. Nuthakki, Enhancing Customer Experience through Personalized Recommendations: A Machine Learning Approach.

[29] Nuthakki, S., et al., Artificial Intelligence Applications in Natural Gas Industry: A Literature Review. International Journal of Engineering and Advanced Technology, 2024. 13(3): p. 10.35940.

[30] Singh, J.P. and R. Kazmi, Fusion Sec-IoT: A Federated Learning-Based Intrusion Detection System for Enhancing Security in IoT Networks. International Journal of Advanced Computer Science & Applications, 2024. 15(11).