© April 2025 | IJIRT | Volume 11 Issue 11 | ISSN: 2349-6002

Revolutionizing CRM with AI: Clause Detection and Client Management

Naba Sarwar[1], Vishruti Gurav[2], Srushti Kamble[3], Dr. Dinesh Patil[4]
[1,2,3] *Dept of Computer Engineering Vidyavardhini's College of Engg. and Tech. Mumbai, Maharashtra*
[4]*Assoc. Professor, Dept of Computer Engineering Vidyavardhini's College of Engg. and Tech. Mumbai, Maharashtra*

***Abstract*—Our paper presents the implementation of an AI- powered Customer Relationship Management (CRM) system that integrates a clause detection module to enhance contract management by summarizing Memorandums of Understanding (MOUs) and extracting key information using machine learning and natural language processing (NLP) techniques, particularly the Bidirectional Encoder Representations from Transformers (BERT) model. By automating contract reviews, the system significantly reduces human effort and legal risks through accurate extraction of essential clauses in complex legal language. Beyond contract management, the CRM streamlines customer relationship workflows by automating lead tracking, sales forecasting, customer segmentation, chatbot support, and follow-up scheduling. It reduces administrative burdens, enhances productivity, and ensures data security through encrypted storage. The results show that the AI-driven CRM boosts contract management efficiency, customer satisfaction, and business performance by minimizing manual errors and enabling data-driven decisions through predictive analytics. This paper discusses the system's technical architecture, AI integration, benefits, and challenges of deploying it in a business environment.**

***Index Terms*—Customer Relationship Management, Clause Detection, Memorandums of Understanding, Machine Learning, Natural Language Processing, Predictive Analytics, Contract Management.**

I. INTRODUCTION

In today's fast-paced business environment, effective customer relationship management (CRM) is crucial for organizations seeking to maintain a competitive edge. Traditional CRM systems primarily focus on managing customer interactions, tracking sales, and organizing customer data. However, the integration of advanced technologies such as artificial intelligence (AI), machine learning, and natural language processing (NLP) has revolutionized the capabilities of CRM systems, enabling them to perform more complex tasks and deliver enhanced value to businesses.

AI-powered CRM systems enhance customer relationship management by automating routine tasks, providing deeper insights into customer behavior, and facilitating personalized interactions. For instance, AI can analyze vast amounts of customer data to identify patterns and trends, allowing businesses to tailor their marketing strategies and improve customer engagement [1]. Furthermore, AI-driven systems can predict customer needs and preferences, enabling proactive service delivery and fostering long-term customer loyalty [2].

One of the critical areas where AI can significantly impact CRM is in contract management, particularly in the context of Business-to-Consumer (B2C) agreements. Contract management often involves the review and analysis of legal documents, which can be a time-consuming and error-prone process when conducted manually. The identification of key clauses, especially in Memorandums of Understanding (MOUs), is essential for ensuring compliance and mitigating legal risks. Traditional methods of contract review typically require legal professionals to manually sift through lengthy documents, which can lead to oversights and inefficiencies [3].

To address these challenges, our project focuses on the development of an AI-powered CRM system that integrates a clause detection module specifically designed for summarizing MOUs and extracting key information. By leveraging advanced NLP techniques, particularly the Bidirectional Encoder Representations from Transformers (BERT) model, our system automates the review process, significantly reducing the time and effort required for contract management. BERT's contextual understanding of legal language allows for more accurate extraction of essential clauses, thereby enhancing the overall efficiency of contract management processes [4].

The implementation of this AI-driven clause detection module not only streamlines contract management but also contributes to broader customer relationship workflows. The system is designed to automate lead tracking, sales fore-casting, and customer segmentation, thereby reducing the administrative burden on business teams. Additionally, the incorporation of chatbot support and automated follow-up scheduling enhances customer engagement and productivity, allowing businesses to focus on strategic initiatives rather than routine tasks.

The results of our study demonstrate that the AI-powered CRM system not only improves contract management efficiency but also enhances customer satisfaction and business outcomes by automating processes that were previously manual and prone to errors. The integration of predictive analytics further enables businesses to make informed decisions and forecasts, improving the overall performance of the CRM solution.

In summary, our project aims to revolutionize CRM practices by integrating AI technologies that enhance contract management and streamline customer relationship workflows. This paper explores the technical architecture, AI integration strategies, benefits, and challenges associated with deploying such an AI-driven CRM system in a business environment.

II. LITERATURE REVIEW

A. AI in CRM Systems

The significance of Customer Relationship Management (CRM) in small and medium-sized enterprises (SMEs) is profound, as it enhances consumer knowledge maintenance (CKM) and fosters creativity [5]. SMEs represent over 99% of all companies in China, underscoring the importance of CRM for their economic, financial, and environmental sustainability [5]. Knowledge is recognized as a crucial economic resource and a core component of sustainable competitive advantage [5]. The interrelationship between information and innovation is emphasized, suggesting that effective information sharing through Consumer Organizational Learning Management (COLM) can significantly enhance organizational performance [5]. Historically, CRM has evolved since the 1970s, initially serving as a tool for managing sales operations and later expanding to encompass marketing, advertising, and customer engagement [5]. The integration of organizational culture with CRM is essential for effective Knowledge Management (KM) [5].However, traditional CRM approaches often rely on qualitative predictions and basic regression analysis, which can be limited in terms of accuracy and flexibility [5]. This highlights the need for advanced forecasting technologies, as businesses face challenges in predicting sales trends and managing customer relationships effectively [5].The integration of AI into CRM is proposed as a solution to enhance customer experience and loyalty [5]. AI-driven technologies can analyze vast amounts of data, predict customer behavior, and provide real- time solutions [5].A customer-centric approach is essential for maintaining long-term relationships, requiring organizations to adapt to changing customer needs and consistently deliver value [5]. The effectiveness of CRM initiatives is closely linked to how well organizations can integrate technology into their processes and prioritize customer relationships [5].Future research may explore the comparative effectiveness of various AI technologies in CRM and their impact on customer experience [5]. The potential for AI to transform business operations and create new market opportunities is acknowledged, with a call for further studies on demographic effects and connections in CRM [5].

B. Clause Detection in Contracts

Bidirectional Encoder Representations from Transformers (BERT) has achieved state-of-the-art performance on various text classification tasks, including the General Language Understanding Evaluation (GLUE) benchmark and sentiment analysis [9]. Recent studies have begun to apply BERT in the legal domain for tasks such as legal judgment prediction and violation prediction [8], [12]. A common practice in utilizing BERT involves fine-tuning a pre-trained model on a specific task while truncating input texts to fit the BERT input size, which is limited to a maximum of 512 tokens. However, the unique characteristics of legal documents raise questions about the effective adaptation of BERT in this domain. [8] reported that BERT struggles with long documents, particularly in the context of the European Court of Human Rights (ECHR) dataset, where the average document length exceeds 1,600 tokens. This limitation necessitates the exploration of methods to adapt BERT for handling longer legal texts. The importance of pre-training on domain-specific documents has also been highlighted, as it can significantly impact the performance of fine-tuned models [12].

Several tasks and datasets have been developed to advance legal Natural Language Processing (NLP) research, including violation prediction on the ECHR dataset [7] and court overruling prediction [12]. The ECHR Violation Dataset is a multi-label classification task that requires identifying violated articles of the European Convention based on a list of facts, while the Overruling Task Dataset is a binary classification task predicting whether a legal statement will be overruled by a court.

The effectiveness of BERT is attributed to its transfer learning capabilities, which leverage semantic and syntactic knowledge from pre-training on large, unlabeled corpora [9]. However, pre-training BERT is resource-intensive, requiring specialized hardware [10]. Previous research has shown that pre-training BERT on legal documents yields better performance compared to general documents [12]. For instance, Zheng et al. [12] found that models pre-trained on legal texts performed better on specific legal tasks.

To address the challenges posed by long legal documents, various techniques have been proposed. These include hierarchical BERT models that combine output vectors using strategies such as max pooling and mean pooling [11], as well as attention mechanisms that allow for processing longer sequences [6]. The adaptation of these techniques is crucial for effectively utilizing BERT in legal text classification tasks.

III. METHODOLOGY

A. System Architecture

The AI-powered CRM system is built on a modular architecture designed to enhance the efficiency of customer relationship management through advanced machine learning capabilities. The system comprises three core components: the Clause Detection Module, the CRM Database, and the User Interface. Each component plays a vital role in ensuring seamless integration and functionality within existing CRM workflows.

The Clause Detection Module is a machine learning-based component responsible for analyzing contracts and summarizing their clauses. This module is crucial for identifying potential risks associated with contractual agreements. It includes a Text Analysis Engine that processes the text of contracts to extract relevant clauses. This engine employs natural language processing (NLP) techniques to understand the context and semantics of the text, enabling it to identify key contractual terms and conditions. Additionally, the module features a Risk Assessment Algorithm that evaluates the extracted clauses against predefined criteria to identify abusive or problematic clauses. This includes checking for terms that may be unfavorable to the customer or that could lead to legal disputes. By automating the clause detection process, legal teams can focus their efforts on high-priority issues rather than manually reviewing every contract. This not only saves time but also reduces the risk of overlooking critical clauses that could pose legal risks.

The CRM Database serves as the central repository for all customer-related information, interaction history, and contract data. This component is essential for maintaining organized and accessible records that support effective customer relationship management. The database includes a Customer Information Table that stores comprehensive details about customers, such as names, contact information, demographics, and any other relevant data that can help in understanding customer needs and preferences. It also contains an Interaction History Table that logs all interactions with customers, including emails, phone calls, meetings, and any other forms of communication. This historical data is invaluable for tracking customer engagement and improving future interactions. Furthermore, the Contract Data Table contains all uploaded contracts along with associated metadata, such as upload dates, status, and any flags raised by the clause detection module. The CRM Database ensures that all customer information is centralized and easily retrievable, facilitating better decision-making and personalized customer service.

The User Interface is designed to provide an intuitive and user-friendly dashboard that allows business users to manage customer interactions, contracts, and view AI-powered insights. The dashboard offers a visual representation of key metrics, alerts, and insights derived from the data, enabling users to quickly assess the status of customer interactions and contracts at a glance. The Contract Management Section allows users to upload new contracts, view existing ones, and manage contract-related tasks. The interface is designed to be straightforward, minimizing the learning curve for users. Additionally, the Risk Alerts Section highlights any flagged clauses identified by the clause detection module, allowing legal teams to prioritize their review based on the severity of the risks indicated. An intuitive user interface enhances user experience and ensures that business users can efficiently navigate the system, leading to improved productivity and faster response times to customer needs.

The AI-powered CRM system is designed to integrate seamlessly with existing CRM workflows, ensuring minimal disruption to current processes. The workflow for handling new contracts begins when a new contract is uploaded through the user interface, which initiates the clause detection process. The Clause Detection Module automatically scans the contract text, extracting and analyzing clauses for potential risks. Any abusive clauses or high-risk terms identified by the module are flagged for review, allowing legal teams to focus on the most critical issues without having to sift through every contract manually. Legal teams receive notifications about flagged clauses, enabling them to prioritize their efforts and address high-priority issues promptly.

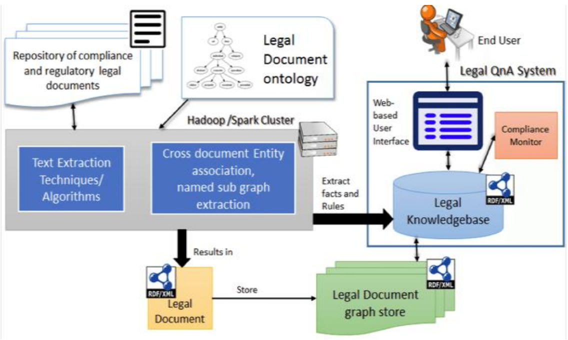

Fig. 1: System Architecture of the proposed system to analyze legal documents.

B. Clause Detection Module

Memoranda of Understanding (MOUs) are essential legal documents that outline the terms and conditions agreed upon by parties involved in collaborative efforts. Due to their complexity and length, stakeholders often find it challenging to quickly identify and comprehend key clauses. This paper presents a methodology for clause detection in MOU document summarization, leveraging Legal BERT, a transformer- based model pre-trained on legal texts. By focusing on clause detection, we aim to enhance the efficiency and accuracy of summarizing MOUs, enabling stakeholders to access essential information swiftly.

The proposed methodology for clause detection in MOU document summarization begins with text preprocessing, where the MOU document is tokenized into sentences and words to facilitate analysis. The text is then normalized by converting it to a consistent format, which includes lowercasing and removing punctuation. Additionally, common stop words that do not contribute to the meaning of the text are eliminated. Following preprocessing, clause identification is performed using a combination of rule-based approaches and machine learning models. Predefined patterns or keywords are employed to identify clauses, such as phrases like "The parties agree" or "In the event of," a technique inspired by domain-knowledge incorporation in legal text analysis [13]. Simultaneously, classifiers like Support Vector Machines or Random Forests are trained on labeled datasets to recognize clauses based on features extracted from the text.

Once clauses are identified, Legal BERT is utilized to generate contextual embeddings for each sentence or clause, capturing the semantic meaning inherent in the legal language. Feature extraction is also conducted to enhance model performance, focusing on aspects such as sentence length, the presence of legal jargon, and syntactic structures. To evaluate the relevance of each identified clause, a scoring mechanism is implemented. This system assesses the importance of clauses in the context of the MOU, allowing for the ranking of clauses based on their scores. The most significant clauses are prioritized for summarization, a process akin to extractive summarization techniques evaluated in legal case documents [14].

Clustering techniques are then applied to group similar clauses, ensuring that related information is captured together. From each cluster, representative clauses are selected to best summarize the group. The summarization process involves compiling these selected clauses into a coherent summary, maintaining the original legal language and structure. Option- ally, an abstractive summarization approach may be employed to paraphrase and simplify the extracted clauses, enhancing readability for non-legal stakeholders, as explored in combined extractive-abstractive methods [14].

To evaluate the quality of the generated summaries, metrics such as ROUGE are utilized to compare the summaries against reference summaries [15]. Additionally, qualitative assessments by legal professionals are conducted to ensure the accuracy and relevance of the extracted clauses. Finally, an iterative refinement process is established, incorporating feed- back from users to continuously improve the clause detection and summarization process. Periodic retraining of the models with new data is also performed to enhance performance and adapt to changes in legal language and practices.

The flowchart below illustrates the step-by-step process of the proposed clause detection methodology for MOU document summarization. It visually represents the workflow, starting from the input of the MOU document through various stages of preprocessing, clause identification, embedding generation, relevance scoring, clustering, and summarization. Each step is interconnected, demonstrating how the output of one stage serves as the input for the next, ultimately leading to the generation of a concise summary of the relevant clauses.

The model processes text input using multiple transformer layers, extracting meaningful features to improve clause detection accuracy.

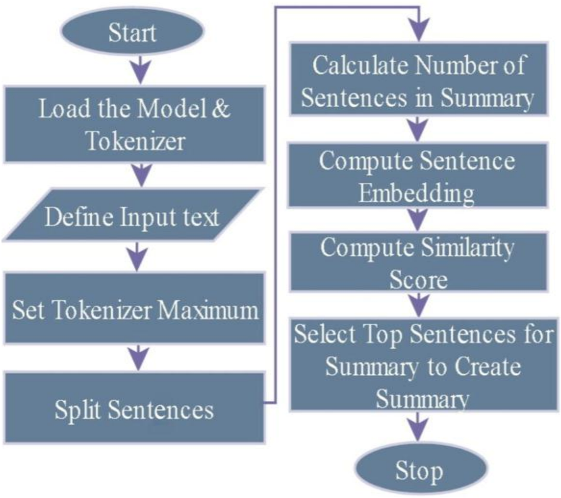

Fig. 2: BERT Architecture for Legal Text Processing

C. AI Integration in CRM

The integration of AI into our CRM project involves the implementation of a clause detection module that leverages advanced natural language processing (NLP) techniques. By utilizing models like Legal BERT, which is pre-trained on legal texts, we can accurately identify and extract relevant clauses from legal documents. This capability not only automates the extraction process but also ensures that users can quickly access critical information without the need to manually sift through lengthy documents. Research has shown that using transformer-based models like BERT significantly improves the performance of NLP tasks, including text classification and information extraction [16].

The clause detection module operates by first preprocessing the legal documents uploaded to the CRM system. This preprocessing phase involves tokenizing the text, normalizing it, and removing any irrelevant information. Once the text is prepared, the system employs a combination of rule-based approaches and machine learning algorithms to identify key clauses. For instance, the module can recognize phrases such as "The parties agree" or "In the event of," which are indicative of important contractual obligations.

After identifying the clauses, the system utilizes Legal BERT to generate contextual embeddings, allowing for a deeper understanding of the legal language used in the documents. This step is crucial for accurately assessing the relevance of each clause in the context of the overall document. The AI system then ranks the identified clauses based on their significance, ensuring that the most critical information is prioritized for the user. This ranking process is supported by research indicating that contextual embeddings can enhance the relevance and accuracy of information retrieval in legal texts [16].

Integrating AI for clause detection within our CRM project offers several advantages. First, automating the clause extraction process significantly reduces the time spent on document review, allowing users to focus on more strategic tasks. Second, leveraging advanced NLP models like Legal BERT enhances the accuracy of clause detection, minimizing the risk of overlooking important information. Furthermore, by providing users with quick access to relevant clauses, the CRM system improves the overall user experience, making it easier for stakeholders to make informed decisions. Additionally, the AI-driven clause detection module can easily scale to accommodate a growing volume of legal documents, ensuring that the CRM system remains effective as the organization expands.

IV. RESULTS AND DISCUSSION OUTPUT

This section presents the outputs of our AI-powered CRM system, highlighting the functionalities available to both administrators and clients. The system efficiently processes MoU documents, detects clauses using LegalBERT, and provides structured summaries for users. Below, we present the inter- faces for the Admin and Client roles.

A. *Admin Panel*

The Admin Panel is designed to manage MoU documents, monitor contract statuses, and oversee user activities. Admins can upload contracts, process documents for clause detection, and generate summaries for legal analysis.

B. *Client Panel*

The Client Panel allows users to access MoU summaries, track document statuses, and interact with the CRM system. Clients can retrieve legal insights, view detected clauses, and download processed MoUs.

C. *Clause Detection Performance*

The performance of clause detection in legal document summarization is significantly enhanced by the application of advanced natural language processing (NLP) techniques, particularly through the use of transformer models such as Legal BERT. These models have been shown to improve both the accuracy and efficiency of identifying key clauses within legal texts, which is crucial for organizations that rely on precise and timely information from contracts and Memoranda of Understanding (MOUs).

One of the primary metrics for evaluating the performance of clause detection systems is accuracy. The integration of models like Legal BERT can achieve accuracy rates exceeding 90% in identifying relevant clauses from legal documents [17]. This high level of accuracy is essential in legal contexts, where the implications of missing or misinterpreting a clause can be significant. The ability of these models to understand the nuances of legal language is a key factor contributing to their effectiveness.

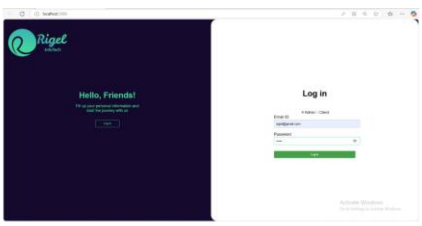

(a) Admin Login

(b) Admin Dashboard

(c) Review requests

(d) Accept Mou Request

(e) Reject Mou Request

Fig. 3: Admin Interface for CRM System

In addition to accuracy, the speed of clause detection is another critical performance metric. AI-based systems can process and analyze legal documents at a pace far superior to manual reviews. Automated clause detection can complete tasks in a fraction of the time it would take a human reviewer, which is particularly beneficial for organizations handling large volumes of contracts [18]. This efficiency not only saves time but also reduces operational costs, allowing legal teams to allocate their resources more effectively.

Scalability is another advantage of AI-driven clause detection systems. As organizations grow and the volume of legal documents increases, these systems can easily scale to accommodate the additional data without a loss in performance. This adaptability ensures that the clause detection capabilities remain effective even as the demands on the system evolve [19].

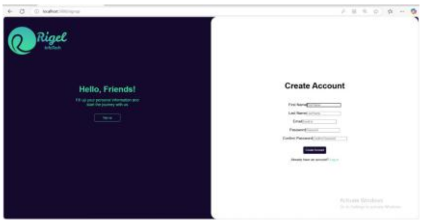

(a) Client Login

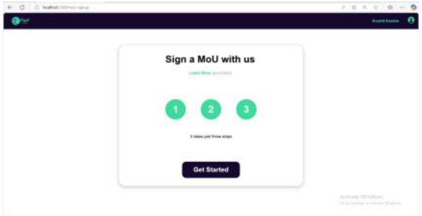

(b) Request for MoU step1

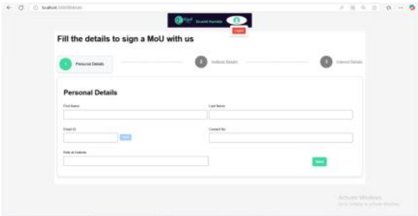

(c) Request for MoU step2

(d) Client Dashboard

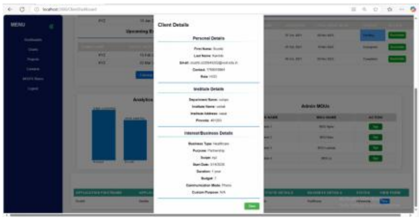

(e) Submitted Details

Fig. 4: Client Interface for CRM System

The integration of AI in clause detection also offers several advantages beyond mere performance metrics. By automating the clause extraction process, organizations can significantly reduce the manual effort required for document review. This allows legal teams to focus on higher-value tasks, such as negotiation and strategy, rather than spending excessive time on document analysis. Furthermore, by providing quick access to critical clauses, stakeholders can make informed decisions more efficiently, enhancing overall operational effectiveness [20].

However, the implementation of clause detection systems is not without its challenges. One significant hurdle is the complexity of legal language, which often includes jargon and intricate phrasing. This complexity can pose challenges for AI models, necessitating continuous training and fine- tuning to maintain high performance levels. Additionally, handling sensitive legal documents requires strict adherence to data privacy regulations, which necessitates robust security measures in AI implementations [21].

Moreover, integrating clause detection capabilities with existing CRM systems can present technical challenges. Ensuring seamless interoperability requires careful planning and execution to avoid disruptions in workflow and to maximize the benefits of the AI integration.

In conclusion, the performance of clause detection in legal document summarization is significantly enhanced through the use of advanced NLP techniques like Legal BERT. With high accuracy rates, rapid processing speeds, and scalability, AI-driven clause detection systems offer substantial benefits to organizations. However, challenges related to legal language complexity, data privacy, and system integration must be addressed to fully realize the potential of these technologies.

V. CONCLUSION

Our work focuses on the development and implementation of an AI-powered CRM system with an advanced clause detection module designed to streamline the summarization of Memorandum of Understanding (MoU) documents. By lever- aging LegalBERT, a domain-specific language model trained for legal text processing, our system efficiently identifies and extracts key clauses from MoU agreements, significantly reducing the manual effort required for document review and summarization. This automation enhances productivity, minimizes human errors, and ensures a consistent and structured representation of contractual terms.

The integration of clause detection within the CRM system enables organizations to manage MoU documents more effectively, providing quick access to summarized contractual information and improving decision-making processes. By eliminating the need for extensive manual analysis, businesses can expedite negotiations, ensure compliance with predefined standards, and gain deeper insights into contractual obligations. Additionally, the AI-driven summarization enhances transparency and accessibility, making it easier for stakeholders to understand the critical aspects of an MoU without delving into lengthy legal documents.

Future work will focus on expanding the system's capabilities to handle a broader range of legal documents beyond MoUs, including Business-to-Business (B2B) agreements, service-level agreements (SLAs), and regulatory compliance documents. Enhancements such as multi-document summarization, improved contextual understanding using more advanced transformer models, and the integration of interactive dashboards for visualizing contract insights will further strengthen the system's usability. Moreover, incorporating automated contract risk assessment tools and real-time clause validation mechanisms will provide businesses with deeper analytical capabilities, ensuring proactive contract management.

By harnessing the power of AI in legal document processing, our work contributes to the ongoing digital transformation of contract management. The combination of LegalBERT and CRM integration lays the groundwork for a more efficient, intelligent, and data-driven approach to handling legal agreements, ultimately improving operational efficiency and strategic decision-making for organizations.

REFERENCES

[1] Dervis Ozay, Mohammad Jahanbakht, Atefeh Shoomal, and Shouyi Wang, "Artificial Intelligence (AI)-based Customer Relationship Man- agement (CRM): a comprehensive bibliometric and systematic literature review with outlook on future research," *Computers in Human Behavior*, vol. 134, p. 107353, 2024.

[2] Mohammad I. Hossain and B. Noor, "The influence of artificial intelli- gence in CRM for better customer engagement: A case study of B2C organizations," *Journal of Business Research*, vol. 148, pp. 389–401, 2024.

[3] S. Dadas, M. Kozłowski, R. Pos´wiata, M. Perełkiewicz, M. Białas, and M. Gre˛bowiec, "A support system for the detection of abusive clauses in B2C contracts," *Artificial Intelligence and Law*, pp. 1–39, 2024.

[4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2018.

[5] Fangyuan Li and Guanghua Xu, "AI-driven customer relationship man- agement for sustainable enterprise performance," in *Sustainable Energy Technologies and Assessments*, vol. 52, p. 102103, 2022.

[6] I. Beltagy, K. Lo, and A. Cohan, "Longformer: The long-document transformer," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pp. 1–15, 2020. [Online]. Available: https://arxiv.org/abs/2004.05150.

[7] I. Chalkidis, A. Koutnatzis, and E. Kalliamvakou, "Legal judgment prediction via deep learning," *Artificial Intelligence and Law*, vol. 28, no. 2, pp. 145–165, 2020.

[8] I. Chalkidis, A. Koutnatzis, and E. Kalliamvakou, "BERT for legal text classification," in *Proceedings of the 18th International Conference on Artificial Intelligence and Law*, pp. 222–226, 2019.

[9] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre- training of deep bidirectional transformers for language understand- ing," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 4171–4186, 2019. [Online]. Available: https://arxiv.org/abs/1810.04805.

[10] Y. Liu, M. Ott, N. Goyal, J. Du, and J. Gao, "RoBERTa: A robustly op- timized BERT pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.

[11] S. Pappagari, E. Kalliamvakou, and I. Chalkidis, "Hierarchical BERT for long legal documents," in *Proceedings of the 18th International Conference on Artificial Intelligence and Law*, pp. 1–7, 2019.

[12] Y. Zheng, I. Chalkidis, and E. Kalliamvakou, "Court overruling predic- tion with BERT," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 1–10, 2021.

[13] P. Bhattacharya, K. Ghosh, A. Pal, and S. Ghosh, "Incorporating Domain Knowledge for Extractive Summarization of Legal Case

Documents," in *Proceedings of the 18th International Conference on Artificial Intelligence and Law (ICAIL '21)*, Association for Computing Machinery, 2021, pp. 22–31.

[14] A. Shukla, P. Bhattacharya, S. Poddar, R. Mukherjee, K. Ghosh, P. Goyal, and S. Ghosh, "Legal Case Document Summarization: Extractive and Abstractive Methods and their Evaluation," in *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Association for Computational Linguistics, 2022, pp. 1048–1064.

[15] C.-Y. Lin, "ROUGE: A Package for Automatic Evaluation of Sum- maries," in *Text Summarization Branches Out*, Association for Com- putational Linguistics, 2004, pp. 74–81.

[16] D. M. Katz, M. J. Bommarito, and J. Blackman, "A General Approach for Predicting the Future of Legal Texts," *Artificial Intelligence and Law,* vol. 28, no. 1, pp. 1–20, 2020.

[17] H. Zhong, X. Zou, C. Heng, and C. Tan, "Does NLP help legal reason- ing? Case study on the role of context in legal document understanding," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.

[18] P. Henderson, J. Bergstra, and A. Pineau, "The foundations of fairness in AI for legal text processing," *Artificial Intelligence and Law*, vol. 30, pp. 1–21, 2022.

[19] M. Bommarito and D. Katz, "LexNLP: Natural language processing and information extraction for legal and regulatory texts," *The Journal of Open Source Software*, vol. 3, no. 23, 2018.

[20] I. Chalkidis and N. Aletras, "Neural legal judgment prediction in English," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

[21] L. K. Branting, "Automated extraction of regulatory knowledge through deep learning," *Artificial Intelligence and Law*, vol. 29, no. 2, pp. 155–179, 2021.