International Journal of Applied Engineering Research ISSN 0973-4562 Volume 21, Number 2 (2026) pp.111-118
 https://dx.doi.org/10.37622/IJAER/21.2.2026.111-118

AI-Enhanced Multilingual Voice-Based Legal Form Automation Platform for Citizens

Dr. Anupama K C, Aadarsh Keshri, Aakriti, Rahul S Kamath, Rajesh M R

Department of Computer Science and Engineering
(IoT & Cybersecurity including Blockchain Technology)
Bangalore Institute of Technology, Bengaluru, India

Abstract :
Legal processes are often difficult for citizens to access due to linguistic diversity, technical terminology, and limited digital familiarity. This paper presents a multilingual, voice-driven legal form automation platform that allows users to complete legal documentation through natural speech interaction. The system combines speech recognition, natural language understanding, intelligent field mapping, and automated document generation to simplify legal workflows. Experimental results indicate improved accessibility, reduced completion time, and enhanced usability.

Keywords : Legal AI, Speech Recognition, Multilingual NLP, E-Governance, Voice Interfaces, Document Automation

1 INTRODUCTION

Access to legal services continues to be a significant challenge, particularly in countries like India where linguistic diversity and varying levels of digital literacy create barriers for citizens. While sectors such as banking, healthcare, and taxation have undergone rapid digital transformation, legal documentation processes remain complex and difficult for non-technical users to navigate. Legal forms often involve unfamiliar terminology, strict formatting requirements, and procedural dependencies, making them challenging to understand and complete correctly.

In conventional legal workflows, individuals frequently rely on intermediaries such as document writers, clerks, or legal professionals even for routine tasks such as affidavit preparation, name change applications, property declarations, and traffic-related appeals. Although legal professionals are essential for handling complex disputes, many basic documentation processes can be simplified through guided digital systems. Errors or incomplete submissions often result in rejection of applications, leading to delays and increased administrative burden.

The multilingual nature of India introduces additional complexity. Many official legal forms are primarily available in English or Hindi, limiting accessibility for individuals who are more comfortable with regional languages. Citizens in rural and semi-urban areas may also face challenges due to limited digital exposure, further increasing dependency on third-party assistance. These factors contribute to increased cost, time, and effort required to access legal services.

Recent progress in artificial intelligence, particularly in speech processing and natural language understanding, has enabled systems to interpret conversational input and convert it into structured information [1, 23, 24]. Such capabilities are especially valuable in legal workflows, where user-provided input must be transformed into formally structured documents. Voice-driven interaction reduces reliance on typing and simplifies user engagement, making systems more inclusive and accessible.

This work presents an AI-powered multilingual legal assistance platform that allows users to complete legal forms using natural speech interaction. Users can provide input in their preferred language, which is converted into text using Automatic Speech Recognition (ASR) [20]. The processed text is analyzed using Large Language Models (LLMs) to identify legal intent and extract relevant entities [1, 25].

To improve accuracy and reliability, the platform incorporates a hybrid validation mechanism that combines rule-based checks with AI-driven contextual validation. This approach helps ensure that generated documents are both complete and logically consistent, reducing the likelihood of rejection during submission. In contrast to traditional legal chatbots that focus primarily on answering queries, the proposed platform supports end-to-end automation, including form filling, validation, document generation, and storage.

The overall architecture is designed using a layered approach consisting of user interaction, client application, processing, storage, and external service layers. This modular design enables scalability, maintainability, and secure handling of sensitive information. The platform also supports multiple input modes, including voice and text, ensuring flexibility for diverse user groups.

1.1 System Interaction Overview

The platform supports voice-based, text-based, and assisted input modes. Voice input is processed through a pipeline that includes noise filtering, normalization, and language detection. The processed data is then analyzed for intent classification and entity extraction, ensuring consistent performance across different input formats.

1.2 Real World Impact and Motivation

The proposed system is intended to assist citizens in completing routine legal processes independently. By providing guided interaction and simplified explanations, it reduces reliance on intermediaries and improves user confidence. The system is designed as an assistive tool rather than a replacement for legal professionals.

From a deployment perspective, a hybrid architecture is adopted that combines cloud-based processing with lightweight local operations. This ensures usability even in environments with limited connectivity. Privacy considerations are addressed through minimal data storage, secure authentication mechanisms, and controlled access to user information.

1.3 Research Contributions

The key contributions of this work include:

- A multilingual voice-driven legal form automation pipeline
- Integration of Large Language Models for intent detection and entity extraction
- A hybrid validation framework combining rule-based and AI-based verification
- Secure and privacy-aware handling of legal data
- A scalable architecture supporting future integration with government systems

These contributions aim to improve accessibility, reduce complexity, and support citizen-centric legal service delivery.

2 LITERATURE REVIEW

The integration of artificial intelligence in legal technology has evolved significantly in recent years, particularly in areas such as document automation, legal information retrieval, and conversational assistance systems. Despite these advancements, many existing solutions still lack comprehensive support for multilingual interaction and end-to-end automation of legal workflows.

Early work in natural language processing demonstrated that transformer-based models are capable of capturing contextual relationships within text [23, 24], allowing meaningful information to be extracted from unstructured inputs. Such capabilities are highly relevant in the legal domain, where user-provided information often needs to be interpreted and structured into formal documentation. These models form the foundation for modern AI-driven legal assistance systems.

Speech-based interaction has also been explored as a means to improve accessibility [5, 16]. Studies indicate that combining speech recognition with language processing can significantly enhance usability, especially for individuals who are not comfortable with typing or navigating complex interfaces. In multilingual environments like India, voice-based systems provide a more natural and inclusive method of interaction.

Several digital legal platforms have been developed to improve access to justice [4, 11] by enabling remote assistance and reducing dependency on physical legal offices. However, many of these systems primarily focus on providing guidance or answering queries, rather than supporting complete document automation. As a result, users often still require manual intervention to finalize legal documents.

Recent research has introduced AI-driven legal chatbots capable of handling user queries [2, 10] and providing basic legal guidance. While these systems improve accessibility, they often lack mechanisms for structured document generation and validation. This limitation highlights the gap between conversational assistance and fully automated legal workflows.

Other approaches have explored retrieval-based and knowledge graph-driven systems [3, 7] to provide legal recommendations. Although these methods enhance information access, they typically require manual verification and are not designed for real-time document creation. This indicates the need for systems that combine intelligent reasoning with structured output generation.

Multilingual speech processing remains an important research challenge [12, 13]. Variations in accents, limited training datasets for regional languages, and domain-specific terminology can affect recognition accuracy. These challenges are particularly critical in legal contexts, where precise interpretation of user input is necessary.

Cloud-based architectures have also been widely adopted in legal automation systems [8, 22] to support scalability, distributed access, and secure storage of sensitive data. Security frameworks emphasize the importance of encryption, authentication, and audit logging to ensure data protection and compliance.

Despite the progress in individual areas such as speech processing, natural language understanding, and legal automation, existing systems often lack integration across these components. Many solutions do not provide seamless workflows from user interaction to document generation and validation.

The proposed work addresses these limitations by combining multilingual speech recognition, AI-driven intent detection, structured entity extraction, hybrid validation mechanisms, and automated document generation into a unified platform. This integrated approach aims to improve accessibility, reduce manual effort, and enable citizen-centric legal service delivery.

3 METHODOLOGY

The proposed platform follows a structured, multi-stage processing pipeline designed to convert unstructured speech input into legally valid and structured document outputs. The methodology emphasizes accuracy in speech interpretation, intent detection, entity extraction, validation, and document generation. Each stage is designed as an independent module to ensure scalability, maintainability, and efficient processing.

3.1 Overall Processing Pipeline

The system operates through the following key stages:

1. User Authentication and Session Initialization
2. Voice Input Capture and Preprocessing
3. Speech-to-Text Conversion
4. Language Detection and Normalization
5. Legal Intent Detection using Large Language Models
6. Legal Entity Extraction and Field Mapping
7. Hybrid Validation Processing
8. Document Generation and Secure Storage

3.2 User Authentication Module

Secure access to the platform is ensured through OAuth-based authentication combined with JSON Web Token (JWT) session management. This approach minimizes risks associated with password storage while maintaining secure communication between client and server. Each authenticated session is validated before allowing access to protected services.

3.3 Speech Input Processing Module

Audio input is captured using browser-based or device-based microphone interfaces. The captured signal undergoes preprocessing steps such as noise reduction, normalization, silence removal, and segmentation. These steps improve clarity and ensure reliable speech recognition across different environments.

The processed audio is then converted into text using Automatic Speech Recognition (ASR) models [20]. Preprocessing combined with contextual correction improves accuracy, especially in multilingual scenarios.

3.4 Language Detection and Normalization

The textual output is analyzed to identify the input language using multilingual detection models. Based on the detected language, appropriate processing pipelines are selected. Normalization techniques are applied to standardize formats such as dates, numbers, and identifiers, ensuring consistency across different inputs.

3.5 Legal Intent Detection using Large Language Models

The normalized text is processed using Large Language Models to determine the user's intent [1, 25]. Instead of relying on predefined templates, the model interprets contextual meaning and classifies the type of legal form required.

The output is structured into a machine-readable format containing:

- Identified legal form type
- Required entities
- Optional supporting information
- Validation constraints

3.6 Legal Entity Extraction and Field Mapping

Relevant entities such as names, addresses, dates, and identification details are extracted from the processed input. These entities are mapped to predefined legal form templates using schema-based mapping techniques. This step ensures that user input is accurately aligned with required document fields.

3.7 Hybrid Validation Engine

To ensure correctness, a hybrid validation mechanism is applied. Rule-based validation checks for mandatory fields, formatting constraints, and predefined rules. In parallel, AI-based validation verifies contextual consistency and logical correctness of the provided information [22].

Combining these approaches reduces errors and increases the likelihood of successful document submission.

3.8 Document Generation Engine

Once validation is completed, the structured data is used to generate legally formatted documents. Template-based rendering ensures that the output adheres to required legal standards. The generated documents are typically exported in PDF format for submission purposes.

3.9 Cloud Storage and Tracking

Generated documents are stored securely in cloud-based systems. Metadata and audit logs are maintained to track document history and user activity. This enables reliable access, traceability, and compliance with data management requirements.

3.10 System Security Approach

Security is implemented across all stages of the pipeline. Key measures include encrypted communication (HTTPS), token-based authentication, secure credential storage, API rate limiting, and audit logging. These mechanisms ensure protection of sensitive legal data.

3.11 Performance Optimization Strategy

To improve efficiency, the system incorporates caching mechanisms, asynchronous processing, and optimized API interactions. The modular architecture allows independent scaling of resource-intensive components such as AI processing modules, ensuring consistent performance under varying workloads.

4 SYSTEM ARCHITECTURE

The proposed platform is designed using a layered and microservices-based architecture to ensure scalability, modularity, and secure processing of legal data. The architecture enables seamless integration of speech processing, AI-driven reasoning, document generation, and cloud-based storage [8].

The architecture is divided into five primary layers, each responsible for a specific set of functionalities:

1. User Interaction Layer
2. Client Application Layer
3. Application Processing Layer
4. Database and Storage Layer
5. External Service Integration Layer

4.1 User Interaction Layer

The User Interaction Layer serves as the entry point for users interacting with the platform. It supports both voice-based and text-based input, allowing users to provide legal information in a flexible manner. Multilingual support is incorporated to ensure accessibility for users from diverse linguistic backgrounds.

This layer focuses on usability by providing:

- Voice-driven input for natural interaction
- Text-based input as an alternative mode
- Real-time feedback during form completion
- Accessible interface design for non-technical users

Voice-based interaction significantly reduces dependency on typing and improves usability, especially for users with limited digital familiarity.

4.2 Client Application Layer

The Client Application Layer is responsible for managing user sessions and handling communication between the user interface and backend services. It ensures smooth data transmission and responsive interaction.

Key responsibilities include:

- Session handling and state management
- Input validation before submission
- Rendering of dynamic user interfaces
- API request handling and response display

Lightweight preprocessing at this layer reduces backend load and improves overall system responsiveness.

4.3 Application Processing Layer

The Application Processing Layer forms the core intelligence of the platform. It integrates multiple services responsible for interpreting user input and generating structured outputs.

4.3.1 Speech Processing Service

Audio input is converted into text using Automatic Speech Recognition models. Preprocessing techniques such as noise filtering and normalization enhance recognition accuracy.

4.3.2 AI Legal Reasoning Service

Large Language Models analyze the processed text to identify legal intent and extract relevant structured data. Contextual understanding enables accurate mapping between user input and legal form requirements.

4.3.3 Validation Engine

A hybrid validation mechanism is applied to verify both structural correctness and contextual accuracy. Rule-based validation ensures required fields are present, while AI-based validation checks logical consistency.

4.3.4 Document Generation Service

Structured data is transformed into legally formatted documents using predefined templates. This ensures that generated outputs meet standard legal formatting requirements.

4.4 Database and Storage Layer

The Database Layer manages storage of user data, legal form data, generated documents, and system logs. A cloud-based NoSQL database is used to support flexible schema design and high scalability.

This layer supports:

- Storage of user profiles and session data
- Legal form data management
- Document metadata and version tracking
- Audit logs for traceability

Cloud-based storage enables high availability, fault tolerance, and distributed access.

4.5 External Service Integration Layer

The External Services Layer connects the platform with third-party services required for functionality and scalability.

Integrated services include:

- Speech recognition APIs
- Large Language Model APIs
- OAuth-based authentication services
- Email and notification services
- Future integration with government legal systems

Using external services reduces infrastructure complexity and accelerates deployment.

4.6 Security Architecture

Security is implemented across all layers to ensure protection of sensitive legal data. The platform incorporates multiple safeguards including encrypted communication, token-based authentication, and controlled access mechanisms.

Key security features include:

- HTTPS-based secure communication
- OAuth authentication
- JWT-based session management
- API rate limiting
- Secure storage of credentials
- Role-based access control

4.7 Scalability and Performance Design

The architecture is designed to support horizontal scalability using a microservices approach. Resource-intensive components such as AI processing modules can be scaled independently based on demand.

Caching mechanisms and optimized API interactions reduce latency and improve response time.

4.8 Reliability and Fault Tolerance

Reliability is ensured through the use of fallback mechanisms, retry strategies, and distributed storage replication. This minimizes system downtime and ensures consistent performance even under high load conditions.

5 ALGORITHMS AND AI MODELS

The platform utilizes multiple AI-driven algorithms to convert unstructured speech input into structured legal document data. The processing pipeline includes speech recognition, intent detection, entity extraction, validation, and document generation.

5.1 Speech Input Processing Algorithm

[H] Speech Input Processing Algorithm User Audio Input Clean Text Output

Capture audio stream Apply noise filtering Normalize audio signal Detect input language Convert speech to text Perform token normalization Return clean text output

5.2 Legal Intent Detection Algorithm

[H] Legal Intent Detection using LLM Clean Text Input Legal Intent Classification

Load LLM model Parse input text Extract semantic context Identify legal domain Classify legal form type Return structured intent data

5.3 Legal Entity Extraction Algorithm

[H] Legal Entity Extraction Processed Text Input Extracted Legal Entities

Load entity extraction model Extract person names Extract addresses Extract dates Extract identification numbers Validate extracted data Return structured entity data

5.4 Legal Field Mapping Algorithm

[H] Legal Field Mapping Algorithm Extracted Entities Completed Legal Form Structure

Load legal form template Map entities to form fields Check mandatory field completeness Generate auto-fill suggestions Return structured form data

5.5 Hybrid Legal Validation Algorithm

[H] Hybrid Legal Validation Algorithm Structured Legal Form Data Validated Legal Document

Check mandatory fields Validate formatting rules Perform AI-based context validation Verify logical consistency validation passes Approve document Generate error feedback

5.6 Document Generation Algorithm

[H] Legal Document Generation Algorithm Validated Legal Form Data Generated Legal PDF

Load legal template Populate template fields Format document layout Generate PDF output Store document in cloud Return document reference ID

6 IMPLEMENTATION

The implementation of the proposed platform focuses on building a modular and scalable system capable of handling real-time speech processing, AI-driven analysis, and legal document generation. A microservices-oriented approach is adopted so that individual components can operate independently while communicating through secure APIs.

6.1 Frontend Implementation

The user interface is developed using modern web technologies to support responsive and interactive workflows. The design prioritizes simplicity and accessibility, enabling users to complete legal forms with minimal effort.

The frontend stack includes:

- React.js for component-based UI development
- Next.js for optimized rendering and routing
- Tailwind CSS for responsive styling
- Web Speech API for capturing voice input

Both voice and text input modes are supported, allowing users to interact with the system based on their comfort level.

6.2 Backend Implementation

Backend services are implemented using a RESTful architecture to handle communication between frontend components, AI services, and storage systems. The backend manages authentication, data processing, and document generation workflows.

Core backend modules include:

- Authentication and session management service
- Speech processing integration layer
- Legal intent detection service
- Document generation module
- Audit logging system

These services are designed to operate independently, improving scalability and maintainability.

6.3 Database Implementation

A cloud-based NoSQL database is used to store user data, legal form information, and generated documents. The flexible schema allows efficient handling of structured and semi-structured data.

The database maintains:

- User profiles and authentication data
- Legal form templates and filled data
- AI processing logs
- Document metadata and history
- Audit trails for system activities

6.4 AI Integration

Artificial intelligence components are integrated through API-based services.

6.4.1 Speech Recognition

Speech input is converted into text using ASR models. Preprocessing techniques improve performance in noisy and multilingual environments.

6.4.2 Language Model Processing

Large Language Models analyze user input to identify intent and extract structured legal information. These models enable contextual understanding of conversational input.

6.4.3 Validation Models

AI-based validation ensures logical consistency and improves the reliability of generated documents.

6.5 Authentication and Security

Security is implemented through multiple layers to protect sensitive data:

- OAuth-based authentication
- JWT-based session handling
- HTTPS encrypted communication
- Secure environment variable management
- API rate limiting and monitoring

6.6 Document Generation

Legal documents are generated using template-based rendering. The system ensures proper formatting and compliance with standard legal document structures. Output files are generated in PDF format and stored securely.

6.7 Performance Optimization

System performance is improved through:

- Caching frequently used responses
- Asynchronous processing of AI tasks
- Optimized API communication
- Efficient database indexing

6.8 Deployment and Testing

The platform is deployed using cloud-based infrastructure with containerized services. Continuous integration and deployment pipelines support regular updates.

Testing includes:

- Unit testing for individual modules
- Integration testing across services
- Performance and load testing
- Security validation
- User acceptance testing

7 RESULTS AND DISCUSSION

The performance of the proposed platform was evaluated across multiple parameters, including speech recognition accuracy, entity extraction performance, document generation reliability, system latency, and overall user experience.

7.1 Experimental Setup

Evaluation was conducted using test cases covering different legal scenarios such as affidavit generation, name change applications, and traffic-related forms. The system was tested under varying network conditions and input types to simulate real-world usage.

7.2 Speech Recognition Performance

Speech recognition accuracy was measured across different languages. The use of preprocessing techniques improved performance in noisy environments.

Observed accuracy:

- English: 96.2%
- Hindi: 94.8%
- Regional languages (average): 92.4%

7.3 Entity Extraction Performance

The system's ability to extract structured information from user input was evaluated using multiple datasets.

Observed results:

- Name extraction: 97.1%
- Address extraction: 93.6%
- Date extraction: 95.4%
- Legal declaration detection: 91.8%

7.4 Document Generation Performance

The document generation module was tested for accuracy and reliability.

- Successful document generation rate: 98.3%
- Validation failure rate: 1.7%

The hybrid validation approach significantly reduced errors compared to rule-only validation systems [22].

7.5 System Latency

The total processing time from input to document generation was measured.

- Speech processing: 1.2 seconds
- AI processing: 1.8 seconds
- Document generation: 0.9 seconds
- Total average time: 3.9 seconds

7.6 System Reliability

The system demonstrated stable performance in cloud deployment environments.

- Uptime: 97.4%
- API success rate: 98.9%
- Database success rate: 99.1%

7.7 User Experience Evaluation

User testing indicated improvements in usability and efficiency.

- Reduction in form completion time: 38%
- User satisfaction score: 4.3/5
- Reduction in submission errors: 41%

7.8 Discussion

The results demonstrate that combining speech recognition, language models, and hybrid validation significantly improves the efficiency [20, 25] of legal document automation. Multilingual support enhances accessibility, particularly for users with limited digital literacy.

Compared to traditional methods, the platform reduces manual effort, improves accuracy, and provides a more user-friendly interaction model. These findings indicate strong potential for deployment in real-world legal service systems.

8 CONCLUSION

This paper presents an AI-powered multilingual voice-enabled platform designed to simplify legal document automation for citizens. By combining speech recognition, natural language understanding, and intelligent validation mechanisms, the platform enables users to complete legal forms through intuitive voice interaction.

The integration of Automatic Speech Recognition and Large Language Models allows conversational input to be transformed into structured legal data with high accuracy. The use of hybrid validation further improves reliability by ensuring both syntactic correctness and contextual consistency of generated documents. As a result, the system reduces errors, minimizes dependency on intermediaries, and improves overall efficiency in legal workflows.

The layered architecture ensures scalability, modularity, and secure handling of sensitive data. Experimental evaluation demonstrates strong performance across key metrics, including speech recognition accuracy, entity extraction accuracy, document generation success rate, and system responsiveness.

Despite these strengths, certain limitations remain. The system relies on internet connectivity for AI processing, and performance may vary for low-resource regional languages due to limited training data. Addressing these challenges will further enhance usability and robustness.

Overall, the proposed platform demonstrates the potential of AI-driven systems in improving accessibility, reducing complexity, and enabling citizen-centric legal service delivery.

9 FUTURE SCOPE

While the proposed platform demonstrates strong performance in multilingual legal automation, several enhancements can further improve its capabilities and real-world applicability.

Future developments can focus on improving contextual intelligence by incorporating user history and interaction patterns. Context-aware systems can provide personalized suggestions, reduce repetitive input, and enhance user experience.

Enhancing speech recognition accuracy for low-resource regional languages remains an important area of improvement. Training domain-specific models using legal vocabulary datasets can improve performance in multilingual environments.

Integration with government legal service platforms represents a significant opportunity for expansion. Direct connectivity with systems such as e-Courts, DigiLocker, and digital identity verification services can enable seamless document submission and verification workflows.

Privacy-preserving AI techniques such as federated learning can be explored to ensure that sensitive user data remains secure while still enabling model improvements. This approach is particularly relevant in legal applications involving confidential information.

Future versions can also include mobile application deployment with offline capabilities, allowing users in low-connectivity environments to access basic features such as speech input and document generation.

Additionally, incorporating document scanning and OCR-based input can enable users to upload and process existing legal documents. Real-time analytics dashboards can also be developed to monitor system performance and user interaction patterns.

In the long term, the platform can evolve into a comprehensive legal assistance ecosystem capable of supporting end-to-end workflows, including case filing, document preparation, tracking, and advisory support.

References

[1] T. Brown et al., “Language Models are Few-Shot Learners,” NeurIPS, 2020.

[2] M. Dhore et al., “BetterCall: AI-Based Legal Assistant,” Computational Law Review, 2023.

[3] S. Jacob et al., “LegalLink: Retrieval-Augmented Legal Assistance,” AI and Law Journal, 2023.

[4] R. Brescia, “Technology and Access to Justice in Rural Communities,” Center for Access to Justice Report, 2018.

[5] D. Jurafsky and J. H. Martin, Speech and Language Processing, Pearson, 2023.

[6] T. Brugger et al., “Multilingual Legal Sentence Boundary Detection,” LREC Conference, 2022.

[7] V. Rodríguez-Doncel and E. Montiel-Ponsoda, “Multilingual Legal Knowledge Graphs,” ISWC Conference, 2020.

[8] NIST, Cloud Computing Security Reference Architecture, 2023.

[9] N. Jain and G. Goel, “AI-Based Legal Assistance Using Rule-Based Systems,” International Conference on Legal Informatics, 2019.

[10] A. R. Kandula et al., “Chatbot-Based Legal Advisory Using NLP,” International Journal of Advanced Legal Studies, 2020.

[11] I. S. Ayebo and T. Ayodele, “Digital Platforms for Pro Bono Legal Services,” Journal of Legal Technology, 2021.

[12] J. Iranzo-Sánchez et al., “Europarl-ST: Multilingual Speech Translation for Legal Domain,” ACL Workshop, 2020.

[13] N. Garneau and O. Bolduc, “Evaluation of Legal Speech Recognition Systems,” Legal Speech Research Lab, 2024.

[14] R. Stern et al., “One Law, Many Languages: Multilingual Legal Reasoning,” EMNLP, 2022.

[15] H. Westermann and J. Savelka, “Multimodal LLMs for Legal Document Understanding,” ICDAR, 2023.

[16] F. A. Habeeb et al., “Speech-Enabled Web Accessibility Interfaces,” W4A Conference, 2021.

[17] O. Isaac and E. Godwill, “Digital Legal Empowerment in Rural Regions,” Journal of Public Interest Technology, 2020.

[18] Y. Xu et al., “Automated Form Understanding Using LayoutLM,” KDD, 2020.

[19] G. Kim et al., “Donut: OCR-Free Document Understanding,” ECCV, 2022.

[20] A. Radford et al., “Whisper: Robust Multilingual Speech Recognition,” ICML, 2023.

[21] P. Liu et al., “Prompt Engineering for Domain-Specific LLMs,” ACM Computing Surveys, 2023.

[22] R. Shokri et al., “Privacy-Preserving AI Systems for Public Services,” IEEE Security and Privacy, 2021.

[23] A. Vaswani et al., “Attention Is All You Need,” NeurIPS, 2017.

[24] J. Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers,” NAACL, 2019.

[25] Y. Gao et al., “Evaluating Large Language Models for Legal Tasks,” AAAI, 2024.

[26] K. Ashley, Artificial Intelligence and Legal Analytics, Cambridge University Press, 2022.

[27] M. Whittaker et al., “AI Accountability in Public Systems,” AI Now Institute Report, 2021.

[28] MongoDB Inc., MongoDB Atlas Documentation, 2024.

[29] Pallets Projects, Flask Official Documentation, 2024.

[30] OpenAI, OpenAI API Documentation, 2024.

[31] Google Developers, OAuth 2.0 Authentication Documentation, 2024.

[32] Google Cloud, Speech-to-Text API Documentation, 2024.