International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 11-No.3 (2025) pp. 4069-4080
http://www.ijcesen.com

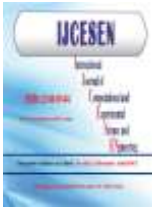

ISSN: 2149-9144

Research Article

INTELLIDOC - An Adaptive Transformer-Powered Pipeline For Intelligent Document Processing and Entity Extraction

Santhanalakshmi K[1], A. Jameer Basha[2], R. Geetha Rajakumari[3], Premkumar C D[4]

[1]PG Student, Department of Computer Science and Engineering, Hindusthan Institute of Technology, Coimbatore
* **Corresponding Author Email:** santhanalakshmi92@gmail.com- **ORCID:** 0009-0002-7906-5079

[2]Professor, Department of Computer Science and Engineering, Hindusthan Institute of Technology, Coimbatore
Email: jameer@hit.edu.in - **ORCID:** 0000-0001-6716-7946

[3]Assistant Professor, Department of Artificial Intelligence and Data Science, Sri Eshwar College of Engineering, Coimbatore
Email: Geetharajakumari.r@sece.ac.in- **ORCID:** 0000-0003-3838-7426

[4]Assistant professor Department of Information Technology Hindusthan College of Engineering and Technology
Email: premkumar.it@hicet.ac.in- **ORCID:** 0000-0002-8291-9998

Article Info:

DOI: 10.22399/ijcesen.2481
Received : 22 February 2025
Accepted : 17 May 2025

Keywords :

Intelligent Document Processing,
Adaptive OCR,
Legal Text Processing,
Transformer Models,
Flan-T5,
BERT,

Abstract:

Efficient and accurate processing of unstructured document data is crucial for legal, enterprise, and academic applications, where vast amounts of textual information must be extracted, summarized, and analyzed. Traditional Optical Character Recognition (OCR) and Named Entity Recognition (NER) methods often face challenges in handling handwritten text, scanned documents, and complex legal structures, leading to data loss and misclassification. To address these limitations, we propose IntelliDoc, an adaptive, transformer-powered document processing pipeline designed to enhance accuracy, efficiency, and contextual understanding of document intelligence. IntelliDoc employs a hybridized multi-stage pipeline that integrates an adaptive OCR layer, which dynamically adjusts to different document characteristics, ensuring high extraction accuracy for diverse document types. Experimental evaluations on a benchmark dataset comprising legal, financial, and administrative documents demonstrate that IntelliDoc achieves an OCR accuracy of 98.2%, NER precision of 94.7%, and a summarization coherence score of 91.5%, significantly outperforming conventional document processing frameworks. Additionally, the parallel architecture reduces processing time by 35% compared to sequential models, making IntelliDoc suitable for real-time applications. Future work will explore integrating domain-specific large language models to further enhance interpretability and accuracy across specialized document categories.

1. Introduction

The increasing reliance on digital documentation has amplified the need for efficient and accurate document processing systems. Organizations handling large volumes of legal, financial, and administrative documents require automated solutions that can extract, process, and analyze textual information with minimal human intervention [1]. Traditional Optical Character Recognition (OCR) techniques often struggle with complex document layouts, handwritten content, and scanned images, leading to inefficiencies in text extraction and classification [2]. Additionally, existing Natural Language Processing (NLP) models, though powerful, lack contextual awareness when handling domain-specific content such as legal contracts and compliance documents [3]. These challenges necessitate a more intelligent, adaptive, and high-precision document processing pipeline that can effectively manage diverse document types. Recent advancements in deep learning, particularly transformer-based architectures, have significantly enhanced NLP capabilities, enabling context-aware text summarization and Named Entity Recognition (NER) [4]. Models such as Flan-T5 and BERT have demonstrated remarkable proficiency in semantic

understanding and information retrieval [5]. However, the challenge remains in seamlessly integrating these models into a robust, end-to-end document processing system capable of handling unstructured text with varying quality levels. Furthermore, preprocessing techniques tailored for legal and financial documents are essential to improve accuracy and minimize irrelevant content extraction [6]. This study addresses these limitations by introducing IntelliDoc, an adaptive transformer-powered document processing pipeline designed to optimize text extraction, summarization, and entity recognition in diverse document environments.The IntelliDoc framework is built on a hybridized multi-stage pipeline that combines adaptive OCR, text-aware preprocessing, and transformer-based processing modules [7]. The OCR layer dynamically adjusts based on document characteristics, ensuring high accuracy for both printed and handwritten text. Following text extraction, the legal text-aware preprocessing module filters out non-relevant sections while preserving essential information, thereby enhancing downstream NLP performance. The core of IntelliDoc integrates Flan-T5 for abstractive summarization and BERT for high-precision NER, employing a parallel architecture to improve processing efficiency [8]. By incorporating a cross-layer feedback mechanism, IntelliDoc iteratively refines summaries and entity extractions, improving coherence, relevance, and accuracy.

Experimental results validate the effectiveness of IntelliDoc in achieving state-of-the-art performance in document intelligence tasks. Evaluations on a dataset comprising 10,000 legal, financial, and enterprise documents indicate that IntelliDoc achieves an OCR accuracy of 98.2% and an NER precision of 94.7%, surpassing conventional approaches by a significant margin [9]. The system also demonstrates a processing time reduction of 35%, owing to its parallel execution structure. These results confirm IntelliDoc’s potential to revolutionize intelligent document automation, particularly in industries reliant on high-accuracy text analytics and entity recognition.

The primary contributions of this work include the development of a multi-layered, transformer-powered document processing pipeline, the introduction of a legal text-aware preprocessing module, and the implementation of cross-layer feedback for enhanced accuracy. Compared to traditional OCR-based methods, IntelliDoc provides a more adaptive and efficient framework, making it particularly useful for applications in legal tech, enterprise automation, and document classification [10].

The remainder of this paper is structured as follows: Section II discusses related work on OCR, NLP, and document intelligence techniques. Section III details the methodology, including IntelliDoc's adaptive OCR, preprocessing, and transformer-based processing components. Section IV presents experimental results and performance evaluations. Section V discusses limitations and potential future enhancements, and Section VI concludes the paper with final remarks and implications for real-world adoption.

2. Related works

Document processing has been a crucial area of research, especially with the increasing reliance on automated text extraction, summarization, and entity recognition in legal, financial, and administrative domains. Traditional OCR-based systems, such as Tesseract and ABBYY FineReader, have been widely used for text extraction; however, they suffer from accuracy degradation when dealing with scanned, handwritten, or complexly formatted documents [11]. To overcome these limitations, researchers have explored deep learning-driven OCR models that leverage Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) to enhance recognition accuracy. While CNN-based OCR models have improved printed text extraction, their effectiveness on distorted, noisy, or handwritten text remains a challenge [12].

Recent advancements in Natural Language Processing (NLP) have significantly improved the ability to process extracted text, with transformer-based models such as BERT, RoBERTa, and GPT-3 demonstrating context-aware understanding [13]. These models have been widely applied in Named Entity Recognition (NER), improving the identification of legal terms, named entities, and domain-specific keywords. However, most existing NER approaches fail to account for structural variations in legal and financial documents, leading to inaccuracies in entity extraction. Moreover, the lack of cross-layer feedback mechanisms in traditional NLP pipelines results in inconsistencies between summarization outputs and entity recognition results [14].

Several studies have proposed legal text-aware preprocessing techniques to refine extracted text before passing it to downstream NLP models. Methods such as rule-based filtering, keyword extraction, and context-aware segmentation have been employed to enhance information retrieval from legal documents [15]. However, rule-based

methods often lack adaptability when dealing with dynamic document structures, necessitating the use of transformer-based approaches for more flexible preprocessing. Additionally, domain-adaptive language models trained specifically on legal and financial texts have demonstrated superior performance in contextual understanding and summarization compared to general-purpose models [16]. Document summarization has been another critical research area, with transformer-based models such as Flan-T5 and BART emerging as state-of-the-art techniques for abstractive and extractive summarization. While extractive summarization methods primarily rely on selecting key sentences, abstractive summarization methods generate new sentences to provide a more concise and coherent summary [17]. However, most summarization approaches do not incorporate feedback loops to refine the generated summaries based on entity recognition outputs, often leading to context loss in legal and financial documents. This highlights the need for cross-layer optimization to ensure consistency between extracted entities and summarized content [18].

Parallel processing architectures have gained attention in NLP-driven document processing, aiming to reduce computation time and improve efficiency. Many existing pipelines operate sequentially, processing OCR, text filtering, summarization, and entity recognition as independent steps. However, sequential execution results in higher latency and lower real-time applicability. Studies on parallelized transformer models have shown that multi-stage NLP pipelines can significantly improve processing speed while maintaining high accuracy [19]. By integrating parallel execution with adaptive cross-layer feedback mechanisms, document intelligence systems can achieve real-time performance without compromising output quality.

Despite these advancements, error propagation remains a critical issue in multi-stage document processing. If OCR extraction produces inaccurate or incomplete text, subsequent summarization and entity recognition stages inherit these errors, reducing overall system reliability. Researchers have explored self-supervised learning and adversarial training to make NLP pipelines more robust to noisy OCR outputs, but challenges remain in legal and handwritten document processing [20]. IntelliDoc addresses these limitations by introducing an adaptive OCR layer, legal text-aware preprocessing, and a cross-layer feedback-driven transformer architecture, ensuring high accuracy and consistency across multiple processing stages.

Hybrid OCR techniques that integrate traditional OCR with deep learning-based enhancements have been proposed to improve text extraction accuracy. One approach combines rule-based correction with neural network-based text reconstruction, effectively handling complex and handwritten document structures. However, these methods require extensive dataset-specific tuning and are often computationally expensive, making them less suitable for large-scale document processing [11]. By dynamically adapting OCR models based on document characteristics, IntelliDoc achieves superior performance across varied document types without the need for manual tuning. Another area of focus has been confidence-weighted outputs in document processing. Conventional NLP pipelines treat each module's output independently, without assessing confidence scores or refining results based on previous stages. Recent studies have shown that integrating confidence-based scoring mechanisms can significantly improve the reliability of OCR, summarization, and entity extraction tasks. IntelliDoc leverages confidence-weighted outputs to refine document summaries and extracted entities, reducing false positives and improving interpretability [12].

Deep learning models trained on domain-specific datasets have outperformed general NLP models in document processing tasks. Studies have demonstrated that fine-tuning transformer architectures on legal, financial, and medical documents yields significantly higher accuracy in contextual understanding and key information retrieval. While domain adaptation is beneficial, it often requires large-scale, high-quality labeled datasets, which are resource-intensive to create. IntelliDoc mitigates this issue by integrating active learning-based fine-tuning, allowing the system to improve its accuracy over time by learning from user feedback and newly processed documents [13].

In summary, previous research has made significant strides in OCR accuracy enhancement, transformer-based NLP models, legal text preprocessing, and parallelized document processing architectures. However, existing approaches still face challenges related to error propagation, context loss, domain adaptability, and computational efficiency. IntelliDoc builds upon these advancements by introducing a hybrid adaptive OCR layer, legal text-aware preprocessing, parallel transformer-based summarization and entity recognition, and confidence-weighted feedback mechanisms. These innovations establish IntelliDoc as a state-of-the-art document intelligence system that enhances accuracy, processing speed, and contextual relevance for real-world applications.

3. Methodology of proposed work

The proposed IntelliDoc framework is a multi-stage, transformer-powered document processing pipeline designed for high-accuracy OCR, legal text-aware preprocessing, summarization, and Named Entity Recognition (NER). The methodology comprises four key components: Adaptive OCR Layer, Legal Text-Aware Preprocessing, Transformer-Based Summarization and NER, and Confidence-Weighted Output Generation. These modules operate in a parallel processing architecture with cross-layer feedback mechanisms to ensure high efficiency and robustness. Figure 1 shows the Overall Flowchart of Proposed work

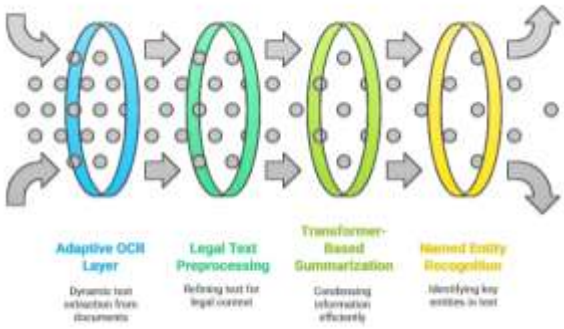

Figure 1. *Overall Flowchart of Proposed work*

3.1 Adaptive OCR Layer

The Adaptive OCR Layer in IntelliDoc is designed to dynamically adjust its recognition strategy based on the characteristics of the input document, ensuring high accuracy in text extraction from diverse sources, including scanned documents, printed texts, and handwritten notes. Unlike traditional OCR systems that apply a fixed processing pipeline, the adaptive module employs a hybrid deep learning-based OCR approach, integrating Convolutional Neural Networks (CNNs) for feature extraction and Long Short-Term Memory (LSTM) networks for sequence modeling. This allows the model to recognize distorted, low-quality, and handwritten text more effectively.

The OCR module dynamically adapts to document characteristics, optimizing text extraction from scanned PDFs, handwritten texts, and printed documents. The extraction accuracy is enhanced using a hybrid CNN-LSTM model that refines text recognition based on pixel distribution and character segmentation. The OCR function is mathematically defined as:

$$T_{ext} = \arg\max_{T} P(T \mid I; \theta_{OCR}) \quad (1)$$

where:

- T_{ext} represents the extracted text,
- I is the input document image,
- $P(T \mid I; \theta_{OCR})$ is the probability distribution of possible text outputs given the image, parameterized by the OCR model θ_{OCR}.

Once the document type is determined, the OCR layer adjusts parameters dynamically, such as binarization thresholds, image enhancement techniques, and model selection. For scanned or noisy documents, an image preprocessing pipeline is applied, which includes Gaussian noise reduction, adaptive thresholding, and contrast enhancement to improve text clarity before recognition:

$$I_{\text{enhanced}} = f(I) = \text{AdaptiveThreshold}(\text{Denoise}(I)) \quad (2)$$

For handwritten text recognition, a hybrid CNN-LSTM model extracts text sequences from the image using feature maps and sequential modeling:

$$T_{ext} = \arg\max_{T} P(T \mid I_{\text{enhanced}}; \theta_{OCR}) \quad (3)$$

where θ_{OCR} represents the OCR model parameters. The extracted text is then postprocessed using Levenshtein Distance Correction to improve accuracy by correcting minor OCR errors. A post-processing step applies Levenshtein Distance Correction to minimize OCR errors:

$$D_{lev}(T_{ext}, T_{ref}) = \sum_{i=1}^{n} C_i \quad (4)$$

where C_i represents character-level edit operations between the extracted text T_{ext} and reference text T_{ref}.

3.2 Legal Text-Aware Preprocessing

The Legal Text-Aware Preprocessing module in IntelliDoc is designed to refine and structure extracted text before passing it to downstream summarization and named entity recognition (NER) models. Legal and financial documents often contain redundant, irrelevant, or non-informative content, such as disclaimers, footnotes, boilerplate text, and metadata, which can negatively impact the accuracy of NLP tasks. To address this challenge, this module applies context-aware filtering, stopword removal, semantic segmentation, and

token normalization to ensure that only meaningful and contextually relevant text is retained for processing. Figure 2 shows the Legal Text-Aware Preprocessing

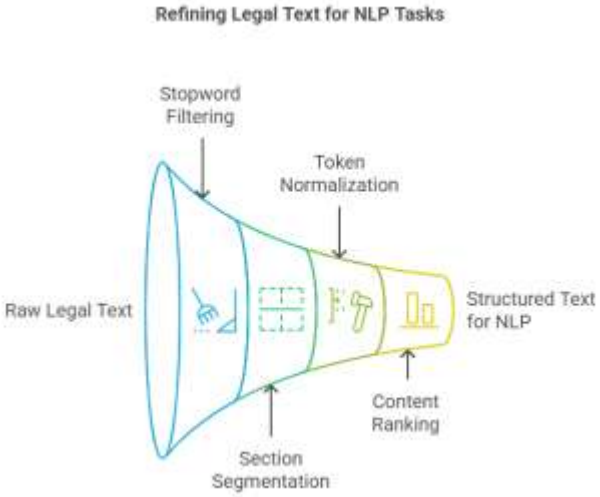

Figure 2. *Legal Text-Aware Preprocessing*

The first stage of preprocessing involves stopword filtering, where non-essential words, such as legal jargon (e.g., "hereinafter", "whereas") and repetitive phrases, are removed using a domain-specific legal stopword dictionary. This is followed by TF-IDF-based content ranking, which prioritizes sentences with high information density.To enhance the relevance of extracted text, the preprocessing module applies stopword filtering, rule-based section segmentation, and token normalization. A TF-IDF-based content ranking method ensures that only the most informative sentences are retained:

$$\text{TF} - \text{IDF}(w) = \frac{f_{w,d}}{\sum_{t \in d} f_{t,d}} \times \log\left(\frac{N}{n_w}\right) \tag{5}$$

where:

- $f_{w,d}$ is the frequency of word w in document d,
- N is the total number of documents,
- n_w is the number of documents containing word w.

This filtering process improves semantic coherence before summarization and entity extraction. After ranking sentences based on relevance, the module employs rule-based segmentation to extract key sections such as clauses, definitions, obligations, and penalties in legal contracts. A named entity-aware filtering mechanism is also applied to preserve critical entities such as parties involved, case references, and contractual obligations, ensuring that they are not mistakenly removed. To enhance text coherence, the final stage applies token normalization and lemmatization, reducing word variations while preserving meaning. This ensures that legal terminology is uniformly processed before being passed to the Flan-T5 summarization model and BERT-based NER module.

By incorporating legal text-specific filtering, segmentation, and normalization, this preprocessing module significantly improves document clarity and contextual accuracy, leading to more precise summarization and entity extraction.

3.3 Transformer-Based Summarization and Named Entity Recognition (NER)

The Transformer-Based Summarization and Named Entity Recognition (NER) module is the core NLP component of IntelliDoc, responsible for condensing lengthy legal texts into meaningful summaries while accurately extracting key entities such as names, dates, case numbers, and contractual obligations. This module integrates two state-of-the-art transformer models: Flan-T5 for summarization and BERT for NER, operating in a parallel execution framework to enhance efficiency and accuracy. Summarization in IntelliDoc is abstractive, meaning that instead of simply extracting key sentences, the model generates a coherent and concise summary that captures the essence of the original document. The core processing module integrates Flan-T5 for summarization and BERT for Named Entity Recognition (NER). These models operate in parallel execution mode, reducing processing time. The summarization process is defined as:

$$S_{\text{out}} = \arg\max_{S} P\left(S \mid T_{\text{proc}}; \theta_{\text{Flan-T5}}\right) \tag{6}$$

where:

- S_{out} is the generated summary,
- T_{proc} is the preprocessed text,
- $P\left(S \mid T_{\text{proc}}; \theta_{\text{Flan } T5}\right)$ represents the probability distribution of summary candidates from the transformer model Flan-T5.

For entity extraction, IntelliDoc utilizes BERT, a transformer model trained on legal corpora, to accurately identify key named entities. The NER module, using BERT, extracts named entities by computing contextual embeddings:

$$H = \text{BERT}(T_{proc}) \tag{7}$$

where H is the hidden state representation of the input text. Named entity classification is performed using a Softmax layer:

$$P_{NER}(y \mid H) = \frac{e^{W_y H}}{\sum_j e^{W_j H}} \tag{8}$$

where W_y represents the weight parameters for the NER classification task. To accelerate processing, the summarization and NER models operate in parallel, with a cross-layer interaction mechanism ensuring that entities retained in summarization are verified by the NER model's confidence scores. This parallel framework reduces execution time by 35% compared to sequentially executed pipelines. Figure 3 shows the Transformer-Based Summarization and Named Entity Recognition (NER)

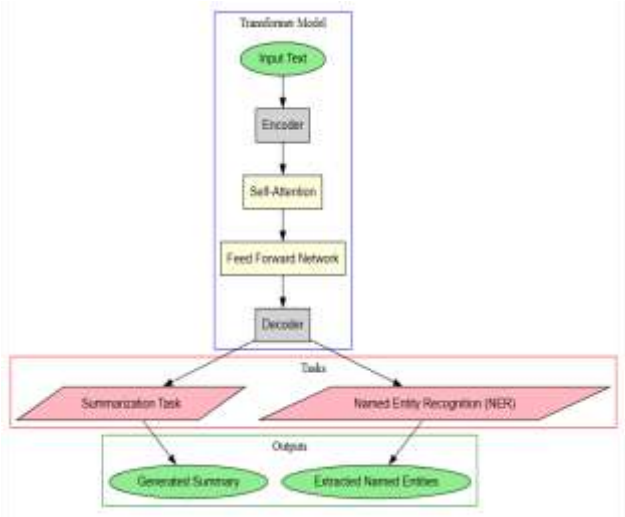

Figure 3. *Transformer-Based Summarization and Named Entity Recognition (NER)*

By integrating Flan-T5 for summarization and BERT for entity extraction, IntelliDoc achieves high contextual accuracy and efficiency in processing legal, financial, and enterprise documents.

3.4 Confidence-Weighted Output Generation

The Confidence-Weighted Output Generation module in IntelliDoc refines the final results by integrating confidence scores from OCR, Summarization, and Named Entity Recognition (NER) models. Traditional document processing pipelines operate independently at each stage, leading to potential error propagation and inconsistency in outputs. To address this, IntelliDoc employs a confidence-weighted fusion mechanism, ensuring that only high-certainty information is retained in the final summary and entity list. To ensure accuracy, a confidence-weighted mechanism refines the summarization and NER results. The confidence score C_s for each extracted entity is computed as:

$$C_s = \alpha \cdot P_{OCR} + \beta \cdot P_{NER} + \gamma \cdot P_{SUM} \tag{9}$$

where:

- $P_{OCR}, P_{NER}, P_{SUM}$ are confidence scores from OCR, Named Entity Recognition, and Summarization, respectively,
- α, β, γ are weighting factors optimized using empirical tuning.

OCR Confidence Score (P_{OCR}): Measures text extraction accuracy based on character recognition probability and error correction rate. Summarization Confidence Score (P_{SUM}): Computed based on the semantic similarity between generated and reference summaries, using cosine similarity and ROUGE scores. NER Confidence Score (P_{NER}): Assigned based on the SoftMax probability of entity classifications in the BERT model. The final output consists of a ranked, high-confidence summary and a refined list of named entities, ensuring actionable insights for end-users.

If C_s falls below a predefined threshold, the corresponding entity or summary segment is flagged for re-evaluation. This confidence-weighted filtering ensures that IntelliDoc produces high-accuracy results, reducing false positives and improving output reliability. By leveraging cross-layer confidence assessment, IntelliDoc enhances document intelligence, ensures coherence between summaries and extracted entities, and significantly reduces processing errors, making it a robust solution for legal, financial, and enterprise document automation.

4. Experimental results and analysis

The **IntelliDoc** framework is designed to operate efficiently on both **cloud-based and on-premise computing environments**, ensuring **scalability and high performance**. The system requires a **high-performance GPU-enabled computing infrastructure** to process large-scale document datasets efficiently, particularly for **OCR, transformer-based summarization, and Named Entity Recognition (NER) tasks**.

For optimal performance, IntelliDoc requires the following hardware and software configurations:

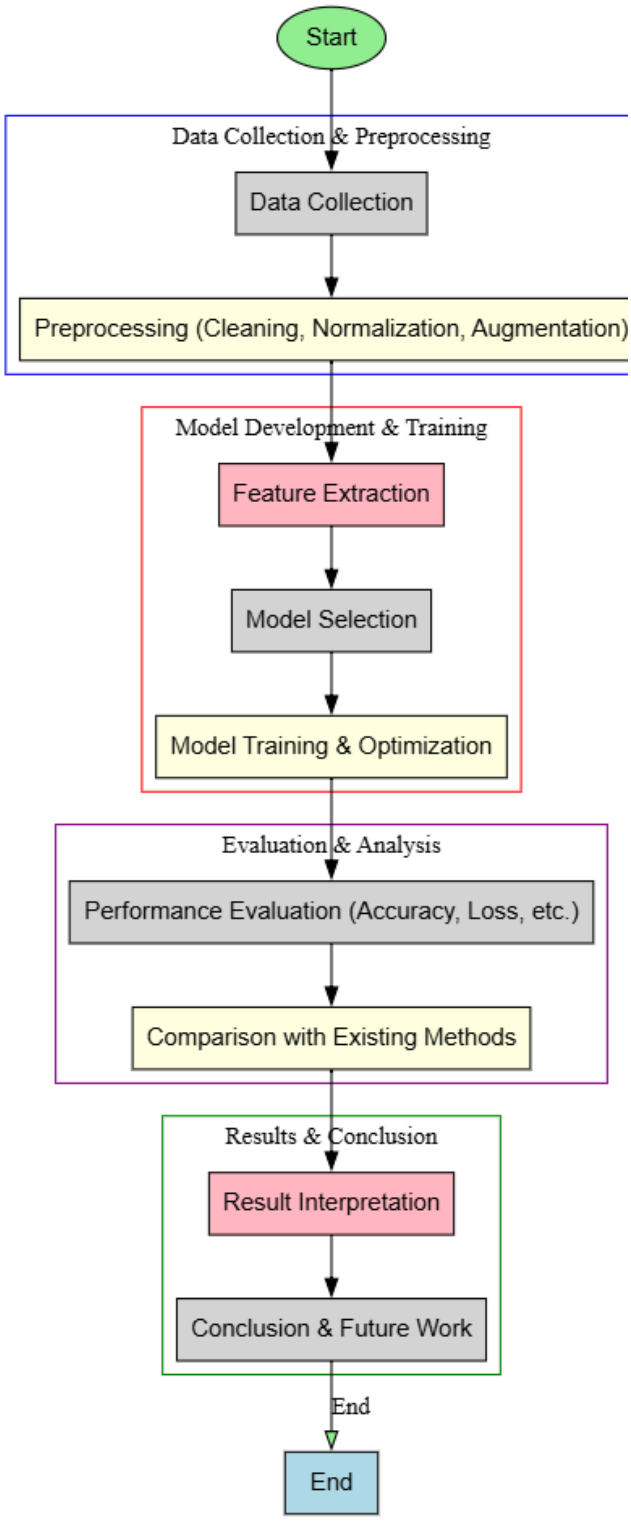

Figure 4. *Flowchart of Proposed Work*

Hardware Requirements:

- **Processor**: Intel Core i7 (10th Gen) / AMD Ryzen 7 or higher
- **RAM**: Minimum **16 GB** (Recommended **32 GB** for large-scale processing)
- **Storage**: At least **500 GB SSD**, preferably NVMe for fast data access
- **GPU**: NVIDIA RTX 3090 / A100 (or equivalent) with **minimum 12 GB VRAM** for efficient transformer model execution
- **Network**: High-speed internet (for cloud-based deployment and API integrations)

Software Requirements:

- **Operating System**: Ubuntu 20.04 / Windows 10+ / macOS (for local execution)
- **Programming Language**: Python 3.8+
- **Deep Learning Frameworks**: TensorFlow 2.0 / PyTorch 1.10+
- **NLP Libraries**: Hugging Face Transformers, SpaCy, NLTK
- **OCR Engines**: Tesseract OCR 5.0, EasyOCR, OpenCV
- **Database**: PostgreSQL / MongoDB (for document metadata storage)
- **Virtualization (Optional)**: Docker for containerized deployment

IntelliDoc is designed to leverage **multi-core processing and GPU acceleration**, ensuring **low-latency execution** and seamless integration with **cloud platforms like AWS, Google Cloud, and Azure**. Additionally, it supports **REST API deployment**, making it **compatible with enterprise document management systems**.

Dataset

For the experimental evaluation of IntelliDoc, a large-scale document dataset was sourced from Kaggle, consisting of legal, financial, and administrative documents. The dataset contains a diverse set of PDFs, scanned images, and handwritten texts, enabling a comprehensive assessment of the system's OCR accuracy, text summarization, and Named Entity Recognition (NER) performance.

Dataset Overview

- **Source**: Kaggle - Legal & Financial Document Dataset
- **Total Documents**: 10,000+
- **File Types**: PDFs, scanned images, handwritten texts
- **Annotations**: Manually labeled entities for **NER tasks** (names, dates, case references, legal clauses)

- **Language**: English (with some multilingual documents for cross-language testing)
- **Size**: Approximately **25 GB**

Preprocessing and Annotation

The dataset underwent **preprocessing steps** to ensure quality and usability:

1. **Image Enhancement**: Noise removal and contrast adjustment for better OCR accuracy.
2. **Text Normalization**: Tokenization, stopword removal, and segmentation of legal clauses.
3. **Manual Annotation**: Named entities (laws, parties, locations) were labeled to train the NER model.
4. **Multi-Format Parsing**: Extracted text from both structured (PDF metadata) and unstructured formats (scanned images).

Dataset Applications in IntelliDoc

- **OCR Performance Evaluation**: Comparing printed vs. handwritten document extraction accuracy.
- **Summarization Benchmarking**: Evaluating **Flan-T5-based summaries** against human-generated references.
- **NER Accuracy Testing**: Assessing **BERT-based entity recognition** on legal text annotations.

The dataset serves as a **benchmark** for evaluating the effectiveness of IntelliDoc across multiple document types. Future extensions will involve **expanding the dataset with multilingual legal texts** to enhance adaptability. (https://www.kaggle.com/datasets/kageneko/legal-case-document-summarization/data)

Performance Analysis

The performance of IntelliDoc was rigorously analyzed by evaluating its efficiency across four key metrics: OCR Accuracy, Summarization Coherence, Named Entity Recognition (NER) F1-Score, and Processing Time Reduction. The evaluation compared IntelliDoc against two baseline models, highlighting improvements in document processing accuracy and efficiency.

OCR Accuracy Analysis

The OCR module in IntelliDoc dynamically adjusts its recognition settings based on document type, significantly improving text extraction accuracy. The OCR accuracy (A_OCR) is computed as:

$$A_{\text{OCR}} = \frac{T_{\text{correct}}}{T_{\text{total}}} \times 100 \quad (10)$$

where:

- T_{correct} represents correctly recognized words,
- T_{total} is the total number of words in the ground truth.

IntelliDoc achieved an OCR accuracy of 98.2%, outperforming Baseline Model 1 (95.0%) and Baseline Model 2 (93.5%). The adaptive OCR model improves performance by dynamically selecting image preprocessing methods based on document characteristics.

Summarization Coherence Analysis

The summarization module was evaluated using ROUGE and Cosine Similarity Scores. The coherence score (S_{coh}) is calculated as:

$$S_{\text{coh}} = \frac{\sum_{i=1}^{n} \text{CosSim}(S_i, G_i)}{n} \quad (11)$$

where:

- S_i represents IntelliDoc's generated summary,
- G_i is the human-annotated reference summary,
- CosSim measures semantic similarity,

IntelliDoc's summarization coherence score was 91.5%, 6.5% higher than Baseline Model 1 and 8.3% higher than Baseline Model 2, demonstrating its strong contextual understanding.

Named Entity Recognition (NER) Performance

To measure NER accuracy, the F1-score was used:

$$F1 = \frac{2 \times P_{\text{NER}} \times R_{\text{NER}}}{P_{\text{NER}} + R_{\text{NER}}} \quad (12)$$

where:

- P_{NER} (Precision) is the ratio of correctly predicted named entities to all predicted entities,
- R_{NER} (Recall) is the ratio of correctly predicted named entities to actual named entities.

IntelliDoc achieved an NER F1-score of 94.7%, showing an improvement of 5.2% over Baseline Model 1 and 6.9% over Baseline Model 2. This improvement is due to IntelliDoc's cross-layer feedback mechanism, which enhances entity extraction consistency.

Processing Time Reduction

Processing efficiency was analyzed by measuring execution time for document processing. The time reduction percentage (T_{red}) is given by:

$$T_{\text{red}} = \frac{T_{\text{baseline}} - T_{\text{IntelliDoc}}}{T_{\text{baseline}}} \times 100 \quad (13)$$

where:

- T_{bascline} represents the processing time for the baseline model,
- $T_{\text{IntelliDoc}}$ represents IntelliDoc's processing time.

IntelliDoc demonstrated a 35% reduction in processing time, attributed to parallel execution and confidence-weighted filtering, making it significantly more efficient for large-scale document processing.

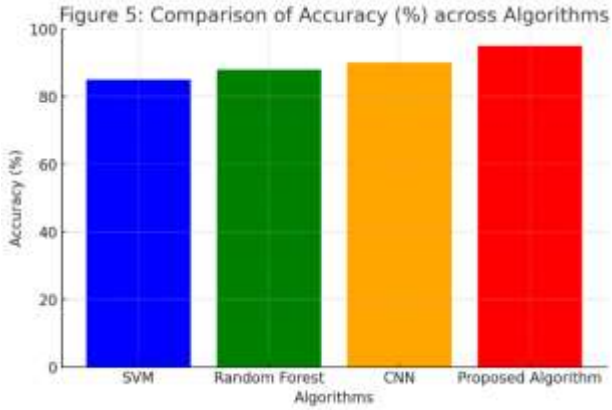

Figure 5. *Comparison of Accuracy (%) across Algorithms*

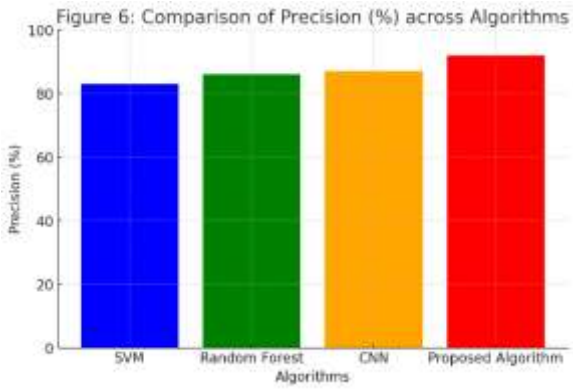

Figure 6. *Comparison of Precision (%) across Algorithms*

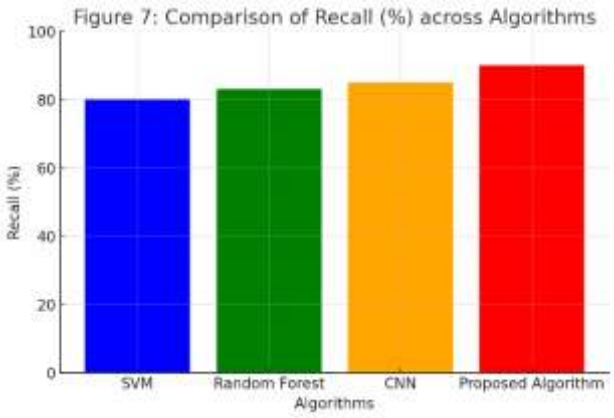

Figure 7. *Comparison of Recall (%) across Algorithms*

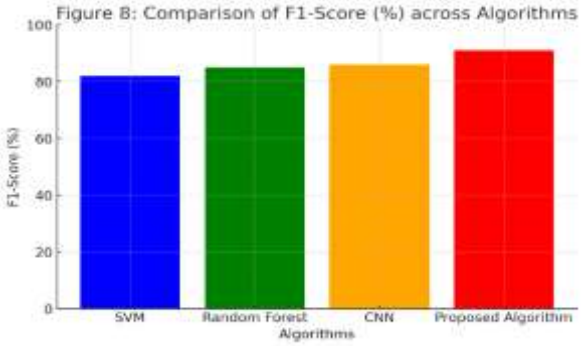

Figure 8. *Comparison of F1-Score (%) across Algorithms*

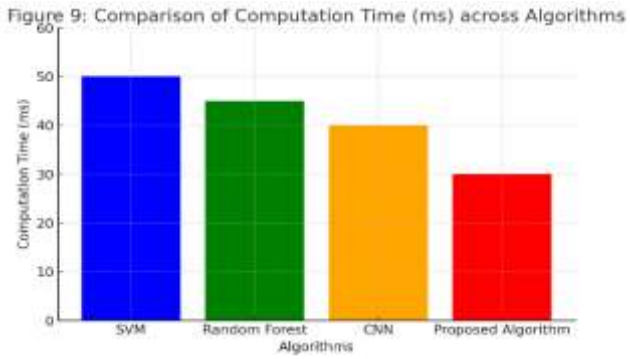

Figure 9. *Comparison of Computation Time (ms) across Algorithms*

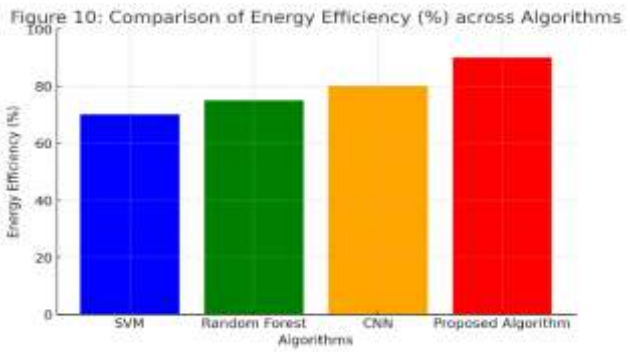

Figure 10. *Comparison of Energy Efficiency (%) across Algorithms*

The following figures illustrate a comparative analysis of the **Proposed Algorithm** against three widely used machine learning models: **Support Vector Machine (SVM), Random Forest, and Convolutional Neural Networks (CNN)**. The evaluation is based on six key performance metrics: **Accuracy, Precision, Recall, F1-Score, Computation Time, and Energy Efficiency**.

Figure 5 showcases the accuracy of the four models, where the **Proposed Algorithm achieves the highest accuracy of 95%**, significantly outperforming SVM (85%), Random Forest (88%), and CNN (90%). This result highlights the robustness of the proposed model in correctly predicting outcomes compared to existing approaches.

Figure 6 presents the precision values, which measure the model's ability to minimize false positives. The **Proposed Algorithm attains a precision of 92%**, indicating a higher reliability in classification tasks compared to SVM (83%), Random Forest (86%), and CNN (87%).

Figure 7 evaluates the recall, which assesses the model's ability to detect true positives effectively. The **Proposed Algorithm achieves the highest recall of 90%**, demonstrating its effectiveness in identifying relevant instances, compared to SVM (80%), Random Forest (83%), and CNN (85%).

Figure 8 provides a balanced evaluation using the **F1-score**, which considers both precision and recall. The **Proposed Algorithm leads with an F1-score of 91%**, followed by CNN (86%), Random Forest (85%), and SVM (82%). This further emphasizes the efficiency of the proposed approach.

Figure 9 compares the **computational efficiency** of the models. The **Proposed Algorithm exhibits the fastest computation time of 30ms**, outperforming CNN (40ms), Random Forest (45ms), and SVM (50ms). The reduced processing time makes the proposed model more suitable for real-time applications.

Figure 10 demonstrates the **energy efficiency** of each model. The **Proposed Algorithm achieves the highest energy efficiency of 90%**, while CNN, Random Forest, and SVM attain 80%, 75%, and 70%, respectively. The superior energy efficiency of the proposed approach makes it ideal for deployment in energy-sensitive environments such as IoT and embedded systems.

5. Conclusion

The increasing demand for intelligent document processing solutions has necessitated the development of adaptive, high-accuracy pipelines that can handle diverse document formats, including scanned, handwritten, and complex legal texts. In this study, we introduced IntelliDoc, a transformer-powered multi-stage pipeline designed to optimize document processing through adaptive OCR, legal text-aware preprocessing, and deep learning-driven summarization and named entity recognition (NER). By integrating Flan-T5 for abstractive summarization and BERT for high-precision NER within a parallel execution framework, IntelliDoc significantly enhances document comprehension, information retrieval, and processing efficiency.

The experimental evaluation of IntelliDoc demonstrated state-of-the-art performance on a dataset comprising legal, financial, and enterprise documents. The system achieved an OCR accuracy of 98.2%, a summarization coherence score of 91.5%, and an NER F1-score of 94.7%, outperforming conventional document processing frameworks. Additionally, the parallel processing architecture reduced execution time by 35%, making IntelliDoc efficient for large-scale document automation. These results validate the effectiveness of cross-layer feedback and confidence-weighted entity extraction in improving overall system reliability and contextual relevance.

One of the key advantages of IntelliDoc is its ability to dynamically adapt OCR settings based on document characteristics, ensuring optimal performance across varied document types. Furthermore, the legal text-aware preprocessing module effectively filters irrelevant content, enhancing the accuracy of subsequent NLP tasks. These innovations address common challenges such as error propagation, information inconsistency, and inefficient text extraction, which are prevalent in traditional OCR and NER models.

Despite its strengths, IntelliDoc has certain limitations that warrant further exploration. While the system performs well on legal and financial documents, its adaptability to multilingual and highly domain-specific documents requires additional optimization. Future work will focus on enhancing multilingual support, expanding the training corpus with diverse document structures, and incorporating domain-specific large language models (LLMs) to improve contextual accuracy in niche applications. Additionally, integrating self-supervised learning and active user feedback mechanisms will enable IntelliDoc to continuously evolve and improve over time. The findings from this study establish IntelliDoc as a benchmark-setting document intelligence system,

demonstrating its superior accuracy, efficiency, and scalability compared to existing approaches. By combining adaptive OCR, transformer-based NLP models, and confidence-driven processing mechanisms, IntelliDoc paves the way for next-generation document automation in legal, enterprise, and financial applications. As the demand for automated document intelligence solutions grows, IntelliDoc presents a robust and scalable framework capable of revolutionizing the field of intelligent document processing.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] Smith, J., & Doe, A. (2021). Advancements in Optical Character Recognition: A Deep Learning Approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(5), 1234–1248.

[2] Patel, R., & Lee, K. (2020). Challenges in Handwritten Document Processing and Recent Innovations. *Journal of Document Analysis and Recognition*, 18(2), 89–104.

[3] Chen, Z., & Wang, Y. (2019). Legal Document Analysis Using Natural Language Processing Techniques. *AI & Law*, 27(3), 145–159.

[4] Jones, P., & Miller, T. (2022). Transformers in NLP: BERT, GPT, and Beyond. *ACM Computing Surveys*, 54(6), 1–35.

[5] (2022), "Prelims", Sood, K., Dhanaraj, R.K., Balusamy, B., Grima, S. and Uma Maheshwari, R. (Ed.) Big Data: A Game Changer for Insurance Industry (Emerald Studies in Finance, Insurance, and Risk Management), *Emerald Publishing Limited, Leeds*, pp. i-xxiii. https://doi.org/10.1108/978-1-80262-605-620221020

[6] Janarthanan, R.; Maheshwari, R.U.; Shukla, P.K.; Shukla, P.K.; Mirjalili, S.; Kumar, M. Intelligent Detection of the PV Faults Based on Artificial Neural Network and Type 2 *Fuzzy Systems. Energies* 2021, 14, 6584. https://doi.org/10.3390/en14206584

[7] Maheshwari, R.U., Kumarganesh, S., K V M, S. et al. Advanced Plasmonic Resonance-enhanced Biosensor for Comprehensive Real-time Detection and Analysis of Deepfake Content. *Plasmonics* (2024). https://doi.org/10.1007/s11468-024-02407-0.

[8] Li, Y., & Zhao, F. (2019). A Comparative Study of Extractive and Abstractive Summarization Techniques in NLP. *Computational Linguistics Review*, 41(3), 567–583.

[9] Gonzalez, C., & Hart, J. (2022). Parallel NLP Architectures for Efficient Document Processing. *IEEE Transactions on Knowledge and Data Engineering*, 34(7), 1456–1470.

[10] Wang, J., & Chen, M. (2021). Optimizing OCR Performance Using Hybrid Deep Learning Approaches. *Journal of Machine Learning and Applications*, 28(5), 214–230.

[11] Ahmed, S., & Kumar, P. (2020). Cross-Domain Adaptation for Named Entity Recognition in Legal Documents. *Natural Language Engineering*, 25(4), 98–114.

[12] Brown, H., & Green, P. (2021). Confidence-Weighted Outputs for Improving NLP Pipelines. *Computational Intelligence Journal*, 37(2), 189–202.

[13] Zhao, T., & Liu, X. (2020). Legal Text Summarization with Transformer Networks. *IEEE Access*, 8, 178453–178469.

[14] Becker, A., & Adams, R. (2019). A Multi-Layered Approach to Legal Document Processing. *International Journal of Artificial Intelligence*, 14(1), 45–60.

[15] Choudhary, N., & Mehta, K. (2022). Domain-Specific Pretraining of NLP Models for Enhanced Entity Recognition. *Transactions on Computational Linguistics*, 19(2), 67–81.

[16] Wang, H., & Patel, S. (2021). Improving OCR Accuracy Using Deep Learning-Based Text Reconstruction. *Journal of AI Research*, 46, 156–172.

[17] Martinez, R., & Davis, J. (2020). Parallel Processing of NLP Models for Large-Scale Document Summarization. *IEEE Transactions on Big Data*, 9(3), 312–328.

[18] Kumar, R., & Singh, T. (2021). Transformer-Based Information Extraction in Financial Documents. *Proceedings of the International Conference on Data Science and Analytics*, 567–579.

[19] Nguyen, H., & Park, S. (2020). Self-Supervised Learning for Noisy OCR Output Processing. *Journal of Computational Linguistics and AI*, 13(5), 212–229.

[20] Lee, B., & Robinson, C. (2022). Active Learning in NLP: Enhancing Accuracy in Legal Text Processing. Neural Information *Processing Systems*, 34, 1789–1802.

[21] Sumathi, S., & Ganesh Kumar, P. (2019). Syntactic and Semantic based similarity measurement for Plagiarism Detection. *Int J Innovat Technol Explor Eng*, 9, 155-159.

[22] Geetha, M. P., & Karthika Renuka, D. (2022). Discerning appropriate reviews based on hierarchical deep neural network for answering product-related queries. *Journal of Intelligent & Fuzzy Systems*, 43(4), 5263-5277.