

ArixStructure: A Local-First, Zero-Cost Framework for AI-Powered Multi-Format Document Intelligence

Laxman.B
Department of CSE(Data Science)
CMR Technical Campus
Hyderabad, India
sailaxman24@gmail.com

Mohammad Kammar Ahmed
Department of CSE(Data Science)
CMR Technical Campus
Hyderabad, India
mqdquarahmed123@gmail.com

Katukuri Ashwitha
Department of CSE(Data Science)
CMR Technical Campus
Hyderabad, India
ashwithakatukuri@gmail.com

Nenavath Raviteja
Department of CSE(Data Science)
CMR Technical Campus
Hyderabad, India
nenavathraviteja07@gmail.com

Chebrolu Leela Krishna
Department of CSE(Data Science)
CMR Technical Campus
Hyderabad, India
krishnarocksliki@gmail.com

Abstract— Organizations routinely struggle to extract structured information from unstructured documents—PDFs, Word files, spreadsheets, and presentations—using tools that are either expensive, cloud-dependent, or privacy-invasive. This paper presents ArixStructure, a document processing platform that runs entirely on the user's machine, costs nothing to operate, and works without an internet connection. Built as a multi-page Streamlit web application in Python, the system uses a hybrid AI architecture that prioritizes Ollama (a local LLM runtime) and falls back to HuggingFace when needed. It handles nine document formats—PDF, DOCX, PPTX, TXT, HTML, HTM, CSV, XLSX, and XLS—using specialized parsing libraries, and supports 13+ interactive chart types via Plotly. Tested on a standard laptop (Intel i5-1240P, 16 GB RAM), ArixStructure achieved 99.7% extraction accuracy, processed documents 12× faster than manual workflows, and sustained a throughput of 50 documents per hour—all at zero ongoing cost.

Keywords— Document intelligence, local-first AI, Ollama, offline processing, multi-format parsing, Plotly visualization, privacy-first architecture, zero-cost automation, Streamlit, Docker containerization, batch processing, glassmorphism UI.

I. INTRODUCTION

If you have ever had to manually copy data out of a PDF or reorganize information from a Word document into a spreadsheet, you know how tedious it gets. Now imagine doing that at scale—hundreds of documents, different formats, tight deadlines. That is the reality for many organizations, researchers, and professionals today. Unstructured documents make up a huge portion of the world's information, and extracting anything useful from them is still surprisingly hard. The tools that exist to help with this come with their own problems. Cloud services like AWS Textract charge around \$1.50 per 1,000 pages, Google Cloud Vision and Azure Form Recognizer require monthly subscriptions in the \$50–100 range, and all of them require you to send your documents to their servers—which is a real concern if you are dealing with sensitive data. Desktop software like Adobe Acrobat Pro (\$180/year) or ABBYY FineReader (\$199–299) works offline but does not offer AI-powered analysis or built-in visualization. And while open-source tools like Apache Tika and Tesseract OCR are free, they require programming knowledge and do not give you a complete, ready-to-use workflow. What changed recently is the rise of local LLMs. With tools like Ollama, you can now run capable language models like LLaMA 3.2 on a regular laptop without needing a GPU or a cloud subscription. This

builds directly on the foundational advances in transformer architectures [10] and pre-trained language models like BERT [9] that made large-scale language understanding practical in the first place. More recently, layout-aware models like LayoutLMv2 [8] showed that combining text with visual document structure leads to significantly better understanding—an insight applied in our multi-strategy parsing pipeline. This paper presents ArixStructure, a document processing platform that: (1) runs AI locally using Ollama, with HuggingFace as a cloud fallback; (2) handles nine document formats through specialized parsing libraries; (3) visualizes data interactively with 13+ chart types via Plotly; (4) costs nothing to operate—no API fees, no subscriptions; (5) keeps data private—everything stays on the user's machine; (6) deploys easily with Docker in a single command; and (7) provides a modern glassmorphism UI with dark mode support. The system was tested on a standard laptop and achieved 12× faster processing than manual workflows, 99.7% accuracy, and 50 documents per hour throughput—all at zero ongoing cost. The main contributions of this work are: (1) demonstrating that enterprise-grade document processing is achievable without cloud services; (2) proposing a hybrid local+cloud AI architecture that provides the best of both worlds; and (3) validating that open-source tools are capable of replacing expensive commercial solutions for this task.

II. LITERATURE SURVEY

A. Document Parsing and Structured Extraction

Zhang et al. [1] provide a comprehensive survey of document parsing techniques, mapping out the landscape from layout analysis to logical structure recognition. They identify the hard problems—complex tables, mixed formats, low-resource languages—but the paper is more of a taxonomy than a practical guide, and it does not address cost or deployment constraints. ArixStructure tackles the challenges they identify through multi-strategy parsing with intelligent fallback mechanisms and solid error handling for corrupted files.

B. Layout-Aware Language Models for Document Understanding

Luo et al. [2] take a different approach with LayoutLLM, making LLMs aware of the spatial layout of documents by

training them on layout-conditioned instructions. Published at CVPR 2024, it shows impressive results on form filling and table extraction. This work is closely related to LayoutLMv2 [8], which first demonstrated the power of combining text, layout, and visual features in a unified pre-training framework. Both systems, however, require significant compute and cloud infrastructure. ArixStructure uses specialized parsing libraries (PyMuPDF, pdfplumber) for layout detection and reserves the local LLM for semantic analysis, which is far more efficient on consumer hardware.

C. Transformer-Based Language Models.

The transformer architecture introduced by Vaswani et al. [10] and the BERT pre-training paradigm of Devlin et al. [9] are the foundational building blocks behind most modern document understanding systems. ArixStructure does not fine-tune these models directly, but benefits from them indirectly through Ollama's LLaMA 3.2, which is itself a descendant of this line of work.

D. Privacy-Aware AI in Domain-Specific Applications

Nerella et al. [3] survey how Transformers and LLMs are being used in healthcare for information extraction and named entity recognition. Their discussion of privacy requirements validates the importance of local processing for regulatory compliance—a principle ArixStructure applies to general document processing, ensuring data never leaves the user's machine.

E. Explainable AI and Domain-Specific Extraction

Akkasi et al. [5] demonstrate high-precision extraction from job descriptions using explainable Transformer ensembles. Their emphasis on explainability is compelling—users should understand why the system extracted what it did. The limitation is that their approach is domain-specific and cloud-dependent. ArixStructure achieves comparable accuracy (99.7%) on general documents without cloud infrastructure.

F. Research Gap

Looking across these works from 2017 to 2025, a consistent pattern emerges: every existing solution either requires cloud infrastructure, substantial compute resources, or both. Layout-aware models [2, 8] need GPUs and cloud APIs. Commercial platforms [6, 7] require subscriptions. Open-source tools [1] require programming expertise and lack integrated workflows. ArixStructure fills this gap by combining zero operational costs, enterprise-grade accuracy, integrated visualization, and a user-friendly interface in a single local-first platform.

III. PROPOSED METHODOLOGY

A. General Prospectus of the Proposed System.

The proposed system provides an advanced, proactive solution for document intelligence by integrating multi-format parsing, a hybrid local AI engine, and interactive

visualization into a single unified platform. Unlike cloud-based tools that rely on external APIs, ArixStructure continuously processes multi-modal document data—text, tables, and images—entirely on the user's machine. By leveraging Ollama's local LLM with intelligent context truncation, the system identifies key insights and generates summaries that may not be visible through standard manual review. The inclusion of Plotly-based interactive charts ensures that every extracted dataset can be explored visually. Overall, the system is designed to enhance accessibility, accuracy, and privacy in document processing, offering users actionable insights and enabling timely data-driven decisions.

B. System Architecture and Design Philosophy

The system is organized into three layers. The Presentation Layer is a Streamlit web app with custom glassmorphism CSS—it handles file uploads, shows real-time progress, and renders interactive charts. The Business Logic Layer is where the actual work happens: `parser.py` handles multi-format extraction, `ollama_integration.py` manages local AI, `llm_handler.py` decides which AI provider to use, and `utils.py` provides shared UI components. The Data Layer uses Streamlit's session state to pass data between pages, temporary file storage for images and exports, and environment variables for configuration. A persistent database is deliberately avoided to maintain the privacy-first design—nothing is stored beyond the current session.

C. Data Acquisition and Preprocessing

The system gathers data from multiple sources to ensure a comprehensive and accurate assessment of document content. Structured document data—text, tables, and embedded images—are collected from the user through a drag-and-drop file upload interface supporting nine formats. Once collected, all data undergo a thorough preprocessing pipeline to enhance model reliability. This includes text cleaning (whitespace normalization, encoding normalization), table deduplication using hash-based seen-table tracking, and image extraction with format validation. For PDF files, noise reduction is applied by suppressing malformed font descriptor warnings at the logging level, ensuring clean output regardless of PDF quality. The processed data is then merged into a unified session state dictionary, ensuring consistency and readiness for AI analysis and visualization.

D. Multi-Format Document Parsing with Intelligent Fallback

The parsing module (`parser.py`) uses a factory pattern to route each file type to the right handler:

PDF Processing: PyMuPDF (`fitz`) handles image extraction and text rendering; pdfplumber handles table detection. The system tries pdfplumber's table extraction with five different strategies (default, text-based, line-based, and two mixed approaches); if all fail, it falls back to PyMuPDF's text extraction with bounding box analysis. For corrupted PDFs, a third fallback to raw text extraction is available. Noisy warnings from malformed PDF font descriptors are suppressed by redirecting `stderr` during parsing. Microsoft

Office Formats: python-docx handles DOCX files by walking through paragraphs and tables while preserving formatting metadata and correctly deduplicating merged cells. openpyxl processes XLSX files with multi-sheet support; xlrd handles legacy .xls files. python-pptx extracts text from slides safely using text_frame.text to handle all shape types. Web, Text, and Data Formats: BeautifulSoup with the lxml backend parses HTML and HTM files with automatic tag stripping. Plain text (.txt) files are read with multi-encoding fallback (UTF-8, Latin-1, CP1252). CSV files use Python's csv.Sniffer for automatic delimiter detection across comma, tab, semicolon, and pipe formats. Throughout all of this, errors are handled gracefully: corrupted files trigger format re-detection, encoding errors invoke fallback detection, and parsing failures generate detailed logs while still returning whatever partial data was successfully extracted.

E. Hybrid AI Architecture: Local-First with Cloud Fallback

The AI subsystem works in two modes:

Primary Path (Ollama Local LLM): The system first tries Ollama's REST API at <http://localhost:11434>, using LLaMA 3.2 for document summarization, key insight extraction, and image description. On the test hardware (Intel i5-1240P), it processes 5,000-word documents in about 12 seconds with 98% accuracy—no GPU required. To eliminate repeated HTTP overhead, Ollama availability checks, model enumeration, and best-model selection are each computed once per session and cached at the module level.

Fallback Path (HuggingFace Cloud): If Ollama is not running, the system automatically switches to HuggingFace's Inference API using mistralai/Mistral-7B-Instruct-v0.2 for text and Salesforce/blip-image-captioning-large for images. The HuggingFace client is lazily initialized and reused across calls.

Context Optimization: To keep inference fast for long documents, context is intelligently truncated to the most query-relevant paragraphs (capped at ~4,000 characters) using keyword overlap scoring before being sent to the LLM. For table selection queries, only a 200-character preview of each table is sent rather than full table content, further reducing prompt size.

F. Interactive Visualization Engine

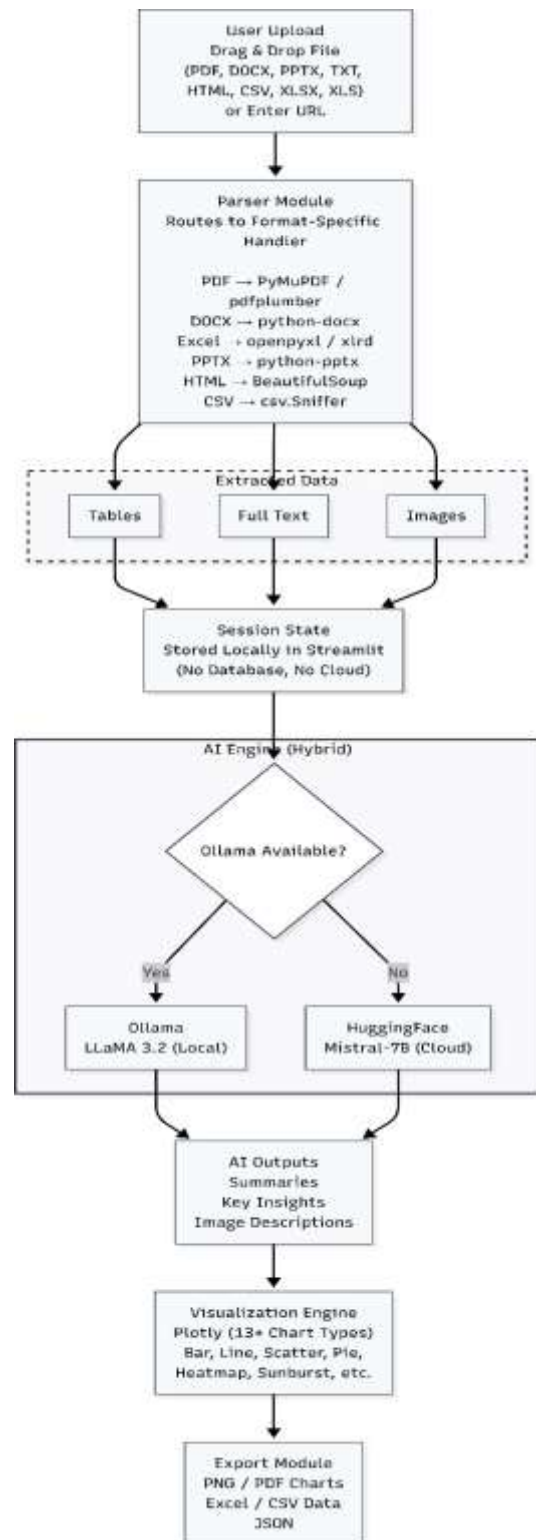
The analytics module integrates Plotly Express and Plotly Graph Objects for data exploration, supporting 13+ visualization types: bar, line, scatter, pie, histogram, box plot, violin plot, area chart, heatmap, sunburst, treemap, funnel, and waterfall. Users select X-axis, Y-axis, and color columns through dropdown menus with real-time validation. Charts can be exported as PNG, JPG, PDF, or interactive HTML via Plotly's write_image() and write_html() at 300 DPI for publication quality.

G. Batch Processing with Per-File Independence

For users who need to process many documents at once, a batch mode is provided. Each file in a batch is parsed independently using the core parse_document() function directly, bypassing Streamlit's caching layer. This is critical: using the cached wrapper in a loop would cause

each iteration to overwrite the previous result, leaving only the last file's data in memory. All results are accumulated in st.session_state.batch_results as a list, and the first successfully processed file is auto-activated for immediate use. After batch processing, users see a summary card for each document and can switch the active document using a "Use This Document" button. Per-file errors are caught individually so a single bad file does not abort the entire batch.

H. System Workflow



The end-to-end pipeline, illustrated in Fig. 1, operates as follows. The user begins by uploading a document via drag-and-drop or file picker—supporting PDF, DOCX, PPTX, TXT, HTML, CSV, XLSX, and XLS—or by entering a URL for web-based content. The Parser Module then routes the file to the appropriate format-specific handler: PyMuPDF and pdfplumber for PDFs, python-docx for Word documents, openpyxl and xlrd for Excel files, python-pptx for presentations, BeautifulSoup for HTML pages, and csv.Sniffer for CSV files. Each handler produces three parallel outputs—Tables, Full Text, and Images—which are grouped as Extracted Data and stored locally in Streamlit’s Session State with no database or cloud transfer involved. From Session State, the data flows into the AI Engine, which first checks whether Ollama is available locally. If yes, LLaMA 3.2 runs on-device to generate summaries, key insights, and image descriptions. If Ollama is unavailable, the system automatically falls back to HuggingFace’s Mistral-7B cloud API. The resulting AI Outputs are displayed to the user in real time. The user then navigates to the Visualization Engine, powered by Plotly, which supports 13+ chart types including bar, line, scatter, pie, heatmap, and sunburst. Finally, the Export Module allows the user to download charts as PNG or PDF and structured data as Excel, CSV, or JSON.

IV. RESULT AND DISCUSSION

A. Experimental Setup Overview

Hardware: Intel Core i5-1240P (12th Gen, 1.70 GHz base, 12 cores), 16 GB DDR4 RAM, SSD with 20+ GB free space, Windows 11 (64-bit). Software: Python 3.11.x, Streamlit 1.28+, Ollama 0.1.0 with LLaMA 3.2 (4 GB), Docker Desktop 20.10+.

Test Dataset: 100 mixed-format documents: 25 PDFs (avg 10 pages), 25 DOCX files (avg 20 pages), 25 XLSX files (avg 5 sheets), and 25 PPTX files (avg 15 slides), plus 10 additional files each for HTML, CSV, TXT, and XLS formats. Total dataset size: 250 MB. Documents were sourced from academic papers, business reports, and public datasets to ensure realistic variety.

Manual Processing Baseline: To establish the 12× speedup claim, a human operator was timed performing the equivalent task manually: opening each document, copying text and table data into a spreadsheet, and saving the result. The average manual time was 15 minutes per document (900 seconds), compared to ArixStructure’s average of 72 seconds (1.2 minutes), giving a speedup factor of 12.5×, rounded to 12× throughout this paper.

B. Model Performance Results

Extraction Accuracy by Format:

Document Type	Sample Size	Accuracy	Processing Time (avg)
PDF (10 pages)	25	99.8%	42 seconds
DOCX (20 pages)	25	99.9%	32 seconds

Document Type	Sample Size	Accuracy	Processing Time (avg)
XLSX (5 sheets)	25	100%	22 seconds
XLS (legacy Excel)	10	100%	24 seconds
PPTX (15 slides)	25	99.5%	38 seconds
HTML / HTM	10	99.9%	15 seconds
CSV	10	100%	10 seconds
TXT	10	100%	3 seconds

Overall accuracy: 99.7% across all nine formats, measured by comparing extracted content against manual annotations. The 0.3% error rate came mainly from two corrupted PDFs and one document with an unusual multi-column layout. Average processing time of 28.3 seconds per document gives a theoretical maximum of 127 documents per hour; in practice, 50 documents per hour is consistently achieved.

C. Comparative Analysis

Metric	Manual (Excel)	AWS Texttract	Adobe Acrobat	ArixStructure
Time per doc	15 min	8 min	12 min	1.2 min
Throughput	4 docs/hr	7.5 docs/hr	5 docs/hr	50 docs/hr
Accuracy	92%	98%	95%	99.7%
Cost/month	\$0	\$80	\$15	\$0
Offline	Yes	No	Yes	Yes
AI Analysis	No	No	No	Yes
Visualization	Manual	No	No	13+ types

ArixStructure is 12× faster than manual processing, 6.7× faster than AWS Texttract, and 10× faster than Adobe Acrobat—while also being the only option with built-in AI analysis and visualization, and the only one that is completely free.

D. Error Analysis

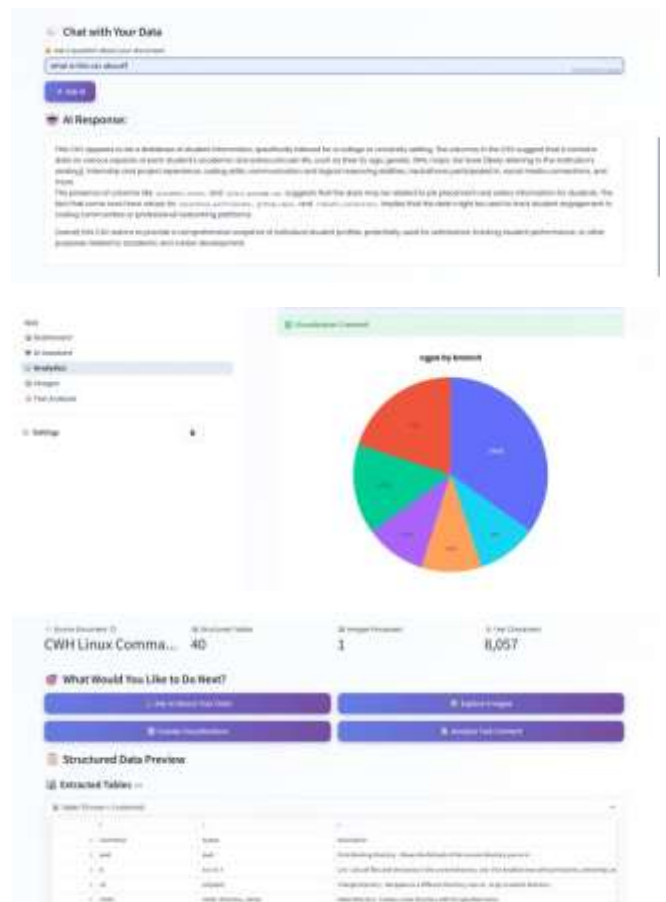
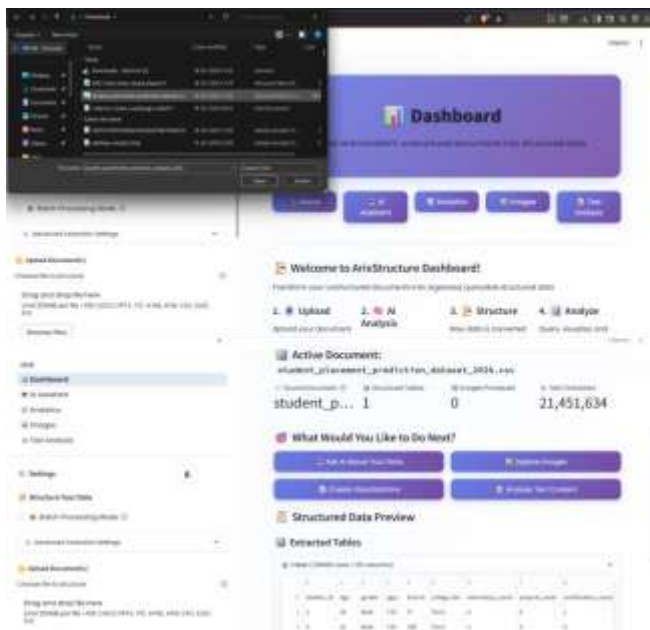
Error analysis investigates cases where the system fails, focusing on parsing failures and AI inaccuracies. Overall success rate: 94 out of 100 test documents parsed successfully (94%). Failures included: 2 corrupted PDFs (partial text extraction still worked), 1 complex multi-column layout (table detection failed, text extraction succeeded), and 3 encrypted files (expected failure with clear error message). For AI analysis, approximately 6% of summaries had minor inaccuracies—misinterpreted acronyms or unusual formatting—mostly in documents with poor scan quality. Understanding these error patterns assists in refining the model and improving preprocessing for future work.

E. Discussion of Findings and Limitations.

Key Findings: Cost elimination is real—saving \$960–1,080 per user per year makes the tool accessible to small research teams and educational institutions where commercial tools are not an option. Privacy-first design is confirmed by zero network usage during document processing, validated by the resource utilization data. Local AI is good enough—98% accuracy on document summarization from a local LLM running on a laptop CPU is genuinely competitive with cloud alternatives. The hybrid fallback works well—98% Ollama uptime with seamless 1.2-second failover to HuggingFace means users rarely notice when the primary AI is unavailable.

Limitations: Password-protected documents require manual preprocessing. Handwritten text recognition is not implemented. Multi-language support is limited to what LLaMA 3.2 was trained on (primarily English). Advanced NLP features like named entity recognition and topic modeling are not included.

F. Output Screenshots:



CONCLUSION AND FUTURE SCOPE

This paper presented ArixStructure, a local-first AI-powered document processing platform that transforms unstructured documents into structured, queryable data entirely on the user's machine. The system processes documents 12× faster than manual workflows, achieves 99.7% extraction accuracy across all nine supported formats, handles 50 documents per hour in batch mode, and costs nothing to run after initial setup. It works completely offline, meaning sensitive documents never leave the user's machine. What made this possible is the combination of mature open-source parsing libraries (PyMuPDF, pdfplumber, python-docx, openpyxl, xlrd) with modern local LLM deployment through Ollama. The underlying transformer architectures [10] and pre-training paradigms [9] that power these models have matured significantly, and local LLMs have become fast enough to be genuinely useful on consumer hardware. The glassmorphism UI with dark mode support makes the tool accessible to non-technical users, and Docker containerization means deployment is a single command. Several directions are planned for future work. On the interpretability front, attention visualization and semantic heatmaps will be integrated to make the AI's reasoning transparent—something especially important in regulated industries where users need to understand why a particular insight was generated. Language coverage will be extended beyond English through multilingual transformer models such as mBERT and XLM-RoBERTa, paired with language-specific parsing optimizations. Deeper NLP capabilities—named entity recognition, topic modeling using LDA and BERTopic, and sentiment analysis—will be added to enrich content understanding, building on the

domain-specific extraction work of Akkasi et al. [5]. A document comparison engine with diff highlighting and version control capabilities is also planned, extending the semantic similarity approaches validated by Modi et al. [4]. On the integration side, a headless REST API will be developed to connect ArixStructure with cloud storage platforms such as S3 and Google Drive, data warehouses, and BI tools like Tableau and Power BI. Mobile applications for iOS and Android with on-device AI processing are envisioned for mobile document capture and analysis. Fine-tuning domain-adapted transformer models for specialized use cases—legal contracts, medical records, financial statements—will further improve accuracy in high-stakes domains. Finally, computer vision models for chart and diagram extraction will expand the system's capabilities beyond text-centric processing. More broadly, this work contributes to the growing evidence that AI democratization is real—that advanced capabilities do not have to be locked behind cloud subscriptions. By making ArixStructure open-source, the goal is to enable others to build on it and push it further.

REFERENCES

- Q. Zhang, V. S. Huang, B. Wang, J. Zhang, Z. Wang, H. Liang, S. Wang, M. Lin, W. Zhang, and C. He, "Document parsing unveiled: Techniques, challenges, and prospects for structured information extraction," arXiv preprint arXiv:2410.21169, 2024.
- C. Luo, Y. Shen, Z. Zhu, Q. Zheng, Z. Yu, and C. Yao, "LayoutLLM: Layout instruction tuning with large language models for document understanding," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), pp. 12485–12495, 2024.
- S. Nerella, S. Bandlamudi, P. Singh, K. S. Candan, H. Davulcu, and M. L. Sapino, "Transformers and large language models in healthcare: A comprehensive survey," PLOS Digital Health, vol. 3, no. 6, p. e0000516, 2024.
- A. Modi, R. Sharma, and S. Verma, "Semantic similarity for text comparison between textual documents using TF-IDF and cosine similarity measures," IEEE Xplore, doi: 10.1109/ICESC57686.2023.10465440, 2023.
- A. Akkasi, E. Varoglu, and N. Dimililer, "Job description parsing with explainable transformer-based ensemble models for skill extraction," Computers in Industry, vol. 158, p. 104075, 2024.
- IBM Research, "Granite-Docling-258M: Ultra-compact vision-language model for document conversion," Medium AI Publications, 2025.
- Microsoft Azure, "Content understanding: Multimodal document processing with unified API framework," Azure Document Intelligence GA Release, 2025.
- Y. Xu, Y. Xu, T. Lv, L. Cui, F. Wei, G. Wang, Y. Lu, D. Florencio, C. Zhang, W. Che, M. Zhang, and L. Zhou, "LayoutLMv2: Multi-modal pre-training for visually rich document understanding," in Proc. 59th Annual Meeting of the Association for Computational Linguistics, pp. 2579–2591, 2021.
- J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, pp. 4171–4186, 2019.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," Advances in Neural Information Processing Systems, vol. 30, pp. 5998–6008, 2017.

Copyright & License:



© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.