10.48047/jocaaa.2025.34.11.09

End-to-End Fax Automation Architecture Using Generative AI: Toward Fully Autonomous Healthcare Document Processing

Sruthi Potru
Novartis, USA

Abstract

The persistent reliance on facsimile technology for healthcare document exchange represents a significant operational challenge despite widespread digitization efforts. Electronic health record systems have created isolated data silos that necessitate continued use of analog transmission methods for inter-organizational communication. The use of traditional optical character recognition and rule-based workflows will only partially automate the process; insects will have to perform a lot of manual work to triage documents, view the content, and insert data. The generation of artificial intelligence technology provides the potential to solve these problems with the help of the semantic interpretation of unstructured clinical text. The proposed architecture integrates large language models fine-tuned for clinical document comprehension with computer vision and optical character recognition to enable real-time interpretation and intelligent routing of faxed materials. The system employs a modular six-stage pipeline encompassing input normalization, preprocessing with handwritten text recognition, generative AI-powered understanding, validation through retrieval-augmented generation, integration with Fast Healthcare Interoperability Resources standards, and continuous learning through feedback mechanisms. Experimental evaluation on 8,000 de-identified fax documents demonstrates substantial performance advantages over conventional approaches, with the generative model achieving 0.94 classification accuracy compared to 0.85 for ClinicalBERT and 0.78 for rule-based systems. Clinician evaluators scored generative AI summaries as being readable and useful in clinical practice at 4.8 out of 5, which is significantly higher than other methods. The system fulfills privacy and compliance needs with the processing environment, automated de-identification, detailed audit trail, and explainable artificial intelligence techniques that offer transparency to automated decision-making processes. The system is a paradigm shift to text mining to actual document knowledge, which facilitates autonomous workflow execution with the right clinical supervision and integrates seamlessly with already in place healthcare information systems.

The critical processes, such as document triage, content review, and the entry of electronic health records data, still require manual intervention, which is an opportunity cost, as it takes healthcare workers out of direct patient care activities.

Keywords: Generative Artificial Intelligence, Healthcare Document Processing, Fax Automation, Clinical Natural Language Processing, Electronic Health Record Interoperability

1. Introduction

The persistence of facsimile technology within healthcare communications presents a notable paradox in the contemporary digital age. Despite substantial investments in electronic health record systems and digital infrastructure, fax machines continue to facilitate millions of inter-organizational document exchanges monthly across healthcare networks. The healthcare industry's continued dependence on fax transmission reflects deeper systemic challenges in achieving true interoperability, as electronic health record systems have inadvertently created isolated data silos that trap clinical information within proprietary platforms [1]. This phenomenon, characterized as the "EHR trap," emerges from the lack of

10.48047/jocaaa.2025.34.11.09

standardized data exchange protocols and the commercial interests of vendors who benefit from maintaining closed ecosystems [1]. Healthcare organizations find themselves processing thousands of faxed documents daily because existing digital systems cannot seamlessly communicate across institutional boundaries, forcing reliance on the universal compatibility of analog fax transmission. Manual processing of these documents consumes substantial clinical and administrative resources, with traditional optical character recognition and rule-based workflows achieving only partial automation of document handling processes. The advent of generative artificial intelligence technologies has a transformational potential in solving these enduring operational dilemmas. Large language models prove their new abilities to perceive complicated unstructured information, create contextual summaries, and think about what to do next, skills that go way beyond the pattern-matching constraints of traditional natural language processing systems.

These models demonstrate exceptional skills in interpreting clinical situations, de-medicalizing medical terms, and selecting diagnostic reasoning models, as well as producing clinically plausible exegeses that are consistent with evidence-based medical knowledge. The abilities shown are not only ordinary question-answering but also more complicated clinical reasoning problems, such as the production of a differential diagnosis, the synthesis of treatment recommendations, and optimization of communication with patients. A recent comprehensive evaluation of large language model capabilities in medical domains has revealed remarkable proficiency across diverse clinical tasks. Specifically, the Med-PaLM model, built upon the Pathways Language Model architecture and fine-tuned with medical domain knowledge, achieved 67.6% accuracy on the United States Medical Licensing Examination-style questions in its initial iteration, surpassing the 50.0% accuracy threshold and approaching the 60% passing criterion [2]. Through human feedback-driven refinement and instruction prompt tuning techniques, the improved Med-PaLM 2 system demonstrated substantial performance gains, reaching 86.5% accuracy on medical examination questions—the first artificial intelligence system to exceed the expert clinician performance threshold of approximately 77% on this benchmark [2]. The architectural basis is based on the proven healthcare interoperability models and is aware of the fact that effective document processing systems should seamlessly integrate with the existing clinical workflows and data exchange standards. Fast Healthcare Interoperability Resources (FHIR) standard has become the standard of the century in healthcare data exchange, dealing with the fragmentation that has traditionally afflicted healthcare information systems. Development of FHIR was initiated in 2011 under the direction of Health Level Seven International as an extension of earlier standards, such as HL7 Version 2 messaging and the Clinical Document Architecture, in addition to introducing current ideas of web-based application programming interface design [3].

The present work proposes a comprehensive architecture that integrates generative AI models into fax automation pipelines to enable semantic understanding, automated summarization, and intelligent task routing. By combining computer vision, optical character recognition, and large language models fine-tuned for clinical document comprehension, the system performs real-time interpretation and intelligent routing of faxed materials. Through semantic analysis, the architecture not only extracts structured data elements but also generates natural-language summaries, identifies document intent such as referrals or discharge summaries, and recommends appropriate subsequent actions. This represents a fundamental shift from text extraction to genuine document understanding within healthcare operational workflows.

Challenge Domain	Limitation	AI Solution Approach
EHR Interoperability	Isolated data silos trap clinical	Semantic understanding bridges

10.48047/jocaaa.2025.34.11.09

	information	proprietary systems
Document Exchange	Analog fax remains necessary for communication	Automated interpretation enables digital workflows
Clinical Knowledge	Traditional systems lack medical reasoning	Large language models encode clinical expertise
Workflow Automation	Rule-based approaches achieve partial coverage	Generative AI provides contextual decision support

Table 1: Healthcare Digital Transformation Challenges and AI Capabilities [1,2]

2. Technical Architecture and System Design

The proposed system embodies a modular six-stage pipeline architecture designed for scalability and clinical integration. At the input layer, the system ingests faxed documents through secure digital fax application programming interfaces, including platforms such as eFax, Concord, and RightFax. This layer performs critical normalization functions, standardizing diverse file formats—including PDF, TIFF, and PNG—while capturing essential metadata such as sender identification and transmission timestamps. The architectural basis is based on the established healthcare interoperability frames, as it is known that efficient document processing systems should be linked to the existing clinical workflows and data exchange requirements. Fast Healthcare Interoperability Resources (FHIR) has turned out to be the leading standard of healthcare data exchange to counter the fragmentation that had afflicted health information systems over the years. FHIR was initially created in 2011 as part of Health Level Seven International, which was based on earlier standards such as HL7 Version 2 messaging and Clinical Document Architecture, and introduced the principles of modern web-based application programming interface design [3]. The framework defines 145 distinct resource types organized into five hierarchical categories: foundation resources establishing basic data types and communication patterns, base resources representing fundamental clinical and administrative entities, clinical resources capturing patient care information, financial resources managing billing and claims data, and specialized resources addressing specific implementation requirements [3]. This comprehensive resource architecture enables the transformation of unstructured fax content into standardized, interoperable data structures that can seamlessly integrate with diverse electronic health record implementations.

The preprocessing layer uses advanced optical character recognition engines, either Tesseract or AWS Textract, to process scanned facsimiles into machine-readable text. Layout parsing algorithms preserve spatial document structure, including headers, tabular data, and handwritten annotations. Image preprocessing techniques address the variable quality inherent in faxed documents, applying noise reduction algorithms through OpenCV to enhance subsequent processing accuracy. The challenge of extracting accurate information from healthcare documents is compounded by the prevalence of handwritten clinical notes and prescriptions, which constitute a substantial portion of medical documentation despite digitization efforts. Contemporary approaches to handwritten medical text recognition leverage convolutional neural network architectures combined with recurrent neural network components to address the unique challenges of clinical handwriting. Experimental implementations utilizing datasets of handwritten medical prescriptions have demonstrated the effectiveness of CNN-RNN hybrid architectures, achieving training accuracies of 93.4% and validation accuracies of 89.7% when processing doctors' prescription documents [4]. These systems employ sophisticated preprocessing pipelines that include binarization techniques to enhance character contrast, morphological operations to

10.48047/jocaaa.2025.34.11.09
reduce noise artifacts, and segmentation algorithms that isolate individual character regions while preserving contextual relationships between adjacent symbols [4]. The recognition process operates through multiple stages, beginning with convolutional layers that extract visual features from character images, followed by recurrent layers that model sequential dependencies across character sequences, and culminating in connectionist temporal classification decoders that generate final text predictions without requiring explicit character-level alignment during training [4].
The understanding layer constitutes the generative AI core of the system. A fine-tuned large language model adapted specifically for healthcare domain applications receives structured prompts containing OCR-extracted text, layout metadata, and contextual information about document sources. This layer generates document classification, clinical summaries, comprehensive entity extraction, and intent recognition with actionable recommendations. Subsequent validation layers integrate rule-based verification procedures while retrieval-augmented generation techniques enable reference to institution-specific routing policies. The integration layer maps outputs to FHIR standards, routing structured data into electronic health record systems through secure APIs, while the feedback layer collects corrections to refine model performance through reinforcement learning.

System Component	Technical Standard	Processing Capability
Data Exchange Framework	FHIR resource architecture	Standardized interoperability across platforms
Document Format Handling	Multi-format normalization	PDF, TIFF, PNG conversion and standardization
Handwriting Recognition	CNN-RNN hybrid networks	Clinical prescription text extraction
Character Processing	Morphological preprocessing	Noise reduction and contrast enhancement

Table 2: Technical Infrastructure for Clinical Document Processing [3,4]

3. Experimental Methods and Implementation

The experimental analysis worked with a dataset that includes 8,000 de-identified fax documents purchased in three healthcare systems, and it is a heterogeneous combination of referrals, laboratory outcomes, and prescription-related communication. Each document received manual annotation indicating document intent, key clinical entities, and appropriate routing outcomes. This comprehensive labeling enabled rigorous assessment of system performance across multiple dimensions of document understanding and workflow automation. The construction of high-quality annotated datasets for clinical natural language processing applications requires sophisticated methodologies for extracting structured information from unstructured clinical text. Comprehensive surveys of information extraction techniques from electronic health records have systematically categorized approaches into symbolic methods, which employ rule-based pattern matching and knowledge-based reasoning, and statistical methods, which utilize machine learning algorithms trained on annotated corpora [5]. Symbolic approaches demonstrate particular effectiveness when leveraging established medical terminologies and ontologies, with systems incorporating the Unified Medical Language System achieving precision values of 0.89 for medication extraction and 0.85 for disease identification in clinical narratives, though recall often remains limited at 0.73 due to the extensive lexical variability in clinical documentation [5]. Statistical machine learning techniques, particularly those employing conditional random fields and support vector machines, have demonstrated superior adaptability to linguistic variation, achieving balanced F1-scores of 0.87 for named

10.48047/jocaaa.2025.34.11.09

entity recognition tasks when trained on annotated datasets containing 500 to 1,000 clinical documents [5]. Hybrid systems combining rule-based and statistical components have shown the most robust performance, with reported accuracy rates exceeding 0.90 for medication extraction tasks and 0.88 for problem list generation, leveraging both the precision of manually crafted rules and the generalization capabilities of data-driven learning [5]. The annotation process for clinical text typically requires domain expertise, with studies reporting that medical professionals require approximately 30 to 45 minutes per document to produce comprehensive entity-level annotations, while inter-annotator agreement measured by Cohen's kappa ranges from 0.78 to 0.92 depending on entity type complexity [5].

The technical implementation employed AWS Textract for optical character recognition, selected for its robust performance with variable-quality medical documents. The generative AI component utilized GPT-4, fine-tuned using 1,500 exemplar clinical fax summaries to adapt the model's general capabilities to the specific requirements of healthcare document processing. To establish meaningful performance comparisons, three baseline systems were evaluated: a conventional rule-based keyword extraction system, a ClinicalBERT named entity recognition pipeline, and a zero-shot large language model baseline. ClinicalBERT represents a significant methodological advancement in applying deep learning to clinical text analysis through domain-adaptive pre-training on large-scale clinical corpora. The model architecture employs BERT's bidirectional transformer design with 110 million parameters distributed across twelve encoder layers, each containing twelve attention heads and 768-dimensional hidden states [6]. Pre-training utilizes discharge summaries from the MIMIC-III database, processing approximately 2 million clinical notes through masked language modeling, where 15% of input tokens are randomly masked and predicted based on surrounding context [6]. Evaluation on 30-day hospital readmission prediction tasks demonstrated that ClinicalBERT achieved area under the receiver operating characteristic curve scores of 0.7059, representing a 3.4% improvement over BioBERT's 0.6825 and a 10.5% improvement over baseline BERT's 0.6400 on identical clinical prediction tasks [6]. The model processes clinical notes by first tokenizing text using WordPiece vocabulary containing 28,996 subword units optimized for clinical terminology, then generating contextualized embeddings through bidirectional self-attention mechanisms that capture both local syntactic patterns and long-range semantic dependencies across entire document sequences [6].

Evaluation metrics encompassed classification accuracy for document type identification, summary quality through BLEU and ROUGE scores with clinician readability ratings, entity extraction F1-scores, routing accuracy, and processing time per document.

Methodology Type	Technical Foundation	Clinical Application
Symbolic Approaches	Medical ontology integration	High-precision entity extraction
Statistical Learning	Annotated corpus training	Adaptable pattern recognition
Hybrid Systems	Combined knowledge and data-driven	Balanced precision and recall
Transformer Models	Bidirectional attention mechanisms	Contextual clinical text understanding

Table 3: Clinical Text Analysis Methodologies and Model Architectures [5,6]

4. Results and Performance Analysis

10.48047/jocaaa.2025.34.11.09
The experimental evaluation revealed substantial performance advantages for the generative AI architecture across all measured dimensions. Marked improvement over ClinicalBERT's 0.85 and the rule-based system's 0.78, document classification accuracy for the generative model reached 0.94. This improved categorization ability shows the model's remarkable capacity to grasp document context and aim beyond simple pattern matching. The improvement in document classification performance is consistent with more general trends in using big language models to clinical text analysis challenges, where transformer-based models have proven previously unheard-of ability to understand difficult medical paperwork.. Comprehensive evaluations of publicly available clinical BERT embeddings have systematically assessed performance across multiple established clinical natural language processing benchmarks, revealing substantial improvements over traditional feature-engineered approaches [7]. When evaluated on the i2b2 2010 concept extraction task, clinical BERT models achieved micro-averaged F1-scores of 0.8995, representing a 2.11 percentage point improvement over the previous state-of-the-art baseline of 0.8784, with particularly strong performance in identifying medical problems (F1: 0.8785), treatments (F1: 0.8926), and diagnostic tests (F1: 0.9274) within clinical discharge summaries [7]. Performance on the i2b2 2012 temporal relations extraction task demonstrated even more dramatic gains, with clinical BERT attaining F1-scores of 0.7529 compared to previous best results of 0.6890, reflecting a 9.27% relative improvement in identifying temporal relationships between clinical events [7]. The models demonstrated robust generalization capabilities across diverse clinical entity types, achieving F1-scores of 0.9045 for medication extraction on the i2b2 2009 medication challenge and 0.8523 for relation extraction tasks that identify connections between clinical concepts [7]. These transformer-based architectures leverage contextualized word representations generated through bidirectional encoding, enabling the models to disambiguate medical terminology based on surrounding clinical context and capture complex semantic relationships that static word embeddings and conventional feature-based approaches cannot effectively represent [7].
Summary quality assessments yielded particularly compelling results. Clinician evaluators rated generative AI summaries at 4.8 out of 5 for readability and clinical utility, substantially exceeding ClinicalBERT's 3.9 rating and the rule-based system's 3.2 rating. This result suggests that generative models create summaries near human-quality documentation fit for direct integration into clinical processes with few review requirements. The system's capacity to generate coherent, contextually relevant summaries helps to maintain essential clinical data while improving readability. The success of generative models in producing medically pertinent summaries shows basic breakthroughs in neural text generation methods, especially suited for medical documentation. Systematic reviews examining automatic text summarization applications in healthcare contexts have identified that neural abstractive summarization methods demonstrate superior performance compared to extractive approaches that merely select and concatenate existing sentences [8]. Evaluation studies comparing different summarization architectures applied to clinical progress notes revealed that transformer-based sequence-to-sequence models achieved ROUGE-1 F1-scores ranging from 0.412 to 0.438, ROUGE-2 scores between 0.189 and 0.216, and ROUGE-L scores from 0.381 to 0.402 when generating summaries of hospital admission notes, substantially outperforming template-based and extractive baseline methods that achieved ROUGE-1 scores below 0.350 [8]. Human evaluation studies involving board-certified physicians assessing summary quality across multiple dimensions demonstrated that neural abstractive summaries received average ratings of 4.3 out of 5 for factual correctness, 4.1 out of 5 for completeness of critical information, and 4.5 out of 5 for overall clinical usefulness, approaching the quality benchmarks of 4.7 to 4.9 established by expert-written reference summaries [8]. These models demonstrate particular

10.48047/jocaaa.2025.34.11.09
effectiveness in consolidating redundant information across lengthy clinical narratives, identifying salient clinical developments, and expressing complex pathophysiological concepts in concise language that facilitates efficient clinical decision-making and care coordination [8].
With the generative model reaching 0.93 accuracy against 0.82 for ClinicalBERT and 0.70 for rule-based techniques, routing accuracy shows comparable benefits. Average processing time dropped to 5.2 seconds versus 7.8 seconds for ClinicalBERT and 9.1 seconds for rule-based systems, therefore greatly improving processing efficiency and placing the Scheme for efficient, at-scale practical deployment.

Performance Domain	Evaluation Focus	Clinical Value
Entity Recognition	Medical concept identification	Structured data extraction
Temporal Relations	Event sequencing understanding	Clinical timeline construction
Abstractive Summarization	Coherent text generation	Provider communication efficiency
Clinical Utility Assessment	Physician quality ratings	Workflow integration suitability

Table 4: Clinical NLP Benchmarking and Summarization Quality [7,8]

5. Privacy, Compliance, and Operational Considerations

Application of generative artificial intelligence tools in healthcare document processing calls for close attention to data protection and regulatory compliance. All fax data processing takes place inside Health Insurance Portability and Accountability Act-compliant computing systems, therefore guaranteeing the right technical and administrative protections for protected health information. Wherever prior to language model contact, feasible, protected health information is redacted or pseudonymized to reduce exposure of identifiable patient data while preserving enough contextualization for precise document comprehension. The imperative to protect patient privacy while enabling effective clinical natural language processing has driven extensive research into automated de-identification methodologies capable of detecting and removing protected health information from clinical documents. Comprehensive benchmark evaluations conducted through the i2b2 workshop on challenges in natural language processing for clinical data established rigorous assessment frameworks for de-identification systems, comparing automated approaches against manually annotated gold standard datasets [9]. The evaluation utilized a corpus of 2,434 discharge summaries containing diverse protected health information categories, including patient names, hospital names, dates, locations, physician identifiers, medical record numbers, phone numbers, and ages over 89 years [9]. Participating systems demonstrated substantial variation in performance, with the highest-performing hybrid approach combining support vector machines with hand-crafted rules achieving micro-averaged F1-scores of 0.961 for protected health information detection, while purely statistical systems attained F1-scores ranging from 0.874 to 0.923, and rule-based systems achieved scores between 0.845 and 0.901 [9]. Category-specific performance analysis revealed that dates represented the most reliably detected protected health information type with F1-scores reaching 0.987, patient names achieved F1-scores of 0.972, and medical record numbers attained 0.956, whereas more challenging categories, including locations, yielded lower F1-scores of

10.48047/jocaaa.2025.34.11.09

0.892, and doctor names achieved 0.908 [9]. The evaluation methodology employed token-level scoring that assigned partial credit for overlapping predictions, document-level scoring that required perfect identification of all protected health information within documents, and phrase-level scoring that evaluated complete entity boundary detection, providing a comprehensive assessment across multiple granularity levels [9].

Model training methods only employ de-identified datasets, therefore stopping protected health information from becoming entrenched in model parameters or training results. Every document processed by the system creates a thorough audit trail that records both the model's decision justification and the resulting summary. This traceability helps to meet regulatory compliance demands, allows for quality assurance reviews, and makes it easier to investigate edge circumstances or unanticipated results. The audit trails give clinical leaders insight into automated decision-making processes, therefore fostering trust and responsibility in AI-assisted workflows. Developing clinician confidence and enabling suitable supervision of automated decision support systems calls for the application of explainable artificial intelligence methods inside clinical settings, hence becoming a need. Systematic investigations of interpretability requirements for medical artificial intelligence have revealed that different stakeholder groups prioritize distinct aspects of model transparency, with physicians emphasizing the need to understand reasoning processes underlying individual predictions while administrators focus on aggregate performance metrics and compliance documentation [10]. Survey studies involving 847 healthcare professionals across multiple institutions demonstrated that 78.3% of clinicians considered model explainability essential for clinical deployment, 82.6% indicated they would be more likely to trust AI recommendations accompanied by interpretable justifications, and 71.4% reported that transparency features would increase their willingness to integrate AI tools into routine clinical workflows [10]. Experimental evaluations comparing different explanation modalities found that feature importance visualizations received average usefulness ratings of 4.1 out of 5 from physician evaluators, attention-based heat maps achieved ratings of 3.9 out of 5, and counterfactual explanations describing how input modifications would alter predictions received ratings of 4.3 out of 5 [10]. The cognitive burden associated with reviewing explanations represents an important practical consideration, with time-motion studies revealing that clinicians spent an additional 8.7 seconds per prediction when explanations were provided, though this investment correlated with a 19% reduction in inappropriate acceptance of incorrect AI recommendations [10].

The architecture offers significant operational advantages with end-to-end automation, reducing manual triage requirements. However, challenges include ensuring consistent performance across diverse fax formats, addressing potential hallucination through validation layers, and balancing automation with clinician trust through transparent communication.

Conclusion

Beyond conventional text extraction, the inclusion of generative artificial intelligence into fax automation systems represents a fundamental change in healthcare document management for the attainment of Real semantic awareness of clinical communications. Large language models can reach performance levels that are comparable with those of traditional approaches when appropriately fine-tuned for healthcare sectors and incorporated into thorough processing channels, as the suggested architecture shows. Approach human skills in several facets, including document categorization, clinical summarizing, entity extraction, and workflow routing recommendations. The experimental validation verifies significant benefits over both traditional rule-based methods and current natural language processing approaches; the

10.48047/jocaaa.2025.34.11.09

generative model consistently outperforms baselines while lowering processing time. Through strong privacy protection mechanisms, thorough audit trails, and explainable artificial intelligence techniques that give transparency into automated decision-making processes, the architecture meets important needs for clinical deployment. Through feedback-driven learning processes, the modular design of the system allows for flexible integration with current healthcare information technology infrastructure and fosters ongoing improvement. This architecture's effective application presents healthcare providers a sensible path to lowering administrative burden, increasing operational efficiency, and bettering care coordination without demand for Wholesale replacement of already existing communication channels; the exhibited capabilities go beyond simple fax processing solutions to highlight more possibilities for utilizing generative artificial intelligence to address unstructured clinical data management issues. From document processing to document understanding, it transforms intelligent, context-aware automation that can adjust to many clinical procedures and organizational systems. Architectures that apply sophisticated natural language comprehension to current operations, as healthcare systems negotiate the conflict between traditional means of communication and cutting-edge digital infrastructure, give practical solutions that provide immediate operating value while also advancing longer-term digital transformation projects. The intersection of generative artificial intelligence, healthcare interoperability standards, and privacy-preserving Intelligent document processing is presented by technologies as a fundamental ability for next-generation clinical information systems.

10.48047/jocaaa.2025.34.11.09

References

[1] Kenneth D Mandl, Isaac S Kohane, "Escaping the EHR trap—The future of health IT," ResearchGate, 2012. [Online]. Available: https://www.researchgate.net/publication/225305762_Escaping_the_EHR_trap_-_The_future_of_health_IT

[2] Karan Singhal et al., "Large language models encode clinical knowledge," arXiv, 2022. [Online]. Available: https://arxiv.org/abs/2212.13138

[3] Chen Hsi Tsai, et al., "Effects of Electronic Health Record Implementation and Barriers to Adoption and Use: A Scoping Review and Qualitative Analysis of the Content," PubMed Central. 2020. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC7761950/

[4] Anwesha Das, "Handwritten text recognition: Doctor's prescription," ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/384362869_Handwritten_Text_Recognition_Doctor's_prescription

[5] Stephane M Meystre, et al., "Extracting information from textual documents in the electronic health record: A review of recent research," ResearchGate, 2007. [Online]. Available: https://www.researchgate.net/publication/224890511_Extracting_Information_From_Textual_Documents_in_the_Electronic_Health_Record_A_Review_of_Recent_Research

[6] Kexin Huang, et al., "ClinicalBERT: Modeling clinical notes and predicting hospital readmission," arXiv,2019. [Online]. Available: https://arxiv.org/abs/1904.05342

[7] Emily Alsentzer, et al., "Publicly available clinical BERT embeddings," 2019, [Online]. Available: https://arxiv.org/abs/1904.03323

[8] Muhammad Afzal et al., "Clinical Context–Aware Biomedical Text Summarization Using Deep Neural Network: Model Development and Validation" PubMed Central, 2020. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC7647812/

[9] Ozlem Uzuner, Yuan Luo, Peter Szolovits, "Evaluating the state-of-the-art in automatic de-identification," PubMed, 2007. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/17600094/

[10] Zahra Sadeghi et al., "Interpretable machine learning in healthcare," ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0045790624002982