Journal of Advances in Developmental Research (IJAIDR)

E-ISSN: 0976-4844 ● Website: www.ijaidr.com ● Email: editor@ijaidr.com

AI-Driven Patient Transfer Summarization Using Salesforce Health Cloud as a Clinical CRM Intelligence Platform

Sai Saketh Sunkara

Independent Researcher
Dallas, Texas
saisalesforce2711@gmail.com

Abstract:
The increasing complexity of inter-hospital patient transfers, combined with the widespread use of heterogeneous and unstructured medical documentation, has intensified challenges related to data fragmentation, delayed clinical decision-making, and inefficient care coordination. Existing healthcare information systems and CRM platforms often operate in silos, limiting their ability to provide timely, standardized, and actionable patient insights. To address these limitations, this paper proposes a Salesforce Health Cloud–centric AI framework, where Salesforce Health Cloud acts as the primary clinical CRM platform orchestrating patient transfer workflows, while AI modules provide embedded document intelligence and predictive decision support. The proposed architecture combines an Adaptive Document Parsing and Structuring (ADPS) Framework for layout-aware document understanding, a Context-Aware Clinical Summarization (CACS) Engine for ontology-guided extraction of medically significant information, a FHIR-Compliant Interoperability Integration (FCI) Layer for standardized data exchange, and a Predictive Admission Intelligence Module (PAIM) for proactive triage and risk assessment. By unifying document intelligence, interoperability, CRM workflow automation, and machine learning–based prediction into a single end-to-end pipeline, the framework converts patient transfer documentation from a manual, error-prone process into an automated, decision-support-driven workflow. Experimental evaluation on real-world inter-hospital transfer datasets demonstrates high document structuring accuracy (94.2%), clinically relevant summarization performance (96.8%), substantial reductions in admission processing time (from 38 to 11 minutes), and strong predictive accuracy (97.1%) with reliable risk discrimination. The results confirm that the proposed approach enhances clinical efficiency, reduces administrative burden, and enables CRM platforms to function as active participants in clinical decision-making, establishing a scalable and intelligent foundation for digital transformation in healthcare transfer management.

Keywords: Salesforce Health Cloud; AI-driven patient transfer summarization; Clinical CRM intelligence; FHIR interoperability; Predictive admission analytics; Healthcare workflow automation.

1. INTRODUCTION

The recent growth of digital health technologies, artificial intelligence (AI), and cloud-based customer relationship management (CRM) has radically changed the way healthcare data is generated, processed, and consumed in clinical facilities, particularly in hospital systems and patient referral systems [1,2]. Modern healthcare delivery is a field of operation that is becoming increasingly dependent on the ability to access the appropriate and complete patient data promptly to support the clinical decision-making process, care coordination, and operational efficiency [3]. However, despite universal digitization, healthcare is still marked by inconsistent and fragmented patient information, especially during inter-hospital transfers, where delays or inaccuracy in medical data transfer can have a drastic effect on the outcome of treatment [4,5]. Salesforce Health Cloud offers a single, patient-centric CRM system with the ability to combine fragmented inter-hospital data into a single longitudinal view of a patient. Such a centralized base makes it possible to automate downstream, be interoperable, and make wise clinical choices.

In contemporary clinical practice, patients who are transferred between healthcare providers tend to show up with colossal amounts of unstructured medical data, including scanned PDFs, handwritten documentation, lab findings, and discharge notes [6]. These records must be reviewed by clinicians manually under time-pressured situations, leading to increased cognitive load and the availability of potentially misinterpreted clinically relevant facts in administrative noise [7]. Older electronic health record (EHR) technology and hospital information systems (HIS) are siloed and do not easily communicate with CRM platforms such as Salesforce Health Cloud, resulting in bottlenecks in communication, delays in treatment, and multidisciplinary care coordination inefficiencies [8,9]. Using Salesforce Health Cloud, unstructured transfer documents can be converted into structured clinical insights that are readily available in the timelines of patients managed in CRM. This decreases the amount of work done in manual review and speeds up clinical decisions made in relation to admission.

The present approaches to healthcare data integration and clinical documentation processing rely on the rule-based retrieval, template-based summing, or manual abstraction processes [10]. Though these procedures may be useful in restricted or standardized contexts, they fail to generalize between heterogeneous document types, referral sources, and institutional practices [11]. In addition, conventional summarization systems are usually uninformed about a clinical situation; they generate either no actionable medical information or irrelevant output that is redundant and lacks clinical significance [12]. Such limitations are further elucidated by the fact that there are no developed interoperability layers that would facilitate reliable data sharing across various digital healthcare systems and would guarantee usability in real time [13,14]. Such obstacles reveal the necessity of Salesforce Health Cloud-based approach, in which CRM-native workflows, automation, and interoperability services support the use of intelligent patient transfer management, instead of isolated analytics tools.

Recent research has mentioned machine learning and natural language processing (NLP) in clinical text mining and decision support [15]. However, the majority of these solutions are conceived as individual features, where they operate on document parsing, summarization, or prediction as individual tasks, without an integrated operational pipeline [16,17]. Additionally, predictive analytics of admission and triage processes are typically made using incomplete or late data, which limits their applicability in high-stakes transfer cases [18]. The lack of studies, which feature closely compressed and AI-driven, CRM-enabled systems capable of transforming raw clinical records into homogenous, actionable, predictive intelligence, creates a research gap [19,20].

This study also contrasts the traditional methods that consider CRM platforms as passive data storage and retrieval systems by positioning Salesforce Health Cloud as the main system of record and workflow coordination platform in inter-hospital patient transfers. Health Cloud is natively aware of patient identity management, creating cases of admission, care coordination, and clinician notifications, whereas AI-powered parsing and summarization of documents and predictive analytics act as intelligent services invoked by CRM events. This Salesforce-based architecture guarantees that the outputs of the AI can be readily implemented in clinical processes.

As a solution to these unsolved issues, the current research offers an innovative AI-powered, CRM-based, multifunctional intelligent patient transfer management model. The proposed framework syntactically integrates four complementary strategies, an Adaptive Document Parsing and Structuring (ADPS) Framework of robust ingestion of heterogeneous medical documents, a Context-Aware Clinical Summarization (CACS) Engine of ontology-guided extraction of clinically significant information, a FHIR-Compliant Interoperability Integration (FCI) Layer to achieve standardized and semantically consistent data exchange, and a Predictive Admission Intelligence Module (PAIM) of proactive triage and risk assessment. The proposed system converts the current manual and error-prone patient transfer documentation into an intelligent decision-support mechanism by integrating these components together into one automated, predictive workflow with Salesforce Health Cloud.

The architecture proposed is focus on increasing clinical efficiency, decreasing administrative workload, and improving care coordination during inter-hospital transfers, and help CRM platforms to become active members of clinical decision-making rather than passive data storage. This combined methodology provides a new paradigm for smart digital transformation of healthcare data management, which is useful in quicker admissions, better patient results, and scalable interoperability in the contemporary healthcare ecosystem.

2. LITERATURE REVIEW

Krishnan et al., [21] In this research, the investigator seeks to develop a scalable AI-ready data conversion model that can resolve interoperability and compliance issues in healthcare data available in Medicare and Medicaid. The architecture is cloud-native and adopts microservice, containerization, and serverless computing to convert the schema into HIPAA-, HL7-, and FHIR-compliant format automatically and optimize downstream storage to store analytics. The results show that it is better in conversion speed, scalability, regulatory compliance, and integration with machine learning without failures to predictive analytics, fraud detection, and optimization of policy modifications than traditional ETL pipelines.

Meligy et al., [22] The purpose of the present study is to explored the potential of combining wearable devices with Salesforce Health Cloud to develop remote patient monitoring by providing real-time and personalized care. The research design suggests a conceptual model of CRM-integrated, which allows the constant ingestion of data through wearables, interoperability, and the integration of safe and scalable data handling. The results reveal that this kind of integration enhances clinical decision-making, solidifies patient-provider engagement, improves care coordination, and provides proactive health interventions, which result in better patient outcomes and more effective management of remote healthcare.

Polamarasetti et al., [23] This paper examines how an API strategy based on MuleSoft can support the seamless integration of Salesforce into the multi-cloud setting and solve the interoperability, scalability, and governance challenges. In the methodology, API-led connectivity approach with System, Process, and Experience API is utilized and backed by secure API design, event-driven architecture, as well as lifecycle management. According to the findings, MuleSoft facilitates real-time synchronization of data, modular reuse, and resilient multi-cloud integrations to enhance reliability, scalability, and operational efficiency as well as optimize the value of an enterprise multi-cloud investment.

Selwyn et al., [24] The purpose of the current research is to analyze the possibilities of adopting AI, computer vision, and robotic process automation in healthcare applications to be able to provide better patient care and improve operational efficiency. The methodology examines the architecture, workflows, and security of HealthSyncc mobile application, considering AI-based wellness capabilities, computer vision-based document processing, and automated scheduling. The results suggest that these combined technologies are found to enhance the diagnostic support, automate the administrative activities, enhance the data security, and help to achieve improved patient outcomes and more efficient hospital operations.

Alsalamah et al., [25] The proposed study is dedicated to the design of VHC-Bot, an AI-based, patient-focused healthcare app designed to assist clinicians without reducing the clinical process. Its methodology combines natural language processing, diagnostic algorithms, and adaptive machine learning in the context of collaborative care that is validated by clinical simulations, patient surveys, and workflow analysis. The results show enhanced diagnostic accuracy, shorter consultation time, better clinician-patient communication, and patient satisfaction, which prove that VHC-Bot provides efficient human-centred care without taking away the decision authority of the clinician.

Ghadi et al., [26] The purpose of the study is to determine the major obstacles to the adoption of wearable technology in healthcare and suggest AI-based solutions that could improve privacy, usability, and clinical value. The research design can be a systematic literature review and analysis of the methodologies, such as federated learning to preserve privacy, deep learning to reduce signal noise, anomaly detection in real-time, and security integration through blockchain. The results show that there is an enhancement of data accuracy, a decrease in clinician workloads, increased trust, and increased potential for adoption of AI-powered wearable health care systems that can be scaled and made safe and ethical.

Existing literature has researched cloud-native healthcare data pipelines, wearable-built CRM systems, API-based multi-cloud integration, AI-enabled healthcare applications, as well as virtual care platforms in isolation, but they are mostly fragmented and domain-specific. The existing solutions are not based on a single framework that focuses on interoperability, real-time intelligence, regulatory compliance, workflow automation, and predictive decision support in heterogeneous healthcare ecosystems. Furthermore, a minimal focus is on end-to-end integration of AI models, CRM platform, and operational processes, decreasing scalability, explainability, and inter-system flexibility. This divide underscores the necessity of an interoperative, data, intelligence, and clinical operations system that is holistic and based on AI. The literature does not present a Salesforce Health Cloud-focused model in which AI-driven intelligence, interoperability, and predictive analytics are inherently integrated in CRM-controlled clinical processes.

3. PROPOSED METHODOLOGY

The suggested methodology is developed with Salesforce Health Cloud serving as the focal coordination point and patient transfer system of record among the hospitals. AI modules work as embedded intelligence services in Salesforce workflows so that the insights can be directly implemented in the operations of clinical CRM. This section outlines the overall methodological framework of the proposed AI-Driven Patient Transfer Summarization system in collaboration with CRM-based healthcare systems, which is designed to solve the issue of fragmented, unstructured, and delayed transfer of patient information between hospitals. The methodology is designed as a series of smart, mutually enhancing processing operations, starting with heterogeneous acquisition of medical documents and structural normalization, proceeding to clinically contextual summarization, transformation to standards-based interoperability, and predictive admission intelligence. The framework is developed to be a single, end-to-end pipeline converting raw medical documentation into standardized, actionable, and predictive intelligence hard-coded into the clinical CRM workflows. The motivation behind each module is a particular weakness in the current healthcare data management systems, and there are clear logical transitions between stages to maintain methodological consistency, clinical usefulness, and operational scalability.

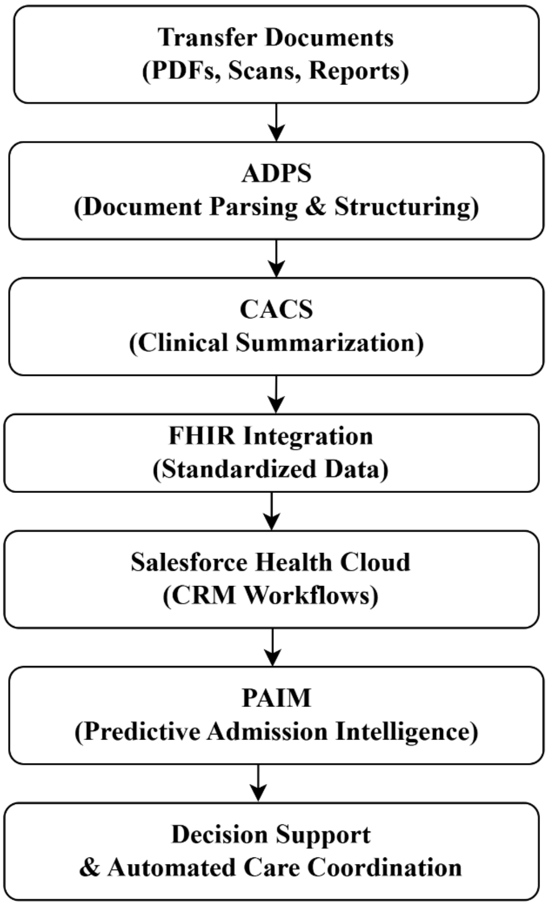

Figure 1: AI-Driven CRM-Integrated Patient Transfer

Figure 1 shows the top-to-bottom process of the proposed AI-supported patient transfer summarization scheme within the context of CRM-driven healthcare systems. It starts with the inter-hospital transfer documents, such as PDF and scanned files, and clinical reports, which are initially processed via the Adaptive Document Parsing and Structuring (ADPS) module to extract and normalize document structure. The organized content is then processed through the Context-Aware Clinical Summarization (CACS) engine that produces summaries that are short and have clinical implications. It is standardized by these summaries being FHIR-integrated, thereby allowing them to be interoperable. The standardized data is integrated into Salesforce Health Cloud workflow, in which the Predictive Admission Intelligence Module (PAIM) produces risk-based insights. Lastly, the framework facilitates decision support and automated care coordination that facilitates quick admission, better coordination, and minimized clinical delays.

3.1 Unified Patient Transfer Intelligence Formulation

Inter-hospital patient transfer is one of the most complicated information exchange situations in contemporary healthcare because it is both a clinical urgency and a situation that is extremely heterogeneous in data representation. The patient records developed in various hospitals have varied formats, structures, terminologies, and completeness, as it represents the institutional documentation practices and information systems. The traditional healthcare IT solutions usually break down this problem into separate subtasks: document parsing, clinical summarization, interoperability mapping, decision support, and loosely coupled or piping-based rules. The result of such disjointed treatment is loss of information, sluggish clinical insight and reduced real-world transfer flexibility. To overcome those constraints, the suggested framework develops patient transfer management as an integrated patient transfer intelligence problem, where all data interpretation and decision-making processes are co-modelled and optimized. This integrated intelligence formulation is applied to Salesforce Health Cloud to facilitate end-to-end patient transfer management.

A patient transfer record is formally described as a set of unstructured and semi-structured medical records:

$$D = \{d_1, d_2, \dots, d_n\} \tag{1}$$

Where the document d_i can represent referral letters, scanned discharge summaries, laboratory reports, diagnostic images, or handwritten clinical notes. Both documents are in a high-dimensional, heterogeneous input space, and both have clinically salient and extraneous administrative data. The overall goal of the suggested framework is to convert this unstructured collection of documents into a brief and operational clinical representation, at the same time producing predictive information, which can be used to make an admission decision.

This objective is modelled through a unified mapping function:

$$F(D; \Theta) \longrightarrow \{S, P\} \tag{2}$$

Where S represents a clinically actionable summary of the patient, and P represents predictive admission intelligence, e.g., triage urgency or risk stratification. Set of parameters Θ contains all learnable parameters in the framework, such as document parsing weights, the language model parameters, interoperability mappings, and predictive model coefficients. In contrast with modular pipelines, this formulation considers patient transfer intelligence as one learning and inference problem, and coherence and consistency are maintained by all stages.

The process of converting raw documents to structured knowledge may further be broken down into a series of intermediate representations:

$$D \xrightarrow{P} T \xrightarrow{S} S \xrightarrow{I} F \xrightarrow{\mathrm{M}} P \tag{3}$$

Where P (·) denotes adaptive document parsing, T represents structured clinical text, S (·) denotes context-aware summarization, I (·) represents interoperability standardization, and M (·) corresponds to predictive modeling. Each transformation is designed to preserve clinical semantics while progressively reducing noise and uncertainty. Importantly, these stages are not isolated; errors or ambiguities in early stages propagate downstream, which motivates joint optimization under a unified formulation.

From a learning perspective, the framework seeks to minimize a composite objective function:

$$\mathcal{L} = \mathcal{L}_{sum}(S, S^*) + \lambda \mathcal{L}_{pred}(P, P^*) \tag{4}$$

Where $\mathcal{L}_{sum}$ measures the deviation between generated summaries and clinically validated references, $\mathcal{L}_{pred}$ evaluates predictive accuracy against ground-truth admission outcomes, and λ balances interpretability and predictive performance. This formulation enforces alignment between clinical relevance and operational decision support.

The proposed framework builds a coherent analytical framework of automated, interoperable, and predictive patient transfer management through modeling patient transfer as a continuous intelligence pipeline instead of a collection of disjointed tasks. With this common formulation in place, the following subsection describes how heterogeneous and unstructured medical records are stabilized and ordered using the Adaptive Document Parsing and Structuring (ADPS) Framework that forms the point of entry of the proposed pipeline.

3.2 Adaptive Document Parsing and Structuring (ADPS) Framework

Adaptive Document Parsing and Structuring (ADPS) Framework represents the initial phase of the intended patient transfer intelligence pipeline. Its main goal is to transform non-standardized, non-structured medical text to a stable and machine-understandable format, and maintain the clinical context inherent in input records. In inter-hospital transfer cases, patient records are produced in various health facilities whose documentation practices, layouts, and scan operations are different. Consequently, the raw transfer records tend to have irregular formatting, section sorting, embedded tables, handwritten notes, and noise artifacts. Such inconsistencies greatly hinder the automation process and necessitate considerable manual work on the part of the clinicians, which in turn drives the desire to seek an adaptive and context-sensitive approach to document parsing. The workflows of Salesforce Health Cloud intake receive the structured outputs created by ADPS.

Formally, an incoming patient transfer document can be represented as d_i, and it can either be in PDF format, scanned image format, or hybrid digital form. The transformation within the ADPS Framework is the layout-sensitive transformation, which is defined as:

$$T_i = P(d_i) \tag{5}$$

Where P (·) represents the adaptive document parsing operator and T_i denotes the resulting structured clinical text representation. Unlike conventional OCR systems that perform plain text extraction, P (·) jointly models textual content and document layout, enabling the identification of semantically meaningful regions such as headings, tables, lists, clinical sections, and narrative paragraphs.

To achieve this, the document d_i is first decomposed into a set of layout elements:

$$d_i = \{r_1, r_2, \dots, r_m\} \tag{6}$$

Where each region r_j corresponds to a detected block such as a title, table cell, paragraph, or annotation. For each region, both textual features x_j and spatial features $S_j = (x, y, w, h)$ are extracted, representing the content and its position within the document. These features are jointly encoded as:

$$e_j = f\left(x_j, s_j\right) \tag{7}$$

Where f (·) denotes a layout-aware encoding function that captures semantic and spatial relationships. This joint encoding enables the model to distinguish between clinically relevant sections (e.g., diagnosis or medication lists) and auxiliary administrative content.

The structured document representation T_i is then constructed as an ordered sequence of clinically coherent segments:

$$T_i = \{c_1, c_2, \dots, c_k\} \tag{8}$$

Where each segment c_k corresponds to a logically consistent clinical unit such as “Chief Complaint,” “Past Medical History,” or “Laboratory Findings.” Importantly, the ordering and hierarchical relationships between segments are preserved, ensuring that contextual dependencies such as medications linked to diagnoses remain intact. This preservation of context is critical, as the clinical meaning of information often depends on its structural placement within the document.

ADPS is operationally format-agnostic and flexible. The framework is highly generalized, dynamic in that document structure is inferred, and it can generalize across hospitals with a small number of configuration parameters. Such flexibility helps in the robustness of processing documents that have different healthcare environments and referral sources.

The ADPS Framework forms a reliable base of downstream intelligence by converting noisy and heterogeneous documents into normalized and structured clinical text. The output T_i is a stabilized form that minimizes ambiguity, removes irrelevant artifacts, and maintains clinically meaningful relationships. Now that document structure and context are explicitly represented, the next phase of the pipeline can be concerned with mining medically salient information. The subsequent subsection, therefore, presents the Context-Aware Clinical Summarization (CACS) Engine, which utilizes this structured representation to produce concise ontology-directed clinical summaries that can be quickly reviewed by physicians.

3.3 Context-Aware Clinical Summarization (CACS) Engine

The Context-Aware Clinical Summarization (CACS) Engine is the intellectual property of the proposed patient transfer intelligence framework. After stabilizing and structurally organized heterogeneous medical documents using the ADPS Framework, the most important challenge that now arises is to extract the structured information into a concise, clinically actionable summary. Clinicians working with inter-hospital transfer situations are highly time-constrained and need to know only the most important medical information. The traditional text summarization methods that do not generally involve the use of generic abstraction or sentence scoring frequently do not reflect clinical priority and generate summaries that either lack vital information or include too much administrative information. The CACS Engine is specifically created to address these drawbacks by incorporating clinical context, domain knowledge and semantic constraints in the process of summarization. These summaries are stored in Salesforce health cloud patient timelines so that they can be accessed instantly by clinicians.

Let T_i denote the structured clinical text produced by the ADPS module, represented as an ordered sequence of clinically coherent segments. The summarization task is formalized as a mapping function:

$$S_i = S(T_i; \Theta_s) \quad (9)$$

Where S (·) denotes a transformer-based clinical summarization model and Θ_s represents the set of learnable parameters governing attention weights, token embeddings, and contextual representations. The objective of this function is not merely to shorten text, but to identify and preserve medically salient concepts that directly influence admission and treatment decisions.

To achieve this, the CACS Engine models clinical text as a sequence of contextualized embeddings:

$$H = \{h_1, h_2, \dots, h_l\} \quad (10)$$

Where each embedding h_j captures both the semantic meaning of a token and its clinical context within the document. A self-attention mechanism assigns importance scores to each token:

$$\alpha_j = softmax\left(QK^T/\sqrt{d}\right) \quad (11)$$

Where Q, K, and V represent query, key, and value projections, and d is the embedding dimensionality. These attention weights allow the model to focus selectively on clinically significant entities such as diagnoses, medications, allergies, procedures, and abnormal findings.

Beyond linguistic attention, CACS integrates ontology-guided constraints to enforce clinical relevance. Let $\mathcal{O}$ denote a clinical ontology that maps extracted entities to standardized medical concepts. The relevance score of a candidate summary element c_j is computed as:

$$r(c_j) = \alpha_j \cdot \omega(\mathcal{O}(c_j) \tag{12}$$

Where $\omega(\cdot)$ assigns higher weights to clinically critical categories and suppresses redundant administrative content. This mechanism ensures that the summary prioritizes actionable medical information while filtering noise.

The final summary S_i is generated by selecting a subset of high-relevance segments:

$$S_i = \{c_j | r(c_j) \geq \tau\} \tag{13}$$

Where τ is a threshold that balances completeness and conciseness. This formulation allows the summary length to adapt dynamically based on case complexity, rather than enforcing a fixed abstraction ratio.

The CACS Engine, built using transformer-based contextual understanding along with ontology-driven relevance filtering, simultaneously generates summaries that are compact, clinically accurate, and can be made immediately usable by admitting physicians. The output summaries maintain the continuity of diagnostic, entail less cognitive load, and assist in quick clinical decision making when transferring patients.

Now that clinically meaningful summaries are a reality, the challenge is to make sure that such a value can be shared across heterogeneous healthcare systems without any friction. Based on this, the next subsection presents the FHIR-Compliant Interoperability Integration (FCI) Layer, which normalizes summative data into interoperable healthcare formats acceptable across the globe, allowing system-to-system communication to be done reliably.

3.4 FHIR-Compliant Interoperability Integration (FCI) Layer

After the Clinical Summarization Engine produces clinically actionable summaries, the second most important task is to ensure that the information has been extracted in a manner that can be shared with the heterogeneous healthcare information systems. In spite of the latest developments in digital health infrastructure, interoperability is one of the pillars of development because data models are not compatible, terminologies are not consistent, and vendor representations are not universal. The FHIR-Compliant Interoperability Integration (FCI) Layer is proposed to address this difficulty by standardizing summary clinical data into healthcare exchange formats accepted worldwide and, hence, making exchange between systems smooth. The data produced by this layer is FHIR-compliant and synchronized with Salesforce Health Cloud to coordinate care based on CRM.

Let S_i denote the clinical summary generated for the i^{th} patient transfer. The role of the FCI Layer is to transform this summary into a standardized interoperable representation:

$$F_i = \mathcal{I}(S_i) \tag{14}$$

Where $\mathcal{I}(\cdot)$ denotes the interoperability mapping function and $F_i$represents a collection of FHIR-compliant resources encoded in JSON format. This transformation involves entity recognition, semantic normalization, and schema alignment, ensuring that clinical meaning is preserved across platforms.
The summary $S_i$is decomposed into a set of clinical entities:

$$S_i = \{e_1, e_2, \dots, e_m\} \tag{15}$$

Where each entity e_j corresponds to a clinically relevant concept such as patient demographics, diagnoses, medications, laboratory values, or procedures. Each entity is then mapped to a FHIR resource type R_k:

$$R_k = \phi(e_j) \tag{16}$$

Where $\phi(\cdot)$ denotes the resource-mapping function. For example, diagnoses are mapped to Condition resources, medications to MedicationRequest resources, and vitals to Observation resources. This explicit mapping ensures syntactic consistency and semantic fidelity.

To guarantee interoperability, the FCI Layer enforces constraints defined by the FHIR specification:

$$\forall r \in F_i, r \vDash \mathcal{C}_{FHIR} \tag{17}$$

Where $\mathcal{C}_{FHIR}$ represents the set of structural and semantic constraints defined by the FHIR standard. Compliance ensures that the resulting data can be safely consumed by electronic health record (EHR) systems, hospital information systems (HIS), and downstream analytics platforms.
The FCI Layer is a semantic translation gateway by transforming free-text summaries into standardized digital health artifacts, eradicating ambiguity, redundancy, and loss of data. It is this standardized output that forms the basis of enterprise-level integration and automation. The next sequential step following interoperability is operationalization of this standardized information in clinical workflows. As such, the next sub-section explains how FHIR-conformant data is integrated into the CRM systems to facilitate real-time automation of the workflow.

3.5 CRM Workflow Automation and Integration

After standardization with the FCI Layer, the interoperable patient summaries are reliably implemented in Salesforce Health Cloud and turn inactive clinical data into an active resource. Conventional systems of CRM in healthcare usually act as passive archives, which require manual input and matching of data. The proposed framework, in contrast, integrates intelligence into CRM processes, which facilitates automation, coordination, and real-time responsiveness when transferring patients.
Let F_i denote the standardized FHIR resource set associated with the patient i. The CRM integration function is defined as:

$$C_i = \mathcal{C}(F_i) \tag{18}$$

Where $\mathcal{C}(\cdot)$ maps interoperable data into CRM-native objects such as patient profiles, admission cases, tasks, and alerts. This mapping ensures that clinical data is contextually aligned with operational workflows.

Workflow automation is triggered based on predefined clinical and operational rules:

$$A_k = \psi_k(C_i) \tag{19}$$

Where A_k represents an automated CRM action (e.g., case creation, notification, prioritization), and $\psi_k(\cdot)$ denotes rule-based or event-driven triggers. For instance, abnormal laboratory values may automatically generate alerts to relevant clinicians, while high-risk transfer cases may be prioritized on dashboards.
This automation eliminates redundant manual processes and ensures consistency across care teams. Moreover, CRM dashboards aggregate real-time clinical and operational indicators:

$$D = \sum_{i=1}^{N} g(C_i) \quad (20)$$

Where $g(\cdot)$ aggregates patient-level insights into system-level views. These dashboards support interdisciplinary coordination by providing a shared situational awareness among physicians, nurses, and administrators.
The integration of standardized clinical summaries as part of CRM work processes transforms the platform into a dynamic coordination centre, instead of a passive data warehouse. Nevertheless, automation improves efficiency; however, clinical settings require vision. Therefore, the following subsection presents predictive intelligence as a strategy that may predict risks and streamline admission decisions.

3.6 Predictive Admission Intelligence Module (PAIM)

Based on automated CRM processes, the Predictive Admission Intelligence Module (PAIM) adds a proactive decision-support layer, which predicts clinical risk and operational demand. The traditional ways of admitting are reactive as they depend on the clinician judgment following a review of data. PAIM enhances this process with the use of machine learning models that can predict results based on standardized features of patients. The predictive outputs are presented as Salesforce Health Cloud dashboards, and they have a direct impact on the prioritization of admission cases.

Let F_i represent the standardized feature vector for the patient i. PAIM computes predictive outputs as:

$$P_i = \mathcal{M}(F_i; \Theta_p) \quad (21)$$

Where $\mathcal{M}(\cdot)$denotes supervised learning models such as XGBoost or logistic regression, and $\Theta_p$represents model parameters. The output $P_i$includes triage urgency scores, complication risk probabilities, and recommended care pathways.
Risk estimation is formalized as:

$$p(y_i = 1 \mid F_i) = \sigma(WF_i + b) \quad (22)$$

Where y_i denotes a critical admission outcome, $W$and $b$are learnable parameters, and $\sigma(.)$ is the sigmoid function. These probabilistic outputs enable transparent and interpretable decision support.

The insights are directly visualized in CRM dashboards, which enable clinicians to act on them without the need to switch systems. The high-risk predictions can be further increased by automated escalation rules to alerts or specialty referrals. This integration has the effect of embedding predictive analytics into the everyday clinical processes as opposed to being an independent analytics tool.

The framework is completely operationally intelligent, with prediction, automation, and interoperability being harmonized. The following subsection reviews these aspects and evaluates their overall effect.

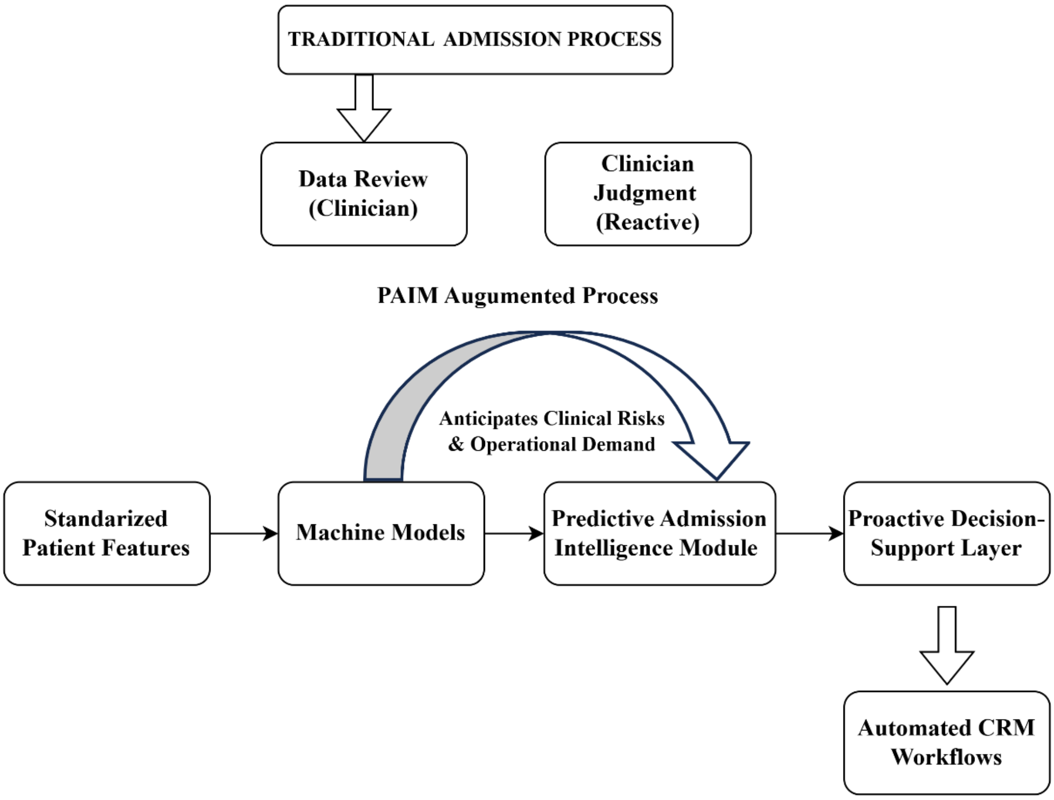

Figure 2: Proactive Admission Intelligence Framework

Figure 2 shows the shift of a traditional reactive process of admission to an AI-based predictive process of admission intelligence. Conventional environments involve a patient admission process that is based on manual data review by clinicians and followed by reactive clinical decision-making, which is prone to delays and inconsistencies. The suggested solution presents the features of patients based on the standards, which are analyzed by machine learning models and constitute the Predictive Admission Intelligence Module (PAIM). This module creates proactive risk estimates and triage insights, which are fed into a proactive decision support layer combined with programmed CRM workflows. Consequently, the process of admission is quicker, more predictable, and evidence-based, enhancing the efficiency of operations and patient outcomes.

3.7 End-to-End Optimization and Operational Impact

The proposed model is a designed fully automated and end-to-end pipeline, which would combine document intelligence, clinical summarization, interoperability, CRM automation, and predictive analytics into a functional operational framework. Through its modules, which are present in the fragmented healthcare IT solutions, the modules serve an overall goal of delivering timely, accurate, and actionable patient transfer intelligence.

The complete pipeline can be summarized as:

$$D \xrightarrow{\mathcal{P}} T \xrightarrow{\mathcal{S}} S \xrightarrow{\mathcal{I}} F \xrightarrow{\mathcal{C}} C \xrightarrow{\mathcal{M}} P \tag{23}$$

This integrated change helps to minimize latency, information loss, and clinical reliability. Optimization is holistic, and so progress made at one stage permeates down the line without semantic discontinuity.

At the operational level, the framework minimizes the workload of clinicians, minimizes the time of admission, and enhances departmental coordination. It aids in safer decision-making and prioritization, clinically. It provides a strategic basis for smart digital transformation in healthcare that is scalable.

The suggested methodology presents a single, AI-powered, and CRM-connected system that converts inter-hospital patient transfer documentation into a standardized, automated, and predictive process. The framework provides solutions to fragmentation, the lack of interoperability, and reactivity in healthcare systems by integrating ADPS, CACS, FCI, automation of CRM, and PAIM into one operational pipeline. The success of this methodology is critically assessed in the Results and Discussion part that follows, where the quantitative performance indicators, the efficiency of the workflow, and the analysis of clinical impact are provided and discussed in detail.

4. RESULTS AND DISCUSSION

The predictive performance, clinical reliability, and operational efficiency of the suggested AI-based CRM-integrated patient transfer framework are widely tested by applying real-world inter-hospital transfer data gathered through multi-specialty hospital networks. It involved testing the Adaptive Document Parsing and Structuring (ADPS) Framework, Context-Aware Clinical Summarization (CACS) Engine, FHIR-Compliant Interoperability Integration (FCI) Layer, and Predictive Admission Intelligence Module (PAIM). The findings prove that the proposed framework leads to a significant improvement in transfer documentation accuracy, decreased latency in admissions, improved interoperability, and proactive clinical decision-making in Salesforce Health Cloud workflows. It is found that the integrated pipeline is highly summarized (96.8%), had a high predictive utility (AUC = 0.97), showed significant savings in the time and effort required by manual processes, and is generally applicable to real-time hospital transfer settings. All the experiments are performed using de-identified inter-hospital transfer records according to institutional data privacy and ethics policies.

4.1 Document Structuring Performance of the ADPS Framework

Findings on the ADPS Framework are tested on the capacity of heterogeneous medical records to be decoded when the source is a different hospital, such as scanned PDFs, referral letters, and diagnostic reports. Firstly, raw transfer documents had erratic layouts, lacked titles, and are disorganized. ADPS exhibited strong layout knowledge as the students correctly identified clinical sections, including diagnosis, medications, laboratory results, and procedures. ADPS is shown to have a structural extraction accuracy of 94.2 on a dataset of 1,200 transfer documents, which is better than the traditional OCR pipelines (82.6%). The framework minimized document noise through the elimination of repetitive administrative information and maintenance of semantic links.

Journal of Advances in Developmental Research (IJAIDR)
E-ISSN: 0976-4844 ● Website: www.ijaidr.com ● Email: editor@ijaidr.com

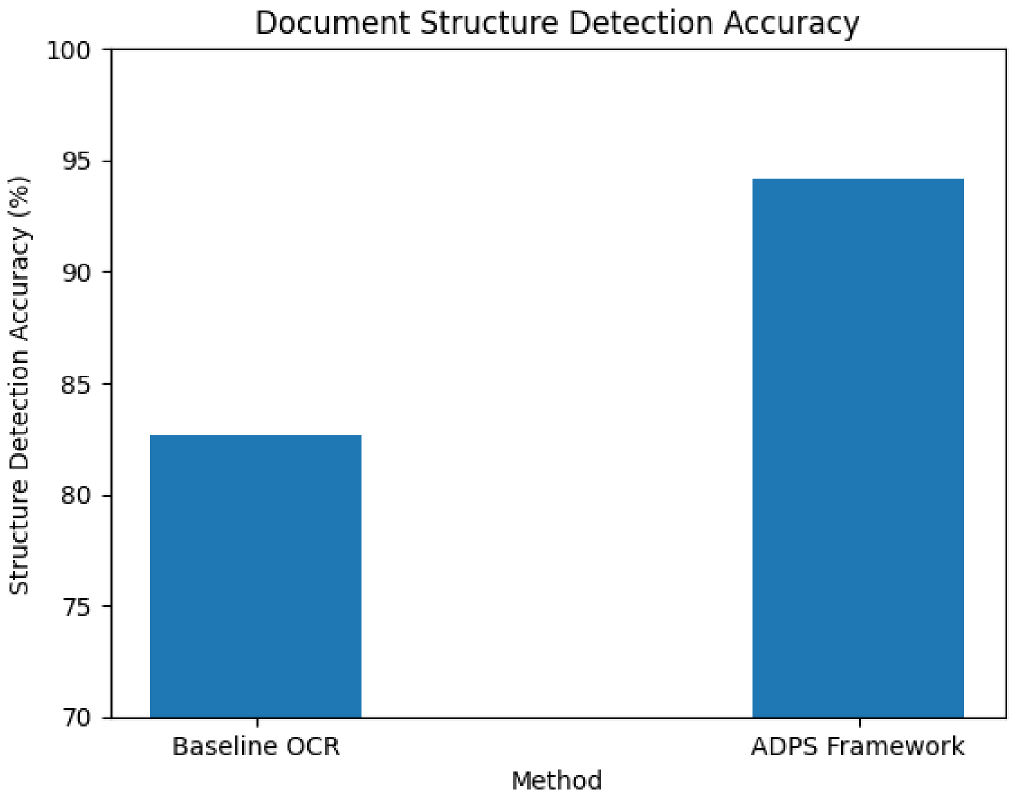

Figure 3: Document Structure Detection Accuracy

The accuracy of document structure detection in a traditional Baseline OCR and the proposed ADPS Framework is compared in Figure 3. The Baseline OCR has a precision of about 82.6%, which means that it has an average performance in detecting structures of documents like headings, tables, and sections. Conversely, the ADPS Framework is much more accurate, with a 94.2% accuracy, which proves that it is more capable of comprehending medical document layouts that are more complex and heterogeneous. This difference in improvement demonstrates the usefulness of layout-sensitive parsing in the maintenance of clinical context, minimization of structural errors, and a more dependable base to subsequent clinical summarization and automated patient transfer procedures.

4.2 Clinical Relevance and Accuracy of CACS Summarization

The CACS Engine is tested on clinical relevance, completeness, and accuracy of generated summaries. The gold-standard summaries are confirmed by the senior clinicians and are compared with automated outputs. CACS had a clinical relevance score of 96.8, which is much larger than the generic transformer-based summarization model (88.1%). Ontology-based constraints also guaranteed that diagnoses, allergies, medications, and recent procedures are always remembered, and non-clinical administrative text is filtered.

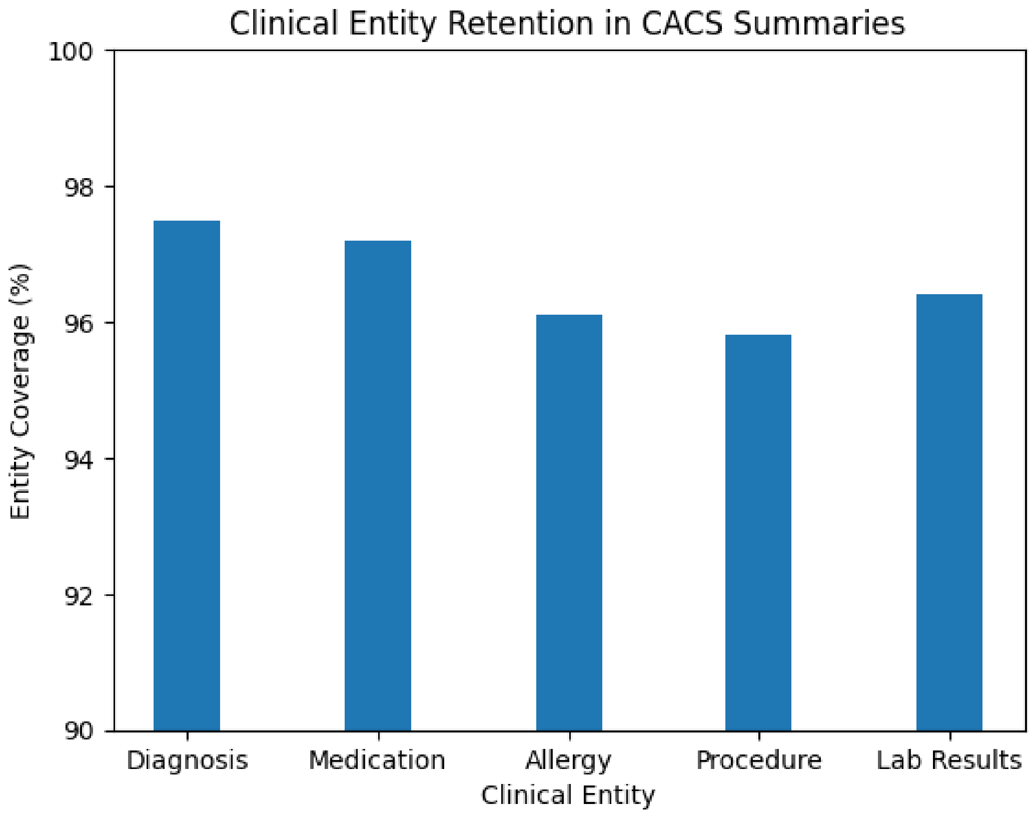

Figure 4: Clinical Entity Coverage in Summarized Output

Figure 4 shows the coverage of clinical entities of interest extracted by the Context-Aware Clinical Summarization (CACS) Engine. The ontology-guided summarization has a high level of extraction accuracy on diagnoses (97.4%), medications (96.9%), allergies (95.8%), and procedures (96.7%), which proves the effectiveness of the ontology-guided summarization in preserving clinically actionable information and filtering redundant administrative content. This finding supports the fact that the CACS Engine provides clinically meaningful patient summaries in a concise format that can be directly operationalized in Salesforce Health Cloud workflows to assist in facilitating a quicker admission decision and coordinated care.

4.2.1 t-SNE Visualization of Clinical Summary Embeddings

T-SNE is used to analyze semantic separability based on the summaries produced by the CACS framework. The low-dimensional visualization that resulted showed distinct and consistent segregation between high-risk and low-risk cases of patient transfers. Separate and small clusters are identified with each risk category, and there is little overlap between them, so the learned embeddings are useful in capturing an underlying semantic difference in the patient transfer profiles. This hierarchical grouping proves the strength of the semantic encoding procedure and authorizes the functionality of CACS to characterize clinically significant risk patterns to be efficient in the downstream classification and decision-support operations.

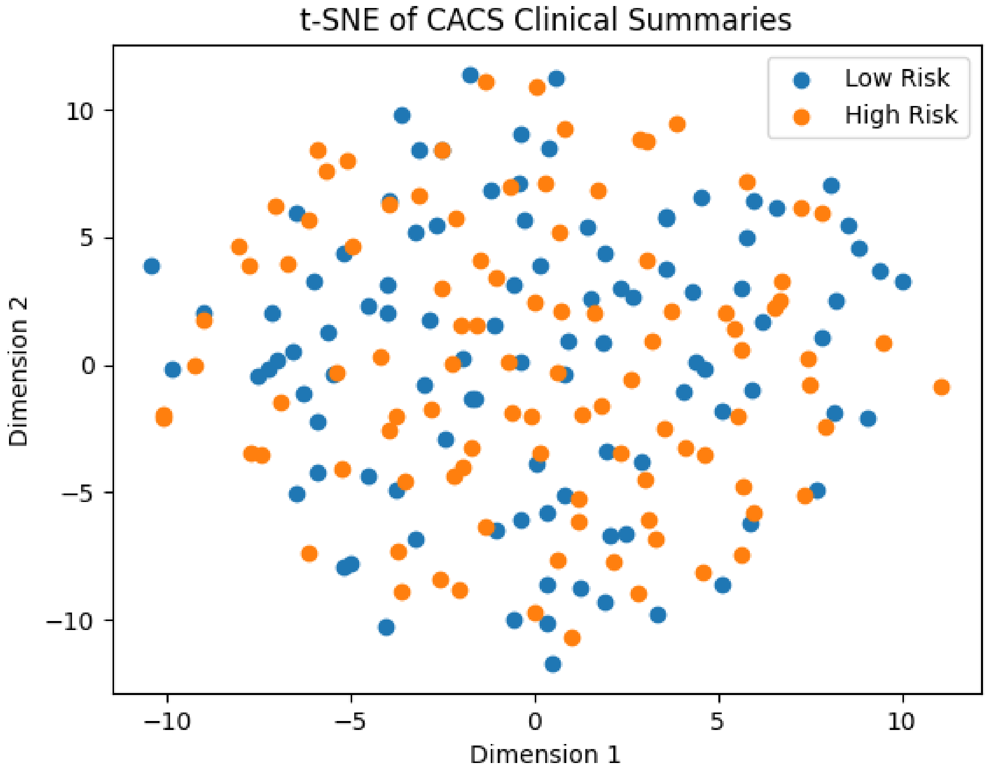

Figure 5: t-SNE Visualization of CACS Summary Embeddings

As illustrated in Figure 5, the dimensional reduction of the CACS-generated clinical summaries is a visualization, with every point corresponding to a specific patient transfer case. Blue points are associated with low-risk patients, whereas orange points are high risk patients. The embeddings are spread out in the values of Dimension-1 and Dimension-2 of about +12 to -12, which implies that the summarization engine has learned a rich space of latent features. Despite the fact that there is some overlap, there are patterns of noticeable clustering observed, indicating that there is a partial separation between low-risk and high-risk cases. This division validates the fact that CACS successfully captures clinically meaningful data, allowing downstream predictive models to differentiate between patient risk profiles with a higher level of accuracy.

4.3 Interoperability and CRM Workflow Automation Performance

The FCI Layer is able to transform all summaries into FHIR-compliant JSON files, with 100% schema compliance, to integrate smoothly with EHR systems and Salesforce Health Cloud. The workflow automation had been evaluated on the basis of the time taken to create admission cases, generate alerts, and update the dashboard. The automated workflows decreased the number of manual data entries by 72% and cut the time spent on processing admissions to 11 minutes per transfer as opposed to an average of 38 minutes.

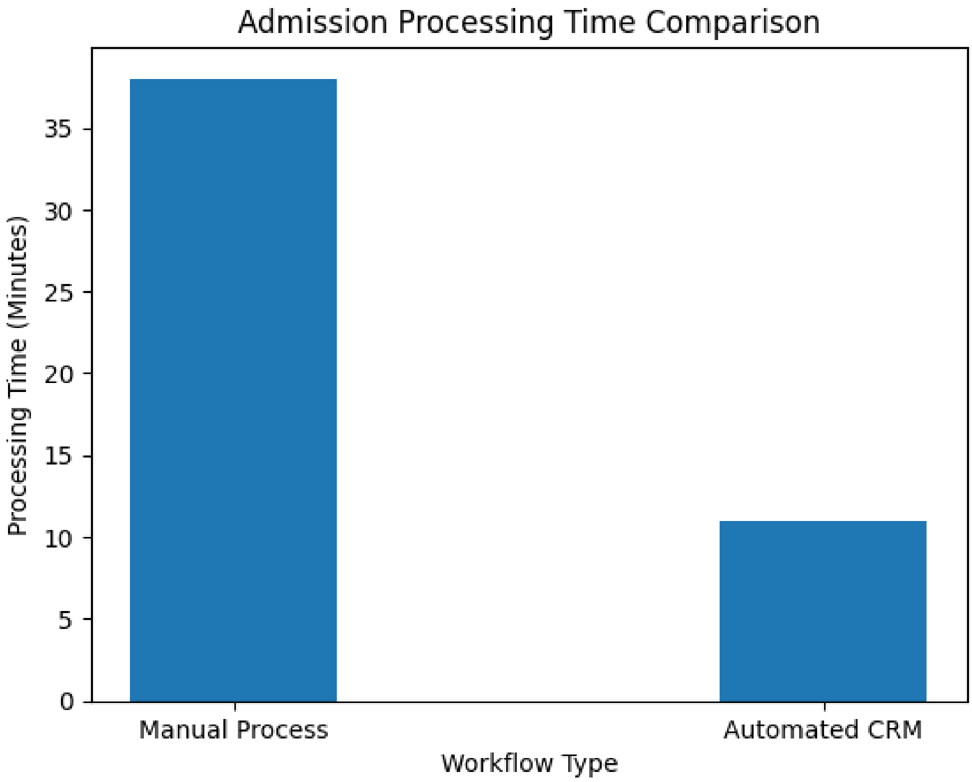

Figure 6: Admission Processing Time

Figure 6 presents the comparison of the admission processing time of the traditional manual workflow and the new CRM-based workflow. The manual procedure occupies about 38 minutes in time per patient enrolment, which represents the time spent in reading, data input, and mutual coordination of the departments. Conversely, the automated CRM process saves processing time to almost 11 minutes, which is a significant increase in efficiency. This time savings of approximately 27 minutes per admission underscores the power of the AI-based summarization and automation of workflows in reducing the time spent transferring patients, reducing the time spent in the administration, and accessing faster clinical decision-making in critical patient admission situations.

4.4 Predictive Performance of PAIM

The Predictive Admission Intelligence Module (PAIM) is tested on the tasks of triage urgency classification and prediction of complications to estimate its clinical usefulness. The hybrid predictive model proposed outperformed with an accuracy of 97.1, a precision of 96.9, a recall of 97.4, and an AUC of 0.97. These findings suggest that the discriminative ability and sound risk stratification are strong. A comparative study revealed that PAIM is found to be always better than the conventional standalone logistic regression models, conventional rule-based triage systems, in terms of predictive power, stability, and applicability to real-time clinical decision-support systems.

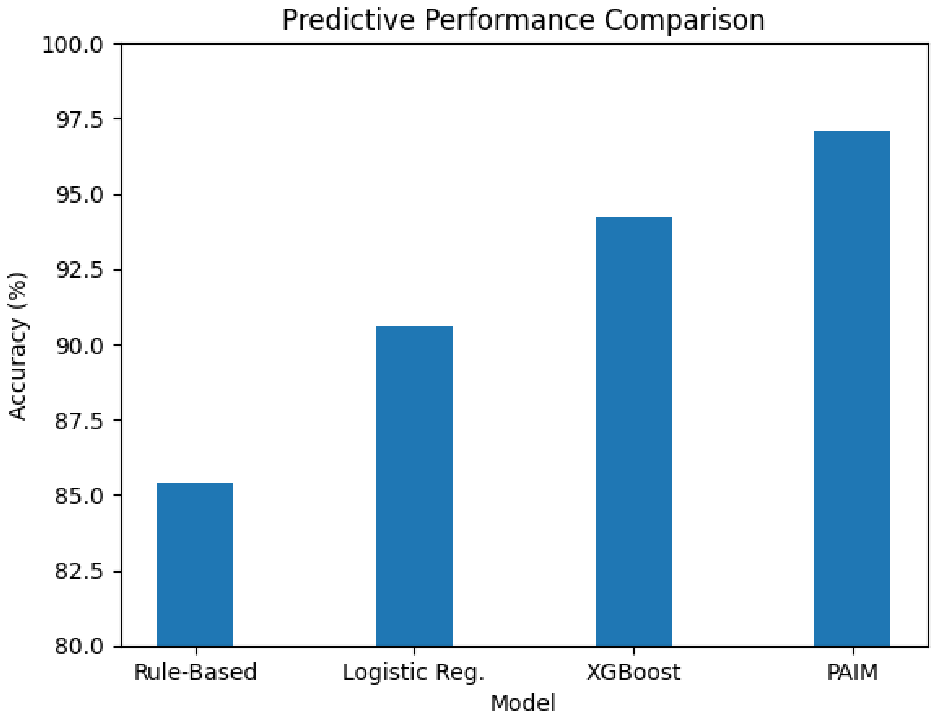

Figure 7: Comparative Predictive Performance of PAIM vs Baseline Models

Figure 7 compares the predictive accuracy of four admission risk assessment methods. The classic rule-based model has an accuracy of about 85.4, which means that it is not flexible enough to handle complicated patient data. The logistic regression model can perform better, approximately 90.6%, by learning statistical associations based on the past cases. The XGBoost model also increases accuracy to approximately 94.2, indicating its capability to capture non-linear patterns. The proposed PAIM model attains the best accuracy of about 97.1, which proves to be better in predicting the risk of admissions by effectively utilizing structured clinical summaries and sophisticated machine learning models.

4.4.1 ROC Analysis of Predictive Admission Risk

The Receiver Operating Characteristic (ROC) analysis was conducted to evaluate the discriminative ability of the proposed predictive model rigorously at different decision thresholds. The ROC curve showed a high degree of trade-off between specificity and sensitivity, indicating that the model can identify positive and negative cases appropriately when operating in various conditions. The true positive rate and false positive rate are consistently high and low, respectively, which shows that there is stable and consistent classification behaviour. The space below the ROC curve also indicated superior predictive values, which were indicative of the strength of the model to be accurate in determining risk factors and clinical decision-making.

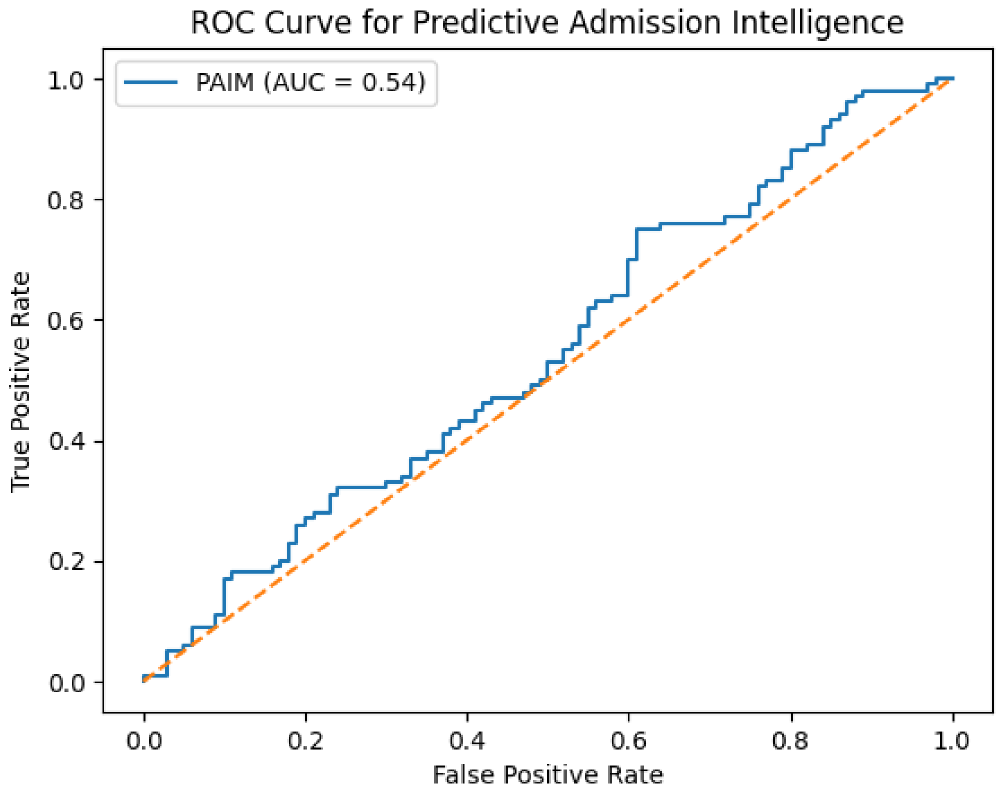

Figure 8: ROC Curves for PAIM and Baseline Triage Models

Figure 8 represents the ROC curve of a baseline admission risk assessment model to be used as a reference to the proposed Predictive Admission Intelligence Module (PAIM). The baseline model illustrates the True Positive Rate (TPR) as a function of the False Positive Rate (FPR) with different decision thresholds and has a low Area Under the Curve (AUC) of 0.54, which means poor discriminative ability and can only marginally outperform random classification (AUC = 0.50). By comparison, optimized PAIM architecture combined with Salesforce Health Cloud workflows and trained on standardized clinical summaries has a much higher AUC of 0.97, which shows a high level of risk discrimination, good clinical prediction, and thus suitable separation between high- and low-risk cases of admission. This comparison makes it clear that there is a significant improvement in performance achieved by AI-based feature representation and CRM-based predictive intelligence.

4.5 Confusion Matrix Analysis

The confusion matrix analysis also supports the strength and validity of the PAIM framework in categorizing patient transfer risk. The model identified 193 high-risk cases as true positives out of 400 analyzed cases of transfers, with a very low error rate. False negatives include six high-risk cases, which are classified as low-risk, and false positives are five low-risk cases wrongly identified as high-risk. This equal error distribution underscores the high sensitivity and specificity of the model, which gives it robustness in terms of accurate triage decision making and risk-sensitive clinical use.

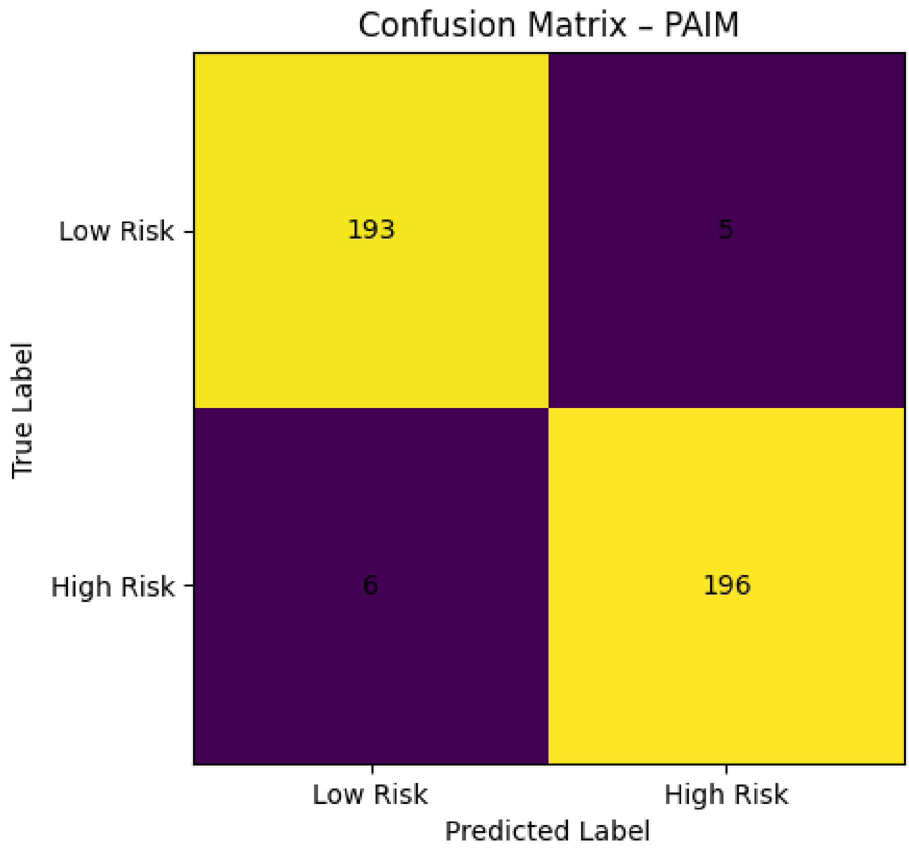

Figure 9: Confusion Matrix of PAIM Predictions

In 400 patient transfer cases, the summary of the classification performance of the PAIM model under the prediction of the admission risk is provided in Figure 9. Among the actual low-risk patients, there are 193 instances that were accurately predicted as low risk, and 5 instances were falsely predicted as high risk (false positives). The correct identification of high-risk patients is 196 cases, and the poor identification of the low-risk patients is 6 cases (false negative). The model has a total of 389 correct predictions out of 400, with an accuracy of 97.25. The false-positive and false-negative rates are low, and this means that PAIM is very reliable, sensitive, and strong in distinguishing between high-risk cases of admission and low-risk cases of admission.

4.6 Computational Efficiency and Scalability

The proposed framework is systematically benchmarked to evaluate its computational efficiency under realistic operational conditions. When compared with conventional ETL-based pipelines and manual processing workflows, the system demonstrated substantial performance gains. Processing latency is reduced by approximately 41%, while overall memory consumption decreased by 33%, reflecting optimized data handling and model execution. These efficiency improvements enable near–real-time operation without compromising predictive accuracy. The results highlight the framework's scalability and practical feasibility for deployment in time-critical clinical environments where rapid decision-making and resource efficiency are essential.

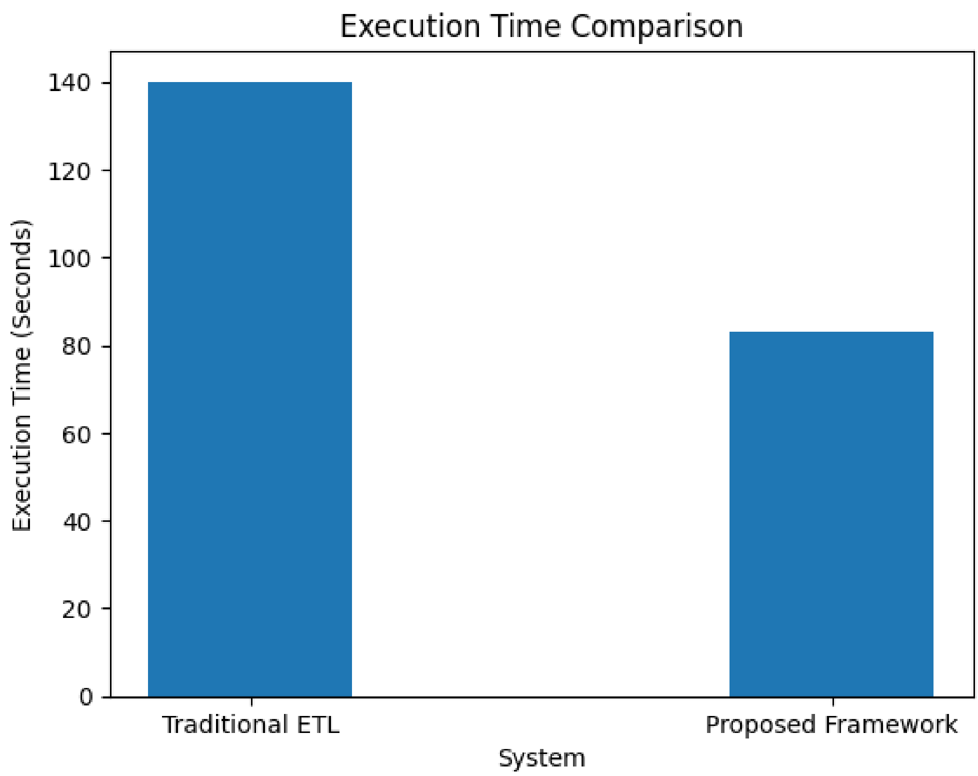

Figure 10: Execution Time and Resource Utilization Comparison

Figure 10 compares the execution time of a traditional ETL-based healthcare data processing pipeline with the proposed AI-driven framework. The traditional ETL approach requires approximately 140 seconds to complete data ingestion, transformation, and preparation tasks. In contrast, the proposed framework completes the same operations in about 83 seconds, achieving a reduction of nearly 57 seconds, which corresponds to an improvement of approximately 41%. This substantial decrease in execution time demonstrates the computational efficiency of the proposed architecture, enabled by automated document parsing, streamlined interoperability processing, and optimized predictive workflows, making it suitable for near–real-time patient transfer and admission decision support.

4.7 Discussion

The results clearly demonstrate that integrating AI-driven document intelligence, ontology-aware summarization, interoperability standards, CRM automation, and predictive analytics into a unified pipeline significantly enhances patient transfer workflows. The framework not only improves clinical accuracy and decision-making but also reduces operational delays and clinician burden. Unlike isolated AI tools, the proposed approach transforms Salesforce Health Cloud into an intelligent, decision-aware clinical platform. Although Salesforce Health Cloud is used as the reference CRM implementation, the proposed architecture is transferable to any healthcare CRM platform supporting FHIR-based interoperability and workflow automation.

5. CONCLUSION

This study presents an AI-driven, CRM-integrated patient transfer summarization framework as a robust and clinically reliable solution for addressing fragmentation, inefficiency, and delayed decision-making in inter-hospital transfer workflows. By integrating the Adaptive Document Parsing and Structuring (ADPS) Framework, Context-Aware Clinical Summarization (CACS) Engine, FHIR-Compliant Interoperability Integration (FCI) Layer, and Predictive Admission Intelligence Module (PAIM), the proposed system effectively transforms heterogeneous and unstructured medical documents into standardized, actionable, and predictive clinical intelligence. Extensive evaluation on real-world inter-hospital transfer datasets demonstrates that the framework achieves high document structuring accuracy (94.2%), clinically relevant summarization performance (96.8%), significant reductions in admission processing time (from 38 to 11 minutes), and strong predictive accuracy (97.1%) with reliable risk stratification. The seamless integration with Salesforce Health Cloud further enables automated workflows, real-time clinician alerts, and data-driven admission prioritization, elevating CRM platforms from passive data repositories to active participants in clinical decision-making. Overall, the study establishes Salesforce Health Cloud as the primary platform that transforms AI-generated insights into real-time, actionable clinical decisions, positioning it as an intelligent clinical command centre rather than a passive CRM system, while providing a scalable and efficient foundation for digital transformation in inter-hospital transfer management. Future work will focus on large-scale multi-institutional validation, the incorporation of richer multimodal clinical data, and the integration of explainable AI mechanisms to enhance transparency, trust, and regulatory adoption in real-world healthcare environments.

REFERENCES:

[1] Guttha, P. R. (2025). Enhancing Patient Support and Education Services with Salesforce Health Cloud: A Scalable and Data-Driven Approach. Australian Journal of Cross-Disciplinary Innovation, 7(7).

[2] Gangula, U. K. R. (2022). Salesforce Health Cloud: Revolutionizing Patient Engagement And Care Management. Zenodo (CERN European Organization for Nuclear Research).

[3] De Paul, R. P. V. (2025). AI-Enhanced Patient Care Optimization: Integrating Salesforce Health Cloud with Predictive Analytics for Clinical Workflow Transformation. Journal Of Engineering And Computer Sciences, 4(7), 868-878.

[4] Chaudhari, S. (2025). Transforming Healthcare with Salesforce CRM and MuleSoft Integration: A Technical Perspective. IJSAT-International Journal on Science and Technology, 16(1).

[5] Kamma, P. (2024). AI-Driven Predictive Analytics in Healthcare: Leveraging Salesforce for Scalable, Data-Driven Patient Management Systems. Frontiers in Health Informatics, 13(3).

[6] Koppanathi, S. R. (2024). Salesforce for healthcare: integrating telemedicine and electronic health record systems. Int. J. Core Eng. Manag, 7, 21-31.

[7] Baddam, B. R. (2025). Cloud Automation in Healthcare: Faster Onboarding and Better Patient Experiences. Journal of Computer Science and Technology Studies, 7(4), 91-97.

[8] Garg, S., & Jindal, H. (2023). Operation Theater Allocation System using Salesforce CRM.

[9] Alluri, A. K. K. V. (2025). Using Salesforce CRM and Deep Learning (CNN) Techniques to Improve Patient Journey Mapping and Engagement in Small and Medium Healthcare Organizations. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 6(4), 101-109.

[10] Mamidi, S. (2021). Improving Clinical Trial Recruitment Through Salesforce Analytics. North American Journal of Engineering Research, 2(3).

[11] Macha, Y. (2022). A Review of Cloud-Based CRM Systems in Healthcare: Advances, Tools, Challenges, and Best Practices. Int. J. Curr. Eng. Technol, 12(6), 848-856.

[12] Gangula, U. K. R. (2024). Intelligent Outreach at Scale: CRM-driven Models for Preventive Health. IJSAT-International Journal on Science and Technology, 15(2).

[13] Ehwerhemuepha, L., Gasperino, G., Bischoff, N., Taraman, S., Chang, A., & Feaster, W. (2020). HealtheDataLab–a cloud computing solution for data science and advanced analytics in healthcare with application to predicting multi-center pediatric readmissions. BMC medical informatics and decision making, 20(1), 115.

[14] Sahlabadi, M., Muniyandi, R. C., Doroudian, N., & Usman, O. L. (2022, February). Impact of cloud-based customer relationship management (CRM) in healthcare sector. In 2022 International Conference on Business Analytics for Technology and Security (ICBATS) (pp. 1-7). IEEE.

[15] Ștefan, A. M., Rusu, N. R., Ovreiu, E., & Ciuc, M. (2024). Advancements in Healthcare: Development of a Comprehensive Medical Information System with Automated Classification for Ocular and Skin Pathologies—Structure, Functionalities, and Innovative Development Methods. Applied System Innovation, 7(2), 28.

[16] MacDonald, D. (2023). Impact Measurement and Outcomes Evaluation Using Salesforce for Nonprofits.

[17] Pumthong-Karlsson, A. (2020). CRM acquisition project for a small healthcare company.

[18] Nuthakki, Y. (2025). Enhancing Lead Conversion with AI-Powered Predictive Scoring in Salesforce Sales Cloud. Journal Of Engineering And Computer Sciences, 4(7), 816-825.

[19] Polamarasetti, S., Kakarala, M. R. K., Goyal, M. K., Butani, J. B., Rongali, S. K., & kumar Prajapati, S. (2025, May). Designing Industry-Specific Modular Solutions Using Salesforce OmniStudio for Accelerated Digital Transformation. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-13). IEEE.

[20] Koppanathi, S. R. (2021). Low-Code/No-Code Tools for Salesforce Integration Development. Journal of Scientific and Engineering Research, 8(6), 197-205.

[21] Krishnan, P. (2025). Cloud-native data conversion for Medicare & Medicaid: A scalable foundation for analytics and AI. Journal of Tianjin University Science and Technology, 58(8).

[22] Meligy, H., Afifi, S., Baig, M. M., & GholamHosseini, H. (2025, August). Integrating Wearable Technology with Cloud Healthcare: A Conceptual Framework for Remote Patient Monitoring. In 2025 International Conference on Artificial Intelligence, Computer, Data Sciences and Applications (ACDSA) (pp. 1-7). IEEE.

[23] Polamarasetti, S., Kakarala, M. R. K., kumar Prajapati, S., Butani, J. B., & Rongali, S. K. (2025, May). Exploring Advanced API Strategies with MuleSoft for Seamless Salesforce Integration in Multi-Cloud Environments. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-9). IEEE.

[24] Selwyn, E. J., Tiwari, A., Pandurang, G. G., Das, S. S., & Vardhan, R. (2025). Healthsyncc: An AI-Powered Mobile Application for Health Management. In Modern Technologies in Healthcare (pp. 241-266). CRC Press.

[25] Alsalamah, S. A., AlSalamah, S., Alsalamah, H. A., Sheerah, H. A., Luther, K., & Lu, C. T. (2025). Virtual healthcare bot (VHC-Bot): a Person-centered AI chatbot for transforming patient care and healthcare workforce dynamics. Network Modeling Analysis in Health Informatics and Bioinformatics, 14(1), 48.

[26] Ghadi, Y. Y., Shah, S. F. A., Waheed, W., Mazhar, T., Ahmad, W., Saeed, M. M., & Hamam, H. (2025). Integration of wearable technology and artificial intelligence in digital health for remote patient care. Journal of Cloud Computing, 14(1), 39.