Real-Time Patient Data and Documentation Access Through Mobile Application: A Secure and Efficient Solution for Hospital Clinicians

Matthew Schembri

Supervisor: Mary Grace Seguna Agius

September - 2025

A dissertation submitted to the Institute of Information and Communication Technology in partial fulfilment of the requirements for the degree of BSc (Hons) in Software Development

Authorship Statement

This dissertation is based on the results of research carried out by myself, is my own composition, and has not been previously presented for any other certified or uncertified qualification.

The research was carried out under the supervision of Ms. Mary Grace Seguna Agius

. . . .

Date   Signature

Copyright Statement

. . . .
Date

. . . .
Signature

Acknowledgements

First and foremost, I would like to extend my sincere gratitude to the instructors and mentors at MCAST, whose guidance and knowledge served as the cornerstone of this dissertation. Special thanks go to Mr. Mark Bruno Olivier at Cursor Ltd and to the staff at the Ministry of Health and Active Ageing for their assistance in providing essential resources. This dissertation would not have been possible without the invaluable contributions of the doctors who willingly contributed their time and expertise during the interview process.

On a personal note, I wish to express my deepest gratitude to my family and friends for their constant encouragement. I am especially thankful to my parents, Joseph and Mary-Ann, for instilling in me a passion for learning and for their steadfast belief in my potential. My brother, Luke, has been a valuable source of support throughout this journey. Finally, I want to offer my heartfelt thanks to my partner, Jeannette, for her remarkable patience and unwavering encouragement. Without their support, this achievement would not have been possible.

Abstract

In today's hospital settings, clinicians often face challenges due to fragmented data systems and inefficient workflows, leading to delays in decision-making and potential risks to patient safety. The absence of an unified, real-time perspective on patient data constitutes a notable and enduring obstacle in healthcare information technology. This dissertation examines the issue by designing, developing and rigorously evaluating a novel mobile application intended as a clinical assistive tool. This research centres on a software artefact that serves as a mobile dashboard, granting clinicians immediate access to patient data. This artefact investigates the application of advanced AI through the integration of a feature that utilises a Large Language Model (MedLLaMA2) alongside the SNOMED CT ontology to produce diagnostic suggestions from unstructured clinical notes.

To guide this practice-oriented research, a Design Science Research (DSR) methodology was adopted, supported by a mixed-methods approach to evaluation. The technical performance of the AI-driven diagnostic suggestion algorithm was assessed through several iterative cycles of quantitative evaluation against a synthetic dataset of 250 clinical cases. Concurrently, a qualitative inquiry was conducted through semi-structured interviews with three practising clinicians to ground the research in the realities of clinical practice.

The quantitative evaluation yielded a critical finding: the diagnostic algorithm, despite multiple architectural enhancements, consistently failed to achieve clinical viability. The final performance metrics of low precision (0.05) and high recall (0.40) revealed a "heuristic ceiling," demonstrating that a loosely-coupled architecture connecting a general-purpose LLM to a knowledge base is insufficient for this safety-critical task.

The qualitative findings offered a significant explanatory framework for the technical results. Clinicians expressed unanimous enthusiasm for the core concept of a centralised patient data dashboard, thereby validating the initial problem statement. They expressed scepticism regarding the AI's diagnostic capabilities, noting its absence of clinical context and the risk of "suggestion overload." their expert intuition that an AI lacking deep contextual understanding would generate "noise" was empirically confirmed by the algorithm's elevated rate of false positives.

This study concludes that the development of such a tool is not merely a technical challenge, but a socio-technical one, where clinician trust, workflow integration and contextual awareness are as important as algorithmic performance. The primary contribution is the robust evidence that simplistic AI pipeline architectures are inadequate for complex clinical reasoning. The research culminates in a set of evidence-based design principles for the final artefact, which firmly prioritises the flawless delivery of a reliable, user-friendly patient data dashboard while positioning its AI capabilities as a transparent, supportive tool for clinical brainstorming.

Keywords: Clinicial Decision Support System (CDSS), Design Science Research (DSR), Mobile Health, SNOMED CT, Large Language Model (LLM), Mixed-Methods Research, Healthcare Interoperability, Artificial Intelligence (AI)

Table of Contents

Authorship Statement **i**

Copyright Statement **ii**

Acknowledgements **iii**

Abstract **iv**

List of Figures **ix**

List of Tables **x**

List of Abbreviations **xi**

1 Introduction **1**
- 1.1 Background and Problem Statement 1
- 1.2 Research Framework and Design 1
- 1.3 Purpose Statement, Research Questions and Objectives 3
  - 1.3.1 Purpose Statement 3
  - 1.3.2 Research Questions 3
  - 1.3.3 Research Objectives 4
- 1.4 Significance of the Study 4
- 1.5 Research Structure 5

2 Literature Review **6**
- 2.1 Data Integration and Interoperability Solutions 6
  - 2.1.1 The Challenge of Fragmented Systems 6
  - 2.1.2 Interoperability Standards and Data Exchange Protocols . . . 9
  - 2.1.3 Data Platforms, Knowledge Graphs and Embeddings 11
- 2.2 Remote/Real-Time Patient Monitoring (IoT & AI) 13
  - 2.2.1 Chronic Disease Management and Ageing Populations . . . 15
  - 2.2.2 Implementation Barriers and Data Security 17
  - 2.2.3 Prospects for Integration with EHR and Clinical Workflows 18
- 2.3 Clinical Decision Support Systems 20
  - 2.3.1 Medication Safety and Prescribing Accuracy 20
  - 2.3.2 Diagnostic Support and Risk Stratification 22
  - 2.3.3 ePrescriptions and Medication Safety 24
- 2.4 Synthesis and Research Gap 25

3 Research Methodology **28**
3.1 Methodological Framework 29
3.2 System Design and Development Methodology 30
3.2.1 Artefact Design and Architecture 30
3.2.2 Development Tools, Technologies and Justification 32
3.2.3 Evaluation 33
3.3 Data and Evaluation Methodology 33
3.3.1 Data Collection and Preparation 33
3.3.2 Qualitative Data: Clinician Interviews 34
3.3.3 Test Data Generation: Synthetic Patient Case Notes 34
3.4 Evaluation Strategy 35
3.4.1 Phase 1: Artefact Technical Validation 35
3.4.2 Phase 2: Gathering and Interpreting Results 36
3.4.3 Phase 3: Ensuring Validity of Results 37
3.5 Ethical Considerations 37

4 Analysis of Results and Discussion **39**
4.1 Quantitative Findings: Performance of the Diagnostic Suggestion Algorithm 39
4.1.1 Cycle 1: Baseline Performance of the Hybrid Model 40
4.1.2 Cycle 2: The Precision-Recall Trade-Off in LLM Prompting 41
4.1.3 Final Iteration: The 'Heuristic Ceiling' and Final Performance 43
4.2 Qualitative Findings: The Clinical Context for a Mobile Assistive Tool 44
4.2.1 Theme 1: A System Under Strain - Inefficiency, Fragmentation and Risk 45
4.2.2 Theme 2: Unanimous Need for a Centralized, Real-Time Dashboard 45
4.2.3 Theme 3: AI as a Cautious Assistant, Not a Confident Authority 46
4.2.4 Theme 4: The Non-Negotiable Requirements - Speed, Reliability and Usability 48
4.3 Discussion: Synthesising Quantitative and Qualitative Findings . . . 49
4.3.1 Triangulating the Data: Why the Algorithm's Performance is a Key Finding 50
4.3.2 Findings in Relation to the Research Questions 51
4.3.3 Limitations and Acknowledgement of the Broader AI Landscape 53
4.3.4 Implications for the Design of the Final Artefact 55

5 Conclusions and Recommendations **57**
5.1 Summary of Research 57
5.2 Principal Findings of the Study 58
5.3 Contribution to Knowledge 60
5.4 Recommendations for Future Research 61
5.5 Concluding Summary 62

List of References **63**

Appendix A Interview Protocol **68**

Appendix B Participant Information Sheet **71**

List of Figures

1.1 Research Onion [1] 2

2.1 Key Barriers to Healthcare Interoperability [2] 8
2.2 The Vision of an Integrated Healthcare Ecosystem [3] 12
2.3 Illustration of SNOMED CT concept relationships in a knowledge graph 13
2.4 artificial intelligence-enabled remote patient monitoring Architectures [4] 15
2.5 The impact of CDSS implementation on patient outcomes [2] . . . 16

3.1 Research Pipeline 29
3.2 Architecture Diagram 31

4.1 Cycle 1 Confusion Matrix of baseline hybrid model. 41
4.2 Cycle 2 Confusion Matrix of revised prompting model. 42
4.3 Confusion Matrix of the final model after re-ranking. 44

List of Tables

2.1 Performance of a machine learning-based CDSS for medication error detection [5] 21

3.1 System Specifications 33

4.1 Evaluation metrics for the baseline hybrid model (Cycle 1). 40
4.2 Evaluation metrics for the revised model (Cycle 2). 41
4.3 Evaluation metrics for the final model with re-ranking (Final Iteration). 43

List of Abbreviations

ADE	Adverse Drug Events
AI	Artificial Intelligence
CDA	Clinical Document Architecture
CDSS	Clinical Decision Support System
CI\CD	Continuous Integration/Continuous Delivery
COPD	Chronic Obstructive Pulmonary Disease
DICOM	Digital Imaging and Communications in Medicine
DSR	Design Science Research
EHR	Electronic Health Records
EMR	Electronic Medical Record
ESB	Enterprise Service Buses
FHIR	Fast Healthcare Interoperability Resources
GDPR	EU General Data Protection Regulation
GPU	Graphics Processing Unit
HIE	Health Information Exchange
HIPAA	Health Insurance Portability and Accountability Act
HL7	Health Level 7
ICU	Intensive Care Unit
IoT	Internet of Things
iPaaS	Integration Platform as a Service
IT	Information Technology
LLM	Large Language Model
MTTR	Mean Time To Repair
NEWS2	National Early Warning Score
SMART	Specific, Measurable, Achievable, Realistic, and Time-bound
SOFA	Sequential Organ Failure Assessment

Chapter 1: Introduction

1.1 Background and Problem Statement

In the current healthcare landscape, clinicians have significant challenges in obtaining timely, comprehensive and precise patient information. The healthcare IT landscape is frequently a fragmented ecosystem of diverse electronic health records (EHRs), laboratory systems and imaging storage devices that lack adequate communication among them. The absence of interoperability hinders continuity of care, delays clinical decision-making and increases the risk of medical errors. The operational and financial challenges of sustaining these siloed systems are considerable, frequently obstructing innovation and consuming resources in the management of legacy infrastructures. This study tackles the pressing requirement for a cohesive secure and efficient solution that grants clinicians immediate access to inpatient data at the point of care, thereby enhancing clinical workflows and patient safety.

1.2 Research Framework and Design

This study is systematically structured according to the established "Research Onion" paradigm defined by Saunders et al. [1]. This paradigm ensures uniformity and justification across all layers of the research design from philosophy to data collection. A graphical representation of this framework, emphasising the principal methodological selections for this study is illustrated in Figure 1.1

The outermost layer, research philosophy, is based on pragmatism selected for its emphasis on practical results in addressing real-world issues. The subsequent layer, the research approach is deductive employing existing theories to inform the formulation of a specific solution. The foundation of the framework is the research strategy which employs a mixed-methods approach.

To implement this mixed-methods approach, the study will utilise a convergent design, as outlined by Creswell [6]. This entails two simultaneous work-streams: (i) a technical stream dedicated to the design and (ii) development of the prototype and a qualitative stream aimed at collecting clinician input via interviews. Chapter 3 will comprehensively outline the methods for both streams.

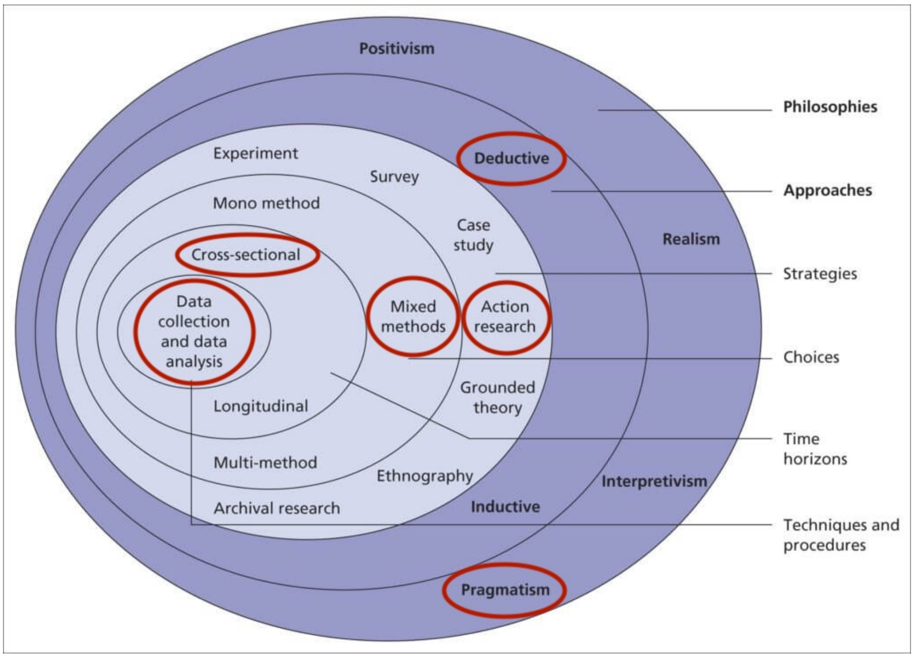

Figure 1.1: *Research Onion [1]*

1.3 Purpose Statement, Research Questions and Objectives

1.3.1 Purpose Statement

The aim of this study is to design, develop and evaluate a mobile medical dashboard for hospital clinicians to address the critical need for optimised real-time access to patient data for hospital clinicians. This project utilises a mixed-methods approach, integrating design science research for the development of a functioning prototype with qualitative inquiry to understand the intricate requirements of clinicians. The primary objective is to develop and validate a user-centred, secure and efficient solution that improves clinical workflows and contributes to better patient care.

1.3.2 Research Questions

RQ1: What are the essential features and workflow improvements that clinicians require from a mobile dashboard to enhance real-time patient data access and documentation?

RQ2: What are the key privacy and security concerns that clinicians have regarding the use of mobile applications in managing sensitive patient data and how can these concerns be overcome?

RQ3: In what ways does real-time patient monitoring through the dashboard influence the speed and effectiveness of clinicians' responses to critical patient events?

1.3.3 Research Objectives

To achieve the aims of this study, a series of structured objectives have been established. The research will begin by critically reviewing the literature on healthcare data management, mobile clinical applications and pertinent standards such as Systematised Nomenclature of Medicine Clinical Terms (SNOMED CT). Subsequently, the project will proceed to design the system architecture and user interface for the mobile dashboard. The core of the practical work will involve the development of a functional artefact utilising react naive, which will integrate features for real-time data access and an AI-driven diagnostic search. Concurrently, the study will conduct semi-structured interviews with clinicians to collect qualitative data regarding their workflow needs. Finally, the research will analyse the qualitative data through thematic analysis, assess the artefact's performance and formulate a series of evidence-based recommendations for future development.

1.4 Significance of the Study

This study offers a significant practical contribution to the field of healthcare technology by developing a user-centred framework for a mobile clinical dashboard. By directly integrating clinician feedback into an iterative design science research process, the research provides a validated model for a system capable of enhancing workflow efficiency, reducing the risk of medical errors, and improving the speed of clinical decision-making. The findings will provide valuable, evidence-based insights for healthcare organisations, software developers and researchers seeking to implement effective and user-accepted mobile solutions in complex inpatient environments.

1.5 Research Structure

This dissertation is structured into five chapters. Chapter 1 has introduced the research problem, the guiding frameworks and the study's objectives. Chapter 2 will provide a comprehensive review of the relevant literature. Chapter 3 will detail the research methodology, outlining the specific procedures for data collection and analysis. Chapter 4 will present and discuss the findings from both the prototype evaluation and the qualitative interviews. Finally, Chapter 5 will conclude the study, summarising the key findings, discussing the limitations and offering recommendations for future research practice.

Chapter 2: Literature Review

The provision of safe and effective patient care increasingly relies on the seamless integration of diverse data sources [3, 7]. This literature review rigorously analyses the current research environment to identify the principle obstacles and suggested remedies to this issue. The analysis commences by examining the issue of fragmented systems and evaluating essential interoperability standards such as Health Level 7 (HL7), Fast Healthcare Interoperability Resources (FHIR) [2, 3]. This research subsequently assesses the function of remote monitoring and AI-driven analytics in enhancing clinical supervision [4,8], before concentrating on the influence of clinical decision support systems (CDSS) on patient safety [5,7,9]. This chapter synthesises distinct domains to underscore the necessity of an integrated clinical dashboard, thereby delineating the primary problem space of this research.

2.1 Data Integration and Interoperability Solutions

2.1.1 The Challenge of Fragmented Systems

Patient data is facing a persistent challenge due to the lack of standardisation in modern healthcare technology. Such technology comprises multiple electronic health records (EHRs), which include laboratory systems and image repositories, commonly referred to as siloed data [3].

Departmental data flow software can be disrupted and hindered by incompatibilities between systems, even within the same institution. This fragmentation impairs the continuity of care, delays clinical decision-making and increases the likelihood of medical errors [3,5,7].

Moreover, external healthcare providers, such as general practitioners, specialists or community clinics, often operate on separate platforms that do not interface with hospital-based systems. This lack of interoperability significantly limits the ability to share patient information in a timely and comprehensive manner, ultimately compromising the effectiveness and efficiency of patient care delivery [3,10].

The growing reliance on big data and real-time analytics further highlights the critical need for integrated patient records. Althati et al. [11] argue that modern machine learning algorithms and real-time streaming platforms perform at their best when supplied with complete and timely datasets. However, the presence of isolated or incomplete data across systems significantly undermines their ability to provide an accurate and holistic view of the patient. Fragmented infrastructures continue to disrupt essential clinical workflows, such as medication ordering and discharge planning, resulting in procedural delays and inefficiencies [2,10]. For instance, when outpatient prescriptions are stored separately from inpatient records, pharmacies encounter difficulties in maintaining consistent medication histories. Likewise, unsynchronised documentation between departments causes delays in perioperative coordination for surgical teams. These operational gaps increase the likelihood of duplicate testing, medication-related errors, and interruptions in the continuity of care [5,7,12].

The financial burden of fragmented healthcare systems is considerable, with studies estimating that redundant tests and adverse drug events caused by disconnected records account for up to 30% of unnecessary spending [12] and 20% of patient incidents [5].

At an operational level, Radwan et al. [2] highlight that maintaining standalone integrations or point-to-point interfaces can absorb up to 25% of a hospital's IT budget, primarily due to ongoing updates, error resolution and manual data reconciliation. This not only places a significant strain on resources but also impedes innovation, as IT teams are often preoccupied with managing legacy systems instead of advancing modern, scalable solutions, as illustrated in Figure 2.1

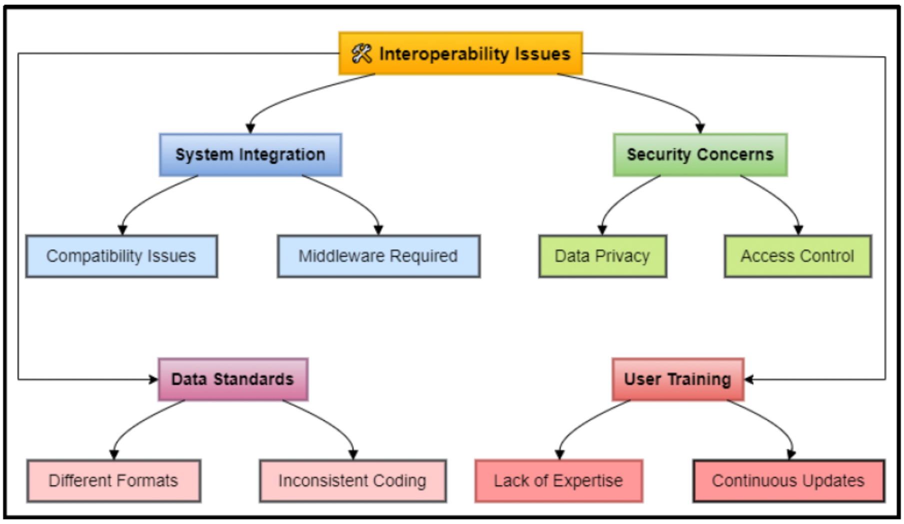

Figure 2.1: *Key Barriers to Healthcare Interoperability [2]*

The adoption of advanced healthcare technologies such as clinical decision support systems (CDSS), remote patient monitoring and real-time analytics is significantly hindered by fragmented data infrastructures. In the absence of a unified, real-time view of patient information, clinicians are forced to compile incomplete datasets from multiple sources, increasing the risk of errors and omissions. Althati et al. [11] contend that integrating data silos is essential to fully realise the potential of machine learning-driven analytics, which depend on comprehensive datasets for accurate prediction. This view is supported by Searle et al. [13], who observe that generating discharge summaries manually from disparate documents is both time-consuming and prone to error, particularly for clinicians operating under strict time constraints.

2.1.2 Interoperability Standards and Data Exchange Protocols

As healthcare organisations strive to integrate data across diverse systems, standards-based interoperability has emerged as a cornerstone for achieving seamless connectivity. The transition from siloed, paper-based records to interconnected digital platforms has been driven by the adoption of interoperability standards that enable heterogeneous systems to communicate effectively. Transitioning from HL7 v2 protocols to FHIR has been shown to reduce data synchronisation latency by 40% to 50% [3]. FHIR promotes interoperability by representing clinical data as modular "resources", such as laboratory results, patient demographics and clinical encounters, thereby minimising the reliance on complex, custom-built interfaces [3]. This resource-based architecture further supports near real-time synchronisation, allowing healthcare providers to access up-to-date patient information rather than outdated snapshots. The immediacy of this data flow enhances clinical decision-making by delivering more relevant and timely information [3].

Nevertheless, the implementation of FHIR alone does not guarantee flawless communication. As noted by Senbekov et al. [10], many institutions continue to rely on a combination of earlier HL7 versions, DICOM (Digital Imaging and

Communications in Medicine), and CDA (Clinical Document Architecture) standards, particularly for imaging. In the absence of robust intermediaries, this heterogeneous mix can impede accurate and efficient data flow. Middleware platforms such as Integration Platforms as a Service (iPaaS) and Enterprise Service Buses (ESBs) are therefore essential to overcoming these limitations. These tools can analyse, validate and transform legacy message formats into contemporary RESTful APIs, bridging structural and semantic gaps between systems.

This architecture is further enhanced through the use of container-based microservices, commonly deployed via platforms such as Docker or Kubernetes. Radwan et al. [2] report that by decomposing monolithic EHR systems into smaller, independent services, hospitals can reduce their mean time to repair (MTTR) for software issues by up to 70%. This modular design allows for more frequent updates without disrupting clinical operations, thereby enabling the rapid deployment of new analytical tools, IoT-driven capabilities and user interface enhancements. Such adaptability strengthens user confidence in the system's stability and reliability.

Ensuring data privacy and adherence to regulations such as GDPR is crucial in any healthcare application. This study does not concentrate on the implementation of comprehensive security measures but recognises the significance of aspects such as encryption and access control in protecting patient data. These factors are essential for the future implementation of mobile clinical systems in practical settings.

In addition, the SMART on FHIR initiative builds upon FHIR's resource-based model to support the development of "plug-and-play" web applications within EHR environments. Early pilot studies have reported 20-25% reductions in development time for specific tools such as medication dosing calculators and risk stratification dashboards when compared to legacy HL7-based integrations [3, 10]. By exposing standardised endpoints and incorporating built-in security layers, SMART on FHIR enables developers to design modular applications that can be rapidly deployed, replaced or upgraded without altering the underlying EHR infrastructure.

Many healthcare organisations use HL7 v2, CDA, DICOM and FHIR concurrently, requiring middleware to manage data transformation, security and message routing. True interoperability often requires combining multiple standards [10, 11]. This strategy, when coupled with containerised micro-services allows institutions to integrate legacy systems while deploying new functionalities.

Collectively, these architectural approaches facilitate fast, accurate and secure data exchange, laying the groundwork for advanced use cases such as real-time analytics, remote monitoring and AI-driven clinical decision support.

2.1.3 Data Platforms, Knowledge Graphs and Embeddings

Despite established interoperability standards, many hospitals are increasingly exploring semantic data models such as knowledge graphs and clinical ontologies to structure patient information for decision assistance and predictive analysis. These models can integrate structured data (e.g., ICD codes, laboratory findings), unstructured text (e.g., clinical notes) and device data into a cohesive, machine-readable format.

A crucial aid in this process is the systematised nomenclature of medicine clinical terms (SNOMED CT), which offers more than 350,000 medical concepts and hierarchical linkages to provide consistent terminology across clinical settings [14]. Utilising SNOMED CT guarantees semantic consistency in the documentation of diseases, symptoms and procedures, hence enhancing subsequent reasoning, alarms and semantic search [15].

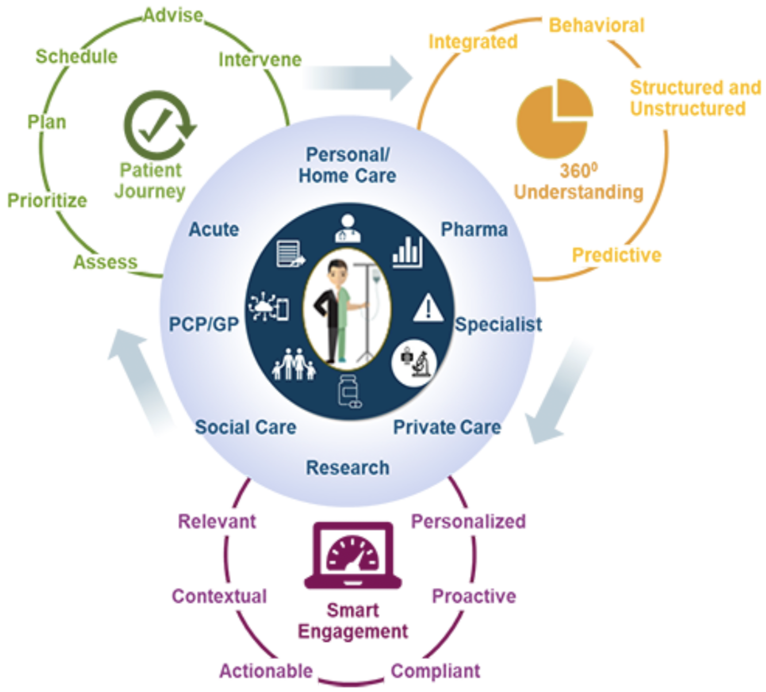

Figure 2.2: *The Vision of an Integrated Healthcare Ecosystem [3]*

Figure 2.3 illustrates the organisation of medical concepts and their semantic relationships (e.g., “is a”, “associated with”) with a SNOMED CT-based knowledge graph, Such representations are useful for supporting clinical decision-making and semantic reasoning, particularly in large-scale systems.

With sophisticated techniques like graph-based embeddings (e.g., Snomed2Vec [15]) and extensive knowledge graphs have demonstrated encouraging outcomes in domains such as patient similarity detection and early warning systems, these approaches are beyond the technical parameters of the current study. Nonetheless, they establish a robust basis for forthcoming improvements to the clinical dashboard especially in the areas of semantic alignment and intelligent diagnostic assistance.

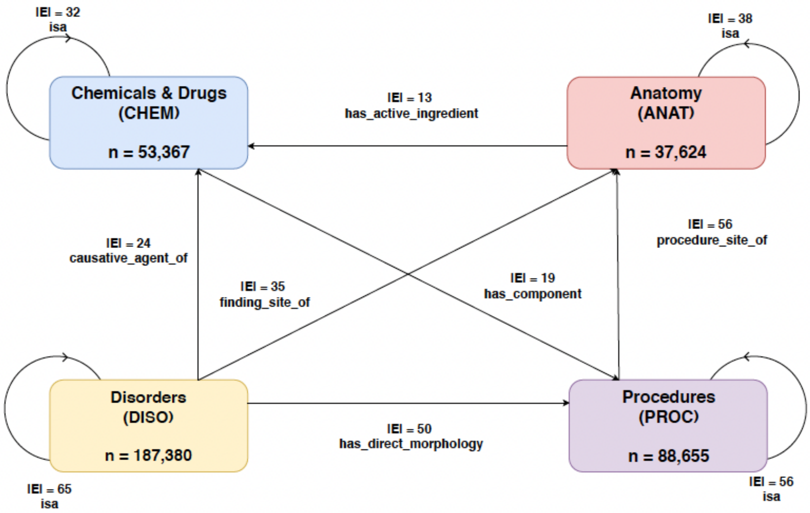

Figure 2.3: *Illustration of SNOMED CT concept relationships in a knowledge graph* [15]

2.2 Remote/Real-Time Patient Monitoring (IoT & AI)

Contemporary healthcare progressively adopts remote and instantaneous patient monitoring to broaden clinical awareness beyond conventional hospital environments [13,16,17]. Utilising the internet of things (IoT) devices, such as wearable, embedded sensors and linked devices healthcare professionals can access constant flows of essential information to facilitate prompt actions [8,10,18]. Parthasarathy et al. [19] describe a prototype that uses microcontroller-driven sensor nodes (like Arduino boards) to collect metrics such as heart rate and body temperature, subsequently transmitting this data through smartphone gateways to a cloud platform for near real-time analysis. Initial feasibility findings indicated a 20% to 30% faster identification of abnormal vital signs, particularly beneficial for step-down units where patients require close monitoring without the intensity of ICU care.

In order to differentiate between correct clinical alerts and minor deviations, Shaik et al. [4] demonstrate how AI-driven techniques might further enhance these monitoring streams. In comparison to threshold-based warning systems, their pilot test revealed that the number of false alarms dropped by 25% using a machine learning model to evaluate ventilation and ECG data, hence reducing "alarm fatigue" among carers. Additionally, contextualised alarms were made possible by incorporating data from the EHR, such as medication history or co-morbidities. The technology prioritises changes in vital signs associated with high-risk profiles over reporting every alteration, freeing up doctors to concentrate on the most urgent situations. The AI-enabled remote patient monitoring architectures are illustrated in Figure 2.4.

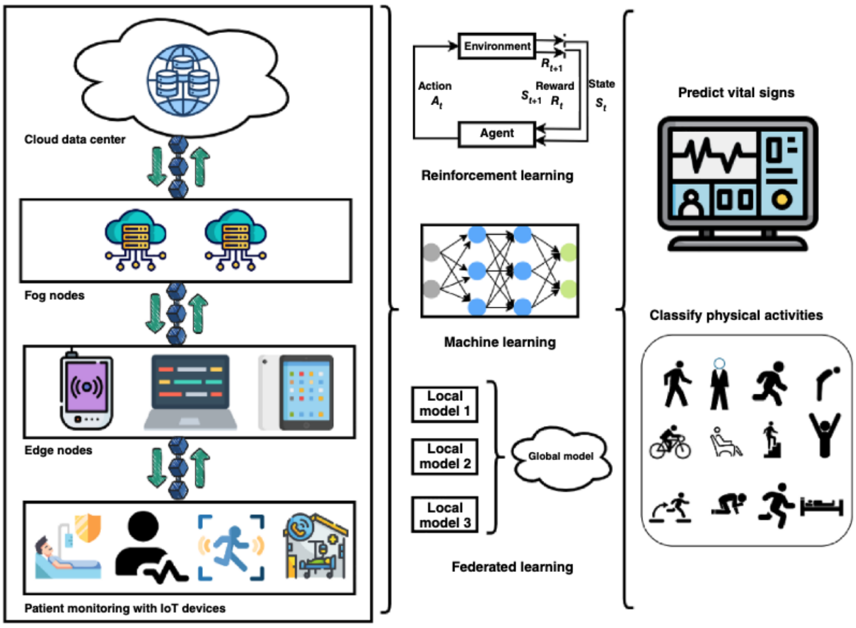

Figure 2.4: *artificial intelligence-enabled remote patient monitoring Architectures [4]*

2.2.1 Chronic Disease Management and Ageing Populations

Remote monitoring has proven to be particularly effective in the treatment of chronic diseases such as diabetes, heart failure and COPD, which together account for a high proportion of healthcare costs. These diseases often require continuous monitoring and care, which can be supported by digital technologies [8]. Gupta [8] describes an IoT+AI platform developed to monitor older adults' daily weight, blood pressure and blood glucose levels at home. This is particularly important given the increasing number of elderly patients and the prevalence of age-related diseases that require consistent monitoring. In a six-month observational study, participants using the system saw a 15% to 20% reduction in hospital admissions, likely due to earlier intervention for subtle warning signs (e.g. sudden weight spikes in heart failure). The ability to recognise subtle warning signs is crucial for effective intervention and can be improved by AI-driven analysis of real-time data [4]. Patient adherence data also feeds into clinical dashboards, allowing providers to investigate medication non-adherence and schedule timely follow-ups or tele-consultations. Such systems help with chronic disease management and support personalised, data-driven and adaptive healthcare [8]. The measurable impact of such systems on key patient outcomes is illustrated in Figure 2.5.

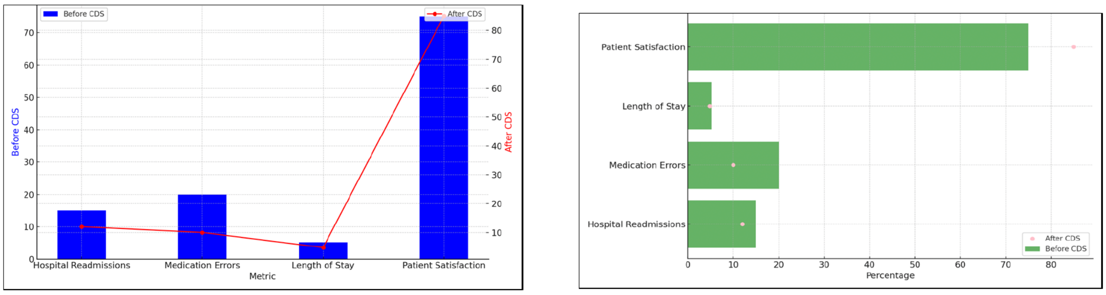

Figure 2.5: *The impact of CDSS implementation on patient outcomes [2]*

By combining real-time streaming with reliable analyses, hospitals can tailor interventions to the profile of each patient. Shaik et al. [4] found that a prediction model for heart failure readmission trained on vital signs and EHR-based comorbidity data reduced 30-day patient readmissions by 10% to 15% compared to standard monitoring. This highlights the importance of integrating multiple data sources for more effective predictive modelling and proactive care [2, 4]. Such findings emphasise that proactive care improves patients' quality of life and relieves pressure on hospital resources, addressing the challenges of rising healthcare costs and the need for efficient use of limited resources [12, 20].

2.2.2 Implementation Barriers and Data Security

Despite these advantages, IOT-based solutions are associated with considerable technical, security and governance challenges. These challenges relate to technical infrastructure, data security, interoperability and regulatory compliance [4], [10], [21]. As noted by Senbekov et al. [10], many hospital IT infrastructures remain outdated, originally optimised for stationary workstations. Adapting these environments to support thousands of mobile, wireless IoT devices requires substantial reconfiguration, particularly with respect to VLAN segmentation, bandwidth, and real-time data flow. Shaik et al [4] emphasise the importance of robust encryption and authentication measures, as wireless medical devices can be hacked or tampered with if these are not in place. The vulnerabilities in medical devices can have serious consequences, such as the possibility of remote manipulation of dosages or devices settings, as seen in some high-profile recalls of infusion pumps or implantable cardiac devices. Such security breaches jeopardise patient safety and undermine trust in digital health solutions [21].

In addition, the proprietary nature of many wearable or sensor platforms makes integration difficult. Manufacturers may use closed APIs or unique data schemas, forcing developers to create custom adapters or rely on middleware that normalises streaming data. This lack of interoperability hinders seamless data exchange and can lead to data silos that limit the effectiveness of integrated healthcare systems [3]. Compliance with regulations such as HIPAA and GDPR also requires robust security measures, including encryption in transit (e.g. TLS/SSL), thorough access controls to prevent unauthorised viewing or manipulation of patient data and secure storage solutions. These measures are essential for both legal compliance and sustaining patient trust. [4], [21]. Shaik et al. [4] reported that transparent communication about privacy and security protocols increased patient confidence in pilot programs; participants appreciated knowing how and why their data was shared. This transparency helps to allay patients' fears and concerns about data security and the impact on the human aspect of healthcare [16], [21].

2.2.3 Prospects for Integration with EHR and Clinical Workflows

The integration of real-time remote monitoring technologies with hospital-based EHRs has the potential to transform clinical workflows. By continuously feeding data from wearable sensors and IoT devices into the EHR, clinicians gain a unified, real-time view of a patient's physiological status, enhancing both decision-making and care coordination. Instead of manually reviewing outputs or navigating disparate systems, healthcare teams can receive contextual alerts, trend analyses and actionable insights directly within their clinical dashboards.

Shaik et al. [4] demonstrate that AI-enhanced remote monitoring systems, when integrated with EHRs, can significantly reduce the cognitive burden on clinicians by filtering data based on individual patient risk profiles. These systems are designed to generate priority alerts only when critical thresholds are exceeded or when variations in vital signs correlate with pre-existing co-morbidities documented in the patient's record. This selective alerting mechanism has been shown to decrease the frequency of false alarms and mitigate alarm fatigue among clinical staff, thereby improving the efficiency and effectiveness of clinical response workflows.

In addition, Radwan et al. [2] report that the integration of AI-generated insights into clinical workflows accelerated discharge planning in test centres by up to 18%. This improvement was achieved by aligning remotely monitored vital signs with discharge readiness criteria defined in the hospital's clinical rules engine. For instance, patients exhibiting steady improvements in oxygen saturation and blood pressure were flagged for early discharge consideration, thereby freeing up critical bed space without compromising patient safety.

Building on this Gupta [8] demonstrates how EHRs can display longitudinal sensor data alongside medication schedules and historical laboratory results. This cross-sectional integration facilitates more comprehensive clinical evaluations, such as identifying correlations between weight gain and the progression of heart failure or detecting post-operative complications at an early stage. Such functionality enables clinicians to better anticipate treatment escalation or recovery trajectories, intervene proactively and deliver more personalised care.

From a workflow perspective, real-time integration significantly reduces the need for manual transcription and frequent vital sign updates by nursing staff, thereby minimising documentation errors and allowing staff to focus more on direct patient care, as noted by Althati et al. [11]. Moreover, real-time data pipelines feeding into the EHR enable the use of automated scoring systems such as NEWS2 and APACHE II, which can be visualised instantly on clinicians' mobile interfaces. Seamless integration also enhances collaboration within multidisciplinary teams. Nurses, physiotherapists, specialists and pharmacists can coordinate their activities through shared real-time dashboards, thereby ensuring continuity of care and reducing delays in interdisciplinary decision-making.

Ultimately, the integration of IoT-based monitoring into EHR systems and clinical workflows represents a fundamental shift towards proactive, continuous and patient-centred care rather than merely a technological advancement. This level of connectivity is critical for enabling hospital mobile dashboards to support bedside clinical documentation, AI-driven alerts and personalised discharge planning.

2.3 Clinical Decision Support Systems

CDSS are computer-based tools that analyse data from various sources to provide prompts and reminders to assist healthcare professionals with decision-making tasks. They are increasingly recognised as essential resources for reducing medication errors and improving prescribing accuracy and safety. Through direct EHR integration, CDSS technologies can identify contraindications or overlaps by instantly comparing current prescriptions with patient allergies, pre-existing conditions and clinical guidelines.

2.3.1 Medication Safety and Prescribing Accuracy

In a machine learning-based assessment of high-risk prescriptions, Corney et al. [5] found that their CDSS can recognise dangerous patterns with high accuracy. In patients with impaired renal function, the tool was particularly successful in detecting dosing irregularities, drug interactions and unauthorised prescriptions, as renal function was a key feature in the model's training [5]. This is consistent with research by other authors suggesting that the inclusion of the CDSS significantly reduces the number of avoidable adverse drug events (ADEs), particularly when combined with renal function alerts and allergy data [7, 9]. For example, one machine learning-based system was able to identify high-risk prescriptions and reduce potential medication errors with 74% precision and 74% recall [5]. In another study, a CDSS for drug allergy management demonstrated 100% accuracy in its evaluations [9]. This suggests that the likelihood of medical intervention would have been significantly higher if the alerts had been issued in real time.

Metric	CDS Alert System	Multicriteria Query	Lumio Medication Algorithm
Recall	0.69	0.66	0.74[a]
Precision	0.54	0.62	0.74[a]
F1 Score	0.61	0.64	0.74[a]
AUCPR	0.56 (95% CI, 0.50–0.62; $P < .00001$)	0.56 (95% CI, 0.51–0.61; $P < .001$)	0.75 (95% CI, 0.70–0.80)[a]
AUROC	0.65 (95% CI, 0.61–0.69; $P < .00001$)	0.68 (95% CI, 0.64–0.72; $P < .0152$)	0.81 (95% CI, 0.78–0.84)[a]

Table 2.1: *Performance of a machine learning-based CDSS for medication error detection [5]*
AUCPR: area under the precision–recall curve; AUROC: area under the receiver-operating characteristic curve; CDSS: clinical decision support; CI: confidence interval. [a]Derived from external validation cohort.

In addition, Gupta [8] emphasises the importance of using SNOMED CT in CDSS frameworks to ensure that clinical terms are consistent and compatible across care settings. Safer prescribing practise is facilitated by the precise mapping of diseases, protocols and pharmacological categories enabled by this semantic integration.

By continuously learning from clinical results, the CDSS equipped with artificial intelligence improves this safety net even further. These systems can provide predictive insights by analysing patterns in groups of patients. For example, they can alert a doctor if a particular drug has caused side effects in patients with similar profiles in the past. According to Shaik et al [4], when machine learning algorithms were combined with CDSS, the number of drugs prescribed dropped by 15% to 20% in a number of healthcare organisations.

All of this data underscores the fact that a strong CDSS serves as a real-time safety buffer against preventable prescribing errors and also guides physicians' decisions. CDSS significantly improves pharmaceutical safety and prescribing accuracy in modern healthcare by automating key tests, recommended evidence-based alternatives and learning from previous results.

2.3.2 Diagnostic Support and Risk Stratification

In high-stress environments such as emergency departments and intensive care units, CDSS are now indispensable for improving diagnostic accuracy and controlling clinical risks [7]. By integrating patient data from laboratory systems, imaging and EMRs, CDSS can synthesise massive data sets and provide immediate alerts or recommendations for diagnosis [2, 5]. In cases of ambiguity or multi-morbidity, where clinicians may miss subtle symptom patterns due to lack of time or cognitive fatigue, this is extremely helpful [22].

To enable triage and differential diagnosis, advanced CDSS platforms utilise AI models such as rule-based engines and neural networks [7, 17]. Corney et al. [5] highlight how the system can evaluate prescriptions by comparing patient-specific variables such as age, renal function and allergy history, significantly reducing diagnostic errors associated with medication-related problems. Sutton et al. [7] provide a comprehensive overview of CDSS, detailing their evolution, functionalities including diagnostics and disease management and the importance of their integration with EHRs. AI-powered CDSS have shown the ability to recognise unusual symptoms in real-life situations, which has led to early research into diseases such as sepsis or heart failure [4, 17].

In addition, the effectiveness of CDSS in risk classification has been studied in detail [5]. These systems can automatically calculate recognised scoring systems, such as NEWS2, by retrieving laboratory results, vital signs and comorbidity indices from EHRs [4, 11]. Automating these assessments not only increases the reliability of clinical judgement but also speeds up the clinical response to patients whose condition is deteriorating [4, 7]. For example, one machine learning-based NEWS2 was able to predict patient deterioration one hour before onset. Separately, integrating risk modelling into CDSS interfaces has been shown to improve the accuracy of predicting outcomes for critically ill patients [4].

Shaik et al. [4] explain how CDSS and AI-assisted monitoring can stratify risk patients with respiratory and heart failure by predicting likely readmission based on historical trends and real-time physiological data. To shift care models from reactive to proactive, these stratified dashboards give clinicians a prioritised list of patients who need urgent care [17]. In addition to improving patient outcomes, hospitals can manage clinicians' workloads more efficiently by focusing their resources on those patients most likely to deteriorate [2].

Visualisation tools integrated into the CDSS interface also improve decision support [7]. These technologies can provide a summary of results in natural language for faster evaluation, graphically represent vital sign patterns and identify significant laboratory anomalies [12, 23]. When such contextual data is easily accessible, clinicians report increased diagnostic confidence, allowing them to differentiate between actual difficulties and expected post-operative changes [22]. Muhiyaddin et al. [22] identified several positive impacts of CDSS on physicians, including improved work efficiency, better personalised care and increased confidence in decision-making.

Finally, the integrated modules of the CDSS platforms for risk stratification and diagnostic support help to close clinical detection gaps. They are indispensable in acute care, as they enable dynamic prioritisation of treatment and promote early intervention [5, 7]. The contribution of CDSS to risk-based triage and diagnostic accuracy is expected to increase as they evolve, particularly as AI models are refined using patient datasets that are increasingly diverse and multimodal [16, 17].

2.3.3 ePrescriptions and Medication Safety

The digitisation of prescriptions known as e-prescribing, has revolutionised medication administration by enhancing safety, precision and interoperability across healthcare environments [12]. CDSS are pivotal in this transformation, especially when coupled with pharmaceutical systems and EHRs [7]. These systems deliver real-time notifications on contraindictions, duplicate prescriptions, drug interactions and dose discrepancies, contingent upon patient-specific variables such as age or renal function [5, 18].

Several studies underscore the efficacy of CDSS in detecting high-risk medications. Corney et al. [5] demonstrated that machine learning based decision support surpassed manual evaluations in identifying potentially unsuitable drugs, particularly in instances of polypharmacy or chronic renal illness. Their algorithm attained a 92% accuracy rate in clinical simulations, mitigating the likelihood of adverse medication effects.

Semantic interoperability is also crucial. Standards like SNOMED CT guarantee uniform terminology and enable automatic substitute recommendations in the presence of allergies, stock shortages or insurance limitations [14, 15]. Algorithms can further enhance dosage by taking into account patient-specific characteristics, which is especially crucial in paediatric or geriatric treatment.

2.4 Synthesis and Research Gap

This literature review highlights a pressing challenge in the development of contemporary healthcare technologies. Although standards such as HL7 FHIR and emerging solutions like IoT monitoring and AI-driven analytics show considerable promise, their practical impact is significantly constrained by persistent issues of data fragmentation and lack of system integration. The evidence consistently shows that even efficient, task-specific AI tools are ultimately limited by siloed data infrastructures, which impede real-time clinical decision-making and compromise patient safety [2, 5, 22].

The architectural priorities of this study are directly informed by these limitations, particularly in relation to advanced functionalities such as e-prescribing. Effective medication management demands seamless data access, real-time safety checks and absolute semantic consistency, the very challenges this research seeks to address. While the literature establishes e-prescriptions as a critical tool for improving prescribing safety and workflow efficiency, the implementation of this feature is beyond the scope of the current artefact due to time and resource constraints. However, the system's underlying architecture is intentionally designed as a robust and extensible platform, ensuring that e-prescribing and automated medication safety systems can be integrated in future iterations without requiring major redesign.

A critical analysis reveals that existing research often mirrors this fragmentation. Studies tend to focus on discrete elements — such as an interoperability protocol [3], a machine learning model [5], or a knowledge graph embedding technique [15] rather than exploring how these components might be cohesively integrated. Consequently, there is a clear research gap in proposing a practical framework that unifies heterogeneous EHR data and real-time IoT streams into a contextualised clinical decision support system (CDSS) embedded directly into clinical workflows [2, 7, 23].

This study aims to address that integration and application gap. Its primary contribution lies in the design and evaluation of a unified clinical dashboard artefact that leverages modern data platforms and AI capabilities to convert integrated real-time data into actionable clinical insights [11, 12, 17]. Rather than starting from scratch, the artefact builds upon existing standards and international best practices — including FHIR, SNOMED CT, and IoT-based monitoring and combines them into a cohesive solution tailored for point-of-care use.

The review has directly informed the design priorities and research strategy of this work. Key takeaways — including the need for semantic interoperability, real-time decision support, and user-centric interfaces were critical in shaping the artefact's architecture. Moreover, a preliminary review of local academic research revealed limited exploration of such integrated systems within the Maltese healthcare context. This gap further justifies the relevance and originality of this study.

By synthesising diverse technologies into a single, context-aware clinical tool, this study offers a practical and locally grounded roadmap for advancing data-driven, proactive, and patient-centred care in hospital environments.

Chapter 3: Research Methodology

This chapter outlines the methodological framework that guided the study. The research adhered to a Design Science Research (DSR) approach, grounded in a pragmatic philosophy. DSR was selected for its suitability in developing innovative technological artefacts that address real-world problems; in this case, the requirement for an extensible, secure and contextually relevant mobile clinical dashboard for inpatient care. A pragmatic philosophy was adopted to prioritise practical solutions and their applicability in clinical settings, aligning with the study's objective of producing a usable artefact rather than solely generating theoretical insights.

The study was structured into three primary phases: (i) methodological foundations, (ii) system development and data collection and (iii) synthesis and analysis of results. This framework facilitated a logical progression, beginning with the establishment of requirements, followed by artefact creation and concluding with its evaluation and refinement. Figure 3.1 summarises the overall research process, with each phase discussed in greater detail in the subsequent section of this chapter.

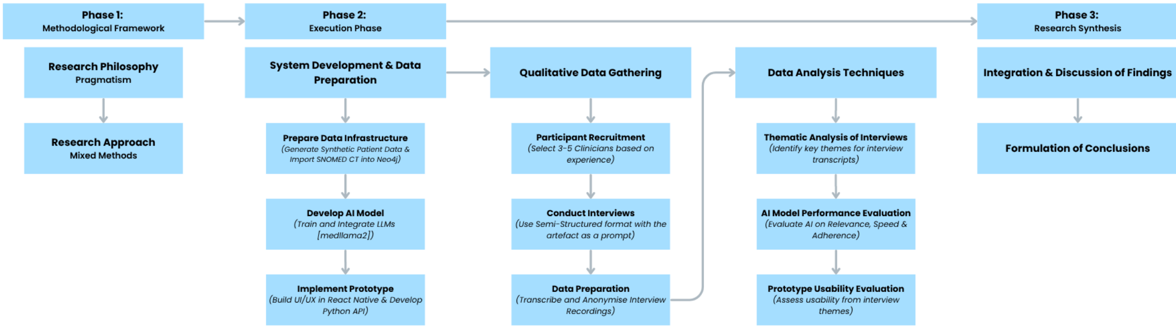

Figure 3.1: *Research Pipeline*

3.1 Methodological Framework

This study was guided by a pragmatic paradigm, as outlined in Chapter 1, which prioritised practical goals and the application of research findings to solve real-world problems. This paradigm aligned with the project's primary aim of creating an effective and functional technological solution for a specific clinical need, rather than focusing solely on theoretical advancements.

A mixed-methods approach was employed to integrate the strengths of technical development and human-centred evaluation. The creation of the mobile clinical dashboard followed a DSR strategy to ensure technical robustness, security and applicability to the inpatient care context. The evaluation phase, however, utilised a qualitative methodology in the form of semi-structured interviews with clinicians. This was deliberately chosen given the exploratory and iterative nature of the artefact; qualitative feedback was deemed more appropriate at this stage than comprehensive quantitative testing. The interviews provided rich insights into clinicians' workflows, priorities, usability preferences and concerns, insights that were essential for refining the artefact. The combined application of DSR and qualitative research ensured that the final solution was not only technically sound but also closely aligned with end-user requirements, thereby achieving the research objectives.

3.2 System Design and Development Methodology

This study utilised the Design Science Research (DSR) methodology to steer the development and assessment of the main research contribution: an innovative mobile clinical dashboard. DSR was particularly appropriate for this study as it is a recognised problem-solving paradigm centred on the development of novel artefacts, in this instance, a technology-driven solution to a real-world healthcare issue. The selection of DSR was further supported by its focus on practical relevance and rigorous evaluation, guaranteeing that the created system was both operational and closely aligned with the requirements of healthcare professionals.

This study's DSR process adhered to a singular, thorough cycle consisting of three fundamental phases: (i) design, (ii) implementation and (iii) evaluation.

3.2.1 Artefact Design and Architecture

The primary artefact is a mobile clinical dashboard aimed at delivering clinicians swift and informed access to diagnostic information and patient data. The system was designed to meet the specifications for data security, real-time performance and clinical accuracy outlined in the literature review (Chapter 2).

The system's architecture, illustrated in Figure 3.2, was composed of three primary layers: a Frontend, a Backend and a Data and AI Layer. The frontend consisted of a mobile user interface designed for intuitive interaction on handheld devices, enabling clinicians to access features efficiently in a dynamic environment. The backend was a robust RESTful API that functioned as the system's central component, overseeing all operational logic, including real-time terminology enquiries and interactions with the AI model. Finally, the Data and AI Layer utilised a dual-database system combined with a locally-hosted Large Language Model (LLM) to facilitate efficient data storage, retrieval and advance search functionalities.

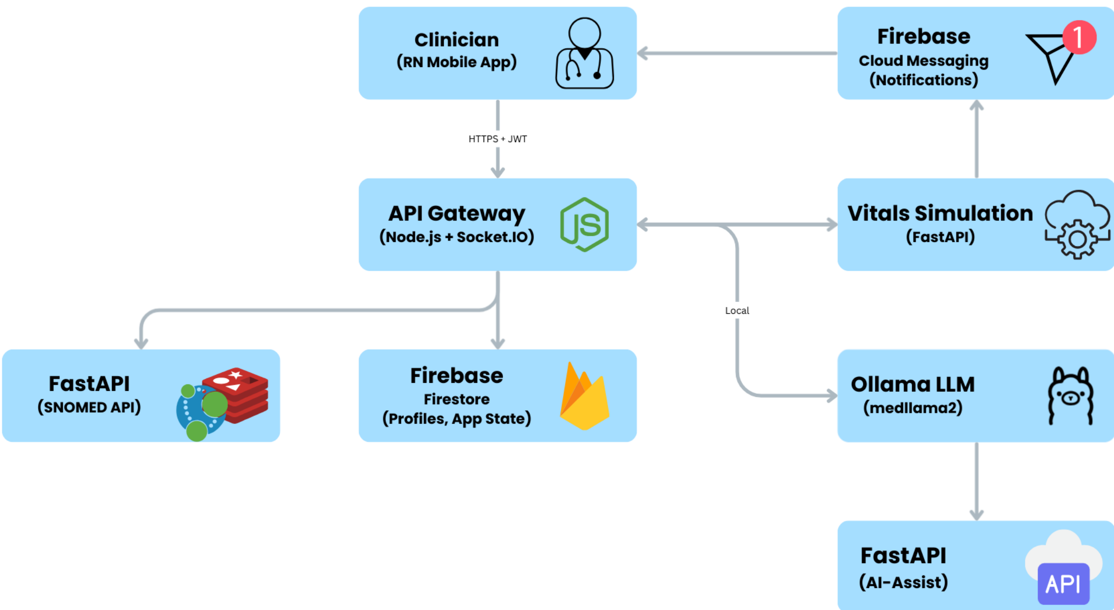

Figure 3.2: *Architecture Diagram*

This architectural design was made to guarantee that all sensitive clinical data is maintained within a secure, self-hosted environment, thereby directly addressing the essential privacy and security requirements of healthcare applications.

3.2.2 Development Tools, Technologies and Justification

The implementation of the artefact involved the selection of a specific suite of modern tools and technologies, with each choice justified by the system's requirements concerning performance, security and scalability.

The backend API was developed using FastAPI, a modern Python web framework selected for its high performance and inherent support for asynchronous operations, which are crucial for handling the application's real-time components. A dual-database strategy was implemented for data management. A Neo4j graph database was utilised for the storage and querying of the complex, hierarchical SNOMED CT terminology, as its graph structure was inherently appropriate for managing interrelated data. This was supplemented by a Redis in-memory cache to ensure low-latency responses for frequently accessed data, thereby ensuring a smooth user experience.

The diagnostic search feature was powered by medllama2, a large languge model (LLM) optimised for the medical field The model was deployed locally using ollama, a framework designed for local execution of LLMs. This self-hosted approach represented a fundamental design choice, ensured that sensitive clinical information was not transmitted to external third-party services.

3.2.3 Evaluation

The last element of the DSR cycle involved the assessment of the artefact. The approaches outlined in Section 3.3 were employed to evaluate the artefact in relation to the study's objectives. Qualitative data obtained from clinicians (Section 3.3.1) was crucial in this phase, offering vital user-centred feedback to validate the design and pinpoint opportunities for future enhancement.

All performance benchmarks for the evaluation were executed on a workstation with the following specs to guarantee the reproducibility of the results:

Table 3.1: *System Specifications*

Component	**Specification**
CPU	AMD Ryzen 7 5800X
RAM	32GB DDR4 at 3200MHz
GPU	Nvidia RTX 3060 (12GB VRAM)
Framework	FastAPI, Neo4j, Redis, Ollama with medllama2

3.3 Data and Evaluation Methodology

This chapter outlines the methods utilised for data collection and preparation, as well as the comprehensive strategy implemented to assess the research artefact. This method integrates synthetic test data generation with qualitative insights from clinicians to facilitate a thorough evaluation.

3.3.1 Data Collection and Preparation

A mixed-methods approach to data was employed, incorporating the creation of a synthetic quantitative dataset for model testing and the acquisition of qualitative data via expert interviews for user-centric assessment.

3.3.2 Qualitative Data: Clinician Interviews

In order to understand the artefact's potential usability, feature significance and workflow integration, qualitative data was obtained through comprehensive, semi-structured interviews. A purposive sampling method was employed to recruit three clinicians, specifically doctors or consultants. Participants were selected with precision based on crucial criteria: direct experience with inpatient care in a hospital setting. This guaranteed that the feedback gathered was particularly pertinent and grounded in the specific context of the inquiry.

Semi-structured interviews with clinicians represent a recognised approach for assessing the utility and usability of emerging clinical technologies [5, 13]. This approach is warranted as it facilitates a detailed examination of the user experience, revealing profound insights into workflow, decision-making and perceived value that are not adequately captured by technical metrics alone.

The interviews followed a structured approach utilising a set of open-ended questions (Appendix A), with the created artefact serving as a prompt to gather specific feedback. All interviews were conducted with informed consent and were audio-recorded and transcribed verbatim to guarantee the accuracy of the data for subsequent analysis.

3.3.3 Test Data Generation: Synthetic Patient Case Notes

A synthetic dataset of patient case notes was created in order to evaluate the AI model's performance quantitatively. Without utilising actual patient data, this was required to establish a controlled and medically realistic test environment.

To guarantee that the case notes accurately depicted a range of clinical settings across many specialities, they were prepared in accordance with accepted clinical case study methods. The AI's capacity to make diagnostic recommendations was tested using this artificial dataset.

3.4 Evaluation Strategy

The assessment of this research was executed employing a threefold approach, as advised for DSR. This methodology guaranteed a comprehensive evaluation by initially verifying the artefact's technical performance, subsequently using it to collect and analyse data to address the research enquiries and ultimately detailing the steps implemented to affirm the validity of those findings.

3.4.1 Phase 1: Artefact Technical Validation

The initial evaluation phase was centred on testing the object to confirm that it functioned as planned from a technical standpoint. This covered the ”Real-Time” and ”Secure” portions of the study.

- Real-time performance: The interference time of the LLM was benchmarked in milliseconds on the specified GPU setup. This quantitative test was created to ensure that the system could produce diagnostic recommendations rapidly enough to be effective in a fast-paced clinical context.
- Security architecture: The suggested secure architecture (including the client-server model and authentication techniques) had been conceptually validated. This was accomplished by discussing the design with the interviewed doc-

tors and requesting input on its effectiveness in resolving their security and privacy concerns.

3.4.2 Phase 2: Gathering and Interpreting Results

The second phase concentrated on collecting and analysing data to address the primary research questions concerning the artefact's effectiveness and utility at the time of evaluation.

- AI model performance: The effectiveness of the AI-driven diagnostic recommendation feature was evaluated through the analysis of the synthetic patient case notes. The LLM processed these notes and its output was assessed based on its clinical relevance, which was determined by comparing the diagnostic suggestions against the expected diagnoses according to the standard SNOMED CT nomenclature.

- User-centred feedback: The transcribed interview data underwent thematic analysis. This process entailed the systematic coding of data to identify recurring patterns, which were subsequently organised into overarching themes. Themes concerning usability, clinician satisfaction and workflow integration were used to evaluate the artefact's overall efficacy from a user-centred perspective, focusing on the "Efficient" dimension of the research. This qualitative approach is known to effectively elucidate user perception in analogous research [13].

3.4.3 Phase 3: Ensuring Validity of Results

The concluding step encompassed actions to ascertain the validity and dependability of the results. The credibility of this work was enhanced by the use of a mixed-methods evaluation. The review integrated quantitative performance parameters (inference time, clinical relevance) with comprehensive qualitative data from expert end-users, resulting in a triangulated and more challenging assessment of the artefact's value than any singular technique could provide. The contrast between the grandeur of the AI's output with recognised clinical standards (SNOMED CT) and the foundation of the qualitative analysis on the first hand experiences of practising doctors guaranteed that the results are both technically robust and clinically pertinent.

3.5 Ethical Considerations

This research was executed in complete compliance with the ethical framework established by MCAST. A mandatory prerequisite for participation was the provision of written informed consent, which was obtained after each participant had reviewed a detailed information sheet (Appendix B). Safeguarding participant anonymity was a key priority; therefore, all transcripts were anonymised using pseudonyms during data preparation. For data security, all digital assets, including audio files were kept in encrypted, password-protected folders with access restricted solely to the researcher. These files were securely and permanently erased at the end of the project lifecycle. Furthermore, all participants were advised of their unconditional right to withdraw from the study at any time, for any reason and a procedure was in place to honour any request for the prompt deletion of their data.

Chapter 4: Analysis of Results and Discussion

This chapter delineates and analyses the key findings of this study, in accordance with the aims specified in Chapter 3. The chapter examines two complementary sources of evidence: the quantitative assessment of the diagnostic recommendation algorithm and the qualitative insights obtained from clinician interviews. Collectively, these findings offer a comprehensive view on the viability, usefulness and ramifications of an AI-driven clinical recommendations system incorporated into a mobile dashboard.

4.1 Quantitative Findings: Performance of the Diagnostic Suggestion Algorithm

The diagnostic algorithm was assessed through a series of iterative development cycles, each aimed at enhancing its capacity to translate unstructured clinical notes into structured SNOMED CT terminology. As Gaudet-Blavignac et al. argued, processing free text is a critical challenge for achieving semantic interoperability in healthcare [14]. The algorithm's performance was evaluated utilising established criteria (precision, recall and F1-score) against a dataset of 250 synthetic patient case notes. The outcomes of these cycles revealed the artefact's initial constraints and subsequent evolution, providing critical insight into its practical applicability.

4.1.1 Cycle 1: Baseline Performance of the Hybrid Model

The initial cycle established a performance baseline for the hybrid model, which integrated a LLM (MedLLaMA2) for concept extraction with a Neo4j graph database containing the SNOMED CT ontology. This design, informed by the work of Agarwal et al. on the *Snomed2Vec* approach, aimed to leverage both the semantic understanding of an LLM and the structured knowledge of a clinical ontology [15].

The model achieved the following micro-averaged metrics (Table 4.1):

Table 4.1: *Evaluation metrics for the baseline hybrid model (Cycle 1).*

Metric	**Value**
Precision	0.11
Recall	0.32
F1-Score	0.17

As illustrated in the confusion matrix (Figure 4.1), the system identified some correct medical concepts but also incorrectly suggested a high volume of irrelevant ones. This predominance of false positives aligned with the low precision score and highlighted a key limitation: an inability to effectively filter unrelated terminology. In a clinical setting, this would undermine clinician trust, as the cognitive burden of reviewing incorrect suggestions could outweigh the benefits of the few accurate detections, a known risk in CDSS design [7].

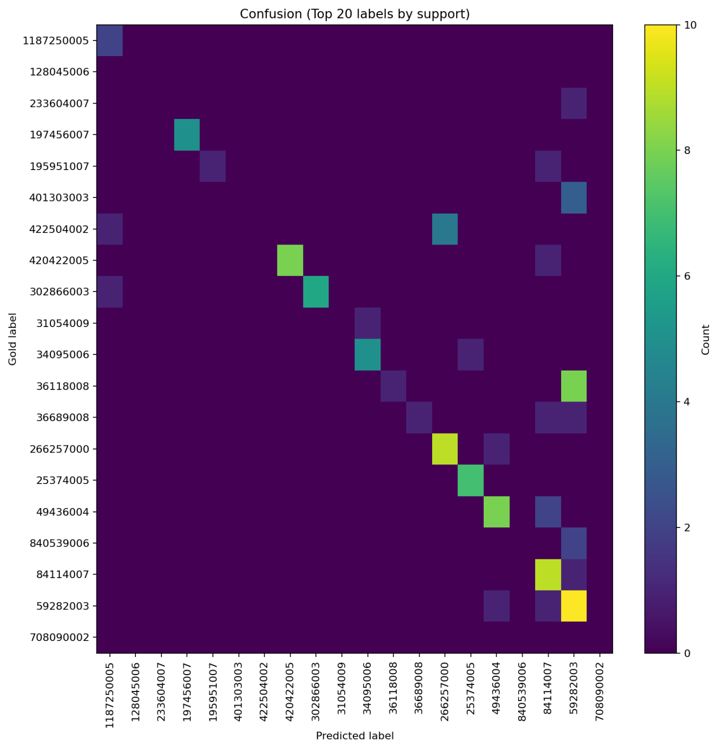

Figure 4.1: *Cycle 1 Confusion Matrix of baseline hybrid model.*

4.1.2 Cycle 2: The Precision-Recall Trade-Off in LLM Prompting

To address the baseline model's shortcomings, the second cycle focused on improving the concept extraction phase by providing the MedLLaMA2 LLM with more detailed and context-aware prompts. The evaluation of the revisited model produced a notably different set of metrics (Table 4.2):

Table 4.2: *Evaluation metrics for the revised model (Cycle 2).*

Metric	Value
Precision	0.05
Recall	0.41
F1-Score	0.10

These results revealed a classic precision-recall trade-off. The enhanced prompts improved the algorithm's ability to identify potentially correct diagnoses, as shown by the recall score increasing from 0.32 to 0.41. However, this gain came at the cost of a sharp drop in precision to 0.05. The system became more sensitive in finding relevant concepts but was also more prone to generating false positives.

Figure 4.2 demonstrates this phenomenon. While more gold labels on the diagonal were identified, the extensive off-diagonal activations revealed the prediction of numerous irrelevant SNOMED CT codes. This visual evidence supported the quantitiative metrics, highlighting that enhanced LLM prompting was insufficient without a robust system for filtering results. As research by Sutton et al. and Muhiyaddin et al. confirmed, an excess of low-quality alerts can lead to "alert fatigue," causing clinicians to ignore the system altogether [7, 22].

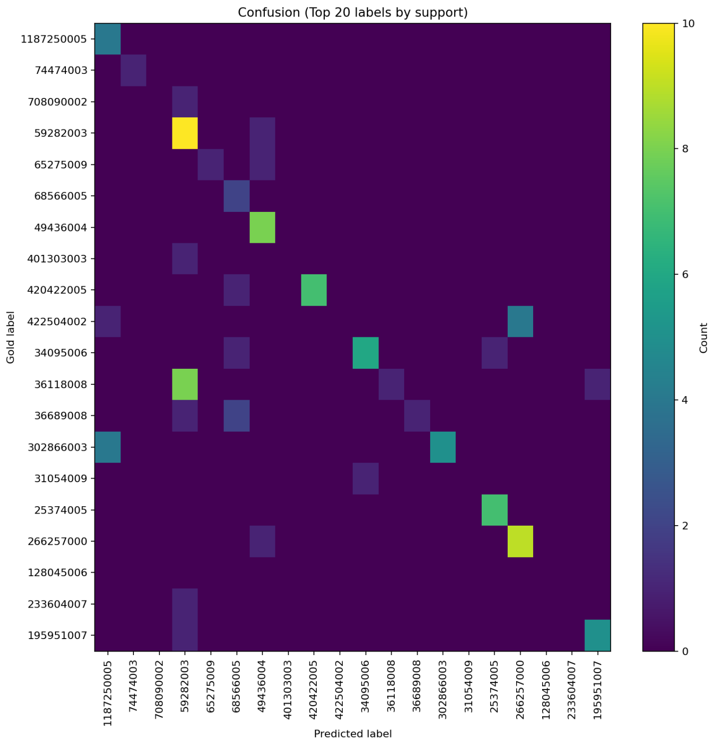

Figure 4.2: *Cycle 2 Confusion Matrix of revised prompting model.*

4.1.3 Final Iteration: The 'Heuristic Ceiling' and Final Performance

The concluding development phase focused on resolving the critical issue of low precision by implementing a feature-based re-ranking strategy. This final model assessed candidates based on LLM confidence, semantic similarity and clinical priors before applying a filtering threshold.

Despite these architectural enhancements, the model's performance did not significantly improve, yielding the following scores (Table 4.3):

Table 4.3: *Evaluation metrics for the final model with re-ranking (Final Iteration).*

Metric	Value
Precision	0.05
Recall	0.40
F1-Score	0.09

This outcome indicated a "heuristic ceiling," a limitation frequently observed in hybrid AI-knowledge graph systems. Without supervised training on clinically validated datasets, the post-hoc heuristics were inadequate to correct the structural noise inherent in the generative model's output. The persistent high rate of false positives (Figure 4.3) posed a significant obstacle to adoption, as clinicians would be unlikely to trust a tool that created additional filtering work. This finding aligns with the broader challenges in AI implementation discussed by Al Kuwaiti et al. and Varnosfaderani & Forouzanfar, where the gap between a model's technical capability and its real-world clinical utility remains a significant hurdle. [16, 17].

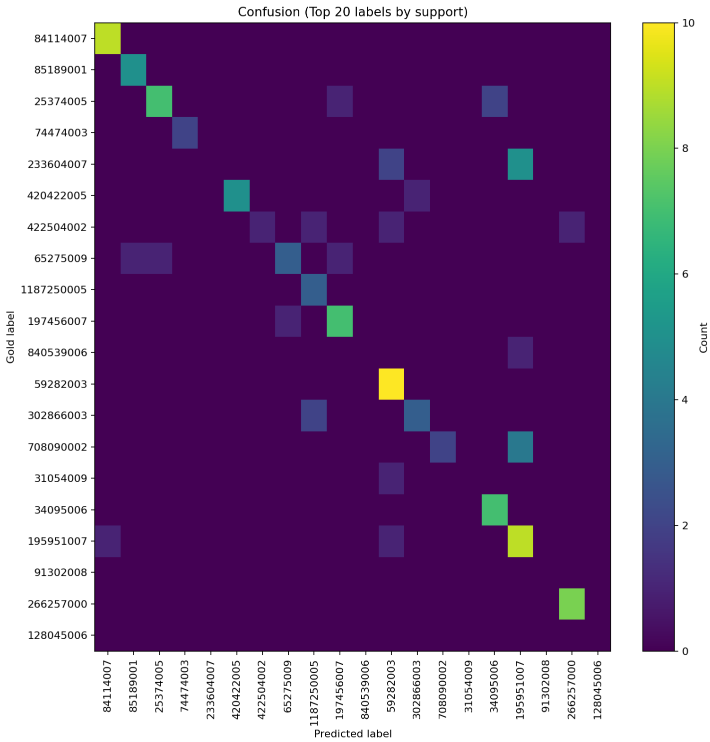

Figure 4.3: *Confusion Matrix of the final model after re-ranking.*

4.2 Qualitative Findings: The Clinical Context for a Mobile Assistive Tool

To complement the quantitative performance evaluation of the artefact, semi-structured interviews were conducted with three practising clinicians. The primary objective was to understand the practical environment of their daily operations, their frustrations with current technologies and their perceptions of a mobile, AI-enhanced clinical tool. This qualitative investigation was essential within a DSR framework, as the value of a technological product is ultimately determined by its utility and acceptance among end-users. As Muhiyaddin et al. asserted, understanding the influence of CDSS on physicians is crucial for its effective deployment [22]. Thematic analysis of the interview transcripts identified four primary themes that define the problem space and elucidate user requirements for the proposed solution.

4.2.1 Theme 1: A System Under Strain - Inefficiency, Fragmentation and Risk

A recurring theme identified in all interviews was the clinical environment's operation under considerable strain, hindered by disjointed information systems, a dependence on outdated or unsuitable tools and manual workflows that contributed to both inefficiency and risk. This finding corresponds with the recognised challenges of data silos and the absence of interoperability, which are considered significant barriers to contemporary, efficient healthcare [2, 3]. Clinicians reported a daily experience of managing a fragmented array of uncoordinated systems. This fragmentation hindered a comprehensive approach to patient care, necessitating the use of various tools such as physical paper files, pagers and multiple, non-integrated EMRs. Doctor 1, for instance, currently utilised four distinct EMRs and characterised their team's reliance on Microsoft Excel for monitoring sensitive patient results as "very dangerous" because of the significant risk of human error. This ad-hoc approach to data management highlights the pressing requirement for robust, integrated data platforms as articulated by Althati et al. [11].

4.2.2 Theme 2: Unanimous Need for a Centralized, Real-Time Dashboard

Despite the complexities and variations in their daily workflows, the clinicians expressed unanimous and enthusiastic support for the core concept of the mobile application: a single, real-time, unified view of their patients. A centralised dashboard that aggregates data from various sources was recognised as an effective solution to the fragmentation issues outlined in the previous theme. This finding is consistent with the vision for integrated healthcare ecosystems articulated by Pushadapu, which utilise real-time data to enhance care coordination and workflow efficiency [3].

All participants communicated the value proposition of the artefact's key feature with clarity and force. Doctor 3 commended the idea of linking all patients in a single, accessible view as "very, very, very, very, very useful" for improving continuity of care and saving important time. Doctor 2 expressed this sentiment, calling such tool "essential" for successfully prioritising patients and developing management regimens. Doctor 1 went on to emphasise its importance in pre-hospital emergency treatment, where quick access to a patient's whole medical history can save lives. A clear consensus emerged among clinicians regarding the identification of the most critical key data points for an at-a-glance dashboard. Essential information required for prompt decision-making encompassed medical and drug history (particularly regarding allergies), current patient vitals and the latest significant lab results (including CBC and CRP). This aspiration for an integrated, real-time flow of vital signs and laboratory data exemplifies the tenets of contemporary remote patient monitoring systems, which seek to furnish clinicians with prompt, actionable information [4, 8].

4.2.3 Theme 3: AI as a Cautious Assistant, Not a Confident Authority

The interviews indicated a varied and sophisticated perspective on the role of AI in diagnostics. Clinicians expressed openness to the potential of AI; however, they categorically opposed the notion of it supplanting their clinical judgement. This perspective is consistent with ethical frameworks that advocate for AI as a tool to enhance, rather than replace, human expertise [21]. This significance of AI was predominantly observed in its capacity to enhance cognitive processes. Doctor 1 utilised AI for brainstorming, generating differential diagnoses and connecting information yet maintained that it cannot fully replace the diagnostic process. This viewpoint positions the optional CDSS as a mechanism for cognitive assistance rather than as an independent entity.

The primary scepticism arose from the notion that AI lacks the ability to apply clinical context effectively. The quantitative findings of this study strongly validated this concern. Doctor 3 evaluated the application as "useful minus the AI" for diagnostic purposes, arguing that a clinician's reasoning is more "filtered" and that an AI might overreact to misleading data, as illustrated by a low SpO2 reading in a patient who is otherwise stable and communicative. The final algorithm's precision score of 0.05 empirically illustrates this issue; its propensity to produce a significant number of false positives backed up the clinicians' expert intuition that the system, in its present form, lacked adequate clinical filtering to be deemed reliable. This lack of contextual reasoning constituted a significant barrier to trust. Doctor 1 emphasised that technology cannot replace the intuitive judgement of an experienced clinician, who is capable of identifying life-threatening conditions through subtle cues that machines may fail to detect. This feedback suggests that for an AI-powered CDSS to be effective, it must demonstrate both accuracy and transparency in its reasoning, a notable challenge emphasised in the literature concerning the application of AI in healthcare environments [17].

4.2.4 Theme 4: The Non-Negotiable Requirements - Speed, Reliability and Usability

For the artefact to be successful adopted in high-stakes clinical setting, the interviews delineated a definitive set of non-negotiable criteria: the tool must be extraordinarily rapid, entirely dependable and intuitively user-friendly. These user-centric requirements reflect known methodologies for effective CDSS implementations, emphasising seamless workflow integration and elevated user acceptance [7]. The inability to fulfil these fundamental functional expectations was recognised as a significant obstacle to adoption, irrespective of the system's technological complexity.

The necessity of speed was a recurring topic. Clinicians worked in a time-sensitive setting where any technological friction could be more detrimental than beneficial. Doctor 2 was clear on this issue, stating that they would "rather opt for the actual physical thing" if the application was delayed. This was framed as a clinical emergency, not just a convenience issue. "The more minutes we lose... the loss of function might vary," as Doctor 1 observed in relation to stroke patients, establishing system delay as a direct influence on patient outcomes. This feedback demonstrated that the mobile artefact must offer data and insights more quickly than existing manual processes in order to be viable.

In addition to speed, the system's absolute reliability was emphasised as a critical factor in clinical safety and professional liability. Doctor 1 raised a significant concern by posing the critical question "If there is a real emergency and the app did not flag it, who will be responsible?" This underscores the significant trust a clinician must invest in such a tool and corresponds with the ethical considerations of accountability in medical AI as discussed by Farhud and Zokaei [21].

Finally, intuitive usability was identified as essential for adoption across a user base with diverse technical skills. Doctor 1 strongly warned against the assumption that all clinicians are tech-savvy, stating, "don't take it for granted, because you're dealing with a doctor... It means that you can create a program with some, complex software... and they will [not] be able to grasp it." This insight is critical, as a tool that is difficult to use will simply be abandoned, a finding consistent with literature on the barriers to the adoption of digital health technologies [10]. The impact of CDSS on clinicians is heavily dependent on its ease of use; a complex interface can increase cognitive load and disrupt workflows, ultimately negating the intended benefits [22]. Therefore, a user-friendly and intuitive interface is not a superficial design choice but a core functional requirement for the artefact's success.

4.3 Discussion: Synthesising Quantitative and Qualitative Findings

The genuine insights of a DSR project arose not from the quantitative or qualitative data in isolation, but from their integration. By integrating the objective performance indicators of the artefact with subjective, contextually rich viewpoints of its intended users, a more profound comprehension of the research problem was attained. This section synthesises the quantitative results from the algorithm's assessment with the qualitative themes derived from the clinicians' interviews. This integration of the "what" (algorithm performance) with the "why" (clinicians' real-world demands and concerns) facilitates a comprehensive study of the artefact's feasibility and yields unambiguous, evidence-based implications for its final design.

This discussion is organised to construct this synthesised argument. The algorithm's particular performance constraints will be closely linked to the issues highlighted by clinicians, illustrating the manual validation of the two data streams. Subsequently, the research questions presented in Chapter 1 will be examined and addressed utilising this consolidated material. The discussion will thereafter examine the study's limitations before finishing with the ramifications of these findings for the design of the final mobile application artefact.

4.3.1 Triangulating the Data: Why the Algorithm's Performance is a Key Finding

This study's primary finding was the strong correlation between the algorithm's quantitative deficiencies and the clinician's qualitative apprehensions. The final precision score of 0.05 served as an empirical validation of clinicians' primary concern regarding the potential for "suggestion overload" from an AI that lacks nuanced clinical context. The algorithm's propensity to generate numerous incorrect suggestions, characterised by a significant number of false positives, exemplified the "noise" cautioned against by Doctor 3. This underscores the necessity for a high-precision, context-aware system for clinicial adoption. This finding aligns with the existing litetature on CDSS where research by Sutton et al. and Muhiyaddin et al. indicates that a high frequency of irrelevant or low-value alerts results in "alert fatigue." This phenomenon causes users to disregard potentially significant information, ultimately diminishing trust and compromising the system's intended function [7,22].

The ongoing challenge to enhance the F1=score beyond the "heuristic ceiling" underscored the significant difficulty of developing a system sufficiently trustworthy to gain clinicians' trust. The qualitative input was clear: the system must be swift, dependable and secure. Doctor 1 incisively enquired, if the application neglected to identify a genuine emergency, "who will be responsible then?". The emphasis on accountability and safety is a pivotal subject in the literature about the ethical ramifications of AI in healthcare, where authors such as Farhud and Zokaei examine the intricate challenges of attributing responsibility in context of automated systems influencing clinical decision-making [21]. The current loosely-coupled architecture's lack of high dependability indicated that a more advanced method is required to fulfil these elevated demands. This corresponds with extensive research in digital health, indicating the effective AI tool implementation necessitates profound integration that transcends basic heuristics to tackle fundamental concerns of trust, safety and accountability [16,17].

4.3.2 Findings in Relation to the Research Questions

The synthesised results from this mixed-methods study provide clear responses to the research questions outlined in Chapter 1. The discussion of each question is enriched by the triangulation of quantitative and qualitative evidence, illustrating the interplay between the artefact's performance and clinicians' real-world needs.

The first research question considered the essential features and workflow improvements required for clinicians to enhance real-time access to patient data and documentation through a mobile dashboard. The findings point to a clear conclusion: a centralised, real-time dashboard that consolidates patient data from fragmented systems is the most essential feature, as consistently highlighted by interviewees. This directly addresses the issue of information silos identified in Theme 1 and in the literature [2, 3]. The principal workflow improvement is the transition from the lengthy process of manual, paper-based data collection to immediate, point-of-care access to a comprehensive patient view. The most critical data elements were identified as medical and drug history, real-time vital signs, and recent significant laboratory results.

The second research question explored the key privacy and security concerns associated with using mobile applications to manage sensitive patient data. Clinicians were notably less concerned with technical safeguards, expressing confidence in conventional protections such as passwords and biometrics. Their primary concern was the clinical reliability and safety of the system. The form of "security" most valued was the assurance that data were accurate and alerts were meaningful. This shifts the security discussion from a purely technical issue to a socio-technical one, emphasising that trust is established not only through encryption and access control but also through the system's demonstrated reliability, a factor critical to the adoption of digital health technologies [10].

The third research question examined how real-time patient monitoring through the dashboard influences the speed and effectiveness of clinicians' responses to critical events. All clinicians agreed that context-aware alerts would significantly improve both response time and effectiveness. However, they emphasised that this benefit is contingent on alert quality. Current systems based on simplistic, inaccurate red flags were criticised for generating unnecessary workload. The ideal system would instead deliver clinically meaningful, context-rich, AI-filtered alerts, advancing beyond basic out-of-range thresholds and aligning with recent literature on remote patient monitoring [4, 8].

4.3.3 Limitations and Acknowledgement of the Broader AI Landscape

The limitations of this project offer crucial context for understanding the results and outline distinct, encouraging paths for future exploration. This study's limitations can be classified into two main categories: methodological and technical.

Methodological Limitations

The main methodological limitations apply to the data utilised in both the qualitative and quantitative aspects of the study.

First, the qualitative findings were based on a small sample of three clinicians working in the Maltese healthcare system. while their insights were extensive and detailed, giving valuable background for this DSR project, a sample of this size is not statistically representative. The highlighted motifs may not apply to all clinical specialities, seniority levels or national healthcare systems. Future research should seek to validate these findings with a wider, more diverse group of clinicians to ensure their broader application.

Second, The algorithm's quantitative assessment was performed on a synthetic dataset. This approach was essential for safeguarding patient privacy and establishing a controlled evaluation environment; however, synthetic data cannot entirely emulate the intricacy, ambiguity and disorder inherent in authentic clinical notes, which frequently include typographical errors, non-standard abbreviations and complex grammatical constructions. Thus, the algorithm's performance measures ought to be seen as a reflection of its capabilities in a regulated environment. Future validation will demand testing the artefact on a substantial, de-identified corpus of authentic patient data, which will involve addressing considerable ethical and data governance concerns.

Technical Limitations

The technical limitations were centred on the choice of the AI model and the overall system architecture.

The adoption of the MedLLaMA2 model was pragmatic design, influenced by the project's limitations, especially the requirement for a locally deployable model to provide complete data privacy and security. This model did not exemplify the pinnacle of medical AI advancements. More robust, large-scale models such as Med-PaLM 2 and MedGemma now exist and have exhibited enhanced performance across several clinical NLP benchmarks. Research by De Vito et al. and Searle et al. illustrates that the selection of model architecture is essential for performance in intricate NLP tasks [9, 13]. Subsequent versions of this research should concentrate on a comparative assessment of these advanced models. This effort must meticulously consider the substantial computational expenses and the data privacy compromises associated with utilising cloud-based APIs, which may be necessary for accessing these larger models.

Moreover, the artefact's "loosely-coupled" architecture, which integrated idea extraction (LLM), candidate retrieval (search) and re-ranking, possessed intrinsic limits. An error in the first LLM generating phase, such as the inability to recognise a crucial concept or the proposal of an irrelevant one, could not be readily rectified in subsequent phases. This error propagation likely led to the observed "heuristic ceiling" in the quantitative results. A more comprehensive strategy, indicating a substantial direction for future research, would involve investigating an end-to-end fine-tuning of a clinical language model. Training a model directly on a collection of clinicial notes and their corresponding SNOMED CT codes could eliminate the fragile intermediate mapping phase, potentially resulting in significant enhancements in precision and overall system reliability.

4.3.4 Implications for the Design of the Final Artefact

The synthesised results from this mixed-methods evaluation yield a coherent and justifiable set of design principles for the final mobile application artefact. The iterative process of development, evalution and user feedback results in a design that is not only technologically feasible but also rooted in the intricate realities of clinical practice. The principal ramifications for the artefact's design are as follows:

1. Prioritise the dashboard as the core feature: The qualitative interviews revealed a clear need for a centralised, real-time patient data dashboard. This feature immediately tackles the most serious pain issue reported by clinicians, system fragmentation and has the most potential for immediate ben-

eficial impact on clinical workflows. As a result, the design of the final artefact must prioritise the faultless, rapid and dependable delivery of this fundamental feature above all else.

2. Position AI as a supportive, transparent tool: The quantitative results, which revealed a "heuristic ceiling" and persistent low precision, directly validate the clinicians' scepticism towards AI as a definitive authority. Consequently, the AI diagnostic feature must be framed as a secondary, supportive tool intended for "brainstorming" or generating a differential diagnosis, not as a primary diagnostic engine. This aligns with the clinicians' expressed wishes for an "assistant" and with the ethical principle of keeping the human clinician "in the loop" [21].

3. Design for precision, trust and usability: The non-negotiable requirements of speed, reliability and usability must be central to the final design. Future development of the AI component must relentlessly focus on improving precision over recall. To build trust, any suggestions shown to the user must be transparent, allowing them to understand *why* a suggestion was made, a key factor for building confidence in any CDSS [7]. Finally, the user interface must be minimalist and intuitive, catering to a user base with diverse technical skills to avoid the risk of the tool being abandoned due to complexity, as warned by clinicians and supported by literature on the adoption of digital health technologies [10, 22].

Chapter 5: Conclusions and Recommendations

The dissertation concludes by synthesising the key findings derived from the Design Science Research (DSR) methodology. The investigation set out to design, develop and evaluate a mobile clinical dashboard intended to address the persistent challenges of data fragmentation and workflow inefficiency in inpatient care. Employing a mixed-methods approach, the work combined the technical evaluation of an AI-powered diagnostic suggestion artefact with a qualitative exploration of clinicians' practical needs.

This closing section begins with a concise summary of the research, revisiting the core problem and the process undertaken to address it. It then highlights the principal outcomes of the study, integrating insights from both the quantitative and qualitative evidence presented in Chapter 4. Attention then turns to the primary and secondary contributions of the work, considering its significance for both academic discourse and clinical practice. Finally, the discussion moves towards future directions, presenting actionable recommendations that stem from the findings and limitations and offering a roadmap for advancing the foundations laid by this thesis.

5.1 Summary of Research

The central problem addressed by this study was the significant challenge faced by hospital clinicians in accessing timely, comprehensive and precise patient information. As established in the literature review, the modern healthcare IT landscape is often a fragmented ecosystem of disparate systems, leading to workflow inefficiencies and potential risks to patient safety. This research sought to address this gap by following a DSR methodology to create and evaluate a novel artefact: a mobile clinical dashboard featuring an AI-driven diagnostic suggestion tool.

The research was conducted via a multi-phase process. An initial literature review identified the principal obstacles and current solutions in healthcare interoperability, remote patient monitoring and Clinical Decision Support Systems (CDSS). This informed the design and development of a functional software artefact. The artefact's core AI component then underwent multiple iterative cycles of quantitative evaluation to assess its efficacy in converting unstructured clinical notes into structured SNOMED clinical terminology. Concurrently, semi-structured interviews were conducted with practising clinicians to gather rich qualitative data regarding their workflows, requirements and perceptions of such a tool. The concluding phase of the research entailed a synthesis of the quantitative and qualitative findings to deliver a comprehensive evaluation of the artefact and its potential clinical utility.

5.2 Principal Findings of the Study

The mixed-methods examination revealed many key conclusions that offer a detailed insight of the issue at hand:

1. **A Clear and Unanimous Need for Data Centralisation:** The qualitative find-

ings unequivocally confirmed that clinicians are burdened by fragmented systems and manual workflows. There was enthusiastic and universal support for the core value proposition a mobile dashboard that consolidates patient data into a single, real-time view. This finding validates the initial problem statement and confirms that the primary utility of the proposed artefact lies in its ability to address system-level inefficiency.

2. **AI Performance Reveals a "Heuristic Ceiling":** The iterative quantitative evaluation of the diagnostic suggestion algorithm revealed a critical technical finding. Despite multiple architectural enhancements, the system persistently struggled with low precision (0.05 in the final run), indicating a "heuristic ceiling." This demonstrates that a loosely-coupled architecture connecting a general-purpose LLM to a knowledge base via post-hoc filtering is insufficient to achieve the high level of accuracy required for a reliable clinical tool.

3. **Clinical Scepticism of AI is Rooted in a Need for Context:** The qualitative data provided a clear explanation for the quantitative results. Clinicians' scepticism towards AI stemmed not from a general distrust of technology, but from a sophisticated understanding of its limitations. They valued AI as a "brainstorming" assistant but were wary of its lack of clinical context; a concern that was empirically validated by the algorithm's high rate of false positives. This highlights that for clinical AI, trust is a function of not only accuracy but also contextual awareness.

4. **Reliability Outweighs Technical Security as a Primary Concern:** While ac-

knowledging the importance of data security, clinicians were far more concerned with the clinical reliability and safety of the artefact. Their questions about accountability in the case of a missed alert reframe "security" within a socio-technical framework, where trust is contingent on proven, dependable performance.

5.3 Contribution to Knowledge

This research offers multiple contributions to the domains of clinical informatics and software development:

The main contribution is the definitive finding that a loosely-coupled architecture for diagnostic suggestions is inadequate for dependable clinical use. This study offers significant evidence for the academic and development worlds by carefully recording recurrent attempts to enhance performance and the eventual "heuristic ceiling." It warns against simplistic implementations and directs future efforts towards more integrated, end-to-end solutions such as fine-tuning, which are essential to address the difficulties of clinical context and precision.

The secondary contributions are dual in nature. This report presents a comprehensive overview of the existing workflow challenges and user requirements within the Maltese healthcare system, providing significant insights for local digital health efforts. Secondly, it illustrates a pragmatic application of the DSR methodology, showcasing how a mixed-methods approach may effectively triangulate technical performance with user-centred feedback to inform the development of a therapeutically pertinent technological artefact.

5.4 Recommendations for Future Research

In light of the findings and constraints of this study, the subsequent recommendations are put of for future research in this domain:

1. **Explore End-to-End Fine-Tuning of Clinical LLMs:** To effectively address the "heuristic ceiling," it is essential to transition from a loosely-coupled architecture. Future research ought to concentrate on optimising a specialised clinical language model by directly utilising a substantial, de-identified corpus of clinical notes paired with the corresponding SNOMED CT codes. This would establish a more comprehensive, end-to-end system and signifies a crucial and essential advancement in this area of research.

2. **Conduct a Comparative Study of State-of-the-Art Models:** More robust, large-scale medical models like Med-PaLM 2 and MedGemma should be compared in future studies. In addition to requiring a comprehensive examination of the related computational and data privacy trade-offs, such a study would offer a more precise baseline for the AI component's performance potential.

3. **Expand the Qualitative Inquiry and Conduct Usability Studies:** To enhance generalisability, the qualitative findings must be supported with a larger and more diverse group of clinicians from various specialities. The next phase in the DSR cycle involves carrying out formal usability studies of the mobile application prototype to collect feedback regarding the user interface and its integration into the clinical workflow.

5.5 Concluding Summary

This dissertation aimed to tackle the essential requirement for enhanced access to patient data in hospital inpatient settings by designing and evaluating a mobile clinical dashboard. The research journey demonstrated that although clinicians robustly confirmed the central notion of a centralised patient data dashboard, the use of AI for diagnostic recommendations presents a significantly intricate difficulty. The principal conclusion of this study is that the effective creation of such a tool is not primarily a technological issue, but a socio-technical one, wherein clinical trust, workflow integration and contextual awareness are equally significant as algorithmic efficacy.

The final artefact, designed according to the concepts and principles established from this research, emphasises the provision of a dependable, user-friendly dashboard while presenting its AI functionalities as a transparent, auxiliary resource for clinical ideation. This paper presents a candid evaluation of the existing landscape and proposes a pragmatic roadmap for the future advancement of technologies that can effectively and securely support clinicians in their essential duties, despite the lengthy journey towards completely autonomous and accurate diagnostic AI.

List of References

[1] Mark NK Saunders, P Lewis, and A Thornhill. Research methods for business students, 03 2023.

[2] Reda Radwan, Mohammad Hossain, Hesham El-Bahkiry, Husham Homeida, Badreldin Ahmed, and Ali Osman. Optimizing electronic health records for enhanced clinical decision support and improved patient care through advanced data integration and machine learning algorithms. *Panamerican Mathematical Journal*, 34, 2024.

[3] Navajeevan Pushadapu, Sr Analyst, Healthpoint Hospital, and Abu Dhabi. Real-time integration of data between different systems in healthcare: Implementing advanced interoperability solutions for seamless information flow. *Distributed Learning and Broad Applications in Scientific Research Annual*, 6, 2020.

[4] Thanveer Shaik, Xiaohui Tao, Niall Higgins, Lin Li, Raj Gururajan, Xujuan Zhou, and U. Rajendra Acharya. Remote patient monitoring using artificial intelligence: Current state, applications, and challenges. *WIREs Data Mining and Knowledge Discovery*, 13:e1485, 01 2023.

[5] Jennifer Corny, Asok Rajkumar, Olivier Martin, Xavier Dode, Jean-Patrick Lajonchère, Olivier Billuart, Yvonnick Bézie, and Anne Buronfosse. A machine learning–based clinical decision support system to identify prescriptions

with a high risk of medication error. *Journal of the American Medical Informatics Association*, 27:1688–1694, 09 2020.

[6] John W Creswell and J. David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. SAGE Publications, 6 edition, 2023.

[7] Reed T. Sutton, David Pincock, Daniel C. Baumgart, Daniel C. Sadowski, Richard N. Fedorak, and Karen I. Kroeker. An overview of clinical decision support systems: benefits, risks, and strategies for success. *npj Digital Medicine*, 3, 02 2020.

[8] Saurabh Gupta. Integrating iot and ai for real-time patient monitoring. *International Journal of Science and Research (IJSR)*, 13:266–270, 12 2024.

[9] Vito De, Gabriele, Ferrucci Filomena, and Athanasios Angelakis. Design and evaluation of a cdss for drug allergy management using llms and pharmaceutical data integration, 2024.

[10] Maksut Senbekov, Timur Saliev, Zhanar Bukeyeva, Aigul Almabayeva, Marina Zhanaliyeva, Nazym Aitenova, Yerzhan Toishibekov, and Ildar Fakhradiyev. The recent progress and applications of digital technologies in healthcare: a review. *International Journal of Telemedicine and Applications*, 2020:1–18, 12 2020.

[11] Chandrashekar Althati, Manish Tomar, and Lavanya Shanmugam. Enhancing data integration and management: The role of ai and machine learning in modern data platforms enhancing data integration and management: The role

of ai and machine learning in modern data platforms enhancing data integration and management: The role of ai and lavanya shanmugam. *Journal of Artificial Intelligence General Science JAIGS*, 2, 2024.

[12] Tayebeh Baniasadi, Mehdi Hassaniazad, Sharareh R Niakan Kalhori, Mehraban Shahi, and Marjan Ghazisaeedi. Optimized patients' length of hospital stay with interventions based on health information technology: A review study. *PubMed*, 271:69–76, 06 2020.

[13] Thomas Searle, Zina Ibrahim, James Teo, and Richard J. B. Dobson. Discharge summary hospital course summarisation of in patient electronic health record text with clinical concept guided deep pre-trained transformer models. *Journal of Biomedical Informatics*, 141:104358, 05 2023.

[14] Christophe Gaudet-Blavignac, Vasiliki Foufi, Mina Bjelogrlic, and Christian Lovis. Use of the systematized nomenclature of medicine clinical terms (snomed ct) for processing free text in health care: Systematic scoping review. *Journal of Medical Internet Research*, 23:e24594, 01 2021.

[15] Khushbu Agarwal, Tome Eftimov, Raghavendra Addanki, Sutanay Choudhury, Suzanne Tamang, and Robert Rallo. Snomed2vec: Random walk and poincaré embeddings of a clinical knowledge base for healthcare analytics. *arXiv:1907.08650 [cs, stat]*, 07 2019.

[16] Ahmed Al Kuwaiti, Khalid Nazer, Abdullah Al-Reedy, Shaher Al-Shehri, Afnan Al-Muhanna, Arun Vijay Subbarayalu, Dhoha Al Muhanna, and Fa-

had A. Al-Muhanna. A review of the role of artificial intelligence in healthcare. *Journal of Personalized Medicine*, 13, 06 2023.

[17] Shiva Maleki Varnosfaderani and Mohamad Forouzanfar. The role of ai in hospitals and clinics: Transforming healthcare in the 21st century. *Bioengineering*, 11:337, 04 2024.

[18] Florence Ranchon, Sébastien Chanoine, Sophie Lambert-Lacroix, Jean-Luc Bosson, Alexandre Moreau-Gaudry, and Pierrick Bedouch. Development of artificial intelligence powered apps and tools for clinical pharmacy services: A systematic review. *International Journal of Medical Informatics*, 172:104983, 04 2023.

[19] Parthasarathy G, Soumya T R, Ramanathan L, and Ramesh P. Improvised approach for real time patient health monitoring system using iot. *Advances in Parallel Computing*, 11 2020.

[20] Paras Malik, Monika Pathania, and VyasKumar Rathaur. Overview of artificial intelligence in medicine. *Journal of Family Medicine and Primary Care*, 8:2328–2331, 07 2019.

[21] Dariush D. Farhud and Shaghayegh Zokaei. Ethical issues of artificial intelligence in medicine and healthcare. *Iranian Journal of Public Health*, 10 2021.

[22] Raghad Muhiyaddin, Alaa A. Abd-Alrazaq, Mowafa Househ, Tanvir Alam, and Zubair Shah. The impact of clinical decision support systems (cdss) on

physicians: A scoping review. *Studies in Health Technology and Informatics*, 272:470–473, 06 2020.

[23] Iacopo Cricelli, Ettore Marconi, and Francesco Lapi. Clinical decision support system (cdss) in primary care: from pragmatic use to the best approach to assess their benefit/risk profile in clinical practice. *Current Medical Research and Opinion*, 38:827–829, 03 2022.

Chapter A: Interview Protocol

The following semi-structured interview questions were used to guide the conversations with clinicians.

Section 1: Workflow Reality (Before Showing Prototype)

1. Can you walk me through how you currently monitor patients during your rounds or shifts? What tools do you use most?
2. What are the biggest constraints or frustrations you encounter when managing patients across multiple wards?
3. How do you decide who to prioritise when handling a full patient list?
4. How are you notified of urgent patient issues today? What do you wish could work better about that process?
5. When a patient's condition changes suddenly, what's your current process for diagnosis? Do you use any tech tools or rely more on experience?

Section 2: Feedback on the Mobile App Concept

("Imagine a mobile app that acts as a bedside companion, giving you a real-time, single-screen view of all your patients – including their vitals, statuses and alerts – helping you stay on top of priorities during rounds. It also integrates AI and SNOMED CT to suggest possible diagnoses based on abnormal vital signs or patient history.")

1. What are your initial impressions of this concept, and does it stand out as useful or concerning?

2. How valuable would it be to see all your patients on one screen with key data? Which 2-3 data points would be most important at a glance?

3. How useful would it be to receive real-time alerts when a patient's condition deteriorates? Would this change how quickly you respond?

4. The app uses AI and SNOMED CT to suggest diagnoses. Would you trust or consider these suggestions in practice? Do you have any concerns?

Section 3: Usability, Integration & Broader Impact (After Showing Prototype)

1. How important is ease of use in a tool like this? What would make or break your willingness to use it?

2. How do you imagine this app working alongside your current EMR or other clinical systems?

3. Do you have any security, reliability or workflow concerns about using a mobile app for patient data?

4. Is there anything missing from what we've described that you would expect or want in a tool like this?

Conclusion

1. Do you have any other thoughts or suggestions we haven’t covered?

2. Would you be open to trying a future prototype?

Chapter B: Participant Information Sheet

Title of Research: Real-Time Patient Data and Documentation Access Through Mobile Application: A Secure and Efficient Solution for Hospital Clinicians

Researcher: Matthew Schembri, BSc (Hons) Software Development

Dear Participant,

You are invited to take part in a research study as part of my final-year dissertation at MCAST. Before deciding, please take the time to read the following information carefully. It outlines the purpose of the study, what your participation involves and your rights as a participant. If anything is unclear or you'd like more information, please feel free to contact me.

What is the purpose of the study?

This study explores how a mobile dashboard application can support clinicians by providing secure, real-time access to inpatient data. The system is designed to enhance clinical workflows through mobile documentation, SNOMED CT integration, and alerts linked to patient monitoring. As a clinician, your input will help evaluate how well the system fits the needs of your day-to-day practice.

Why have I been invited?

You have been selected because of your professional experience in inpatient care and clinical decision-making. Your perspective will provide valuable insights into the usefulness, usability and potential impact of this tool.

What will I be asked to do?

You will be invited to participate in a semi-structured interview lasting approximately 20-30 minutes. Interviews may be conducted online or in person, depending on your preference and availability. Questions will focus on the app's usability, relevance and potential for clinical integration.

Do I have to take part?

Participation is entirely voluntary. You may choose not to answer any question and are free to withdraw from the study at any time without providing a reason. If you withdraw before your data is anonymised, it will be deleted. Once anonymised, it cannot be linked back to you.

Will my information be kept confidential?

Yes. All data will be handled under ethical and data protection standards. Interviews will be audio-recorded (with your consent), transcribed and anonymised. Your name and any identifiable details will be removed. Only the researcher will have access to the raw data, which will be stored securely and encrypted.

What are the risks or benefits?

There are no known risks in taking part. While you may not receive direct benefit, your feedback will contribute to the development of more effective digital tools for clinical environments.

What will happen to the results?

The anonymised results will be included in my final dissertation submitted to MCAST. With your permission, they may also contribute to future academic publications or presentations. You may request a copy of the final dissertation upon completion.

Who is organising this research?

This research is part of my bachelor's degree in software development under the supervision of Ms. Mary Grace Seguna Agius at the Institute of ICT, MCAST.

Contact for further information:

If you have any questions or wish to discuss any part of the study, please contact me at: matthew.schembri.d102517@mcast.edu.mt

Thank you for considering taking part in this study.

Kind regards,

Matthew Schembri
MCAST Software Development Student