Faculty of Engineering of the University of Porto

Framework development to support AI-based medical technology innovation

José Miguel Palha de Araújo dos Santos

Dissertation for the degree of

Master in Molecular Bioengineering

Supervisor: Daniel Marques de Vasconcelos
Co-supervisor: Marta Vranas

June 24, 2025

Resumo

A indústria MedTech tem sido crucial para aumentar a longevidade e qualidade de vida da população, bem como a eficiência dos cuidados de saúde. Estima-se que atinja um valor global de 629 mil milhões de dólares em 2025, com perspetivas promissoras impulsionadas por novas tecnologias, especialmente as que utilizam dados médicos. No entanto, o setor enfrenta dificuldades na disseminação de inovações, refletidas na elevada taxa de insucesso das startups. Estas barreiras agravam-se no caso de tecnologias com base em inteligência artificial (IA), dificultando a transferência tecnológica e impacto.

Para responder às questões "Quais os principais desafios que dificultam a adoção de tecnologias médicas com base em IA?" e "Que boas práticas podem apoiar o desenvolvimento de um framework orientador?", este estudo seguiu uma abordagem qualitativa, exploratória, e sequencial combinando entrevistas, não estruturadas e semiestruturadas, com revisão da literatura focada nas barreiras e promotores da inovação MedTech com base em IA. O primeiro grande contributo é uma análise abrangente dos principais desafios à inovação MedTech com base em IA, acompanhada de boas práticas para os ultrapassar. Estes foram organizados em oito dimensões-chave: 1) Identificação da Necessidade; 2) Acesso a Dados; 3) Viabilidade Técnica e Clínica; 4) Conformidade Regulatória; 5) Comercialização; 6) Implementação Clínica; 7) Ética, Justiça e Sustentabilidade; e 8) Transparência. Entre as boas práticas identificadas, a validação local e monitorização destacam-se como essenciais para a integração responsável de IA na saúde.

Com base nestes resultados, o segundo grande contributo é o VAMI Framework (Value-driven AI MedTech Innovation), criado para orientar inovadores no desenvolvimento e implementação eficaz de soluções médicas com base em IA. Este foi validado através de entrevistas com especialistas em inovação em saúde, cujo feedback confirmou a sua relevância e sugeriu melhorias, como uma representação mais clara da natureza cíclica do processo. O VAMI Framework foi ainda analisado através do caso de estudo CINDERELLA, um projeto de MedTech com base em IA que ajudou a ilustrar e testar o framework num contexto real. O projeto aplicou com sucesso várias das práticas recomendadas, incluindo a constituição de uma equipa multidisciplinar e o desenvolvimento da solução como uma ferramenta de apoio à decisão, para complementar, e não substituir, os profissionais de saúde. No entanto, revelou também desafios, nomeadamente no planeamento regulatório, que o VAMI Framework pode ajudar a antecipar e mitigar, reduzindo o risco de gastos e atrasos no caminho até ao mercado.

Esta dissertação pretende fornecer orientações práticas aos inovadores da área MedTech, através da introdução de um framework que identifica desafios críticos e estratégias para soluções médicas com base em IA. O estudo reconhece, contudo, limitações como a natureza qualitativa e exploratória, a reduzida amostra de entrevistas e o curto período de investigação. Investigações futuras devem validar o framework junto de stakeholders mais diversos, incluindo grandes empresas MedTech e instituições de saúde, assim como testá-lo em projetos reais, permitindo uma iteração e refinamento rápidos.

Palavras-chave: Inovação MedTech, Inteligência Artificial, Inovação MedTech baseada em IA, Dados de Saúde, Cuidados de Saúde Baseados em Valor

Abstract

The MedTech industry has played a vital role in improving populations' longevity and quality of life, as well as the efficiency of healthcare delivery. It is projected to reach a global market value of $629 billion by 2025 and shows exciting prospects as new technologies, particularly those leveraging medical data, enter the sector. However, the industry struggles to disseminate innovations, as reflected in the high failure rate of MedTech startups. This is further complicated by the challenges specific to artificial intelligence (AI)-based technologies, making successful technology transfer and real-world impact even more difficult to achieve.

To answer the research questions, "What main challenges impair the adoption of AI-based medical technologies?" and "What best practices can support the development of a guiding framework?", this study employed a qualitative exploratory sequential approach. It combined unstructured and semi-structured interviews with a literature review focusing on the barriers and enablers of AI-based MedTech innovation. The first major contribution of this dissertation is a comprehensive analysis of the core challenges in AI-based MedTech innovation, accompanied by best practices to overcome them and promote real-world impact. These were organised in eight key dimensions: 1) Need Identification; 2) Data Access; 3) Technical and Clinical Feasibility; 4) Regulatory Compliance; 5) Commercialisation; 6) Clinical Implementation; 7) Ethics, Fairness, and Sustainability; and 8) Transparency. Among the identified best practices, local validation and continuous performance monitoring emerged as essential safeguards for the responsible deployment of AI in healthcare.

Building on these insights, the second key output is the VAMI Framework (Value-driven AI MedTech Innovation), designed to guide innovators in effectively navigating the development and implementation of AI-based medical solutions. The framework was validated through interviews with healthcare innovation experts, whose feedback confirmed its relevance and suggested refinements, such as better representation of the cyclical nature of the innovation process. The VAMI Framework was further examined through the CINDERELLA case study, an AI-based MedTech project that helped illustrate and test the framework in a real-world setting. The project successfully implemented several recommended practices, including building a multidisciplinary team and developing the solution as a decision-support tool to complement, not replace, clinicians. However, it also revealed challenges, namely in regulatory planning, that the VAMI Framework could help anticipate and mitigate, reducing the risk of costly, time-consuming setbacks on the path to market.

This dissertation aims to provide actionable guidance for MedTech innovators by introducing a practical framework that outlines the most pressing challenges and corresponding strategies for the effective development and implementation of AI-based medical solutions. However, the study acknowledges several limitations, including its qualitative and exploratory nature, a small interview sample, and a short research period. Future research should focus on broader validation of the VAMI Framework with a more diverse set of stakeholders, including large MedTech corporations and healthcare institutions, as well as piloting the framework in live projects to enable rapid iteration and refinement.

Keywords: MedTech Innovation, Artificial Intelligence, AI-based MedTech Innovation, Healthcare Data, Value-based Healthcare

Acknowledgments

“Inspiration exists, but it has to find you working.” — Pablo Picasso

Comecei o meu percurso académico incerto sobre o que queria e, passados 5 anos, continuo sem saber. Mas tenho muita sorte, porque daquilo de que não tenho dúvidas é que vivi anos incríveis, enriqueci-me pessoal e academicamente, construí relações especiais e guardo memórias que, com certeza, não irei esquecer.

A todos os membros do SAL, por me terem acolhido, pelo tanto que me ensinaram, pela companhia e pelas palavras de encorajamento. Em especial, aos meus orientadores e professores, Daniel Vasconcelos e Marta Vranas, pelo voto de confiança, pelos recursos e oportunidades que me disponibilizaram, e pelo imenso apoio que me prestaram durante todo o percurso. Aprendi muito, senti-me desafiado e, mesmo quando o caminho se tornava mais nublado, senti-me confiante de que se iria fazer luz. A todos, não só pela orientação no trabalho, mas pelas conversas e incentivos a nível pessoal e profissional.

Aos meus colegas e amigos, pelos muitos bons momentos, conversas, partilhas de sonhos e desabafos... porque percorrer este caminho sem vocês não teria tido a mesma graça. Em especial, à Filipa e ao Francisco, pelas longas horas de estudo, pela paciência e pela motivação que acompanharam o nosso percurso. À família que fiz em Erasmus, pelas aventuras que partilhámos e pelas amizades que ali nasceram.

Aos meus irmãos, porque é muito reconfortante saber que tenho, e terei sempre, com quem contar, conversar e em quem me inspirar. E, em especial, aos meus pais, por me terem proporcionado tantas e tão boas experiências. Desde pequeno, incentivam-me a dar o meu melhor e tranquilizam-me, garantindo que isso chega e que os resultados, mais cedo ou mais tarde, hão de vir.

A todos vocês e a tudo o que me aconteceu,

O meu obrigado.

Index

RESUMO 1

ABSTRACT 3

ACKNOWLEDGMENTS 5

LIST OF FIGURES 8

LIST OF CONTENTS 9

ABBREVIATIONS 10

1. INTRODUCTION 11

2. THEORETICAL BACKGROUND AND KEY CONCEPTS 14

2.1 Medical Technology Innovation 14
2.2 Value-Based Healthcare (VBHC) 15
2.3 AI-driven Medical Technology Innovation 15
2.3.1 Real-world Success Cases of AI-based Medical Technology Adoption 17
2.4 Challenges and Enablers in AI-driven Medical Technology Innovation 18
2.4.1 Data Accessibility 19
2.4.2 Clinical Implementation 21
2.4.3 Ethics and Fairness 23
2.4.4 Regulation 24
2.5 Frameworks for AI-driven Medical Technology Innovation 26
2.5.1 Stanford's Biodesign Innovation Process: A Framework for Healthcare Technology Development 26
2.5.2 Limitations of Traditional Frameworks 28
2.5.3 New Frameworks Targeting the Application of AI to Healthcare 28

3. METHODOLOGY 30

3.1 Research Question, Objectives and Methodology 30
3.2 Research Design 32
3.3 Data Collection 33
3.4 Data Analysis 35

4. RESULTS AND DISCUSSION 37

4.1 The AI-based MedTech Innovation Landscape 37
4.2 Challenges and Best Practices in AI-based MedTech Innovation 38
4.2.1 Need Identification 39
4.2.2 Data Access 40
4.2.3 Technical and Clinical Feasibility 43
4.2.4 Regulatory Compliance 46
4.2.5 Commercialisation 49
4.2.6 Clinical Implementation 51
4.2.7 Ethics, Fairness and Sustainability 53
4.2.8 Transparency 55
4.3 Development of the VAMI Framework to Support AI-based MedTech Innovation 56
4.3.1 Guiding Questions of the VAMI Framework 56
4.4 Positioning the VAMI Framework within Broader Innovation Pathways: Stanford's Biodesign 66
4.5 Validation of the VAMI Framework 68
4.5.1 Framework Scope 68
4.5.2 Insights by Component 69
4.5.3 Structural Considerations and Final Reflections 73

5. CASE STUDY ON THE CINDERELLA PROJECT 75

5.1 The Problem 75

5.2 The CINDERELLA Solution 75
5.3 Aligning the VAMI Framework with the CINDERELLA project 76
5.4 Case Study Conclusions and Key Insights 81
6. SUSTAINABLE DEVELOPMENT GOALS 82
7. LIMITATIONS AND FUTURE IMPROVEMENTS 84
8. CONCLUSIONS 86
9. REFERENCES 88
APPENDIX I 95

List of Figures

Figure 1 Schematic representation of data collection, data management and processing, and healthcare applications. 16
Figure 2 Challenges to overcome for AI-based solutions' successful clinical application. ... 19
Figure 3 The three phases of the Stanford Biodesign process for healthcare technology development - Identify, Invent, and Implement - are shown, each comprising two specific stages. Key activities associated with each stage are outlined below..... 27
Figure 4 Flowchart of the qualitative exploratory sequential research design 31
Figure 5 Diagram of the methodology to analyse the results and discuss the findings. 36
Figure 6 Simplified overview of the AI-based MedTech innovation process. The diagram outlines key phases from need identification to outcome monitoring, highlighting the iterative nature of development through feedback loops. 38
Figure 7 A combined approach to support value-driven AI-based MedTech innovation. The guiding questions of the proposed VAMI Framework are aligned with the Biodesign phases - Identify, Invent, and Implement. 67

List of Contents

Table 1 Number of interviews conducted per dimension of the AI-based medical solution innovation process, separated by work setting. 34
Table 2 Key challenges and best practices for AI-based medical technology innovation: Need Identification. 39
Table 3 Key challenges and best practices for AI-based medical technology innovation: Data Access. 41
Table 4 Key challenges and best practices for AI-based medical technology innovation: Technical and Clinical Feasibility. 44
Table 5 Key challenges and best practices for AI-based medical technology innovation: Regulatory Compliance. 47
Table 6 Key challenges and best practices for AI-based medical technology innovation: Commercialisation. 49
Table 7 Key challenges and best practices for AI-based medical technology innovation: Clinical Implementation. 51
Table 8 Key challenges and best practices for AI-based medical technology innovation: Ethics, Fairness & Sustainability. 53
Table 9 Key challenges and best practices for AI-based medical technology innovation: Transparency. 55
Table 10 Thesis contribution to selected UN Sustainable Development Goals (SDGs). The table outlines relevant SDG targets, the corresponding contributions of this work, and suggested indicators for measuring impact. 83
Table 11 Unstructured interviews conducted to explore the field of AI in healthcare and identify key challenges and best practices in AI-based MedTech Innovation. 108
Table 12 Semi-structured interviews conducted to validate the proposed VAMI Framework 112

Abbreviations

AI	Artificial Intelligence
AI Act	Artificial Intelligence Act
ALTAI	Assessment List for Trustworthy Artificial Intelligence
API	Application Programming Interfaces
CADe	Computer-aided Detection
DICOM	Digital Imaging and Communications in Medicine
EHDS	European Health Data Space
EU	European Union
FDA	Food and Drug Administration
FICO	Fair Isaac Corporation
FL	Federated Learning
FRIA	Fundamental Rights Impact Assessment
GDPR	General Data Protection
HER	Electronic Health Record
HIPAA	Health Insurance Portability and Accountability Act
HL7 FHIR	Health Level Seven Fast Healthcare Interoperability Resources
IP	Intellectual Property
ISO	International Organisation for Standardisation
IVD	In Vitro Diagnostics
IVDR	In Vitro Diagnostics Regulation
KPIs	Key Performance Indicators
MedTech	Medical Technology
MDR	Medical Device Regulation
ML	Machine Learning
NHS	National Health Service
OECD	Organisation for Economic Co-operation and Development
OMOP CDM	Observational Medical Outcomes Partnership Common Data Model
PROMs	Patient-reported Outcome Measures
R&D	Research and Development
RCT	Randomised controlled trials
SaMD	Software as a Medical Device
SaaS	Software-as-a-service
SDGs	Sustainable Development Goals
SDQ	Smart Data Query
SNOMED CT	Systematised Nomenclature of Medicine - Clinical Terms
US	United States
VBHC	Value-based Healthcare
WGS	Whole-genome Sequencing
WHO	World Health Organisation
XAI	Explainable Artificial Intelligence

1. Introduction

Healthcare innovation has been a cornerstone of medical progress, playing a central role in the dramatic rise in human life expectancy and quality of life over the past century [1]. Today, health systems face growing pressures, such as an ageing population, increasing chronic disease burden, and rising healthcare costs. Moreover, the global pandemic exposed weaknesses in healthcare systems, such as inefficiency and inequity, which reinforced the need for scalable and sustainable care strategies [2]. In this context, artificial intelligence (AI) has emerged as a key enabler of healthcare innovation. Its ability to process vast amounts of health data and identify complex patterns makes it particularly valuable, supporting clinical tasks such as monitoring, diagnosis, and treatment optimisation. As a result, governments, payers, regulators and providers are being increasingly driven towards AI-based medical solutions to face rising competition and deliver high-quality, cost-effective care that aligns with value-based healthcare principles [2]. However, the success of digital health innovation remains threatened by barriers to its development and adoption, including concerns around data privacy, fairness, as well as the overall trustworthiness of AI systems [3].

As healthcare systems accelerate their digital transformation, the increase in data collection and availability, paired with advances in computer science, have paved the way for artificial intelligence to transform the field of medical practice [4]. AI tools, designed to simulate human cognitive functions – such as learning and problem-solving [2] - are increasingly being integrated into every stage of the medical care cycle, from drug discovery and disease diagnosis to decision-support systems. These applications aim to accelerate the development of new treatments, enable earlier diagnosis, enhance clinical decisions, and ultimately improve patient outcomes. Moreover, by improving system efficiency, they also enable healthcare professionals to better manage growing workloads and focus on human interaction, facilitating the transition towards patient-centred care [2],[4].

Despite growing enthusiasm for AI in healthcare, the development of advanced AI tools does not guarantee their integration into and adoption within healthcare practice [4]. Research shows that many implementation efforts lack a well-defined motivation grounded in real clinical needs, a critical factor for successful adoption. A recent scoping review by Sharma *et al.* pointed out that over half of the forty-five analysed studies on AI in healthcare failed to justify why the solution was introduced, revealing a broader disconnect between technological solutions and the problems they aim to address [4]. This misalignment has drawn criticism from experts within the field, concerned that the hype surrounding AI is diverting developers' attention away from the iterative, needs-driven development processes required for successful innovation. As Nigam Shah, Stanford Health Care's AI vetter-in-chief, explains: *"A lot of what is being created is formed backwards. Rather than follow what is known as a biodesign process - where need dictates solutions - the AI buzz has many researchers reaching blindly for the next great thing… There's a lot of poorly designed hammers looking for nails."* [5]. Additionally, the scoping review highlighted that much of the literature overlooks the practical challenges of AI implementation, with key steps, such as clinical trust, usability, and regulatory alignment, often left unaddressed [4]. Without a clear understanding of these barriers, there is

a significant risk that AI's potential benefits for patients, healthcare providers, and health systems will remain unrealised. Moreover, it raised awareness of the fact that few studies apply structured frameworks to guide the development and deployment of AI tools in real-world clinical settings, limiting the field's ability to generate replicable, sustainable impact [4].

This lack of structured, needs-driven methodologies has prompted greater attention to guiding innovation frameworks. One such example is the Stanford Biodesign framework, which takes a needs-driven approach, focusing on patients' real-world needs, to guide MedTech innovators through the process of developing medical products, outlining key challenges and best practices to mitigate risk and drive successful implementation and adoption of technologies with real-world impact [6]. This innovation process has found success in leading teams to achieve better patient health outcomes through the development of novel products by following a comprehensive roadmap for identifying unmet clinical needs, brainstorming, prototyping, and implementing potential solutions [6]. However, despite the completeness of the Biodesign framework, the healthcare sector and MedTech industry have undergone significant changes since the framework's last revision, published in the 2015 book *Biodesign: The Process of Innovating Medical Technologies 2nd Edition* [7]. These changes, and particularly the rise in artificial intelligence-driven approaches, bring unforeseen challenges unique to the development of AI-based medical technologies, including issues related to data access, algorithmic bias, or the need for continuous model monitoring.

Driven by the research questions *"What main challenges impair the adoption of AI-based medical technologies?"* and *"What best practices can support the development of a guiding framework?*", the goal of this dissertation is to highlight the key challenges and best practices regarding the development and implementation of AI-based medical technologies. Moreover, it aims to propose a structured framework to guide this innovation process, promoting innovation across healthcare settings and ensuring real-world impact.

During my academic journey, I became increasingly aware of how little research actually translates into applications that return real-world value. Through courses, such as *"Innovation and Management in Bioengineering",* and discussions around the commercialisation of academic technologies, I came to recognise the critical role that structured frameworks play in bridging this goal, guiding the development of new technologies and supporting their adoption and impact. Adding to that, I came to understand that the path of medical innovation from lab to impact is filled with barriers, from the misalignment with clinical needs to weak commercialisation and implementation strategies. Driven by a strong interest in medical innovation, I am particularly motivated by the opportunities offered by data-driven and AI-based tools to reshape healthcare, making it more personalised, efficient and accessible. I also argue that the success of such innovations relies not only on the technology itself but also on project planning and stakeholder involvement.

This study is driven by the belief that artificial intelligence has the potential to not only transform healthcare systems but to expand access to quality care and shift the paradigm from treatment to prevention. However, this promise will remain unrealised unless the barriers to

development and adoption are carefully addressed. To bridge this gap between technological potential and real-world clinical and societal benefit, this study aims to help by laying a foundation for more effective AI-based MedTech innovation. By guiding the creation of solutions that are both clinically relevant and implementable, the goal is to extend quality care to different environments, improve patient outcomes and enhance the experience and efficiency of healthcare delivery.

To this end, the study draws on insights from literature, reports, media, and expert interviews to raise awareness of the key challenges faced by AI-based MedTech innovation and best practices to overcome them. The final output is a framework to support the development and implementation of AI-based medical technologies, offering guiding questions and best practices to promote successful clinical integration and positive real-world impact.

The Theoretical Background and Key Concepts Chapter provides a literature review on medical technology, artificial intelligence, technology transfer and innovation, the AI-based medical innovation landscape, and the challenges specific to AI-based medical technologies. The Methods Chapter describes the research methodology, which followed a qualitative exploratory sequential approach, based on unstructured and semi-structured interviews. The Results and Discussion Chapter presents and discusses the results of the interviews, comparing them with the literature review findings. This Chapter a) illustrates the AI-based MedTech innovation landscape, b) identifies the main challenges and best practices for AI-based medical technology innovation in healthcare settings, c) proposes the VAMI Framework: a framework for Value-driven AI MedTech Innovation, d) explores how the proposed framework can complement broader innovation pathways: Stanford's Biodesign, and e) validates the proposed framework based on experts' grounded recommendations. The Case Study on the CINDERELLA Project Chapter presents a case study in which the AI-based MedTech Project CINDERELLA is evaluated through the lens of the framework proposed in the previous chapter. The case study constitutes a practical implementation of the framework, analysing whether CINDERELLA aligned with the best practices proposed by the framework. In case of misalignment, recommendations based on the framework are provided. The Sustainable Development Goals Chapter explores how this work aligns with the United Nations' Sustainable Development Goals, highlighting its potential contributions to health equity and innovation. The Limitations and Future Improvements Chapter acknowledges this study's limitations and suggests directions for future research. The Conclusion Chapter concludes this dissertation by summarising the main contributions and implications of the study.

2. Theoretical Background and Key Concepts

2.1 Medical Technology Innovation

Innovation can be broadly defined as the implementation of novel solutions which, motivated by commercial and/ or social motives, generate value, harness new technology, and drive competitive advantage and economic growth [8]. In healthcare, innovation has significantly contributed to increased longevity, a better quality of life, and improved healthcare service delivery [1], [9]. While healthcare innovation may also take non-technological forms, such as new care delivery models or public health policies like national screening programs, technological innovation remains a major driver of progress in this field. It encompasses the invention, development, and diffusion of new tools, systems or devices [9], expanding the possibilities for medical treatment and diagnosis.

According to the World Health Organisation (WHO), healthcare systems around the world are under increasing pressure due to factors such as ageing populations, workforce shortages, and financial constraints [10], [11], [12]. The WHO projects a global shortfall of 11 million health workers by 2030, predominantly in low- and lower-middle-income countries, posing serious challenges to achieving universal health coverage [12]. These systemic pressures have amplified the urgency of innovation, not only to improve patient outcomes and system efficiency but also to reduce the burden on overstretched healthcare workers through automation and optimised workflows.

In this context, the medical technology (MedTech) sector plays a crucial role. It is currently amongst the most dynamic industries, ranking second in global patent applications [13]. Its fast pace is reflected in the short product lifetime, with many devices being replaced within 18 to 24 months by superior innovations. Driving this fast-moving landscape is a strong commitment to research and development (R&D), with the average global R&D investment rate (R&D spend as a percentage of sales) estimated to be around 8% in 2024 [13].

While increasingly advanced technologies enable new and effective solutions to complex industry problems, they also intensify competition and raise the bar for companies. Meanwhile, traditional models of care continue to struggle with growing demand, rising costs for payers and governments, and the consistent delivery of high-quality care, particularly due to the growing burden of ageing populations and prevalence of chronic diseases [14].

In response, healthcare is undergoing a fundamental shift towards new models of care focused on delivering higher value, rather than simply increasing volume.

2.2 Value-Based Healthcare (VBHC)

Traditional fee-for-service models have dominated healthcare systems. Under these models, providers such as doctors, healthcare facilities, and clinics are reimbursed based on the volume of services they deliver, regardless of whether they achieve meaningful health improvements. While this approach promotes service provision, it often results in inefficiencies and a disconnect between healthcare costs and patient outcomes. Over time, pressures from an ageing population, rising costs, and unmet patient needs have raised concerns among stakeholders, rendering such models unsustainable.

In response, value-based healthcare (VBHC), introduced by Porter & Teisberg in 2006 [15], emerged as a compelling alternative. It aims to align healthcare delivery with improvements in patients' outcomes relative to the cost incurred [16]. This model redefines value as the ratio of patient-relevant outcomes to the total cost of care. Rather than rewarding volume, VBHC incentivises effectiveness, aiming to deliver better care at lower costs while enhancing patient experience.

Implementing VBHC involves multiple stakeholders, including patients, providers, payers, and policymakers, collaborating to optimise health outcomes in a cost-efficient manner [16]. The approach seeks to improve clinical results, personalise and coordinate care, reduce reliance on long-term interventions, and enhance the experiences of both patients and clinicians. Collectively, these goals aim to foster more sustainable health systems [16].

However, the shift from fee-for-service to value-based models presents significant challenges. For instance, the implementation of VBHC requires structural change, including the establishment of standards for measuring outcomes and data infrastructure capable of tracking both clinical and patient-reported outcome measures (PROMs). These changes require significant resources, robust digital ecosystems, and interoperability in line with regulations such as the General Data Protection Regulation (GDPR) [17].

In this context, artificial intelligence-driven technologies are central to facilitating this transition into VBHC. These tools support real-time outcome tracking, predictive care models, and the delivery of more personalised and efficient care [17].

2.3 AI-driven Medical Technology Innovation

Artificial Intelligence (AI) has come a long way from a concept in science fiction to an essential tool utilised in big tech companies like Apple, Amazon, or Google [18]. It is now set to revolutionise healthcare, with promising applications across the full spectrum of care, from research to clinical decision-making [19]. Figure 1 (source: [18]) illustrates the landscape of AI integration in healthcare, from data collection to healthcare applications.

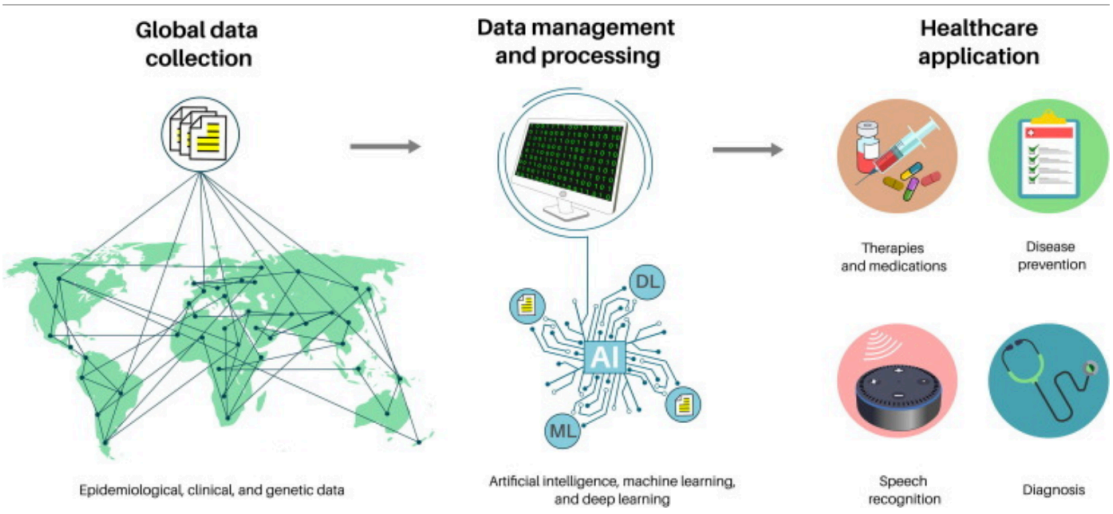

Figure 1 Schematic representation of data collection, data management and processing, and healthcare applications.

At its core, AI is the science and engineering of making intelligent machines, which follow algorithms or a set of rules, to mimic human cognitive functions, such as learning and problem-solving [2]. A key strength of AI lies in its ability to recognise and learn patterns within large, multidimensional datasets – often referred to as big data – and to autonomously refine its performance over time through exposure to new data (continuous learning) [2]. Within AI, machine learning (ML) describes the development of algorithms whose performance improves through experience. Deep learning (DL) is a subset of ML that uses multi-layered neural networks to simulate the complex decision-making power of the human brain [20].

The ability of AI to make sense of vast and multidimensional health data has profound implications for healthcare. It enables the development of tools that address previously inaccessible problems, boost operational efficiency, and deliver more personalised, timely, and effective patient care. In the MedTech sector, this has led to a surge of AI-enabled innovations, including diagnostic tools, clinical decision support systems, and patient monitoring devices. These technologies assist healthcare professionals by enhancing decision accuracy, reducing cognitive burden, and automating routine tasks, freeing up time for more complex care delivery.

The potential of AI to transform healthcare systems at scale is underscored by recent findings. An impact study commissioned by MedTech Europe and conducted by Deloitte estimated that AI-driven innovations could save 400,000 lives annually, generate €200 billion in savings each year, and free up to 1.8 billion working hours, equivalent to having 500,000 additional full-time healthcare professionals [21]. These figures highlight not only the clinical promise of AI but also its macroeconomic and workforce implications, reinforcing its central role in future-ready, sustainable healthcare systems.

2.3.1 Real-world Success Cases of AI-based Medical Technology Adoption

I. AI as a tool to combat the COVID-19 pandemic

The COVID-19 pandemic exposed critical weaknesses in healthcare systems worldwide and highlighted the need for more efficient, outcome-oriented approaches to delivering care. It also demonstrated the transformative potential of data-driven tools, accelerating the adoption of AI across the healthcare sector [18].

One notable example is Pfizer's use of AI to support the development and implementation of its COVID-19 vaccine. During clinical trials, vast amounts of data must be cleaned and analysed before scientists can make critical decisions. Traditionally, this data cleaning process would take over 30 days. However, by implementing a machine learning tool called Smart Data Query (SDQ), Pfizer reduced the process time to 22 hours, accelerating decision-making without compromising data integrity. This use of AI allowed for faster validation of trial results, ultimately enabling the vaccine to be distributed on time [22].

II. AI in Genome Sequencing

Patients with rare or undiagnosed diseases often face lengthy diagnostic journeys, involving multiple clinical visits, high healthcare costs, and prolonged uncertainty. Whole-genome sequencing (WGS) can help identify genetic variants associated with these complex conditions. However, standard workflows typically take weeks to complete and are not part of routine care, delaying critical treatment decisions [23].

To overcome this, Stanford University and its partners developed an ultra-rapid WGS workflow that leverages AI to accelerate the diagnostic process. Using Oxford Nanopore Technologies' PromethION48 sequencer, the system generates real-time electrical signals as DNA strands pass through protein nanopores embedded in a membrane. These signals are decoded into nucleotide sequences (A, T, G and C) using deep learning models powered by NVIDIA GPUs, significantly reducing the time needed. The workflow set a world record by delivering a genetic diagnosis in just 7 hours and 18 minutes [24]. In one clinical case, this innovation enabled physicians to rapidly diagnose a three-month-old infant experiencing unexplained seizures by identifying a critical genetic variant within hours of admission, whereas a standard gene panel took two weeks and failed to detect the condition [25].

III. AI in Cancer Prevention

Colorectal cancer is the second leading cause of cancer-related deaths worldwide [26]. However, screening significantly improves patient outcomes, with over 90% of patients being cured if the disease is caught at an early stage [27]. Colonoscopy is the gold-standard screening method, enabling the detection and removal of precancerous polyps. Despite its effectiveness, the procedure's success can vary significantly depending on the physician's skill, focus, and fatigue, leading to inconsistent detection rates [28].

To address this, Medtronic – one of the world's leading MedTech companies – developed GI Genius™, the first AI-powered, real-time computer-aided detection (CADe) system for colonoscopy. This intelligent endoscopy device integrates seamlessly into standard procedures, without altering the workflow or increasing withdrawal time. GI Genius™ analyses every frame in real time, identifying and flagging suspicious lesions that might otherwise be missed. The tool has a 99.7% sensitivity rate and 1% false positives, significantly enhancing the accuracy and consistency of colorectal cancer screening [29].

While notable advances have been made in specific applications, the widespread integration of AI across care pathways, diseases and healthcare settings remains limited by a range of systemic, technical, and ethical barriers that must be addressed to unlock its transformative potential [2].

2.4 Challenges and Enablers in AI-driven Medical Technology Innovation

A fundamental issue blocking the successful deployment of innovations in healthcare is the disconnect between technological development and clinical needs. This is often fuelled by a tech-push mindset, where excitement around emerging technologies results in innovations that are not grounded in validated real-world needs and local context [2]. Artificial intelligence is a clear example of this trend. As Stanford Health Care's AI Vetter-in-Chief, Nigam Shah, notes about AI, "there's a lot of poorly designed hammers looking for nails" [5]. This lack of alignment results in tools that often perform well in experimental settings but fail to translate into clinical practice.

While this misalignment is a critical barrier, it is only one part of a broader set of challenges which pave the journey toward widespread implementation of AI in healthcare. These range from technical limitations to ethical, regulatory and societal concerns [30]. Figure 2 (source: [31]) illustrates this complex landscape by outlining key stages and barriers in the adoption of clinical AI. Although the figure focuses on medical imaging, the structure exemplifies challenges common across various AI applications in medical technology innovation.

This section explores some of the challenges most frequently covered in the literature, as well as the emerging strategies to overcome them.

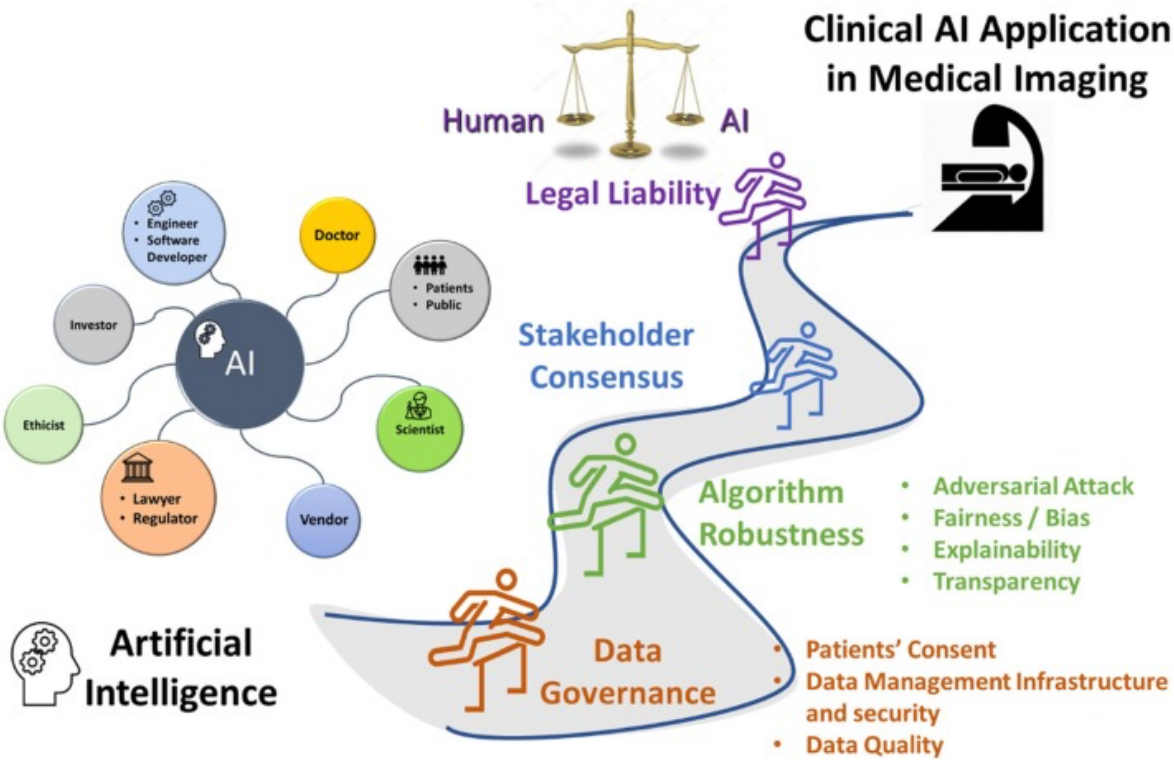

Figure 2 Challenges to overcome for AI-based solutions' successful clinical application.

2.4.1 Data Accessibility

AI's performance in clinical settings is heavily reliant on access to large, high-quality datasets [31]. While increased data availability could accelerate innovation, as seen in sectors where data is more readily accessible [32], it also raises serious concerns regarding patient privacy, data misuse, and ethical oversight [33].

Privacy and Security

Medical data is highly sensitive and typically protected by legal frameworks that grant patients ownership or control over how their information is used [33]. A high-profile case illustrating the potential downfalls is the Royal Free NHS Foundation Trust's unauthorised transfer of 1.6 million patient records to Google's DeepMind, which bypassed patient consent and drew significant public and regulatory backlash [34]. These concerns are not merely hypothetical. For instance, insurance providers may attempt to exploit health data to influence coverage decisions, and pharmaceutical or biotech firms have been known to offer financial incentives for individuals' genetic information [35]. These examples underscore the need for robust data governance and enforceable legal protections to safeguard patient interests.

Cybersecurity is another critical concern. Healthcare organisations often lag behind other sectors in terms of cybersecurity investment, making them particularly vulnerable to breaches [36]. These vulnerabilities further intensify resistance to data-sharing initiatives and promote the use of technical solutions such as anonymisation and pseudonymisation to allow data use while protecting identities. Anonymisation permanently removes identifying information, whereas pseudonymisation allows re-identification under strict controls [37]. However, both approaches face issues with scalability, maintaining data utility, and preventing re-identification in practice [38].

A promising privacy-preserving alternative in data sharing is federated learning (FL), which enables model training across decentralised datasets without transferring patient data between institutions. Instead, only model parameters or gradients are shared, allowing for privacy-conscious collaboration while enhancing data exposure [32], [39]. Studies have shown that FL can match the performance of models trained on centralised data and often outperform models trained on isolated local datasets [40]. Nevertheless, FL is not without risks. Even though patient data is not directly shared, the process of exchanging model updates can potentially leak sensitive information or allow reverse engineering of the original data [39]. Thus, FL remains an evolving solution that demands further research and refinement.

Another emerging approach is the use of synthetic data, generated algorithmically to mimic the statistical properties of real-world medical datasets [41]. This strategy offers a potential route to address privacy concerns, facilitate regulatory approvals, and reduce costs. However, it introduces new risks. Synthetic datasets can inadvertently encode or amplify biases, misrepresent certain populations, or, even if derived from de-identified data, still pose re-identification risks if not rigorously validated [42]. To mitigate these risks, transparency about data generation processes, regular audits, and a clear understanding of synthetic data's limitations are essential.

Standardisation and Interoperability

A persistent challenge in healthcare AI is the lack of standardisation across data sources, systems, and institutions. Medical data is often fragmented and inconsistently formatted, even within the same organisation, making integration and analysis difficult [32]. Variability in data acquisition methods further contributes to this issue. For example, blood oxygen saturation can be measured via pulse oximetry or arterial blood gas analysis, each yielding results in different units (percentage vs. partial pressure) and levels of precision. These inconsistencies affect not only data comparability but also how information is interpreted and stored. Tabular data is especially prone to such discrepancies, given differences in format, structure, and context.

To address these issues, data harmonisation is essential. This involves reconciling disparate data sources into compatible, standardised formats to enable meaningful comparison and analysis [43], [44]. Several international standards and data models support these efforts:

- The Observational Medical Outcomes Partnership Common Data Model (OMOP CDM) provides a standardised structure and content model specifically designed for organising observational health data [45].

- Systematised Nomenclature of Medicine - Clinical Terms (SNOMED CT) offers a comprehensive multilingual clinical terminology that facilitates consistent coding of clinical concepts [46].

- Digital Imaging and Communications in Medicine (DICOM) standardises the handling and storage of medical images and related information [47].

- Health Level Seven Fast Healthcare Interoperability Resources (HL7 FHIR) facilitates the electronic exchange of healthcare information, enabling interoperability between disparate healthcare systems [48].

Together, these standards and data models, developed and maintained by international organisations and consortia such as the International Health Terminology Standards Development Organisation (for SNOMED CT), aim to reduce data fragmentation and ensure interoperability across clinical, research, and administrative applications.

Nonetheless, harmonisation is not without its drawbacks. The process is resource-intensive and complex, potentially leading to oversimplification, loss of contextual nuance, and challenges in aligning data variables across health systems [43]. While harmonisation enhances data usability and cross-institutional analysis, it is not equivalent to creating a complete or universally coherent dataset and remains challenging to scale effectively.

2.4.2 Clinical Implementation

Real-World Performance

Despite significant technical progress, the clinical integration of AI remains limited. One of the primary barriers is the challenge of generalisation, where AI models that perform well during development often underperform when deployed in new or different clinical environments [32]. One common reason for this is overfitting, where algorithms learn unimportant or incorrect associations between patient features and outcomes, especially when influenced by too many variables. This leads to unreliable predictions when applied outside the original data used for training [32]. Other major contributors to performance degradation include population drift, changes in patient demographics or disease prevalence over time and technological drift, such as changes in medical devices or data acquisition methods. These drifts can significantly reduce model accuracy in real-world settings, justifying that practices such as local validation and continuous performance monitoring remain essential to maintain model relevance and reliability [49].

Return on Investment

Evaluating AI's return on investment in healthcare is difficult, partly due to the limited empirical evidence supporting its clinical effectiveness. Most AI research remains confined to controlled, retrospective studies rather than real-world environments. Moreover, randomised controlled trials (RCTs), the gold standard for medical evidence, are often poorly suited to assess dynamic, continuously learning AI systems [50]. This evidence gap diminishes stakeholder confidence and slows adoption.

Workflow Integration and Interpretability

For AI tools to deliver clinical value, they must integrate seamlessly into existing workflows without creating additional burdens for healthcare professionals [32], [51]. Tools that are disruptive, unintuitive, or time-consuming are unlikely to be adopted, especially in high-pressure clinical environments where time and reliability are critical. Resistance from healthcare professionals is common and often stems from fears of job displacement, reduced professional autonomy, and limited understanding of AI's capabilities [31], [52]. These concerns are further compounded by unresolved issues around data privacy, accountability, and the absence of clear legal frameworks [31].

Patients, too, tend to trust human clinicians more than algorithmic systems, making transparent communication essential for public trust and acceptance [52]. Misunderstandings or inflated expectations can result in either rejection of AI tools or overreliance on them inappropriately.

Interpretability

Closely tied to trust and workflow integration is the issue of interpretability. Clinicians are less likely to rely on AI outputs they cannot explain or understand, especially when those outputs influence diagnoses or treatment decisions. Patients, similarly, are less inclined to accept recommendations from incomprehensible algorithms.

To address this, the field of explainable AI (XAI) seeks to make AI outputs more transparent and interpretable [31]. A common XAI method is the use of attention or saliency maps, which highlight the regions of an input (e.g., a medical image) that contributed most to the model's decision. Unfortunately, these visualisations are often inconsistent and can even highlight irrelevant features [53], resulting in XAI being an immature field characterised by inconsistent standards and limited clinical integration. And while many interpretability techniques have been developed in academia, they are rarely implemented in clinical practice [53].

Overcoming the barriers to AI integration in clinical practice requires coordinated, multi-stakeholder action. AI development must be closely aligned with the practical needs and workflows of clinicians, patients, administrators, and policymakers to ensure relevance and usability [31]. Investment in infrastructure, workforce training, and long-term capacity-building programs is equally essential to support effective and ethical implementation.

Transparent communication and inclusive public dialogue are needed to build trust and foster realistic expectations about AI's capabilities and limitations [32]. Furthermore, explainable AI tools should not be treated as add-ons but rather as integral components of trustworthy and accountable AI systems. Their integration into clinical workflows can enhance both usability and acceptance, paving the way for more responsible and sustainable adoption in healthcare settings.

2.4.3 Ethics and Fairness

The integration of AI into healthcare raises critical ethical challenges, including opaque data collection and processing practices, the "black box" nature of many neural networks, inherent bias in training data, and AI hallucinations—false or misleading outputs that are difficult to trace or explain [54].

Bias

Bias is the most prevalent and persistent ethical concern in medical AI. All datasets carry some degree of bias, shaped by factors such as race, ethnicity, gender, and socio-economic or environmental conditions [33]. This can skew outcomes and exacerbate existing inequalities. Importantly, bias doesn't only stem from data. It can also arise from flawed design, poor-quality inputs, or a lack of diversity among development teams, whose implicit assumptions and cultural prejudices and misconceptions can become embedded in the models themselves [32].

When AI systems are trained on such skewed data or shaped by homogenous perspectives, they tend to underperform for underrepresented populations, reinforcing systemic disparities. A U.S. 2019 study by Obermeyer exemplifies this problem: a widely used algorithm, affecting millions of patients, systematically underestimated the needs of Black patients. Because the algorithm predicted healthcare costs instead of actual illness, and because Black patients typically receive less healthcare funding than White patients, it wrongly assigned lower risk scores to sicker Black patients. Correcting this bias would have nearly tripled the number of Black patients receiving additional care, from 17.7% to 46.5% [55].

These findings underline the necessity of using large, inclusive, and representative datasets in training AI systems, along with robust validation processes that guard against algorithmic bias. Moreover, promoting fairness in AI requires not just technical fixes but also systemic changes in how healthcare tools are designed, tested, and deployed.

To support ethical development, several practical tools have been introduced. Notably, the Assessment List for Trustworthy Artificial Intelligence (ALTAI) provides a structured self-assessment checklist for developers to ensure their systems align with ethical principles such as fairness, transparency, and human oversight [56]. Similarly, the Fundamental Rights Impact Assessment (FRIA) helps evaluate how high-risk AI systems might affect individuals' fundamental rights, especially those from marginalised groups [57]. These tools serve as early checkpoints for identifying ethical risks and reinforcing fairness in design before models are deployed in clinical practice.

Accountability

Another major ethical issue involves accountability, specifically, determining who is responsible when an AI system makes a harmful or incorrect decision. This question is particularly urgent in clinical contexts, where errors can have life-threatening consequences. Traditional medical ethics demand a clear chain of responsibility, but AI complicates this: if developers cannot anticipate or explain a model's behaviour, can they be held morally or legally responsible? And if clinicians rely on outputs they don't fully understand or control, are they to blame for poor outcomes? [32], [58].

Many AI systems, especially those based on deep learning, are difficult, if not impossible, to interpret. These are often referred to as opaque models or "black box" systems, where the decision-making process is not transparent to human users. In contrast, transparent models, such as decision trees or linear regression, provide outputs based on understandable, traceable logic. While transparent models may offer less predictive power in some complex tasks, they are often preferred in clinical settings due to their interpretability, which supports accountability and trust.

A 2021 report by Fair Isaac Corporation (FICO) found that 65% of companies using AI cannot explain how their models arrive at specific predictions [59]. This "black box" problem creates serious legal and ethical complications, eroding trust among both clinicians and patients. While some argue that clinicians already rely on tools (e.g., MRI or CT) whose mechanisms are not widely understood, this comparison overlooks the need for accountability in decision-making - unlike those tools, AI often contributes directly to diagnostic or treatment choices [32].

These challenges to AIs' practical implementation highlight a critical need for them not only to perform accurately but to do so in a transparent and accountable way.

2.4.4 Regulation

The regulation of AI in healthcare is complex and rapidly evolving across jurisdictions [60]. While both the European Union and the United States acknowledge the need to ensure safety, effectiveness, and ethical compliance in AI-driven medical technologies, their regulatory approaches diverge in scope, structure, and maturity.

The EU has adopted a proactive and prescriptive approach, most notably through the Artificial Intelligence Act (AI Act), which establishes stringent pre-market obligations, particularly for high-risk AI systems [61]. In parallel, developers must also comply with the Medical Device Regulation (MDR) and In Vitro Diagnostic Regulation (IVDR), a regulatory layering that can introduce overlapping and, at times, conflicting requirements [62]. The 2024 MedTech Europe Report underscores the urgency of aligning these frameworks, particularly regarding definitions of high-risk AI and harmonised conformity assessment procedures [63].

In contrast, the US lacks a dedicated AI-specific legal framework. Instead, regulation is primarily overseen by the Food and Drug Administration (FDA), which employs a risk-based

premarket review process for AI-enabled medical devices. Recognising that traditional regulatory paradigms are unsuitable for adaptive AI technologies, the FDA is increasingly emphasising post-market oversight. Initiatives such as the Predetermined Change Control Plan aim to address the unique challenges of continuous learning systems [64].

This regulatory divergence presents both opportunities and challenges. While it creates uncertainty for global developers navigating inconsistent requirements, it also reflects differing national priorities. The EU emphasises legal harmonisation and safety, whereas the US promotes innovation and regulatory agility. As highlighted by the Organisation for Economic Co-operation and Development (OECD), coordinated international policy efforts are essential to unlock AI's potential, protect patient data, and ensure equitable access to innovation [65].

The regulatory landscape is further complicated by the dynamic nature of AI technologies. Frequent model updates and retraining can quickly render static assessments obsolete, making it difficult for regulators to provide timely approvals and for developers to maintain compliance [60]. This underscores the need for adaptive, future-facing regulatory mechanisms.

Key Regulatory Frameworks

In the European Union,

- General Data Protection Regulation (GDPR):

The GDPR is the EU's primary legal framework for data protection and privacy. It applies to any organisation, regardless of location, that processes the personal data of EU citizens or residents or offers them goods and services. It establishes strict requirements for how personal data, especially sensitive health data, is collected, processed, and stored. It enforces transparency, consent, data minimisation, and individual rights over personal information. The GDPR entered into force on 14 April 2016 and became enforceable on 25 May 2018 [66].

- Artificial Intelligence Act (AI Act):

The EU's AI Act is the first comprehensive legal framework on AI, aiming to promote responsible innovation while minimising risk. It provides AI developers and deployers with clear requirements and obligations for specific AI applications while also aiming to minimise administrative and financial burdens. For instance, it imposes stricter obligations on high-risk AI systems, such as those used in medical diagnostics, including requirements around risk mitigation, data quality, transparency, and human oversight [30], [61]. The Act came into force on 1 August 2024 and will become fully applicable by 2 August 2026 [61].

- European Health Data Space (EHDS):

The EHDS is the EU's first common data space dedicated to healthcare. It facilitates secure access, sharing, and reuse of electronic health data across the Union. Its goal is to empower individuals with control over their health information while also supporting medical research,

innovation, and evidence-based policymaking [67]. The EHDS Regulation was published on 5 March 2025 and entered into force on 26 March 2025 [67].

In the United States,

- Health Insurance Portability and Accountability Act (HIPAA):

HIPAA remains the primary legislation governing data privacy and security in US healthcare. Although not specific to AI, it provides foundational protections for personal health information and sets the framework for data use in AI-enabled tools.

Bringing AI-driven healthcare products to market entails considerable regulatory navigation, time, and investment. While robust oversight is critical to ensuring safety and public trust, overly rigid or fragmented frameworks risk stifling innovation and deterring beneficial technologies. A balance must be achieved between ensuring ethical integrity and allowing for regulatory flexibility. As stated in a 2024 Roche report, "The future of AI in healthcare is a matter of collaboration, transparency and fairness" [54]. To foster meaningful innovation, regulatory systems must evolve to uphold strong safeguards while remaining adaptive to technological progress and the dynamic nature of AI in healthcare [32].

2.5 Frameworks for AI-driven Medical Technology Innovation

As seen in the previous chapter, and unlike traditional medical solutions, AI tools must contend with data dependency, high technical standards, regulatory uncertainty, and ethical scrutiny. In this context, there is an urgent need for robust frameworks that can guide innovators through the full innovation lifecycle, from identifying clinical needs to achieving real-world adoption.

2.5.1 Stanford's Biodesign Innovation Process: A Framework for Healthcare Technology Development

One of the most influential and widely adopted frameworks in the medical technology field is the Stanford Biodesign Innovation Process. First introduced at Stanford University in 2000 [68], this evolving framework has since expanded globally, advising or inspiring over thirty educational programs across North America, Europe, Asia, Africa, and Australia. Its international reach highlights its growing impact on shaping medical innovation worldwide [69].

The Biodesign Innovation process takes a needs-driven approach to MedTech innovation, structured around three core phases – identify, invent, and implement – commonly referred to as the "3 Is" [6]. By prioritising patients' needs and working back towards a solution, it contrasts with more traditional technology-first approaches and has proven effective in guiding the development of impactful healthcare products and services [6]. A prime example is the iRhythm Zio® patch, where the Biodesign approach was instrumental in breaking down the complex challenge of patient cardiac monitoring into manageable needs, leading to a simple,

patient-friendly wearable patch alongside a sophisticated cloud-based algorithm for data analysis and scaling [70]. Since its launch, iRhythm has analysed data from approximately 8 million patients, generating over 1.5 billion hours of curated heartbeat information and supporting more accurate and timely diagnoses [71].

Figure 3 (source: [73]) illustrates the three phases of Stanford Biodesign's framework, outlining two specific stages and key activities for each. Phase one – Identify - focuses on recognising unmet clinical needs through an iterative screening process that evaluates potential impact, stakeholder involvement, and system efficiency [72]. Phase two – Invent – is centred on developing lead concepts using brainstorming, prototyping, and refinement methods. Solutions are assessed across dimensions such as feasibility, intellectual property, business models, and regulatory pathways [72]. Lastly, phase three – Implement – involves designing a go-to-market strategy, addressing funding, regulatory approval, reimbursement, and commercialisation [72]. This process is cyclical and iterative, as new information and results often require returning to prior stages [73].

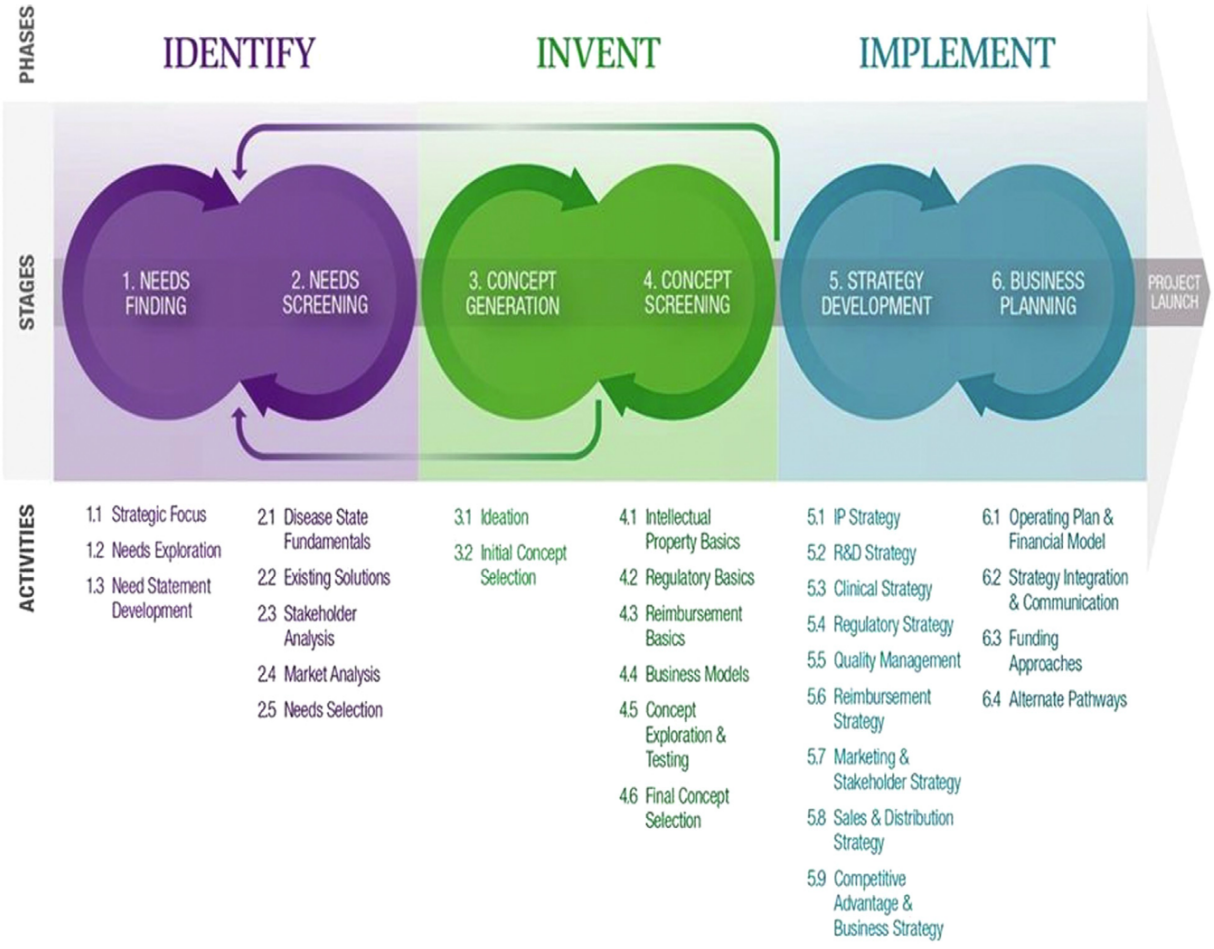

Figure 3 The three phases of the Stanford Biodesign process for healthcare technology development - Identify, Invent, and Implement - are shown, each comprising two specific stages. Key activities associated with each stage are outlined below.

2.5.2 Limitations of Traditional Frameworks

While Stanford's Biodesign framework remains a cornerstone in MedTech innovation, it was introduced in 2000, at a time when AI and digital health technologies had not yet emerged as central components of medical practice. As such, it does not fully address the specific and multifaceted challenges posed by AI-based medical technologies. Notably, it lacks guidance on critical aspects such as algorithmic bias, data governance, dynamic model validation, and the integration of AI into complex healthcare workflows.

Recognising these emerging demands, Stanford expanded its educational offerings in 2023 with two courses: *Biodesign for Digital Health* [74] and *Building for Digital Health* [75]. These are not updates to the core Biodesign framework but hands-on training opportunities. The former guides students in applying the Biodesign innovation process to digital health challenges, focusing on identifying unmet needs and developing solutions that incorporate digital technologies [74]. The latter offers computer science students practical experience developing digital health applications in collaboration with faculty and clinicians on real-world projects [75].

Moreover, Stanford has developed Spezi, an open-source software designed to support the technical development of digital health applications [76]. While Spezi offers valuable infrastructure for developers, it functions as a technical enabler rather than a conceptual extension of the Biodesign process itself.

Despite these important initiatives, the original Biodesign framework and its foundational documentation have not yet been systematically updated to reflect the unique demands of AI-driven innovation, and continue to fall short in providing comprehensive guidance for innovators working at the intersection of healthcare and AI.

2.5.3 New Frameworks Targeting the Application of AI to Healthcare

To address this, a growing number of frameworks have been developed in recent years across academic, institutional, and industry settings. These efforts aim to support the responsible design, validation, and deployment of AI in healthcare. Broadly, these can be divided into two categories: those that address discrete aspects of AI systems, such as ethics or regulation, and those that attempt to offer more comprehensive, end-to-end guidance.

Among the more targeted approaches, multiple have focused on the ethical layer. These include the previously discussed ALTAI (see 2.4.3.), developed by the European Commission to provide developers with a self-assessment tool grounded in established ethical and legal standards, and FRIA (see 2.4.3.), which focuses on fairness and inclusion, advocating for equitable AI deployment. Stanford has also recently released the FURM (Fair, Useful, and Reliable AI Models) framework, encouraging developers to prioritise AI models that are Fair, Useful, and Reliable within clinical environments [77].

Further ethics-oriented frameworks continue to emerge. Dykstra *et al.* (2025) proposed an institutional framework aimed at promoting equitable, ethically governed ecosystems for AI-

augmented care, with a strong emphasis on patient-contextualised data access and multi-domain integration [78]. Similarly, Abujaber *et al.* (2024) presented an ethics framework for healthcare AI research, structured around core ethical principles, implementation guidelines, and operational standards to ensure integrity in the research process [79].

In terms of clinical translation, You *et al.* (2025) introduced a framework that focuses on the validation of AI systems through a clinical trial–informed framework. It introduces a four-phase approach—Safety, Efficacy, Effectiveness, and Monitoring—that mirrors the rigour of drug development and aims to prevent unintended risks while improving transparency and real-world applicability [49].

In contrast, some recent frameworks take a more comprehensive view, seeking to guide the entire AI innovation process. The *FUTURE-AI* framework was developed by an international consortium of 117 interdisciplinary experts, including AI scientists, clinicians, ethicists, and social scientists, collaborating over two years, to establish six guiding principles: Fairness, Universality, Traceability, Usability, Robustness, and Explainability [80]. Its focus on usability and integration helps bridge the gap between high-level ethical aspirations and practical implementation needs. Similarly, the Michigan Health & Hospital Association's AI framework outlines considerations across five key areas: patient outcomes, workflow integration, transparency, safety, and bias mitigation [81].

The increasing number of such frameworks - many published in 2024 and 2025 - signals the growing recognition of AI's complexity and the need for structured, domain-specific guidance. However, despite their contributions, most remain either conceptual or limited in scope. For instance, a recent systematic review by Shiferaw *et al.* (2024) found that while existing AI guidelines often emphasise ethics, reproducibility, and transparency, they tend to fall short in several areas: meaningful engagement of stakeholders such as patients and clinicians, environmental sustainability considerations, and the provision of detailed implementation workflows [82]. Furthermore, many current frameworks do not sufficiently account for real-world complexities such as model drift, reimbursement models or institutional readiness.

To conclude, while many challenges have been thoroughly documented in the literature, some of which are explored in detail in Section 2.4, others remain poorly understood or entirely unaddressed. Moreover, the growing number of available frameworks, combined with significant variability in their scope and quality, underlines the absence of a comprehensive, integrated approach that can guide stakeholders through the full development and implementation pathway of AI-based medical technologies. This highlights a continued need for an operational framework, grounded in stakeholder input, that effectively bridges ethical principles, technical robustness, and real-world feasibility. Such a framework should also integrate both academic foundation and industry expertise, supporting innovators across the entire lifecycle of AI-driven healthcare solutions.

3. Methodology

3.1 Research Question, Objectives and Methodology

A preliminary literature review was conducted from February 2025 to April 2025 to explore the context of the topic and understand key concepts. Keywords (MedTech Innovation, Artificial Intelligence Healthcare Innovation, Biodesign, Artificial Intelligence in MedTech, Artificial Intelligence Regulation, Artificial Intelligence MedTech Reimbursement, MedTech Business Strategy, MedTech Technology Transfer, Artificial Intelligence MedTech Adoption and Implementation) were used to search platforms such as PubMed, ScienceDirect, and Google Scholar. Additional sources were also targeted, including journal articles and reports from relevant stakeholders. Given the large volume of results, titles and abstracts were screened to assess relevance to the research focus. Priority was given to recent publications, highly cited papers, and sources that addressed AI-specific challenges, best practices, and digital innovation frameworks within the healthcare context.

Following the literature review, a qualitative and interview-based approach was adopted to identify and explore best practices for the development of AI-based medical solutions. Qualitative methods utilise meanings, experiences, and perspectives of individuals and communities to understand a phenomenon, especially when conventional sources of information are insufficient or partial to support a comprehensive understanding of the context or geography [83]. Through the identification and validation of key challenges and best practices in AI-based medical technology innovation, this dissertation aims to develop a framework to facilitate this process and enhance its real-world impact. Hence, this study's target research questions are:

What main challenges impair the adoption of AI-based medical technologies?

What best practices can support the development of a guiding framework?

To this end, the adopted methodology aimed to identify the key challenges regarding AI-based MedTech innovation as well as best practices. Interviews with professionals with diverse backgrounds, including MedTech business leaders, clinical advisors, data scientists, and regulatory experts, documented practical experiences and perspectives. These interviews helped identify challenges, success factors, best practices, and opportunities in the field. Based on the interview findings, this study proposes a framework to support MedTech businesses in developing and implementing AI-based MedTech innovations, ensuring their meaningful real-world impact. The flowchart in Figure 4 depicts the steps and methods followed to achieve the final objective; phases 1 to 4 shown in the flowchart will be described in detail in Section 3.2, Research Design.

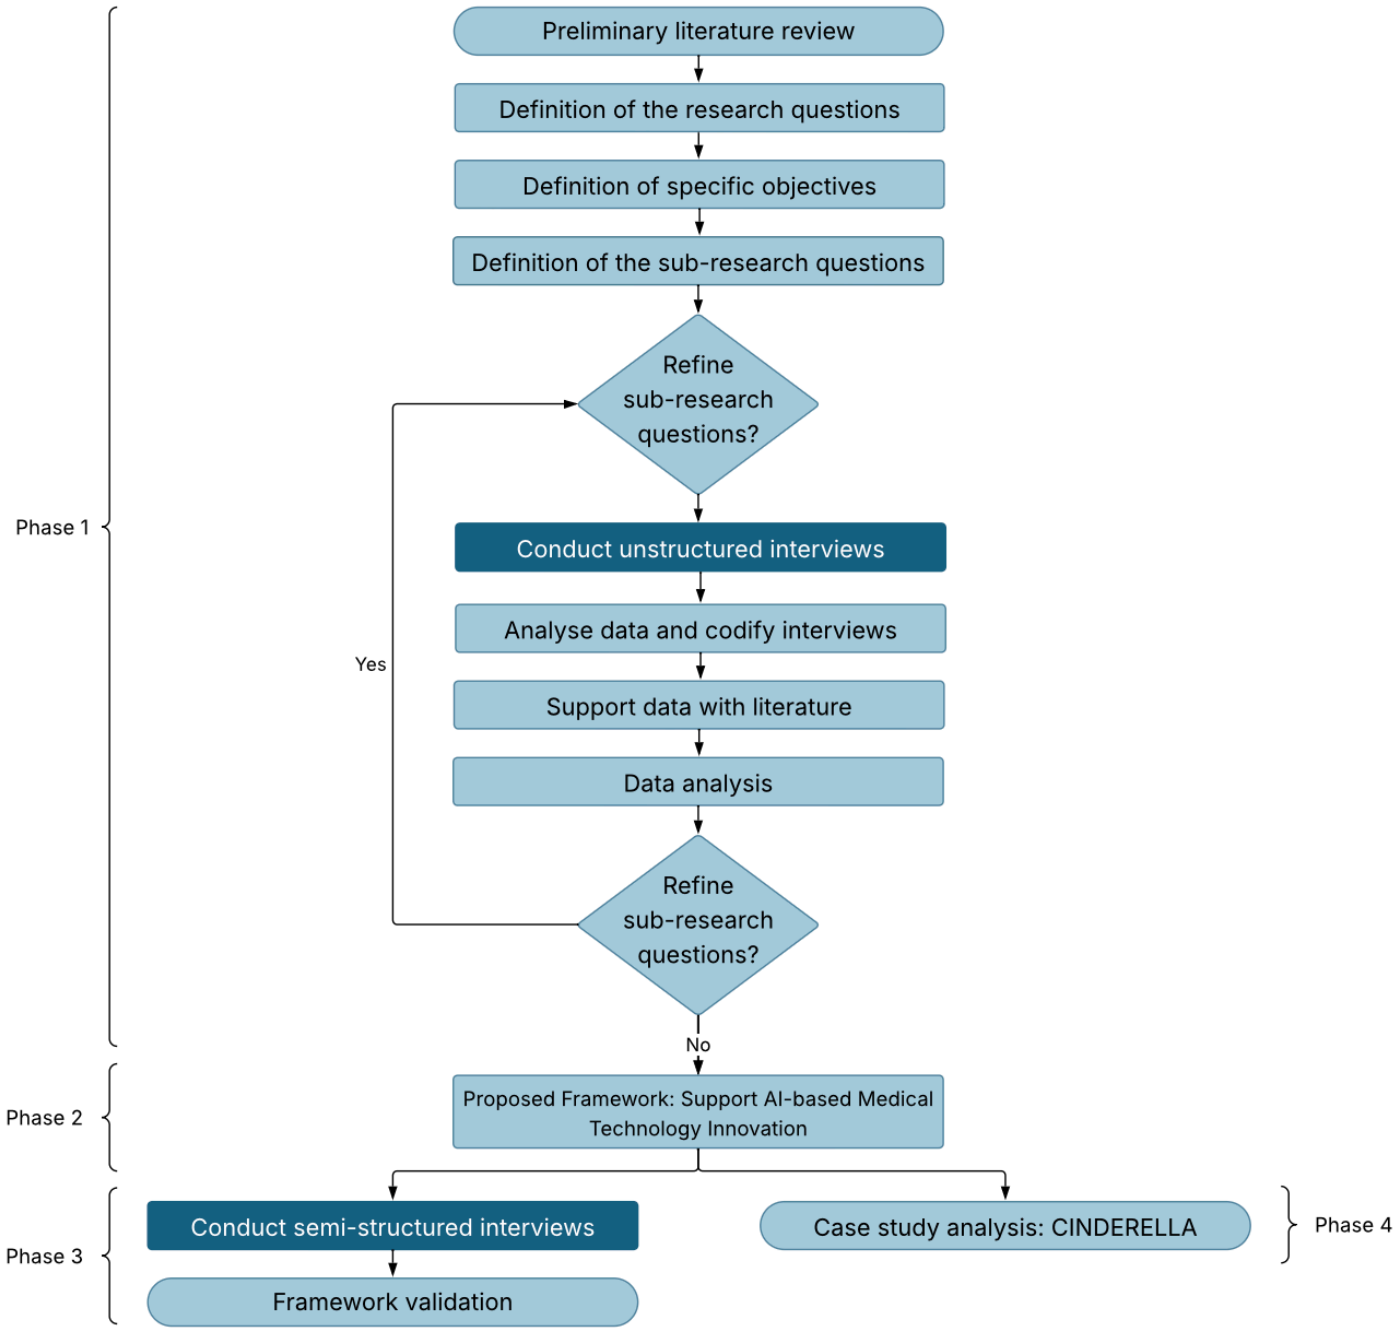

Figure 4 Flowchart of the qualitative exploratory sequential research design

3.2 Research Design

This study employed an exploratory sequential research design, drawing methodological inspiration from the works of Kathleen M. Eisenhardt and Robert K. Yin [84], [85]. The approach began with a comprehensive literature review to inform the development of research questions and guide data collection strategies. Two rounds of interviews, one unstructured and one semi-structured, were conducted with MedTech industry experts to capture diverse stakeholder perspectives and support both exploration and validation. Particular attention was given to selecting appropriate interview formats and structuring them to elicit in-depth insights. Qualitative data were then coded and analysed, with emergent themes compared against existing literature to ensure analytical rigour and contextual relevance. The full structure of the research design is illustrated in Figure 4, which outlines the integration of literature and empirical findings across each phase of the study.

The first phase combined a preliminary literature review with a series of unstructured interviews conducted with participants working with AI across various healthcare domains. Inclusion criteria and a detailed breakdown of the sample are outlined in Section 3.3. These open-ended conversations, conducted with minimal guidance [84], were designed to broadly explore the topic, refine the research focus, and identify initial themes, namely, the key challenges and best practices in the field. Following the interviews, a targeted, follow-up literature review was conducted to deepen the understanding of these emerging insights, verify their relevance, and enrich the thematic analysis.

The second phase involved using the challenges and best practices identified earlier to structure the proposed framework - VAMI: A Framework for Value-driven AI MedTech Innovation, aiming to support the development and implementation of AI-based medical solutions. The framework highlights key barriers and promotes the adoption of best practices throughout the innovation process.

The third phase involved a series of semi-structured interviews, guided by a flexible question set. This phase explored the findings from phase one, identified knowledge gaps, and served to validate the findings and proposed framework. Interviewees were purposefully selected for their experience across the full spectrum of AI-driven medical technology innovation. This phased and adaptive approach ensured that the findings were informed by both the existing literature and real-world experience from the field [101].

In the fourth and final phase, a case study analysis was conducted to complement and contextualise the findings from the interview-based and literature-informed phases. This phase aimed to evaluate the applicability and relevance of the VAMI Framework by examining a real-world example of AI-driven medical technology innovation. The CINDERELLA Project, a Horizon Europe initiative focused on AI-supported clinical decision innovation, was selected for its multidisciplinary scope and integration of AI in the development of a medical tool.

3.3 Data Collection

This research work followed a five-step data collection process to investigate challenges and best practices in AI-based medical technology innovation, validate the proposed framework, and apply it to a real-world case study. Data was collected at each step to progressively inform and validate the study's focus.

Preliminary Literature Review - Data was first collected through a broad review of existing literature and relevant secondary sources, including peer-reviewed articles, stakeholder reports, press releases, and expert commentary. This review aimed to establish a theoretical foundation, contextualise the topic, identify key concepts, and define the research questions and objectives.

Unstructured Interviews - Following the preliminary literature review, unstructured interviews were conducted with eleven professionals working with AI technologies across diverse sectors: a molecular modelling researcher, a professor of bioinformatics, a specialist in secure and privacy-preserving data methods, a PhD candidate in AI ethics, a pharmaceutical data scientist, a product owner in digital health, an innovation director, an application specialist, a machine learning engineer, a data protection officer, and a regulatory affairs director. Participants were selected to capture a broad range of perspectives on key areas of MedTech innovation, including data access, model development, regulation, and adoption. The sample included experts with experience in both European and American markets, two regions recognised by WIPO's 2024 Global Innovation Index as leading global innovation hubs [86]. Interviews were open-ended and exploratory, serving to expand on initial findings from the literature and to clarify the key challenges and best practices associated with AI-based medical technology development and adoption.

Support Data with Literature - A comprehensive literature review was then conducted to support and refine the findings that emerged from the unstructured interviews. This step involved a more rigorous selection of recent articles, reports, news, and other sources of complementary information, ideally published after 2020, to reflect current developments in AI-based MedTech innovation strategy, regulation, etc. This step helped ensure the relevance of the previously findings, while also informing the development of the proposed framework.

Semi-Structured Interviews - With the framework developed and grounded in both empirical and secondary data, a third round of data collection was conducted through semi-structured interviews. Three additional professionals were selected for their direct involvement in the development and deployment of AI-based medical technologies: the CEO and Founder of Amparo, a company revolutionising prosthetics (Portugal); a Radiologist, Head of Clinical Development at CEREBRIU, and Founder of the RAIT Innovation Unit (Denmark); and the Chief Innovation Officer at Complear and Managing Director of Pi Ventures (Portugal). These interviews followed a flexible guide and focused on validating the findings and proposed framework, while also gathering recommendations and lessons learned.

Throughout the study, data on the CINDERELLA Project were collected from various private and public sources. These included the project's official website, published videos and

webinars, social media channels, publicly accessible deliverables, and confidential internal documentation. To complement these materials and better understand the project's structure, objectives, and progress, additional information was gathered through informal conversations and interviews with individuals actively involved in the project. These included one software developer at INESC TEC, one researcher from the same institution who also assumed a project management role, and one business expert from CANKADO. Together, these participants contributed valuable insights into key aspects of development and commercialisation.

All interviews were conducted confidentially and professionally, with participant consent and respect for privacy. While face-to-face interviews were preferred, written responses were accepted when necessary. Interviews lasted approximately 60 minutes. In total, fourteen participants contributed to the study, offering diverse insights across roles, regions, and phases of the innovation process. The triangulation of literature and interview data enhanced the reliability of the findings and supported a robust discussion of results [101][84]. A summary of the main insights and recommendations from each interview is included in Tables 11 and 12 (Appendix I).

Table 1 presents the distribution of interviews conducted across different dimensions of the AI-based medical technology innovation process, categorised by participants' work settings.

Dimension	Setting	Number of interviews
Comprehensive Expertise	Industry	3
Data Access & AI Model Development	Academia	3
	Industry	2
Clinical Implementation & Market Adoption	Academia	1
	Industry	2
Business Strategy	Academia	0
	Industry	1
Medical Device Regulation & Reimbursement	Academia	0
	Industry	1
Ethics & Fairness	Academia	1
	Industry	0

Table 1 Number of interviews conducted per dimension of the AI-based medical solution innovation process, separated by work setting.

3.4 Data Analysis

Insights from the preliminary literature review helped define the central research question, establish key concepts, and identify stakeholder profiles for the unstructured interviews. To better navigate the breadth and complexity of AI-based medical technology innovation, the central research question was divided into four sub-questions:

What is the current AI-based MedTech *innovation landscape?*

What are the main challenges and causes of failure in the development and transfer of AI-based medical solutions?

What are the key best practices that contribute to successful AI-based MedTech *innovation?*

What factors not accounted for in Stanford's original Biodesign framework are relevant to the development of AI-based solutions?

The unstructured interview data was analysed to refine the study's objectives, ensure alignment with both theoretical and practical perspectives, and inform the development of the proposed framework. A thematic coding approach [87] was applied to this data to systematically identify recurring concepts across diverse professional perspectives. This helped distil complex, often fragmented insights into structured, comparable themes. Thematic coding was used in the Results and Discussion Chapter, specifically Sections 4.1 ("The AI-based MedTech Innovation Landscape") and 4.2 ("Challenges and Best Practices in AI-based MedTech Innovation"). Key concepts such as "need identification," "data access," "technical and clinical feasibility," and "regulatory compliance" were coded directly from the interview content and grouped into broader analytical themes like "Challenges and Best Practices," which shaped the structure and focus of each section.

Data from the semi-structured interviews was summarised and analysed to validate the framework and incorporate key recommendations from healthcare innovation experts. Similarly, qualitative data from the CINDERELLA Case Study was analysed, focusing on points of alignment, misalignment, and practical insights. Thematic coding was not applied to these sections (4.3-4.5 and Chapter 5) due to their more targeted and confirmatory nature.

A custom diagram (figure 5) was created to visually represent this analytical process and the resulting structure of the research. It uses a colour-coded system: darker blue indicates the guiding research questions; medium blue corresponds to the key themes addressed in the Results and Discussion Chapter (Sections 4.1– 4.5); and the lightest blue represents the Case Study Chapter 5.

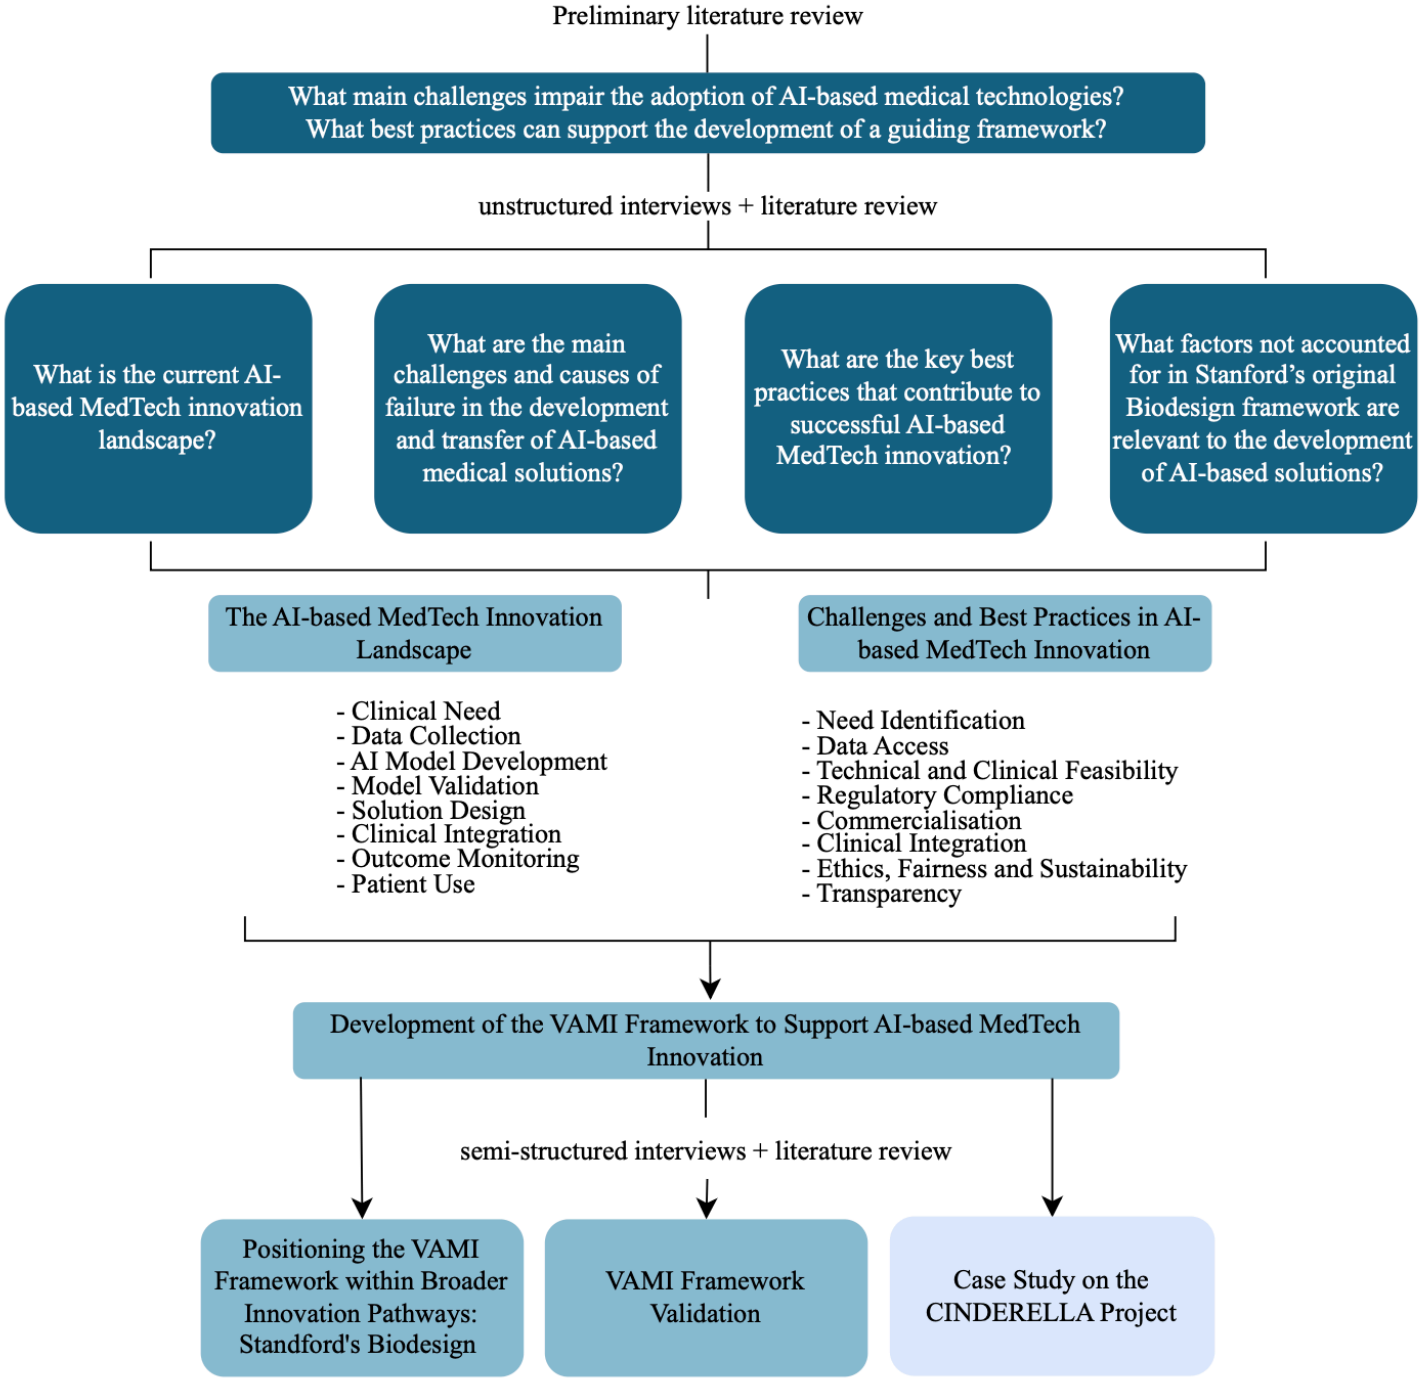

Figure 5 Diagram of the methodology to analyse the results and discuss the findings.

4. Results and Discussion

This dissertation aimed to identify the key challenges and best practices in AI-driven MedTech innovation and propose a framework to guide the development and adoption of such tools. To achieve this, unstructured and semi-structured interviews were conducted with fourteen experts with different backgrounds, including MedTech business leaders, clinical advisors, data scientists, and regulatory experts. The interviewees were based in different geographical locations between the EU and the US, with the majority of the EU having experience with the US in their respective fields of expertise.

This chapter presents and discusses the results from the interviews and literature review, aiming to address the two driving research questions: *"What main challenges impair the adoption of AI-based medical technologies?"* and *"What best practices can support the development of a guiding framework?"* The findings build on existing literature and recommendations while incorporating fresh insights from professionals' experiences, providing an updated perspective on the topic. To address the complexity of this question and the diversity of stakeholders involved, it was further broken down into four specific research objectives.

This chapter is structured to address the four research questions outlined in the study. Section 4.1 explores the current landscape of AI-based MedTech innovation process, responding to the first question. Section 4.2 examines the main challenges and common causes of failure, as well as key best practices for success, thereby addressing the second and third research questions. Based on the insights gathered, Section 4.3 presents the proposed framework to support the development and implementation of AI-based medical technologies. Section 4.4 explores how the proposed Framework can complement Stanford's Biodesign innovation process, thus addressing the fourth research question. Finally, Section 4.5 discusses the relevance and validity of the proposed Framework, drawing on expert feedback gathered specifically for this purpose.

4.1 The AI-based MedTech Innovation Landscape

AI-based MedTech innovation is a complex and highly iterative process, involving a series of interconnected stages, from the initial identification of unmet clinical needs to the real-world deployment and continuous outcome monitoring of a solution. While this process may be seen as sequential, in reality, it is inherently cyclical, with frequent feedback loops that allow for reassessment and refinement.

This section aims to provide a simplified yet structured overview of that landscape. It draws from both academic literature and insights from the interview conducted in this study. Although the framework presented here (illustrated in Figure 6) does not capture every step of innovation or the full extent of its iterative nature, it is intentionally designed to distil the process to a level that is more comprehensible and useful, particularly in identifying key innovation phases.

Figure 6 outlines the main phases of AI-based MedTech development, starting with the identification of a clinical need, followed by data collection to support AI model development and validation. Next comes solution design and integration into clinical workflows, culminating in patient use and the ongoing monitoring of real-world outcomes. Moreover, feedback from later stages, especially during validation, clinical use, and outcome monitoring, can trigger iterations across earlier stages, ensuring that the solution remains safe and clinically relevant.

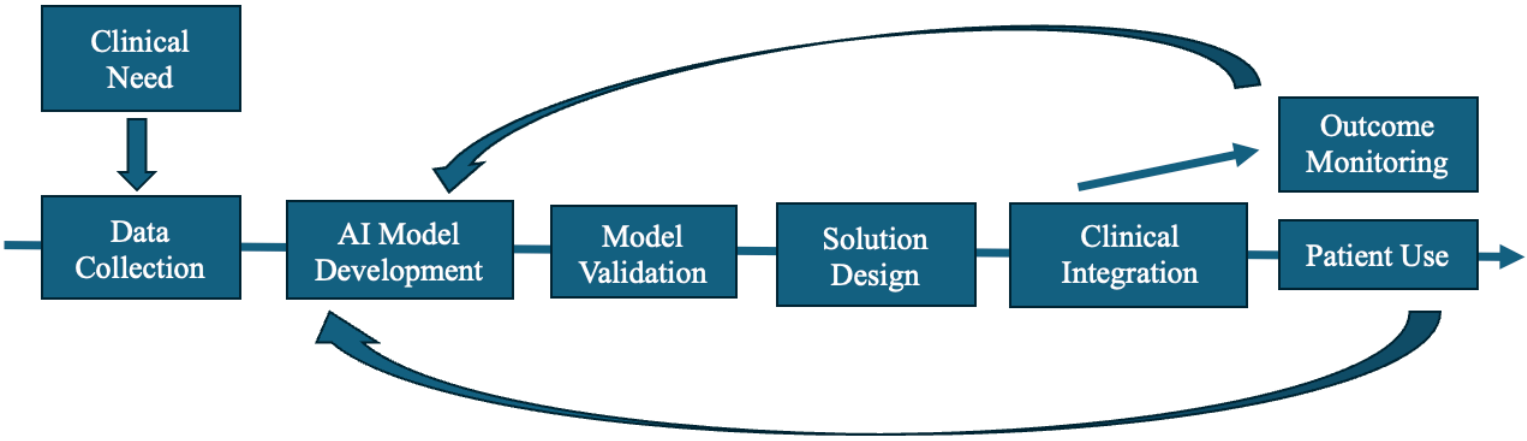

Figure 6 Simplified overview of the AI-based MedTech innovation process. The diagram outlines key phases from need identification to outcome monitoring, highlighting the iterative nature of development through feedback loops.

By presenting this structured yet flexible landscape, this section provides a foundation for understanding the broad path of AI-based MedTech innovation and mapping out its critical challenges.

4.2 Challenges and Best Practices in AI-based MedTech Innovation

Unstructured interviews and a subsequent literature review identified the key challenges faced by AI-based medical technologies from development to implementation, and best practices employed to achieve real-world impact. It is important to note that the best practices are not linked to specific individual challenges but are presented collectively, as many address multiple, interrelated barriers. For instance, in Section 4.2.2 Data Access, the best practice "Explore local healthcare facilities datasets, available data banks, open datasets, and data sets" applies to challenges such as "Data fragmentation across silos" and "Bias due to non-representative datasets". This information is summarised in Tables 2 to 9.

The best practices outlined aim to remain general and broadly applicable, to provide guidance for innovators regardless of regional context, whether in the EU, US, UK, or other health systems. However, the European context was chosen to illustrate practical

considerations and examples throughout this chapter. This regional lens provides concrete references without limiting the generalisability of the proposed insights.

4.2.1 Need Identification

Challenge	Explanation	Potential Solutions/ Best Practices proposed by this study
Mismatch between clinical needs and AI solutions	AI development is often driven by innovation goals, rather than validated clinical needs, leading to low clinical relevance or adoption.	o Identify, analyse, and engage stakeholders (e.g., patients, healthcare providers, insurers, regulators, and engineers) early to identify, validate, and rank real-world needs. o Develop differentiation in the market by analysing existing solutions and identifying gaps. o Get early feedback from relevant stakeholders, to guide product design, ensuring critical requirements are met and facilitating later adoption.

Table 2 Key challenges and best practices for AI-based medical technology innovation: Need Identification.

Among the key barriers to AI development and implementation, the mismatch between AI-based solutions and real-world clinical needs was one of the most mentioned in the interviews. Several interviewees described how solutions are often designed based on technological enthusiasm and not grounded in validated clinical challenges. Duarte Rodrigues, Director of Innovation at Neuroes, emphasised that this disconnect is a primary reason why many AI-based medical solutions fail to achieve real-world adoption. He underlined the importance of early feedback loops with stakeholders, including physicians, patients, and health administrators, to ensure that product development aligns with practical workflow requirements [Duarte Rodrigues, Appendix I]. This concern has been reflected in recent literature. In a 2023 article by Stanford Medicine, Nigam Sha, Stanford Health Care's AI Vetter-in-chief, alerted to this problem, describing how many AI tools are "poorly designed hammers looking for nails" [5].

This challenge is not limited to technical compatibility, but also to the broader understanding of the systems in which these technologies are intended to be deployed. Maria Loureiro, machine learning engineer at Loka, a technology company specialising in AI solutions for healthcare, offered a practical example involving Sword Health's attempt to deploy an AI-based triage tool for INEM, Portugal's emergency medical service. Although the tool successfully improved the efficiency of patient allocation to hospitals, INEM's

infrastructure was not prepared to handle this increase in output. The resulting misalignment led to the discontinuation of the project [Maria Loureiro, Appendix I] [88]. This case showcases how even technically robust solutions can fail when misaligned with healthcare delivery realities, ultimately resulting in wasted resources and missed opportunities.

To mitigate this, interviewees proposed actionable strategies to bridge the gap between innovation and implementation. These included systematically identifying and ranking unmet real-world needs, mapping existing solutions, and ensuring continuous dialogue with stakeholders throughout the development process. Together, these practices point to a more grounded and sustainable path for AI adoption in healthcare.

4.2.2 Data Access

Challenges	**Explanation**	**Potential Solutions/ Best Practices proposed by this study**
Insufficient data quality or completeness for reliable model development	Clinical datasets often contain missing, inconsistent, or inaccurate information, which limits their usefulness for training safe and effective AI models. High-quality, complete data is rare and resource-intensive to curate.	o Explore local healthcare facilities datasets, available data banks, open datasets, and data spaces. o Utilise federated learning to access decentralised data while maintaining data privacy. o Utilise synthetic data, data augmentation, and inferred data to completement gaps or increase limited datasets. o Consider data anonymisation or pseudo-anonymisation to access local data while complying with data privacy regulations. o Use data harmonisation techniques to integrate heterogenous datasets, increasing the amount of accessible high-quality data for AI training. o Develop data-efficient training methods to train models without reliance on large datasets.
Privacy concerns and data protection laws	Strict regulations (e.g. GDPR) restrict data sharing.	
Data fragmentation across silos	Data is often stored in disconnected systems or institutions, limiting access and use.	
Difficulty in harmonising datasets	Integrating and standardizing data from multiple sources can be challenging due to differences in formats, terminologies, and structures. This lack of harmonization hinders the ability to build comprehensive and reliable models.	

Bias due to non-representative datasets	Datasets may underrepresent certain populations, leading to biased model performance.	○ Ensure diverse and representative data to mitigate model bias. ○ Establish multiple partnerships to access multiple different datasets. ○ Use standardised metadata schemes (*e.g., DataCite, Dublin Core*) to ensure consistent data description. ○ Deposit open datasets in repositories that support machine-readable licenses and interoperability (*e.g. Zenodo*). ○ Align with FAIR principles (Findable, Accessible, Interoperable, Reusable).
Lack of data interoperability	Inconsistent data formats, terminologies, and metadata standards across institutions hinder integration, aggregation, and reuse of datasets for AI model development.	
Balancing local-specific data with generalisable datasets	Data from one institution may not reflect broader populations, affecting model generalisation.	
High cost of creating large datasets	Creating and curating high-quality datasets requires substantial time, expertise, and financial investment.	
Commercial restrictions on data sharing	Institutions or companies may restrict data use for proprietary or competitive reasons.	

Table 3 Key challenges and best practices for AI-based medical technology innovation: Data Access.

Data is a key determinant in the success of AI technology, shaping not only model performance but also its ability to generalise to real-world settings and deliver trustworthy, unbiased results. As previously discussed, many AI tools underperform outside of controlled environments, often due to being trained on limited or non-representative datasets. However, access to complete and high-quality data remains a major challenge in the field. A 2023 systematic review by Ahmed *et al.* analysed 50 studies and found that data-related barriers, specifically data quality, access, and dataset size, were the most frequently cited challenges, appearing in 42 of the studies reviewed [89]. More than a theoretical issue, the interviews with professionals in the field confirmed that this has been a persistent barrier to the development of AI-based solutions. Artur Rocha, coordinator of the Human-Centered Computing and Information Science research centre at INESC TEC, whose research focuses on privacy-preserving computing and big data, highlighted that data access challenges remain one of the primary obstacles to progress, particularly due to privacy constraints [Artur Rocha, Appendix I].

This barrier stems from a combination of societal and systemic challenges. European regulatory standards, unlike more permissive frameworks in regions such as the U.S., impose strict conditions on data access. Vasco Dias, Data Protection Officer at INESC TEC, emphasised that regulatory guidance often creates complex and time-consuming approval processes which, while essential for protecting individual privacy, can hinder innovation [Vasco Dias, Appendix I].

To note that Healthcare systems themselves contribute to the problem, as existing practices and fragmented infrastructures obstruct the creation of unified datasets. Inês Campos, a specialist in clinical interoperability at ByMe, explained how siloed systems within healthcare facilities make it difficult to access and aggregate data consistently, significantly limiting its utility for AI development [Inês Campos, Appendix I]. In addition, concerns around data bias and representativeness were raised by João Matos, a PhD researcher in AI ethics at Oxford University, who warned that datasets lacking demographic diversity risk producing models that perpetuate existing inequalities [João Matos, Appendix I].

To address these issues, several best practices have been explored in the literature, with some widely accepted while others are more contested. This variation was also evident within the interviews, where experts expressed differing views on the real-world applicability of these approaches. Amongst the most frequently mentioned solutions, federated learning enables healthcare facilities and institutions to collaboratively train models without transferring sensitive data, as is often considered a promising way forward. However, as shown by Rieke *et al.* (2020), some challenges can impede its implementation, such as system heterogeneity and communication overhead [39]. This technical complexity and significant resource demands were also emphasised by the interviewees. João Fonseca, product owner and data engineer at Promptly Health, added that, given that the infrastructure required for training the algorithms resides within the healthcare facility, these institutions are often responsible for the costly set-up [João Fonseca, Appendix I].

Other data-centric approaches, whose real-world applicability remains debated, include the use of synthetic data, data augmentation, and inferred data techniques. While these methods can help compensate for limited datasets or fill data gaps, several interviewees expressed concern over their ability to capture real-world complexity. Artur Rocha pointed out that while these techniques have potential, they must be rigorously validated before being used in clinical AI development [Artur Rocha, Appendix I]. Moreover, their utility and reliability can vary significantly depending on the specific clinical domain and data type involved. For instance, in the field of immunology, synthetic data may, in some cases, be more reliable than real-world data due to the high levels of noise and inter-individual variability in biological immune responses. By contrast, in imaging-based applications such as radiology, synthetic or augmented data may fail to reproduce the subtle, context-dependent visual cues necessary for accurate diagnosis, potentially leading to misleading results or biased model training [Artur Rocha, Appendix I]. These domain-specific differences highlight the importance of tailoring data strategies to the particularities of each clinical problem, rather than applying generalised assumptions about the effectiveness of synthetic or augmented datasets.

Regarding privacy-preserving techniques, data anonymisation and pseudonymization are often suggested as means to mitigate legal and ethical risks in handling sensitive health data. However, their practical application faces different challenges. Vasco Dias highlighted that even pseudonymized data continues to be treated as legally sensitive, and the narrow legal boundaries, combined with the risk of high penalties, often deter their use in healthcare [Vasco Dias, Appendix I]. This concern is reinforced in the literature, with a 2024 article by Gadotti et al. pointing out that traditional de-identification techniques, such as pseudonymization, offer limited protection against re-identification attacks, especially in the age of big data [90]. The authors argue that modern approaches, including anonymisation and synthetic data, while promising, still require rigorous auditing to ensure robustness against privacy breaches.

Despite ongoing discussions about the practicality and real-world potential of various data strategies, some approaches received more consistent support across the interviews. These include establishing partnerships to access diverse data sources, aligning data strategies with evolving legal frameworks such as GDPR, the AI Act, and FAIR principles, and applying data harmonisation techniques to improve interoperability and dataset usability. Another trend that emerged is the tendency for AI innovators to focus their efforts or relocate to regions with less restrictive data policies, such as the United States. According to insights from Martim Quintas e Sousa, a data scientist at MoonLake Immunotherapeutics, the European Union's stringent data privacy regulations often limit access to clinical datasets, therefore impeding AI development, whereas the US operates under a more permissive, free-market data model where companies can access large volumes of anonymised, cleaned, and standardised data, facilitating research and innovation [Martim Quintas e Sousa, Appendix I].

4.2.3 Technical and Clinical Feasibility

Challenges	**Explanation**	**Potential Solutions/ Best Practices proposed by this study**
Ensuring sufficient accuracy for medical use	Models must meet high clinical accuracy thresholds (e.g. sensitivity, specificity) to be safe and effective.	○ Validate the model's technical feasibility against the required medical accuracy, safety, and regulatory standards. ○ Assess the IT infrastructure required for deployment, including data pipelines, interoperability across systems, computing resources, and cloud integration capabilities. ○ Adopt internationally recognised data formats and ontologies (*e.g., OMOP, DICOM, SNOMED CT, HL7 FHIR*).
Limited robustness and poor generalisability	Models may fail when deployed in new settings or with slightly different inputs.	
Limited model interpretability	Black-box models reduce trust and complicate	

	regulatory approval and clinical use.	○ Enhance interoperability to facilitate AI integration and doctors' work, increasing adoption. ○ Simplify the model to increase interpretability, reduce computational demands, and streamline integration into existing workflows.
Model performance variation across demographics	Performance may degrade across patient subgroups due to training bias.	
Lack of standardised data formats	Absence of standards like DICOM, HL7 FHIR hinders development and integration.	
Model performance degradation due to data or setting drift	Over time, clinical data may shift, reducing model performance.	
Incompatibility with healthcare facility infrastructure	Outdated systems, poor interoperability, or limited computing power hinder AI deployment.	
Difficulty ensuring regulatory safety and effectiveness	Developers must align models with evolving safety and efficacy criteria.	

Table 4 Key challenges and best practices for AI-based medical technology innovation: Technical and Clinical Feasibility.

To be considered viable in clinical practice, AI systems must meet exceptionally high standards of accuracy, robustness, and safety. As Martim Quintas e Sousa pointed out, even substantial improvements over human performance may not be sufficient: "If a radiologist misses 25% of the cases, and the AI model 5%, it is still not good enough" [Martim Quintas e Sousa, Appendix I]. This reflects the strict demands of clinical environments, where even marginal errors can carry significant ethical, legal, and health consequences. In such settings, given that the integration of AI technologies entails considerable changes in workflows, training, and accountability mechanisms, these investments are only justified when the benefits clearly outweigh the associated risks and disruptions.

Despite these barriers, a key driver behind AI adoption remains its potential to mitigate critical workforce shortages, particularly in high-demand specialities, such as radiology. By automating routine tasks and redistributing workload, AI tools can help alleviate systemic strain and enhance healthcare delivery efficiency, an especially valuable proposition in resource-constrained systems.

Nonetheless, clinical feasibility depends not only on technical performance and perceived utility, but also on the system's ability to function reliably across diverse, real-world contexts. A key challenge in this regard is the so-called "myth of generalisability" of AI in healthcare. A 2024 article by Futoma *et al.* highlights how performance can decrease across settings due to variability in clinical practices over time and across health systems, patient demographics and genetic diversity, differences in hardware and software used for data capture, as well as broader social, environmental, and cultural factors [91].

A notable real-world example is Google's AI tool for detecting diabetic retinopathy. While the model demonstrated over 90% accuracy in laboratory conditions, it failed to deliver consistent results when deployed across 11 rural clinics in Thailand. Operational challenges such as poor internet connectivity, inadequate infrastructure, and inconsistent clinical workflows impacted accuracy and usability. For instance, only two clinics had darkened rooms necessary for capturing high-quality retinal images. These constraints led to delays in image uploads and reduced the number of patients screened daily from 200 to 100. Even the feature of immediate diagnostic feedback became counterproductive, as nurses were unprepared to manage follow-up care on the spot, which added pressure to their routine and led them to discourage patient participation [92], [93].

A similar issue emerged within the U.S. healthcare system. A 2021 study by Wong et al. evaluated the performance of a widely implemented proprietary sepsis prediction model and found that it underperformed significantly when deployed across different hospital settings [94]. Although multiple hospitals had adopted the model, its predictive accuracy was inconsistent and often poor, demonstrating that even within the same national healthcare system, variability in patient populations, data recording practices, and care delivery environments can undermine model generalisability. Moreover, this study reflects a broader problem: the widespread adoption of proprietary AI models in healthcare has outpaced proper evaluation, due in part to flexible regulation and easy integration. Many of these models are deployed based on confidential, non-peer-reviewed claims that may not accurately reflect real-world outcomes [94]. This reinforces the urgent need for independent validation, continuous monitoring, transparency, and formal guidance from medical organisations to ensure that AI systems remain effective and safe across diverse clinical environments.

These concerns were strongly echoed by the professionals interviewed. João Matos highlighted the risks of using unrepresentative training data, stressing that it compromises both clinical safety and fairness, potentially reinforcing system bias. He alerted to the need for local validation and ongoing monitoring [João Matos, Appendix I]. Inês Campos pointed to the lack of interoperability in clinical infrastructures, which lack standardised formats and communication between departments. Since this fragmentation hinders the implementation of AI solutions between institutions and can affect their performance, an early assessment of local IT environments is essential, along with the adoption of interoperability standards like OMOP CDM and HL7 FHIR [Inês Campos, Appendix I].

Another major technical challenge is limited interpretability. Explainable AI (XAI) is frequently proposed to bridge this gap, aiming to increase models' transparency and users' trust.

Martim Quintas e Sousa noted that explainability is not only vital for trust-building but often a prerequisite for regulatory approval, as clear reasoning behind model outputs is required for certification [Martim Quintas e Sousa, Appendix I]. In contrast, João Matos offered a more sceptical perspective, arguing that while he recognised the conceptual value of XAI, its current explanation methods may not be meaningful to clinicians and could even reinforce existing biases. [João Matos, Appendix I]. These opposing perspectives underscore the ongoing debate surrounding the practical utility of XAI in clinical settings. A 2023 review by González-Alday highlights that while XAI has the potential to improve transparency and trust in AI systems, its actual effectiveness in clinical practice remains uncertain, emphasising the need to establish standardised evaluation metrics and ensure alignment with the needs of clinicians [95].

4.2.4 Regulatory Compliance

Challenges	**Explanation**	**Potential Solutions/ Best Practices proposed by this study**
Constantly evolving AI-specific medical device regulations	The regulatory landscape for AI tools (e.g., EU AI Act, MDR, GDPR) is still evolving and often varies across regions, creating uncertainty for developers and stakeholders.	o Assess whether the AI tool qualifies as a medical device under applicable medical device regulations and determine its classification based on software classification rules (*e.g., MDR Rule 11 in the EU*) o Identify applicable AI-specific regulatory frameworks and assess requirements related to transparency, risk management, and traceability (*e.g., AI Act in the EU*). o Embed data protection and privacy principles into design and deployment in accordance with relevant protection data laws (*e.g., GDPR in the EU*). o Align with applicable health data governance and cybersecurity frameworks to ensure secure, interoperable, and resilient data exchange (*e.g., EHDS in the EU*). o If the AI system is classified as a medical device, select an appropriate conformity assessment body, considering
Uncertainty regarding AI classification	AI tools may be classified differently depending on the jurisdiction, with some tools being classified as medical devices, others as wellness products, and others falling under unique classifications based on the AI Act, which can affect the regulatory requirements.	
Overlap and complexity in regulatory requirements	AI tools may be subject to multiple overlapping regulations, including the MDR, AI Act, GDPR, and EHDS, which can complicate compliance efforts.	
Lack of clear guidance on AI-based medical	The certification process for AI-based medical devices is often unclear, particularly in terms of AI-specific software under the AI	

devices classification.	Act and MDR, creating ambiguity for manufacturers.	cost, language, and regional interpretations of applicable regulations *(e.g., Notified Bodies under EU MDR)*. o Continuously monitor regulatory developments across applicable frameworks to remain aligned with evolving compliance obligations *(e.g., MDR, AI Act, GDPR, EHDS in the EU)*.
Misalignment between AI development and regulatory approval	Failing to integrate regulatory milestones early can delay or block deployment.	
Selecting the right notified body for certification	Choosing the appropriate Notified Body for certification is critical, but the process can be complex, involving considerations like language, costs, and the specific expertise required for AI medical devices.	
Rigorous cybersecurity and data governance requirements	Medical AI tools must comply with evolving cybersecurity regulations (e.g. GDPR, Cyber Resilience Act), as failure to ensure data security can compromise certification, safety, and public trust.	

Table 5 Key challenges and best practices for AI-based medical technology innovation: Regulatory Compliance.

Regulatory compliance is a critical component in the development and deployment of AI-based medical technologies. Beyond a bureaucratic hurdle, regulatory frameworks fundamentally shape how AI tools are designed, validated, and introduced into clinical practice. Recent literature has highlighted the difficulty of applying static, traditional regulatory structures to fast-evolving, complex AI systems [96]. This difficulty can hinder innovation and limit the real-world impact of new technologies, highlighting the need for adaptive guiding frameworks [97]. These concerns were echoed in several of the interviews conducted for this study and discussed in depth with Célia Cruz, a regulatory strategy expert at Complear, and Vasco Dias [Célia Cruz, Appendix I] [Vasco Dias, Appendix I].

One of the most frequently cited challenges is the dynamic, complex, and fragmented nature of the regulatory landscape. Developers must navigate multiple overlapping legal frameworks that vary depending on the region. Schmidt et al. (2024) highlighted that regulation specific to AI remains nascent and limited. Current regulations consist of a mix of data protection, medical, technological, and human rights policies that collectively form a basic

regulatory environment for AI in healthcare [98]. This fragmentation creates substantial compliance difficulties, which limit innovation.

In the European context, developers must navigate several key legislative instruments. These include the Medical Device Regulation (MDR), which determines whether an AI system qualifies as a medical device and under what classification; the General Data Protection Regulation (GDPR), which governs personal data handling and user rights; the European Health Data Space (EHDS), which addresses data sharing and secondary use of health data; the AI Act, which imposes obligations related to transparency, risk classification, and human oversight; and cybersecurity frameworks such as the NIS2 Directive, which mandates robust security and incident response measures. As Célia Cruz noted, aligning these regulatory layers, each governed by distinct legal entities with their own interpretations and enforcement practices, poses a significant challenge for developers aiming to bring compliant and scalable AI solutions to market [Célia Cruz, Appendix I].

Another key concern is the uncertainty surrounding the classification of AI systems as medical devices, stemming from the flexible and context-dependent nature of their intended use [60], [96], [99]. Accurate classification is essential to determine applicable regulatory obligations, ensure appropriate clinical oversight, and avoid delays or legal ambiguity during market entry. This challenge was echoed in interviews. Duarte Rodrigues noted that some developers deliberately frame their solutions as "general wellness products" to avoid being classified as medical devices and thus bypass the need for clinical trials and more rigorous certification processes [Duarte Rodrigues, Appendix I]. Célia Cruz identified Rule 11 under the MDR as a common entry point for software-based tools and stressed the strategic importance of defining the intended use in ways that, where possible, avoid falling under Class IIa or higher categories that trigger more demanding clinical and technical validations. She also highlighted discrepancies in how Notified Bodies interpret classification rules across EU member states, making the choice of Notified Body and jurisdiction a critical strategic decision [Célia Cruz, Appendix I].

To manage this regulatory complexity, both Vasco Dias and Célia Cruz recommended a structured, step-by-step approach. In the European context, this begins by determining whether the AI system qualifies as a medical device under the MDR, often using Rule 11 as a guiding criterion. This classification should be followed by a thorough assessment of AI-specific obligations introduced by the AI Act. Simultaneously, developers must ensure compliance with the data protection standards set out by the GDPR and the EHDS, while also designing systems that meet applicable cybersecurity requirements. [Vasco Dias, Appendix I] [Célia Cruz, Appendix I]. For developers targeting international deployment, additional frameworks must be considered, for example, in the United States, regulations such as HIPAA (for data protection) and the FDA's Software as a Medical Device.

4.2.5 Commercialisation

Challenges	Explanation	Potential Solutions/ Best Practices proposed by this study
Ambiguity in ownership of data, model, and final product	Collaboration between healthcare facilities, companies, and developers creates complex ownership issues.	○ Evaluate AI-specific intellectual property strategies, including patents, design rights, copyright, trademarks, and trade secrets. ○ Use APIs for integration to protect proprietary AI models while enabling interoperability. ○ Leverage insurance coverage to boost clinical adoption. ○ Enable data-driven reimbursement models by demonstrating AI's clinical and economic impact using real-world data. ○ Explore commercialisation strategies: sell the product, at different regulatory stages, to big pharma companies (have pipelines established for clinical dissemination); offer a subscription-based service with continuous monitoring and updates; license the technology to third parties; open-source the model with a structured dissemination plan. ○ Explore other potential applications, market gaps, for the developed AI models. ○ Enable safe and regulatory compliant data sharing through data banks and data spaces. ○ Implement continuous deployment to keep models monitored and updated.
Choosing appropriate IP protection strategy	Developers must decide between different intellectual property protection strategies depending on their business model.	
Constant need for re-validation and re-training	Real-world performance requires ongoing evaluation and updates, increasing business model complexity.	
Difficulty in adapting AI models to different clinical settings	Local practices, patient populations, and infrastructure vary significantly.	
Uncertainty around AI business models	Developers must choose between different business models, each with pros/cons.	
Unclear reimbursement landscape	AI-based digital tools still lack clear or consistent reimbursement pathways. This affects adoption, especially if out-of-pocket costs fall on patients or facilities.	

Table 6 Key challenges and best practices for AI-based medical technology innovation: Commercialisation.

Commercialising AI technologies is essential for translating innovation into real-world impact, yet it presents a distinct set of challenges. A key concern raised by interviewees, and widely recognised in the literature, is the ambiguity surrounding ownership and the complexity

of navigating intellectual property (IP) rights. As noted by Poddar et al. (2024), the involvement of multiple stakeholders, such as healthcare facilities, software developers, and commercial partners, in the development of a single solution, combined with the dynamic nature of AI (marked by frequent retraining and updates), makes it difficult to apply traditional IP protection rights like patents. This can significantly delay technology transfer and broader adoption in healthcare [100].

To address this, developers should carefully consider alternative IP protection strategies, such as utility models, design patents, and copyright [100]. Trade secrets, for example, are often preferred for AI software to avoid disclosing sensitive algorithmic details that patents would make public. However, this approach may conflict with the regulatory challenges discussed earlier, particularly around the need for algorithmic transparency in healthcare, where explainability and accountability are crucial to building trust.

Another practical solution is the use of application programming interfaces (APIs) to balance IP protection rights with interoperability. By only exposing the model's output while keeping its internal architecture confidential, APIs enable secure external integration without compromising proprietary knowledge. This approach is exemplified by the CINDERELLA project, where the AI models developed by INESC TEC are integrated into the patient-facing app owned by CANKADO via an API (for further detail, see Chapter 5: CINDERELLA Case Study).

Another persistent barrier involves transferring AI models across healthcare settings. AI tools developed and validated in a specific clinical setting or regional context often require re-validation to ensure safety and legal compliance elsewhere, limiting scalability. João Matos highlighted this issue, warning that the lack of local validation and post-deployment monitoring could lead to serious outcomes [João Matos, Appendix I]. This concern was previously discussed in Section 4.2.4: Technical and Clinical Feasibility and justifies the need for commercialisation strategies to consider continuous deployment.

Reimbursement uncertainty presents another commercialisation challenge. Unlike pharmaceuticals or conventional medical devices, AI tools do not always fit into established reimbursement pathways. One promising strategy is the use of real-world evidence to demonstrate both clinical value and economic impact. This can support data-driven reimbursement models and facilitate negotiations with insurers and healthcare providers, thereby increasing adoption.

As for commercialisation pathways suited to AI, multiple strategies were discussed. Duarte Rodrigues, for example, noted that many small companies choose to sell or license their AI solutions at different regulatory stages, such as after obtaining CE marking - the certification indicating compliance with European health, safety, and environmental protection standards. This allows larger, better-resourced MedTech firms to scale and distribute the product [Duarte Rodrigues, Appendix I]. Another common strategy is a software-as-a-service (SaaS) model, offering AI tools on a subscription basis with ongoing updates and monitoring. This approach

is used by companies such as Complear [Maria Loureiro, Appendix I] and Promptly [João Fonseca, Appendix I]. Additional strategies include open-sourcing models with structured dissemination plans to maximise societal impact or adapting the technology to new clinical applications.

4.2.6 Clinical Implementation

Challenges	Explanation	Potential Solutions/ Best Practices proposed by this study
Low clinical usability and risk of disrupting medical practice	Poor user interface (UI), lack of workflow integration, or unclear outputs hinder clinician adoption. AI tools may disrupt clinical routines or decision-making, introducing friction or reducing efficiency.	○ Introduce AI as a decision-support tool to facilitate doctor-AI interaction, meet regulatory compliance, and keep responsibility on the doctor. ○ Assess real-world constraints within the medical setting, including existing workflows, available resources, budget limitations and user-friendliness. ○ Ensure healthcare facility readiness by deploying the required IT infrastructure, including data pipelines, digital systems, databases, and cloud infrastructure. ○ Implement clinicians/ patients training and follow-up to facilitate adoption. ○ Tailor the value proposition to address the specific needs and priorities of key stakeholders such as pharmaceutical companies, healthcare providers, insurers, and regulatory bodies. ○ Leverage explainable AI to increase transparency, facilitating user acceptance and regulatory compliance. ○ Work on a user-friendly interface.
Medical systems limitations	Many healthcare facilities lack the infrastructure to host or run advanced AI tools.	
Lack of resources for clinical readiness	Adoption may require investment in IT systems, data processes, or digital literacy.	
Difficulty building trust in AI	Patients and clinicians are often resistant to change and struggle to trust decisions made by AI.	
Enhanced burden to clinical workflows and systems	AI tools may increase rather than reduce workload if not well integrated.	

Table 7 Key challenges and best practices for AI-based medical technology innovation: Clinical Implementation.

Successfully integrating AI into clinical environments requires alignment with clinical workflows, infrastructure readiness, and user trust. The professionals interviewed in this study repeatedly underscored the practical barriers that hinder implementation, ranging from outdated medical IT systems and limited resources to the risk of disrupting established clinical routines.

From an infrastructural perspective, several participants stressed the importance of assessing real-world healthcare facility constraints early in the development cycle. These include workflow patterns, available staffing, IT and energy limitations, and cost considerations. Inês Campos highlighted that many hospitals still rely on fragmented systems incapable of supporting real-time data integration. She also noted that in Portugal, significant paper-based processes remain common, particularly outside imaging-heavy specialities like radiology and gastroenterology, which tend to be more digitally mature [Inês Campos, Appendix I]. Ensuring healthcare capacity by deploying digital infrastructure, such as data pipelines, cloud services, and database systems, is therefore critical to the sustainable use of AI tools. Duarte Rodrigues further added that cultural readiness and openness to innovation are also often underestimated barriers. For example, he noted a marked difference between northern European countries, which tend to be more receptive to new technologies, and more conservative systems like Portugal's, where institutional inertia can slow adoption [Duarte Rodrigues, Appendix I]. These differences should be carefully considered when selecting the target country or healthcare system for AI deployment.

Low clinical usability emerged as another key concern, with challenges such as non-intuitive interfaces and unclear outputs introducing friction into clinical routines. To address this, user interfaces must be simple, relevant, and tailored to clinical workflows. As Inês Campos explained, healthcare professionals are more likely to adopt tools that are intuitive and easily integrated into their daily routines [Inês Campos, Appendix I]. Another contributor to low clinical usability is the lack of interpretability. As João Matos noted, many AI tools operate as "black boxes," making it difficult for clinicians to trust, understand, or justify the model's decisions [João Matos, Appendix I]. Explainable AI (XAI) is often promoted as a solution to increase transparency and regulatory acceptance, though opinions on its practical effectiveness vary. This ongoing debate was previously discussed in Section 4.2.3 on Technical and Clinical Feasibility.

Another alternative to reduce adoption resistance is the introduction of AI as a decision-support tool, rather than a replacement for clinical judgment. This strategy maintains the clinician's authority, facilitates regulatory compliance, and builds trust. This was suggested by Duarte Rodrigues [Duarte Rodrigues, Appendix I], and is exemplified in the CINDERELLA project (see Chapter 5), where the system is presented as a support tool to assist patients and physicians in choosing between locoregional therapies. Yet, this strategy can also result in unintended consequences. A 2023 study by Jabbour et al. illustrated how AI systems can unintentionally increase the burden on clinical workflows, even when designed to support diagnostic accuracy [101]. While the AI tool improved diagnostic confidence, it also introduced cognitive overload when explanations lacked clinical clarity or relevance, especially in high-

pressure settings like emergency departments [101]. These findings raise a broader concern that even supportive AI systems may compromise safety and efficiency if poorly integrated into decision-making processes.

These examples reinforce the need for user-centred design and careful alignment with existing workflows in AI implementation strategies. Interviewees also stressed the importance of tailoring the AI solution's value proposition to different stakeholders, as healthcare facilities, pharmaceutical companies, insurers, and regulators all have unique priorities and risk tolerances. Integrating AI tools requires considering all these interests, from improving efficiency and reducing diagnostic errors to demonstrating economic value. This approach is essential for successful clinical integration.

Lastly, training emerged as a critical enabler of adoption. Inês Campos noted that digital tools often fail not due to poor performance, but because users are unfamiliar with their capabilities or wary of their risks [Inês Campos, Appendix I]. Providing education and hands-on support for both clinicians and patients can significantly lower resistance and improve engagement with AI technologies.

4.2.7 Ethics, Fairness and Sustainability

Challenges	**Explanation**	**Potential Solutions/ Best Practices proposed by this study**
Bias leading to discriminatory outcomes	Unfair results may harm already prejudiced populations.	○ Validate AI locally, monitor and retrain. ○ Mitigate bias through diverse datasets and fairness audits. ○ Launch pilots strategically with clear KPIs to demonstrate real world impact. ○ Build multidisciplinary teams with AI expertise do drive development. ○ Where required by law, conduct impact assessments for high-risk AI systems (*e.g., FRIA under the EU AI Act*). ○ Apply ethical frameworks and self-assessment tools to support the development of responsible and trustworthy AI systems (*e.g., the EU's ALTAI framework*).
Environmental impact of model training and usage	Large-scale models may have high carbon footprints.	
Difficulty in assessing tangible real-world benefits	After deployment, it's often unclear whether AI tools actually improve outcomes, reduce workloads, or save time. Healthcare facilities struggle to measure tangible impact, making long-term adoption and investment harder to justify.	

Table 8 Key challenges and best practices for AI-based medical technology innovation: Ethics, Fairness & Sustainability.

The ethical deployment of AI in healthcare depends on its ability to promote fairness, safeguard patient rights, and operate sustainably. As João Matos emphasised, AI systems are particularly prone to replicating existing societal and systemic inequalities when trained on biased or non-representative data [João Matos, Appendix I]. These biases often stem from the underrepresentation of specific groups, by ethnicity, gender, geography, or socioeconomic status, due to factors such as unequal access to healthcare or biased medical instruments (e.g. oxygen saturation monitors that misread darker skin tones [102]). Without targeted safeguards, AI tools risk reinforcing existing health disparities.

Fairness also requires protecting patient privacy. As Artur Rocha highlighted, while privacy protections and informed consent can complicate AI development, they must not be ignored, especially in healthcare settings [Artur Rocha, Appendix I]. These challenges underscore the need to embed ethical and legal expertise into AI development teams from the beginning.

To mitigate bias and promote fairness, this study recommends using diverse datasets, conducting local validation, implementing fairness audits, and maintaining ongoing model monitoring. Tools such as the European Commission's Assessment List for Trustworthy Artificial Intelligence (ALTAI) and the Fundamental Rights Impact Assessment (FRIA), required by Article 27 of the AI Act for high-stakes healthcare applications [57], offer structured frameworks to identify and reduce risks [103].

Beyond ethics and fairness, sustainability emerged as a critical but often overlooked issue. João Matos noted the high environmental cost of training advanced AI models, which consume significant energy and rely on rare materials like silica [João Matos, Appendix I]. These costs are rarely considered, even though they carry serious implications, particularly in publicly funded health systems. Sustainable development should therefore include adopting more efficient model architectures and evaluating environmental trade-offs alongside performance metrics.

Another concern raised by Pedro Ferreira, Professor and Researcher of Bioinformatics at the University of Porto, is the difficulty of demonstrating real-world impact, which is key to justifying both the ethical and economic costs of AI [Pedro Ferreira, Appendix I]. Therefore, pilot projects should be strategically designed with clear, measurable KPIs to assess patient outcomes, system efficiency, and return on investment. Frameworks such as Stanford's FURM can support this effort by aligning AI evaluation with functional, usable, reliable, and measurable impact.

Finally, ethical AI development must be deeply interdisciplinary. Several interviewees stressed the importance of collaboration across clinical, technical, legal, and ethical domains to anticipate potential problems and ensure alignment with societal values.

4.2.8 Transparency

Challenges	Explanation	Potential Solutions/ Best Practices proposed by this study
Transparency	A lack of clarity in how AI systems operate can undermine trust, hinder accountability, and hide risks.	o Focus on transparency from the beginning of the project. o Establish clear AI goals. o Use variable (non-bias) data. o Share model training methodology and potential bias. o Clearly present AI model limitations. o Explain models' decision-making logic. o Continuous auditing and performance reporting. o Understandable and accountable AI decision-making.

Table 9 Key challenges and best practices for AI-based medical technology innovation: Transparency.

Transparency is a core requirement for building trustworthy AI tools in healthcare. Without clear and transparent development and implementation practices, such as access to the data and methodologies underlying model development, the validity of AI models and their potential biases cannot be adequately assessed. This raises ethical concerns, as opaque systems may unknowingly perpetuate discrimination or exacerbate existing disparities.

In a study, Obermeyer et al (2019) exposed racial bias in a widely used healthcare algorithm designed to identify and help patients with complex health needs. As mentioned previously (see Section 2.4.2), the algorithm systematically underestimated the needs of Black patients because it used healthcare spending as a proxy for health, a biased metric influenced by the inequalities in access to healthcare [55]. This case illustrates how a lack of transparency in model design can hide critical flaws, leading to unfair and potentially harmful outcomes.

Moreover, lack of transparency is not only critical for trust and regulatory compliance but also for scientific integrity [104]. João Matos emphasised this by warning that opaque models are often reported to outperform transparent ones, which may reflect not genuine superiority but publication bias and selective reporting [105], [106]. This highlights the importance of robust documentation, open sharing of methodologies, and strict adherence to reporting guidelines to prevent manipulation and ensure reproducibility [João Matos, Appendix I].

Ultimately, transparency must be embedded throughout the development lifecycle, not just as a compliance requirement, but as a commitment to trust, safety, and scientific integrity. To this extent, this study highlights several key practices aiming to support transparency from the

earliest stages of development: defining clear AI goals, using representative datasets, documenting training methods, clearly presenting model logic and limitations, and maintaining accountability through continuous auditing and performance reporting.

4.3 Development of the VAMI Framework to Support AI-based MedTech Innovation

The AI-based MedTech innovation landscape presented in Section 4.1, together with the challenges outlined in Section 4.2, provides a comprehensive foundation for understanding the complexity of AI-based MedTech innovation. These challenges highlight critical barriers that developers face throughout the development and deployment of AI-based medical technologies. However, given their scope and complexity, developers may struggle to address them in a structured, actionable, and timely way.

To support this process, a question-based framework – VAMI: A Framework for Value-driven AI MedTech Innovation – was developed. It organises the challenges into guiding questions, helping innovators anticipate and mitigate risks throughout the innovation journey. By framing challenges as questions and pairing them with relevant best practices identified in Section 4.2, the framework aims to provide clarity and direction, thereby increasing the likelihood of achieving real-world impact.

Identifying and validating real-world needs is a prerequisite for any successful technology innovation project. Although the "mismatch between clinical needs and AI solutions", discussed in Section 4.2.1, remains a critical issue, it is not exclusive to AI and is well recognised across MedTech innovation in general. As the framework proposed in this work is tailored to the specific requirements and challenges of AI-based medical technologies, need identification is intentionally excluded from its scope. Instead, the framework is intended to complement broader innovation processes that already address this foundational stage. This will be further detailed in Section 4.4.

The VAMI Framework consists of nine key questions, each highlighting a critical stage or decision point in the AI-based MedTech development lifecycle. While the questions are numbered from 1 to 9 in a temporal flow, some address downstream issues that should be considered early on. This anticipatory approach enables innovators to plan ahead, make informed decisions, and avoid investing in concepts that may prove unfeasible later.

4.3.1 Guiding Questions of the VAMI Framework

Question 1 addresses the fundamental requirements for AI development: data. Specifically, it asks: *"Is there sufficient high-quality, accessible, and unbiased data for AI model training?"*. It stems from the challenges in Section 2.2.2 (Data Access), such as privacy constraints and biases in data collection. However, this question goes beyond technical performance. It serves as a feasibility check, assessing whether a validated clinical need can be addressed using AI. An AI solution targeting a clinically relevant healthcare problem for which the required data is

unavailable, inaccessible, or biased might not be viable. Thus, Question 1 is critical to determine both the suitability of AI for a given need and the practicality of development, helping innovators avoid costly misalignments early in the process.

Question 2 addresses the challenge of infrastructure readiness within healthcare settings, asking: *"Does the healthcare facility have the required infrastructure, and how can it be adapted for AI integration?"* This question builds on critical issues that can render a project unviable, such as system incompatibility with hospital IT infrastructure, as discussed in Section 4.2.3 (Technical and Clinical Feasibility), and the lack of resources or organisational capacity to support implementation, outlined in Section 4.2.6 (Clinical Implementation). Addressing this question early enables innovators to assess the feasibility of deployment, anticipate required adaptations, and reduce the risk of delays or failure due to inadequate clinical environments.

Question 3 concerns the high technical standards required of AI systems in healthcare and asks: *"Does the AI model meet the required medical accuracy, safety, and usability standards?"* It builds on several challenges from Section 4.2.3 (Technical and Clinical Feasibility), as well as the misalignment between AI development and regulatory approval highlighted in Section 4.2.4 (Regulatory Compliance). This question prompts developers to critically assess key model parameters while ensuring that design choices align with regulatory expectations. Failing to address these criteria early may result in costly rework or even prevent the product from entering the market.

Question 4 focuses on the complexities of intellectual property and product ownership in the context of AI. It asks: *"What strategy is most suitable to protect intellectual property and ownership?"* As discussed in Section 4.2.5 (Commercialisation), AI-based solutions often involve multiple stakeholders and data sources, raising questions about data ownership, IP rights, and contractual responsibilities. This question encourages innovators to establish clear, proactive strategies for IP protection and stakeholder alignment to avoid disputes and support long-term commercialisation.

Question 5 targets the evolving and region-specific nature of AI regulation. It asks: *"Which regulatory requirements apply to the AI-based medical technology, and how can compliance be ensured?"* This question is grounded in the regulatory challenges presented in Section 4.2.4 (Regulatory Compliance), particularly the complexity of classifying AI tools and the uncertainty around compliance paths. Although examples are drawn from EU legislation, the best practices associated with this question are relevant across jurisdictions. This question promotes early consideration of regulatory requirements, helping to reduce approval delays and avoid costly redesigns.

Question 6 addresses the commercial dimension of AI-based medical innovation, asking: *"Which business strategies will maximise the real-world impact of the AI-based medical technology?"* It encompasses challenges such as solution generalisability, uncertain reimbursement pathways, and the management of diverse exploitable assets, as discussed in Section 4.2.5 (Commercialisation). This question supports innovators in tailoring their value

propositions for strategic partners and exploring a variety of business models to optimise impact.

Question 7 focuses on the alignment of AI solutions with clinical workflows and practices, asking: *"Can the AI technology integrate feasible clinical workflows and add value to healthcare professionals?"* This question stems from challenges in Section 4.2.6 (Clinical Implementation), including clinician resistance and low clinical usability. Addressing this question helps innovators design AI systems that enhance rather than disrupt care delivery, increasing adoption likelihood and clinical impact.

Question 8 examines ethical and fairness concerns in AI development and deployment. It asks: *"How can AI medical technologies be developed and deployed fairly and ethically?"* Based on insights from Section 4.2.7 (Ethics, Fairness and Sustainability), this question addresses the particular issue of algorithmic bias and unintended consequences for vulnerable populations. It encourages innovators to embed ethical principles and fairness audits into their processes, contributing to equitable healthcare outcomes and societal trust.

Question 9 considers long-term sustainability, asking: *"What other best practices should be considered for the sustainable use of AI in MedTech?"* This question reflects challenges discussed across Section 4.2.7 (Ethics, Fairness and Sustainability), particularly the difficulty in assessing tangible return on investment and the environmental impact of AI solutions. It encourages a forward-looking perspective, prompting innovators to consider practices to ensure lasting value.

Proposed Framework

VAMI - A Framework for Value-driven AI MedTech Innovation

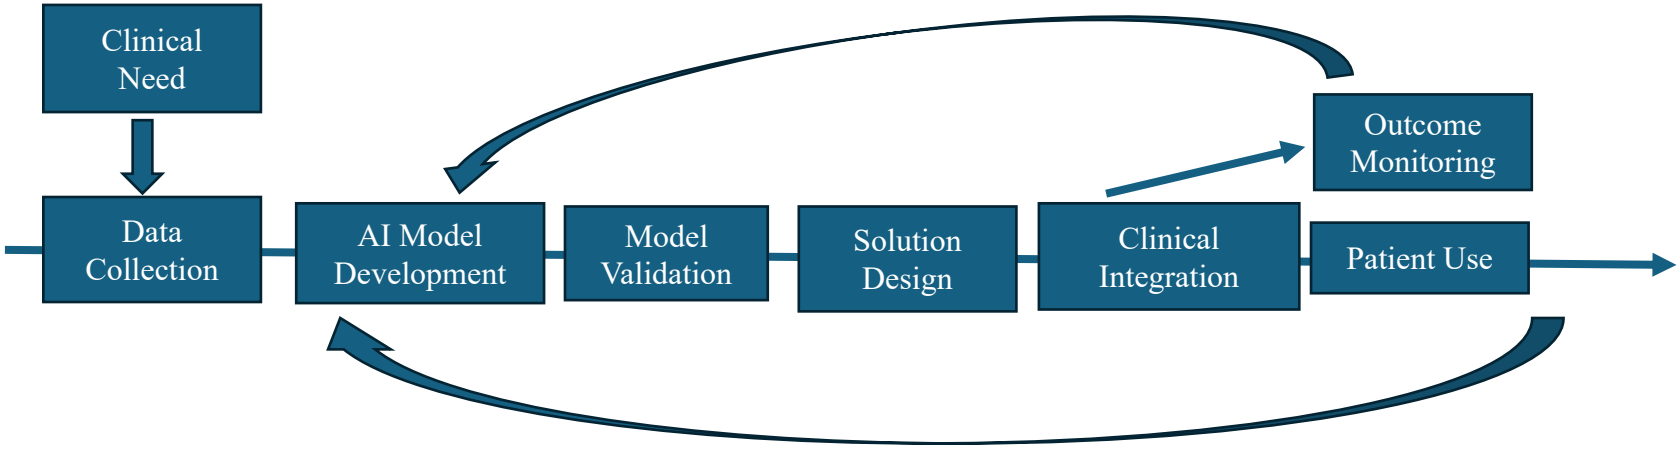

Data Collection

1. Is there sufficient high-quality, accessible, and unbiased data for AI model training?

Challenges	Insufficient data quality or completeness for reliable model development Privacy concerns and data protection laws	Data fragmentation across silos Difficulty in harmonising datasets Lack of data interoperability	Bias due to non-representative datasets Balancing local-specific data with generalisable datasets	High cost of creating large datasets Commercial restrictions on data sharing

Best Practices	Explore local healthcare facilities datasets, available data banks, open datasets, and data spaces. Establish multiple partnerships to access different datasets. Consider data anonymisation or pseudo-anonymisation to access local data while complying with data privacy regulations. Use data harmonisation techniques to integrate heterogeneous datasets, increasing the amount of accessible high-quality data for AI training. Utilise federated learning to access decentralised data while maintaining data privacy. Utilise synthetic data, data augmentation, and inferred data to complement gaps or increase limited datasets. Develop data-efficient training methods to train models without reliance on large datasets. Use standardised metadata schemes (*e.g., DataCite, Dublin Core*) to ensure consistent data description. Deposit open datasets in repositories that support machine-readable licenses and interoperability (*e.g. Zenodo*). Align with FAIR principles (Findable, Accessible, Interoperable, Reusable) Ensure diverse and representative data to mitigate model bias.

Clinical Integration

2. Does the healthcare facility have the required infrastructure, and how can it be adapted for AI integration?			
Challenges	Incompatibility with healthcare facility infrastructure	Medical systems limitations	Lack of resources for healthcare facility readiness

Best Practices	Assess real-world constraints within the medical setting, including existing workflows, available resources, budget limitations and user-friendliness. Assess the IT infrastructure required for deployment, including data pipelines, interoperability across systems, computing resources, and cloud integration capabilities. Adopt internationally recognised data formats and ontologies (*e.g., OMOP CDM, DICOM, SNOMED CT, HL7 FHIR*). Enhance interoperability to facilitate AI integration and doctors' work, increasing adoption.

Model Validation

3. Does the AI model meet the required medical accuracy, safety and usability standards?

Challenges	Ensuring sufficient accuracy for medical use	Limited model interpretability	Lack of standardised data formats	Difficulty ensuring regulatory safety and effectiveness
	Limited robustness and poor generalisability	Model performance variation across demographics	Model performance degradation due to data or setting drift	Misalignment between AI development and regulatory approval

Best Practices	Validate the model's technical feasibility against the required medical accuracy, safety, and regulatory standards. Simplify the model to increase interpretability, reduce computational demands, and streamline integration into existing workflows. Leverage explainable AI to increase transparency, facilitating user acceptance and regulatory compliance. Work on a user-friendly interface.

Solution Design

4. What strategy is most suitable to protect intellectual property and ownership?		
Challenges	Choosing appropriate IP protection strategy	Ambiguity in ownership of data, model, and final product
Best Practices	Evaluate AI-specific intellectual property strategies, including patents, design rights, copyright, trademarks, and trade secrets. Use APIs for integration to protect proprietary AI models while enabling interoperability.	

Solution Design

5. Which regulatory requirements apply to the AI-based medical technology, and how can compliance be ensured?			
Challenges	Constantly evolving AI-specific medical device regulations	Overlap and complexity in regulatory requirements	Selecting the right Notified Body for certification
	Uncertainty regarding AI classification	Lack of clear guidance on AI software certification	Rigorous cybersecurity and data governance requirements
Best Practices	Assess whether the AI tool qualifies as a medical device under applicable medical device regulations and determine its classification based on software classification rules (*e.g., MDR Rule 11 in the EU*) Identify applicable AI-specific regulatory frameworks and assess requirements related to transparency, risk management, and traceability (*e.g., AI Act in the EU*).		

	Embed data protection and privacy principles into design and deployment in accordance with relevant protection data laws (*e.g., GDPR in the EU*). Align with applicable health data governance and cybersecurity frameworks to ensure secure, interoperable, and resilient data exchange (*e.g., EHDS in the EU*). If the AI system is classified as a medical device, select an appropriate conformity assessment body, considering cost, language, and regional interpretations of applicable regulations *(e.g., Notified Bodies under EU MDR)*. Continuously monitor regulatory developments across applicable frameworks to remain aligned with evolving compliance obligations *(e.g., MDR, AI Act, GDPR, EHDS in the EU)*.

Clinical Integration

6. Which business strategies will maximise the real-world impact of the AI-based medical technology?		
Challenges	Constant need for re-validation and re-training Uncertainty around AI business models	Difficulty in adapting AI models to different clinical settings Unclear reimbursement landscape
Best Practices	Tailor the value proposition to address the specific needs and priorities of key stakeholders such as pharmaceutical companies, healthcare providers, insurers, and regulatory bodies. Leverage insurance coverage to boost clinical adoption. Enable data-driven reimbursement models by demonstrating AI's clinical and economic impact using real-world data. Explore commercialisation strategies: sell the product, at different regulatory stages, to big pharma companies (have established pipelines for clinical dissemination); offer a subscription-based service with continuous monitoring and updates; license the technology to third parties; open-source the model with a structured dissemination plan.	

	Explore other potential applications and market gaps for the developed AI models. Enable safe and regulatory-compliant data sharing through data banks and data spaces.

Clinical Integration

7. Can the AI technology integrate feasible clinical workflows and add value to healthcare professionals?			
Challenges	Low clinical usability and risk of disrupting medical practice	Difficulty building trust in AI	Enhanced burden to clinical workflows and systems
Best Practices	Introduce AI as a decision-support tool to facilitate doctor-AI interaction, meet regulatory compliance, and keep responsibility on the doctor. Implement clinician/ patient training and follow-up to facilitate adoption. Implement continuous deployment to keep models monitored and updated. Ensure healthcare facility readiness by deploying the required IT infrastructure including data pipelines, digital systems, databases, and cloud infrastructure.		

Patient Use

8. How can AI medical technologies be developed and deployed fairly and ethically?	
Challenges	Bias leading to discriminatory outcomes
Best Practices	Validate AI locally, monitor and retrain.

	Mitigate bias through diverse datasets and fairness audits. Where required by law, conduct impact assessments for high-risk AI systems (*e.g., FRIA under the EU AI Act*). Apply ethical frameworks and self-assessment tools to support the development of responsible and trustworthy AI systems (*e.g., the EU's ALTAI framework*). Emphasise transparency from the beginning of the project by establishing clear AI goals, using variable (non-biased) data, sharing the model training methodology and potential bias, clearly presenting the AI model limitations, explaining the models' decision-making logic, implementing continuous auditing and performance reporting, and ensuring understandable and accountable AI decision-making.

9. What other best practices should be considered for the sustainable use of AI in MedTech?		
Challenges	Environmental impact of model training and usage	Difficulty in assessing tangible real-world benefits
Best Practices	Launch pilots strategically with clear KPIs to demonstrate real-world impact. Build multidisciplinary teams with AI expertise to drive development.	

4.4 Positioning the VAMI Framework within Broader Innovation Pathways: Stanford's Biodesign

The VAMI Framework is designed to stand as a practical and actionable tool for developers navigating the complex landscape of AI-based medical innovation. By highlighting key challenges in AI-based MedTech innovation and offering actionable strategies to address them, it is a flexible tool for innovators in different contexts.

However, while the VAMI Framework offers value independently, its scope is deliberately narrower than full-cycle innovation methodologies. For instance, it does not address critical early-stage activities such as need identification and clinical validation. To this end, the VAMI Framework can benefit from being embedded within broader innovation frameworks that cover these foundational stages.

One such innovation framework is Stanford's Biodesign process, which has long provided a structured, needs-driven approach to medical technology development. As discussed in Section 2.5.1, despite Biodesign being well-suited for identifying unmet clinical needs and developing viable solutions, it is not specifically designed to address the distinctive challenges associated with artificial intelligence.

To illustrate how the VAMI Framework can complement such broader approaches, this section presents a mapping between the proposed framework's nine guiding questions and the three phases of the Biodesign process. For instance, Questions 1 and 2 fall under Phase 1 – Identify, emphasising the importance of early alignment with data availability and system realities. Questions 3 and 4 are placed in Phase 2 – Invent, addressing concept design, validation and IP protection. Questions 5 to 9 are integrated into Phase 3 – Implement, focusing on commercialisation, adoption, and long-term sustainability. This alignment is illustrated in Figure 7 (adapted from [73]), which shows how the VAMI Framework can be layered into an existing innovation pathway, creating an enhanced process tailored for AI-based MedTech innovation.

It is important to note that both the VAMI Framework and Stanford's Biodesign process are iterative and cyclical, meaning that innovators may revisit earlier phases as new insights emerge throughout development.

Ultimately, while integration with broader frameworks can increase the robustness of innovation efforts, the VAMI Framework was not developed as an extension of Biodesign or any other specific model. Its primary goal is to guide developers through AI-specific challenges, providing a flexible, modular resource that can stand on its own or be used alongside established methods, depending on the context.

VAMI Framework and Biodesign:

A Combined Approach to AI-Based MedTech Innovation

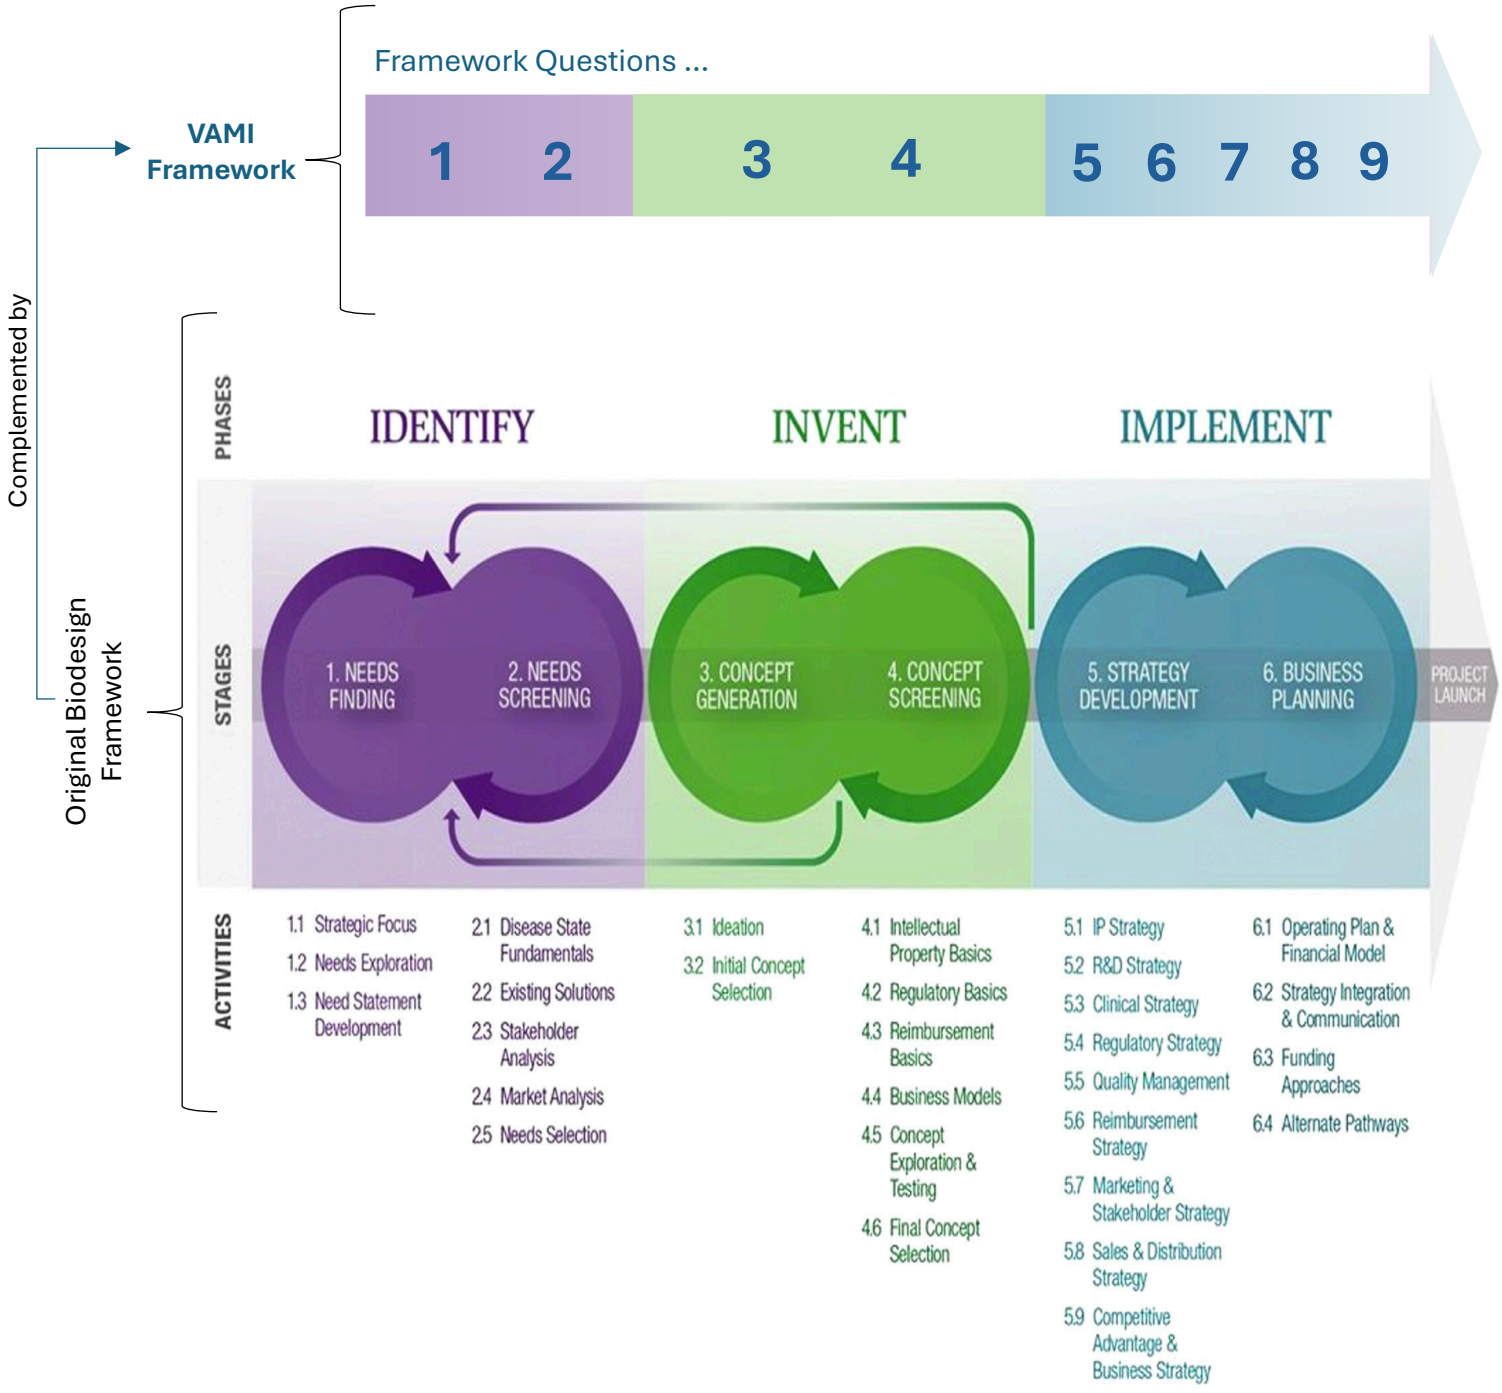

Figure 7 A combined approach to support value-driven AI-based MedTech innovation. The guiding questions of the proposed VAMI Framework are aligned with the Biodesign phases - Identify, Invent, and Implement.

4.5 Validation of the VAMI Framework

To assess the practical relevance and robustness of the VAMI framework, feedback was gathered through in-depth discussions with three professionals working at the forefront of AI-driven MedTech innovation:

- Frederico Carpinteiro, CEO at Amparo (Portugal), brings extensive expertise in the biomedical sector, with a focus on improving patient outcomes and accessibility in prosthetics through innovative smart technologies.
- Michel Nemery, Neuroradiologist, Head of Clinical Development at Cerebriu (Denmark), Chairman of the Department of Radiology at Herlev and Gentofte Hospital, Co-chairman at the Board of Radiology of the Capital Region of Denmark, and founder of Denmark's first needs- and user-driven clinical Innovation Unit, brings a strong clinical, strategic, and innovation-driven perspective.
- Miguel Amador, Partner and Chief Innovation Officer at Complear and Managing Director at Pi Ventures, brings deep expertise in MedTech innovation, with a focus on AI regulation and strategic implementation within healthcare ecosystems on both European and American markets.

Their complementary perspectives**,** spanning entrepreneurship, clinical practice, regulatory expertise and strategic development, helped validate the framework's structure and content. Their input also helped highlight key topics and identify areas for improvement. A summary of their insights and recommendations is provided in Table 12 (Appendix I).

4.5.1 Framework Scope

All three experts recognised the overall relevance and necessity of a specialised tool like the VAMI framework for guiding the development of AI-based medical solutions.

Nemery emphasised that identifying and validating clinical need is crucial for success. He highlighted the importance of having a deep understanding of user needs and the clinical context. Amador also reinforced this concern, underscoring that a strong alignment between clinical needs and MedTech development is essential for ensuring adoption. Following the same logic, Carpinteiro positioned the VAMI framework as a practical complement to Stanford's Biodesign methodology, suggesting it is most effective after needs have been identified but before implementation, bridging the gap between problem identification and solution deployment.

This feedback collectively validates the framework's relevance while also suggesting that future iterations may benefit from more explicitly incorporating a step focused on clinical and stakeholder requirement validation prior to AI-specific considerations. Although this dimension is addressed in the broader study, it currently sits outside the core structure of the VAMI framework, given its specific focus on challenges unique to AI-based innovation.

To validate the internal components of the VAMI framework, namely the nine guiding questions, associated challenges and proposed best practices, each question was discussed during the expert interviews. It is important to note that while Frederico Carpinteiro and Michel Nemery engaged directly with the framework in a question-by-question format, reinforcing, challenging or suggesting refinements to its elements. Miguel Amador adopted a more thematic approach. As such, his insights are integrated primarily into Sections 4.5.1 (Framework Scope) and 4.5.3 (Reflections and Structural Considerations), and do not appear in Section 4.5.2. The following section presents feedback from the two former experts.

4.5.2 Insights by Component

Question 1: Is there sufficient high-quality, accessible, and unbiased data for AI model training?

This question was strongly validated by Frederico Carpinteiro and Michel Nemery as a foundational concern in AI-driven MedTech innovation.

Carpinteiro observed that this question is frequently answered in the negative, suggesting that a non-AI-driven solution may sometimes be more appropriate. In particular, he highlighted the lack of data standardisation and interoperability, especially across EU institutions, as a critical structural barrier, aligning closely with the challenges highlighted in the VAMI Framework. His feedback also pointed to a possible refinement: if data is available, the framework could more explicitly demand a reassessment of stakeholder-specific needs, as concerns and expectations often differ when AI is involved.

Nemery also addressed the importance of this question, raising a concern regarding whether dataset bias can ever be fully resolved due to the deep-rooted biases in healthcare systems. Despite this, he reinforced the value of mitigation strategies already reflected in the framework, such as the inclusion of large, diverse datasets, along with continuous data access and regular monitoring for representativeness. This feedback further highlights the importance of treating data acquisition and assessment as a continuous process, as addressed in the framework's feedback loop between outcome monitoring and data collection.

Question 2: Does the healthcare facility have the required infrastructure, and how can it be adapted for AI integration?

Both experts affirmed the relevance of this question, highlighting infrastructure readiness as central to successful AI deployment.

As noted in the framework, Nemery observed that hospitals often operate in siloes, with slow decision-making processes that hinder the scaling of innovative technologies. He also cautioned that institutional leadership buy-in must be handled carefully, as senior stakeholders have more to lose and are prone to resist disruptive change. To address this, he advocated for "plug-and-play" solutions that integrate seamlessly into existing workflows and use cases, maximising clinician uptake while minimising disruption. While the VAMI Framework already promotes user-friendly design, this feedback suggests that including "plug-and-play"

approaches as best practices could enhance its applicability. He also pointed out that, contrary to academic or urban hospitals, which are subject to more intricate decision processes, rural hospitals may be more open to innovation.

From a complementary standpoint, Carpinteiro stressed that technical infrastructure alone is not enough. Without proper training, frontline staff may input data incorrectly, which compromises AI reliability. This "human-in-the-loop" challenge emerged as a critical bottleneck, suggesting that operational realities must be considered alongside technical feasibility. This challenge is not currently covered by the VAMI Framework and could be considered in future refinements.

Question 3: Does the AI model meet the required medical accuracy, safety and usability standards?

While this question targets key performance thresholds and high standards set for AI-based medical tools, the experts focused instead on the importance of a cyclical, iterative development process that incorporates real-world validation early and continuously.

Acknowledging that needs and conditions often evolve during development, Carpinteiro recommended a structured two-stage validation process: an initial phase to confirm clinical and stakeholder requirements before development begins to ensure alignment with real-world needs, and a subsequent post-development phase to evaluate how well the final solution meets those initial goals and to determine any necessary adjustments or improvements. This “clinical and stakeholder requirements assessment” is not currently addressed in the VAMI Framework. Given its potential to guide development in a more grounded way, it is a valuable addition for future iterations.

Nemery reinforced the need for post-market performance monitoring, including regular model updates to maintain clinical relevance as data and workflows evolve. His feedback validates the VAMI Framework’s cyclical design. Moreover, in the context of solution design, he recommended using a "moonshot" exercise - envisioning a perfect solution - as a method to set aspirational goals for design. This exercise can help ensure that usability and safety standards are not only met but also guided by a vision of what optimal performance could look like.

Together, these insights validate the framework’s emphasis on iterative development, while also indicating that ongoing stakeholder alignment and visionary thinking can strengthen real-world applicability.

Question 4: What strategy is most suitable to protect intellectual property and ownership?

This question received limited attention across the interviews. While the significance of intellectual property (IP) was acknowledged, the limited emphasis suggests that it may not be a dominant concern in AI innovation or may often be embedded in broader development

processes, such as commercialisation strategy. However, given the small number of interviews conducted, it is not possible to draw strong conclusions on this point.

Question 5: Which regulatory requirements apply to the AI-based medical technology, and how can compliance be ensured?

The importance of this question was strongly validated by both experts, who emphasised regulatory clarity as a key enabler of successful and timely MedTech innovation, especially in AI.

Carpinteiro underscored the importance of explicitly addressing data protection compliance from the beginning, while Nemery stressed that clearly defining the AI system's intended use early on simplifies regulatory navigation and prevents misclassification, particularly if the system goes beyond decision support. He further noted that early regulatory clarity aids in commercialisation, particularly when partnering with larger industry players. These insights directly support the best practices proposed under this question in the VAMI Framework.

Overall, their feedback reinforces the idea that regulation should be integrated strategically into early planning and not approached as a final obstacle.

Question 6: Which business strategies will maximise the real-world impact of the AI-based medical technology?

This question, which addresses business viability and the broader strategies needed to sustain adoption, was validated by both experts, with a focus on regulatory and reimbursement considerations.

Carpinteiro advised innovators to consider reimbursement pathways early, as the presence or absence of established codes can significantly accelerate or delay implementation. Nemery focused on defining KPIs linked directly to measurable clinical outcomes, arguing that these metrics should guide both development and adoption.

These insights reinforce the framework's inclusion of a business-oriented question while also illustrating its connection to other areas, particularly the alignment between technical success and business sustainability discussed later in Question 9.

Question 7: Can the AI technology integrate feasible clinical workflows and add value to healthcare professionals?

This question was highlighted by both experts, who emphasised that usability is not a secondary consideration but a critical success factor for AI adoption in healthcare.

Carpinteiro stressed that user-centric design is essential from the outset, with features such as interpretability, transparency, and ease of use embedded from the beginning. These priorities echo and expand concerns raised in Question 2 regarding integration into existing workflows.

Nemery highlighted that AI tools should be positioned as assistive technologies, designed to support rather than replace clinicians. He also reaffirmed the value of the "plug-and-play" functionality to facilitate seamless clinical integration. These inputs align and validate the best practices already proposed in the VAMI Framework.

Beyond these suggestions, Nemery proposed three strategies for AI solutions to increase clinical value while reducing barriers to adoption:

A. Target low-complexity or high-volume tasks: By focusing AI development on areas that do not require deep clinical expertise, such as initial triage, developers can produce tangible, rapid benefits. For example, in stroke triage, most patients turn out not to have a stroke. Therefore, an AI tool that can rapidly and reliably flag these cases could significantly reduce unnecessary hospital stays, increasing hospital savings.

B. Improve on the weakest alternative rather than the best: Rather than aiming to outperform the best clinical experts, AI can offer value by being more consistent than the weakest link in the chain, such as an overworked junior doctor on a night shift. This redefinition of success can lower the performance threshold for impactful deployment.

C. Prioritise timeliness over precision gains: In many clinical scenarios, speed can be more valuable than perfect accuracy. For instance, if an AI tool can provide a diagnosis of comparable quality but in a shorter timeframe, this can lead to earlier interventions, reduced complications, better outcomes, and hospital long-term savings.

These strategies are not currently addressed in the VAMI Framework, but given their potential to expand the scope of clinical utility in real-world settings, including them in future iterations could be valuable.

Question 8: How can AI medical technologies be developed and deployed fairly and ethically?

This question was strongly supported by both experts, particularly Nemery, who highlighted the real-world challenges of fairness and ethical deployment.

He illustrated this through a case at his hospital, where an AI system trained on a different population failed during local validation, dropping from the promised 90% accuracy to just 14%. This example underscores the risk of training models on idealised datasets and the importance of local validation. Nemery further cautioned that technology drift, such as changes in data generation methods (e.g., new imaging equipment), can impact performance just as much as population drift.

These risks align perfectly with the VAMI Framework's highlighted challenges and recommended best practices, such as local validation, ongoing audits, and transparency. Additionally, both experts called for the inclusion of multidisciplinary teams, bringing together legal, clinical, and social perspectives to ensure fairness is embedded throughout both development and deployment. This recommendation also supports the framework's guidance under Question 9, concerning long-term sustainability.

Question 9: What other best practices should be considered for the sustainable use of AI in MedTech?

The framework's final question, regarding sustainability, was discussed with a focus on both outcome measurement and team dynamics.

Carpinteiro cautioned against overly technical evaluations of value, arguing for attention to ROI and net impact. Nemery echoed this concern, stressing that clinical and economic KPIs must guide long-term evaluations. He also reiterated the importance of multidisciplinary collaboration, not only as a mechanism for fairness, as discussed under Question 8, but also as a means of maintaining long-term alignment with institutional priorities and patient outcomes. These insights validate the VAMI Framework's emphasis on continuous monitoring, stakeholder alignment, and team dynamics as key enablers of sustainable AI adoption.

4.5.3 Structural Considerations and Final Reflections

A key structural critique of the VAMI Framework is that, if followed strictly as a step-by-step process, it may not adequately reflect the cyclical and iterative nature of AI-driven MedTech innovation. As discussed in earlier chapters, the development of such technologies is rarely linear, with requirements evolving, and validation often occurring in parallel with design.

To address this, the framework should be understood not as a rigid sequence of steps but as a flexible tool to be consulted at the beginning of a project and revisited regularly throughout its lifecycle. Amador strongly echoed this point, noting that critical aspects, such as ethical oversight (e.g., FRIA), dataset audits, or transparency, are often introduced too late in projects, despite being essential from the start. He stressed that these principles are not "steps" but continuous requirements that must guide the entire process.

Additionally, although the importance of early and continuous stakeholder engagement is emphasised in this study's discussion, it should be more explicitly embedded in the framework itself. As proposed by Carpinteiro, a new step of "Identifying Clinical and Stakeholder Requirements" could be introduced between the stages of clinical need identification and data assessment. This would help ensure that stakeholder insights are translated into clear, actionable technical specifications, bridging the gap between problem definition and solution design.

In addition, Amador brought attention to the often-overlooked importance of following international standards throughout the innovation process. He warned that while several AI-

specific standards are emerging, such as those developed by the International Organisation for Standardisation (ISO), they are still underutilised in practice. Some are poorly known or insufficiently integrated into medical technology innovation, while others conflict with established standards for traditional medical devices (like the MDR), potentially leading to redundancy or conflict. His insights underscore that standards are essential to ensure uniformity, safety, and interoperability, but also that innovators need practical tools to navigate and apply these standards as the regulatory landscape evolves. This reinforces the role of the VAMI Framework as a means to promote structured alignment with best practices and compliance expectations.

Amador further observed that several best practices discussed in the framework are mixed with regulatory obligations. For example, while local validation is considered a best practice, it is not legally required, unlike post-deployment monitoring, which is mandated under regulations like the EU AI Act. He clarified that AI solutions intended for general deployment must be trained on representative datasets (a requirement), but they are not obliged to undergo local validation after deployment, even though performance drift may justify it. Therefore, future iterations of the VAMI Framework could benefit from a clearer distinction between minimum legal obligations and additional best practices that exceed regulatory requirements.

Recurring themes identified throughout the framework validation interviews suggest key enablers of successful AI-based medical technology innovation:

- Ensuring a validated clinical need through early and meaningful stakeholder alignment, with a focus on challenges that can be effectively addressed using data-driven approaches.
- Establishing multidisciplinary teams that bring together technical, clinical, legal, social, and ethical expertise to support comprehensive and context-aware innovation.
- Integrating solutions into existing clinical workflows and healthcare structures to minimise disruption and avoid the need for major systemic changes.
- Conducting local validation, continuous fairness monitoring, and regular model updates to support trust, adoption, and long-term relevance across diverse settings.
- Defining measurable KPIs and outcome metrics to evaluate return on investment, clinical impact, and overall system value throughout the solution lifecycle.
- Adopting relevant standards at every stage of development and deployment, while remaining critical of their limitations.

Ultimately, the validation exercise confirms that the VAMI Framework addresses core challenges in AI-specific MedTech innovation. The expert feedback reinforces its utility as a practical and flexible guide and supports its positioning as a valuable complement to broader methodologies like Biodesign.

To further validate the proposed VAMI Framework and provide insights into its practical application, a real-world case study was included. The selected case is the CINDERELLA Project, an initiative funded by Horizon Europe - the European Union’s flagship research and innovation programme.

5. Case Study on the CINDERELLA Project

The CINDERELLA Project is developing a new AI-based medical tool to support locoregional treatment decisions for breast cancer patients: the *CINDERELLA APProach*.

The *CINDERELLA APProach* integrates a web-based healthcare platform with AI-powered tools designed for clinical settings. Its core objective is to enhance shared decision-making between patients and clinicians by offering more reliable and personalised predictions of the aesthetic outcomes of locoregional therapies. By visualising potential post-treatment appearances, the platform helps patients form realistic expectations, reduce decision-related stress, and feel more confident in their choices. Ultimately, the CINDERELLA project aims to improve patient satisfaction with their treatment results, therefore supporting long-term quality of life.

As a large-scale, multi-stakeholder initiative, the CINDERELLA project reflects many of the real-world challenges and opportunities associated with developing and implementing AI in healthcare. Its inclusion as a case study helps ground the proposed framework in a real-world context.

5.1 The Problem

In 2022, breast cancer was diagnosed in over 2.3 million individuals worldwide [107]. Thanks to early detection and treatment advances, 10-year survival rates now exceed 80 % in many high-income countries [108]. As survival improves, attention has shifted toward optimising post-treatment quality of life (QoL), particularly aesthetic and functional outcomes, which significantly affect physical health, self-esteem, and psychosocial well-being.

Breast cancer treatment typically involves locoregional approaches, often through surgery combined with radiotherapy, aiming to control the disease within the breast and surrounding area. This is the primary form of treatment for early-stage breast cancer, meaning that optimising these procedures has the greatest potential to benefit the growing number of patients diagnosed with low- or intermediate-risk disease [109]. However, there are multiple procedures under the scope of locoregional treatment, and while these options often offer equivalent survival rates, they can lead to significantly different aesthetic outcomes and long-term QoL.

Currently, patients are presented with multiple surgical choices, ranging from mastectomy - full breast removal - to breast-conserving procedures. However, standardised and objective comparative data on aesthetic outcomes and QoL are limited. Decision conversations typically rely on the surgeon’s experience and subjective examples, leaving patients to independently seek information from external sources. This situation creates stress, inconsistency, and misalignment with patient expectations.

5.2 The CINDERELLA Solution

The *CINDERELLA APProach* is the main technological solution developed within the CINDERELLA project, resulting from the collaboration between its two core technology

partners: INESC TEC and CANKADO. It combines (1) AI models developed by INESC TEC, which analyse patients' biometric features to retrieve comparable past cases and generate predictive visualisations of likely aesthetic outcomes, with (2) user-facing digital platforms developed by CANKADO, including a web-based interface for clinicians and a mobile app for patients. By presenting pre- and post-operative images and providing evidence-based information on treatment options and associated risks, the *CINDERELLA APProach* supports shared decision-making in breast cancer locoregional treatment, helping align expectations and ultimately improving the experience for both patients and clinicians.

As an AI-based medical solution designed for clinical application, the *CINDERELLA APProach* presents a relevant case to evaluate the presence of the identified innovation challenges and the suitability of the proposed VAMI Framework's best practices.

5.3 Aligning the VAMI Framework with the CINDERELLA project

In this chapter section, the CINDERELLA project is assessed through the lens of the VAMI Framework. Specifically, the objective is to examine whether the challenges outlined in Section 4.2 have emerged during CINDERELLA's development, how they were addressed, and to what extent the VAMI Framework's recommended practices proved to be effective.

As of the current stage of writing (12 June 2025), the *CINDERELLA APProach* remains in the clinical trial phase, having successfully reached its recruitment target of 1000 patients. The project has not yet entered commercialisation, and several anticipated challenges remain unaddressed. In these cases, the VAMI Framework's guidelines and best practices will serve to anticipate potential barriers and propose best practices for addressing them in future stages.

1. Is there sufficient high-quality, accessible, and unbiased data for AI model training?

The development and training of CINDERELLA's AI algorithms required access to large volumes of high-quality, representative, and unbiased data. As discussed in previous sections, this is essential not only for achieving high model accuracy but also for ensuring consistent performance across diverse populations and clinical settings. In *CINDERELLA APProach's* case, the necessary data consisted of pre- and post-operative photographs of patients' torsos.

To facilitate access to such data, the project partnered with five clinical institutions in Portugal, Germany, Poland, Italy, and Israel. These partners were responsible for patient engagement, data collection, and secure handling of sensitive visual information. Data harmonisation was not a problem, as the data collected, such as biometric measurements, are generally standardised clinical parameters. The geographic distribution of the clinical partners also introduced valuable population diversity, helping to reduce bias and improve the generalisability of the dataset.

Establishing partnerships for data access aligns with the VAMI Framework's proposed best practices. However, even if the solution proves effective during clinical trials, this study

underscores the need for local validation and continuous fairness monitoring following deployment. Regular model updates should also be planned to ensure sustained performance, fairness, and adaptability over time.

2. Does the healthcare facility have the required infrastructure, and how can it be adapted for AI integration?

The *CINDERELLA APProach* is designed as a web-based tool, intentionally requiring minimal adaptation from healthcare facilities. It does not need major changes to physical infrastructure, nor does it rely on specific IT systems or high energy requirements. Clinicians can access the platform via standard web browsers, and patients can engage with it through a mobile application, ensuring ease of integration into existing digital environments.

As a cloud-based solution, the approach relies on stable and reliable network infrastructure to ensure smooth operation and timely data exchange. While cloud architecture enables scalability and remote access, it also requires robust cybersecurity measures to safeguard sensitive patient data and ensure compliance with data protection regulations. This aligns with the regulatory best practices of the proposed framework, particularly emphasising the importance of adhering to relevant health data governance and cybersecurity standards.

To further support the workflow, the *CINDERELLA APProach* can optionally integrate the Pink Robot, a commercial hardware add-on developed for the automatic, standardised capture of patient photographs. While this device can streamline clinical processes and improve the consistency of image quality, its use is not essential for the platform's core functions. Healthcare facilities that decide to adopt the PINK robot should consider several factors, including the need for physical space, compatibility with existing IT, and the costs associated with purchasing, installing, and maintaining the equipment.

For continuous use beyond the clinical trial, routine data sharing with the technology provider, such as uploading new patient images, can be performed directly through the web platform, eliminating the need for new data pipelines or complex integrations.

3. Does the AI model meet the required medical accuracy, safety and usability standards?

The AI models at the core of the *CINDERELLA APProach* have been subject to ongoing testing and refinement by the development team at INESC TEC. As the clinical trial progressed and new patient data became available, the models were continuously retrained and re-evaluated to enhance performance. Ultimately, the clinical trial will determine whether the models meet the required standards of accuracy, robustness, and safety under real-world conditions.

In parallel, both the patient-facing app and the clinician-facing web platform were designed with a strong focus on usability. Their interfaces are being continuously refined based on user feedback collected during the clinical trials. To further enhance trust and usability, the *CINDERELLA APProach* incorporates explainable AI methods, enabling clinicians and patients to better understand how predictions are generated, therefore supporting more

informed and transparent decision-making. This use of user-friendly interfaces and adoption of explainable methods aligns with the VAMI Framework's proposed best practices.

4. Which strategy is most suitable to protect intellectual property and ownership?

The *CINDERELLA APProach* is built from different components, such as the AI algorithms and trained models created by INESC TEC, and the web-based platform developed by CANKADO, which the consortium opted to connect via an API. This structure is in line with the VAMI Framework' proposed best practices and enabled each contributor to retain ownership over their respective technologies while facilitating interoperability.

Moreover, the consortium has largely adopted an open-source-oriented exploitation strategy. This includes publishing the methodologies used to train and validate the models, as well as releasing algorithmic code under permissive licenses. While this approach promotes transparency, collaboration, and scientific advancement, it limits IP protection strategies and may pose challenges in maintaining a competitive advantage.

To support the commercialisation of the *CINDERELLA APProach,* IP protection can be secured, for instance, through trademark registration to protect the brand and support its future adoption and market translation. Furthermore, confidential know-how accumulated in the innovators' teams will contribute to the real-world impact of the project's outcomes.

In addition to technology ownership, data rights represent a critical aspect of the project's exploitation strategy. Healthcare institutions, which contribute high-quality patient data required for AI model training, retain *sui generis* rights over this information. To recognise their role, ensure compliance with data governance regulations, and incentivise ongoing collaboration, the consortium proposed sharing a percentage of future royalties or commercial revenues with these institutions.

5. Which regulatory requirements apply to the AI-based medical technology, and how can compliance be ensured?

During development, CINDERELLA made it a priority to comply with important data protection regulations, particularly the General Data Protection Regulation (GDPR) and emerging guidelines under the European Health Data Space (EHDS). These efforts reflect the data governance best practices outlined in the VAMI Framework, which stress secure, interoperable, and compliant data flows from the earliest stages of development.

However, the regulatory classification of the *CINDERELLA APProach* as a medical device has not yet been completed. This represents a significant gap at the current stage, as classification will define the applicable conformity requirements and determine the appropriate regulatory pathway. According to the VAMI Framework, initiating the classification process early is essential to avoid delays, manage risk, and align the development process with regulatory expectations. Delayed classification can result in time- and resource-intensive adjustments later, particularly when selecting a Notified Body and preparing the necessary technical documentation.

According to the proposed framework, the classification process should begin by consulting the Medical Device Regulation (MDR), with specific attention to Rule 11, which concerns software used for diagnostic or therapeutic decisions or for monitoring physiological processes. The final classification will depend heavily on how the intended purpose of the tool is framed. For instance, if the *CINDERELLA APProach* is clearly positioned as a decision-support system that assists but does not replace clinical judgment, it may be eligible for a lower risk Class I classification. This could significantly reduce the burden of conformity assessments compared to Class IIa or higher classifications.

In parallel, the solution must also be evaluated under the EU Artificial Intelligence Act, which introduces additional obligations based on risk category and intended use. If classified as a high-risk AI system, the *CINDERELLA APProach* would be subject to further requirements regarding transparency, human oversight, and post-market monitoring.

Should the tool be classified as a medical device, selecting an appropriate Notified Body becomes critical. Beyond costs and timelines, this decision should also consider the Body's experience with software-based medical devices, regional interpretations of EU regulations, and language accessibility, especially in multinational consortia.

Finally, it is essential that the project team continues to monitor evolving EU regulations, including the MDR, AI Act, GDPR, and EHDS, to ensure the solution remains compliant and ready for certification, deployment, and potential commercialisation.

6. Which business strategies will maximise the real-world impact of the AI-based medical technology?

The CINDERELLA consortium has not yet defined a comprehensive exploitation strategy, as it will depend significantly on the outcomes of its ongoing clinical trials. Nonetheless, preliminary planning has been carried out with the aim of balancing the interests of all major stakeholders and maximising the project's potential real-world impact. Following the VAMI Framework, several business models could be considered.

One potential direction, towards which the project is currently progressing, is the commercialisation of the *CINDERELLA APProach* through a licensed Software-as-a-Service (SaaS) model targeting healthcare facilities that offer breast cancer locoregional treatments. This would require a well-defined reimbursement strategy, and one promising route could involve data-driven reimbursement models, where real-world data is used to demonstrate the clinical benefit, operational efficiency, or patient-reported outcomes. To support this, local validation efforts and ongoing performance monitoring should be prioritised to ensure fairness, generalisability, and long-term value.

In this context, a tech-based commercial entity such as CANKADO plays a central role as the project's designated exploitation partner. The initial focus is expected to be on early adopters, including CINDERELLA's clinical trial partners and healthcare institutions already integrated with CANKADO's certified platform. CANKADO's extensive experience in digital therapeutics and international healthcare markets, along with the technical maturity of its

solution, strengthens the project's alignment with real-world healthcare systems and enhances the feasibility of future commercial deployment.

Additionally, non-commercial pathways could also be considered. As discussed in Question 4, the consortium could adopt an open-source exploitation strategy, complemented by a robust communication and dissemination plan. This would involve increasing visibility and accessibility of already published models and code, as well as releasing currently undisclosed components where appropriate. Such an approach would enhance the project's scientific reach and facilitate broader uptake.

Finally, another impactful strategy could include curating and sharing the anonymised image database of pre- and post-operative patient photos, under strict data privacy protections. By granting access to this unique dataset, the project could significantly contribute to future research and the development of more robust and fair AI models for healthcare.

7. Can the AI technology integrate feasible clinical workflows and add value to healthcare professionals?

The *CINDERELLA APProach* is designed to integrate smoothly into existing clinical workflows without requiring major infrastructural or procedural changes. As a decision-support tool, it aims to complement rather than replace clinicians, thereby minimising resistance and supporting adoption by enhancing and not disrupting daily practice. This choice is aligned with the VAMI Framework's recommendations for promoting adoption, as well as maintaining clinicians' authority and responsibility over the outcome.

Furthermore, the proposed framework recommends that solution deployment should include targeted training for both clinicians and patients, ensuring proper use and acceptance. Additionally, as previously covered, continuous deployment practices, including performance monitoring and model updates, are also recommended.

8. How can AI medical technologies be developed and deployed fairly and ethically?

The CINDERELLA consortium has prioritised fairness and transparency from the outset. Ethical development was guided by established frameworks such as FRIA and ALTAI, which helped ensure responsible AI design and implementation practices, and are recommended in the proposed framework.

Looking ahead, ethical deployment will require maintaining performance across diverse populations and clinical settings. To address this, in line with the VAMI Framework and as previously covered, the implementation of local validation and continuous performance monitoring to detect and mitigate issues like model drift or underperformance is strongly suggested.

9. What other best practices should be considered for the sustainable use of AI in MedTech?

To ensure the long-term sustainability of AI solutions like the *CINDERELLA APProach*, the VAMI Framework suggests the launch of targeted pilot studies with clearly defined Key

Performance Indicators (KPIs). Relevant KPIs might include reductions in time spent on repetitive tasks, increased clinician–patient interaction time, improvements in workflow efficiency, or direct and indirect cost savings. Tracking such indicators can help assess the real-world impact of the solution and quantify its added value.

5.4 Case Study Conclusions and Key Insights

So far, the CINDERELLA Project has proactively aligned with several best practices outlined in the VAMI Framework. These include assembling a multidisciplinary team that combines AI expertise (INESC TEC), business and implementation know-how (CANKADO), and clinical partnerships with access to real-world data and patient populations, enabling the development of a diverse and robust dataset.

Although a preliminary exploitation strategy for the *CINDERELLA APProach* has been outlined, a comprehensive strategy is still missing. This is a crucial aspect that should ideally be addressed early in the development process. Without a clear roadmap, decisions such as publishing source code or sharing key components openly may unintentionally weaken the project's competitive positioning and reduce its ability to attract follow-on investment or sustain go-to-market efforts. The importance of forward planning is especially pronounced in the context of AI-based medical tools, which face not only technical and integration challenges but also complex adoption dynamics, as discussed in Sections 4.2 and 4.5.

Another gap emerged on the regulatory front. As the analysis suggests, the consortium’s activities appear not to have fully prioritised certain regulatory steps outlined in the VAMI Framework, such as determining the appropriate medical device classification or engaging early with a Notified Body. These elements, while complex, carry significant financial and procedural implications. Addressing them earlier in the development process could help mitigate uncertainty and reduce the risk of delays or unexpected costs as the solution approaches market readiness. The VAMI Framework highlights the value of proactive regulatory alignment, encouraging teams to view compliance not as a final hurdle but as an integral part of strategic planning and design.

To mitigate the risk of stagnation and maximise real-world impact, further alignment with VAMI recommendations is encouraged. In particular, advancing local validation and continuous model monitoring will be critical to maintaining performance, ensuring fairness, and avoiding unintended harm to patients.

6. Sustainable Development Goals

This thesis aims to address pressing societal challenges and contribute to a more equitable, responsible, and sustainable future. In this context, this chapter analyses the potential impact of this research in relation to the United Nations Sustainable Development Goals (SDGs) [110]. Table 10 presents the SDGs most relevant to this work, identifying the specific targets addressed, outlining the thesis's contributions, and proposing performance indicators and metrics to assess its alignment with each goal.

SDG	Targets	Contribution	Performance Indicators and Metrics
3: Good Health and Well-being	3.8	The proposed framework supports the development and adoption of AI-based medical technologies. In doing so, it can enhance access to diagnostic and therapeutic tools and improve the overall availability of quality essential health-care services.	○ Number of AI-based MedTech solutions deployed ○ Patient outcome metrics (e.g., diagnosis time, accuracy)
	3.d	This study identifies key barriers in the development of AI-enabled medical tools. By addressing these, it promotes the integration of predictive AI systems into health infrastructure, enabling earlier detection and more proactive healthcare responses.	○ Number of predictive AI systems ○ Reduction in response time to emerging health threats
4: Quality Education	4.4	This study explores the current landscape, challenges, and best practices in AI-based MedTech innovation. It can serve as a learning tool for bioengineers, MedTech developers, and students seeking to develop skills in value-driven, responsible AI for healthcare.	○ Adoption or reference of the study or framework in academic projects, course materials, or professional training programs.
9: Industry, Innovation and Infrastructure	9.5	This study provides a foundation for continued scientific research in AI-driven healthcare innovation. The proposed framework aims to guide the sustainable and ethically aligned development of MedTech solutions, helping advance the technological capacity of the healthcare sector.	○ Number of research initiatives addressing the AI-based MedTech innovation process ○ Healthcare sector adoption rate of new AI-based medical technologies

10: Reduced Inequalities	10.3	The study draws attention to the risk of AI systems amplifying existing healthcare biases. The proposed framework promotes local model validation, continuous performance monitoring, and transparency as best practices to ensure fairness and equity across institutions and populations.	○ Detection and correction of biased performance in AI systems across populations ○ Implementation of local validation protocols
12: Responsible Consumption and Production	12.5	The proposed framework encourages AI-driven efficiency in healthcare, helping to avoid unnecessary diagnostic procedures or treatments.	○ Reduction in redundant or low-benefit medical interventions with unbalanced return on investment
	12.6	The study presents concrete examples of biased or inequitable AI-based MedTech systems. It advocates for ethical development guidelines that help MedTech companies adopt fairer, more inclusive, and sustainable innovation practices.	○ Inclusion of equity and sustainability guidelines in AI MedTech innovation
16: Peace Justice, and Strong Institutions	16.7	This study explores how AI-based systems influence clinical decision-making. It emphasizes the importance of designing these technologies as support tools that assist, rather than replace clinicians, preserving their authority and professional judgment.	○ Number of AI-based clinical support systems developed or adopted with a human-in-the-loop design approach
	16.10	The study highlights the critical role of health data in AI development while reinforcing the importance of privacy, informed consent, and adherence to data protection regulations (e.g., GDPR). It supports a balance between innovation and rights protection.	○ Inclusion of ethical data-use guidance in MedTech development pipelines

Table 10 Thesis contribution to selected UN Sustainable Development Goals (SDGs). The table outlines relevant SDG targets, the corresponding contributions of this work, and suggested indicators for measuring impact.

7. Limitations and Future Improvements

This dissertation offers an exploratory, qualitative work into the challenges and enablers of AI-driven medical technology innovation. However, several limitations should be acknowledged to contextualise its findings and identify opportunities for future research.

First, the research draws from a relatively small sample of interviewees and was conducted within a limited timeframe, which may affect the generalisability of the findings. While the expert-informed approach provides valuable insights, future studies could benefit from a broader and more diverse participant pool, incorporating a wider range of stakeholder perspectives, including those of private payers, regulators, and patients, to enrich the understanding of sector dynamics, challenges, and best practices. Additionally, extending the geographical scope could provide a stronger foundation for actionable, context-specific innovation recommendations.

Second, while the thematic coding analysis allowed for the extraction of key insights across interviews, the coding and interpretation process is inherently subjective and may still be subject to bias. Although reflexivity and transparency were prioritised throughout this process, future research could incorporate dedicated qualitative analysis tools, such as MaxQDA, to enhance rigour, traceability, and support the identification of underexplored themes, informing the selection of additional interview profiles for supplementary data collection.

Third, even though the proposed VAMI Framework's current form has been validated through expert feedback, further refinement is needed to increase its applicability and robustness. Notably, the validation interviews, while valuable, posed certain methodological challenges. The broad scope and ambition of the framework's questions made it difficult to explore all components with equal depth during time-constrained interviews. Moreover, this validation approach heavily depended on each expert's current interests and focus. Future work should consider allocating more time or adopting more targeted approaches for framework validation, as well as include influential actors that were not strongly represented in the interview sample, such as big tech companies operating in the health space or highly data-centric MedTech firms.

Another limitation concerns the application of the VAMI Framework to the CINDERELLA case study, which occurred while the project was already underway rather than at its inception. As such, its use primarily served to reflect on ongoing activities and identify areas where earlier alignment might have added value. While this limited the opportunity to test the framework's role in early-stage risk identification, it provided valuable insights into its applicability in mid- to late-stage projects. Notably, some of the best practices embedded in the framework were informed by lessons learned during the CINDERELLA project itself. While this alignment enriched the practical grounding of the framework, it also raises concerns about confirmation bias and limited generalisability. In addition, only a limited number of CINDERELLA project partners were involved, although those selected had in-depth knowledge of the project's core technological and business elements.

To address these limitations, a prospective, real-time pilot study applying the VAMI Framework from the beginning of an AI-based medical technology innovation project is recommended. This approach would enable assessment and refinement of the framework based on its practical ability to support decision-making, shape innovation strategies, and mitigate risks.

Finally, it is important to recognise that many of the topics explored in this thesis are still rapidly evolving. Some reports, studies, and policy documents referenced in this study were published only shortly before its completion, particularly in the domains of AI governance and healthcare regulation. This underscores the importance of ongoing research to ensure that frameworks and strategies remain aligned with evolving legal, technological, and clinical landscapes.

8. Conclusions

The aim of this dissertation was to explore the challenges involved in developing and deploying AI-based medical technologies and to propose a framework of best practices that can guide their development, adoption, and ultimately increase real-world impact.

To address this goal, the research was guided by two central questions: *"What main challenges impair the adoption of AI-based medical technologies?"* and *"What best practices can support the development of a guiding framework?"* A mixed-methods approach was adopted, combining a literature review with qualitative insights gathered from expert interviews. In the first round, eleven unstructured interviews were conducted with professionals from a variety of fields, including data science, product management, and regulatory affairs, all with experience in AI healthcare innovation across both European and American contexts. The objective was to understand the key barriers, causes of failure, as well as enabling factors in the development and adoption of AI-based medical technologies. In the second round, semi-structured interviews were conducted with three experts with extensive experience in AI-based medical technology innovation. These interviews served to validate the proposed framework, highlight areas of focus, and suggest improvements for future iterations.

The study identified a wide range of challenges, including limited access to high-quality data, difficulties in clinical integration, highly strict model performance standards, and broader issues around social acceptance and trust. Among these, two major barriers emerged: the lack of validated clinical needs and restricted access to reliable, diverse data. Several best practices were also highlighted through both literature and interviews, including the importance of need-driven innovation, early stakeholder engagement, and ongoing model validation protocols. These findings confirmed that successful AI MedTech innovation requires a comprehensive understanding of technical, clinical, regulatory, and societal dimensions.

To address these findings, this dissertation proposed the VAMI Framework, a set of nine guiding questions designed to help innovators navigate these multidimensional challenges. Each question is paired with best practices derived from both academic and practical sources. To support its practical use, the framework was mapped onto the three phases of the Stanford Biodesign process - Identify, Invent, and Implement - providing a familiar reference for MedTech innovators. Specifically, Questions 1-2 align with the Identify phase, focusing on data and problem definition; Questions 3-4 align with the Invent phase, covering solution design and intellectual property protection; and Questions 5-9 align with the Implement phase, addressing regulatory, business, and sustainability issues. It is important to note that both frameworks are inherently iterative and not strictly linear.

The validation interviews provided additional perspective, reinforcing the relevance of the proposed framework and confirming the importance of many of its guiding questions, highlighted challenges, and suggested best practices. For example, they underscored the significance of local validation, continuous model monitoring and building multidisciplinary teams. At the same time, the interviews revealed areas for improvement. One key concern was that the framework could unintentionally imply a step-by-step approach, when in fact several

best practices, such as regulatory preparation, must be considered from the outset rather than postponed as later-stage checkpoints. Experts also suggested that future iterations of the framework should include a dedicated step for the systematic collection and validation of stakeholder and clinical requirements, which are likely to evolve throughout development. Additionally, explicit integration of existing medical device standards could further strengthen the framework by offering a structured way to address compliance and risk management from early stages.

The relevance of the proposed framework was further illustrated through the CINDERELLA case study. This real-world project successfully implemented many of the VAMI Framework's best practices, including building a multidisciplinary team and using explainable AI to facilitate clinical adoption. However, it also highlighted the risks of insufficient strategic planning. For instance, the early release of open-source components may have impacted future commercialisation options. While this may not completely prevent commercial success, it introduces limitations that could have been mitigated through early planning. Interestingly, intellectual property protection was not directly raised during the validation interviews. The case study also highlighted the importance of early regulatory planning. The project did not fully consider key future steps, such as medical device classification, resulting in potentially costly and time-consuming implications. These oversights underscore how frameworks like VAMI can have a critical role in supporting early, proactive planning and preparation to reduce project risks and enhance the likelihood of real-world impact.

One of the most critical contributions of this study is the emphasis on local validation and continuous deployment, including ongoing performance monitoring and model updating. Until regulatory frameworks fully evolve to answer the urgent concerns around generalisability, model bias, and performance degradation caused by population or technology drift, these practices must remain key safeguards to ensure ethical, fair, and safe implementation of AI in clinical settings.

By proposing the VAMI Framework, this dissertation aims to better equip innovators to navigate the complex process of developing AI-based medical solutions. The ultimate objective is to support the effective, safe, and equitable translation of these technologies into real-world clinical practice, delivering value not only to healthcare providers but most importantly, to patients.

9. References

[1] S. Flessa and C. Huebner, ‘Innovations in Health Care—A Conceptual Framework’, *Int. J. Environ. Res. Public. Health*, vol. 18, no. 19, p. 10026, Sep. 2021, doi: 10.3390/ijerph181910026.

[2] J. Bajwa, U. Munir, A. Nori, and B. Williams, ‘Artificial intelligence in healthcare: transforming the practice of medicine’, *Future Healthc. J.*, vol. 8, no. 2, pp. e188–e194, Jul. 2021, doi: 10.7861/fhj.2021-0095.

[3] M. Mitchell and L. Kan, ‘Digital Technology and the Future of Health Systems’, *Health Syst. Reform*, vol. 5, no. 2, pp. 113–120, 2019, doi: 10.1080/23288604.2019.1583040.

[4] M. Sharma, C. Savage, M. Nair, I. Larsson, P. Svedberg, and J. M. Nygren, ‘Artificial Intelligence Applications in Health Care Practice: Scoping Review’, *J. Med. Internet Res.*, vol. 24, no. 10, p. e40238, Oct. 2022, doi: 10.2196/40238.

[5] M. Conley, ‘How Stanford Medicine is capturing the AI moment’, Stanford Medicine Magazine. Accessed: May 23, 2025. [Online]. Available: https://stanmed.stanford.edu/translating-ai-concepts-into-innovations/

[6] J. Y. X. Chua, E. M. Kan, P. P. Lee, and S. Shorey, ‘Application of the Stanford Biodesign Framework in Healthcare Innovation Training and Commercialization of Market Appropriate Products: A Scoping Review’, *J. Med. Syst.*, vol. 48, no. 1, p. 44, Apr. 2024, doi: 10.1007/s10916-024-02067-x.

[7] P. G. Yock, *Biodesign*, 2nd ed. West Nyack: Cambridge University Press, 2015.

[8] S. Singh and Y. and Aggarwal, ‘In search of a consensus definition of innovation: a qualitative synthesis of 208 definitions using grounded theory approach’, *Innov. Eur. J. Soc. Sci. Res.*, vol. 35, no. 2, pp. 177–195, Apr. 2022, doi: 10.1080/13511610.2021.1925526.

[9] I. of M. (US) C. on T. I. in Medicine and A. C. Gelijns, ‘Medical Technology Development: An Introduction to the Innovation-Evaluation Nexus’, in *Modern Methods of Clinical Investigation: Medical Innovation at the Crossroads: Volume I*, National Academies Press (US), 1990. Accessed: May 23, 2025. [Online]. Available: https://www.ncbi.nlm.nih.gov/books/NBK235486/

[10] M. W. van den Hoed, R. Backhaus, E. de Vries, J. P. H. Hamers, and R. Daniëls, ‘Factors contributing to innovation readiness in health care organizations: a scoping review’, *BMC Health Serv. Res.*, vol. 22, p. 997, Aug. 2022, doi: 10.1186/s12913-022-08185-x.

[11] World Health Organization, *How does population ageing affect health system financial sustainability and affordable access to health care in Europe?* 2025.

[12] ‘Health workforce’. Accessed: May 23, 2025. [Online]. Available: https://www.who.int/health-topics/health-workforce

[13] MedTech Europe, ‘MedTech Europe’s Facts & Figures 2024’, MedTech Europe. [Online]. Available: https://www.medtecheurope.org/resource-library/medtech-europes-facts-figures-2024/

[14] D. R. Pai, ‘Complexities of Simultaneously Improving Quality and Lowering Costs in Hospitals’, *Int. J. Health Policy Manag.*, vol. 12, p. 7442, Oct. 2022, doi: 10.34172/ijhpm.2022.7442.

[15] D. J. van Staalduinen, P. van den Bekerom, S. Groeneveld, M. Kidanemariam, A. M. Stiggelbout, and M. E. van den Akker-van Marle, ‘The implementation of value-based healthcare: a scoping review’, *BMC Health Serv. Res.*, vol. 22, p. 270, Mar. 2022, doi: 10.1186/s12913-022-07489-2.

[16] E. Teisberg, S. Wallace, and S. O'Hara, 'Defining and Implementing Value-Based Health Care: A Strategic Framework', *Acad. Med.*, vol. 95, no. 5, pp. 682–685, May 2020, doi: 10.1097/ACM.0000000000003122.
[17] R. A. Rosalia, K. Wahba, and N. Milevska-Kostova, 'How digital transformation can help achieve value-based healthcare: Balkans as a case in point', *Lancet Reg. Health - Eur.*, vol. 4, p. 100100, Apr. 2021, doi: 10.1016/j.lanepe.2021.100100.
[18] M. L. Leite *et al.*, 'Artificial intelligence and the future of life sciences', *Drug Discov. Today*, vol. 26, no. 11, pp. 2515–2526, Nov. 2021, doi: 10.1016/j.drudis.2021.07.002.
[19] K. Akyüz, M. Cano Abadía, M. Goisauf, and M. T. Mayrhofer, 'Unlocking the potential of big data and AI in medicine: insights from biobanking', *Front. Med.*, vol. 11, p. 1336588, 2024, doi: 10.3389/fmed.2024.1336588.
[20] 'What Is Deep Learning? | IBM'. Accessed: May 26, 2025. [Online]. Available: https://www.ibm.com/think/topics/deep-learning
[21] E. Biundo, A. Pease, M. de Groote, T. d'Argent, and E. de Schaetzen, 'The Socio-Economic Impact of AI in Healthcare', Delloite, Oct. 2020. [Online]. Available: https://www.medtecheurope.org/resource-library/the-socio-economic-impact-of-ai-in-healthcare-addressing-barriers-to-adoption-for-new-healthcare-technologies-in-europe/
[22] 'MRNA Artificial Intelligence in Vaccine Development'. Accessed: May 26, 2025. [Online]. Available: https://www.pfizer.com/news/articles/how_a_novel_incubation_sandbox_helped_speed_up_data_analysis_in_pfizer_s_covid_19_vaccine_trial
[23] R. Glaubitz *et al.*, 'The cost of the diagnostic odyssey of patients with suspected rare diseases', *Orphanet J. Rare Dis.*, vol. 20, p. 222, May 2025, doi: 10.1186/s13023-025-03751-y.
[24] K. Powell, 'World Record-Setting DNA Sequencing Technique Helps Clinicians Rapidly Diagnose Critical Care Patients', NVIDIA Blog. Accessed: May 23, 2025. [Online]. Available: https://blogs.nvidia.com/blog/world-record-genome-sequencing-parabricks/
[25] J. E. Gorzynski *et al.*, 'Ultrarapid Nanopore Genome Sequencing in a Critical Care Setting', *N. Engl. J. Med.*, vol. 386, no. 7, pp. 700–702, Feb. 2022, doi: 10.1056/NEJMc2112090.
[26] 'Colorectal cancer'. Accessed: May 26, 2025. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/colorectal-cancer
[27] H. Brenner, M. Kloor, and C. P. Pox, 'Colorectal cancer', *Lancet Lond. Engl.*, vol. 383, no. 9927, pp. 1490–1502, Apr. 2014, doi: 10.1016/S0140-6736(13)61649-9.
[28] D. A. Corley *et al.*, 'Adenoma detection rate and risk of colorectal cancer and death', *N. Engl. J. Med.*, vol. 370, no. 14, pp. 1298–1306, Apr. 2014, doi: 10.1056/NEJMoa1309086.
[29] 'GI Genius™ Intelligent Endoscopy Module'. Accessed: May 23, 2025. [Online]. Available: https://www.medtronic.com/en-us/healthcare-professionals/products/digestive-gastrointestinal/gastrointestinal-artificial-intelligence/gi-genius-intelligent-endoscopy-module.html
[30] 'Artificial Intelligence in healthcare - European Commission'. Accessed: Jun. 04, 2025. [Online]. Available: https://health.ec.europa.eu/ehealth-digital-health-and-care/artificial-intelligence-healthcare_en
[31] S. N. Saw and K. H. Ng, 'Current challenges of implementing artificial intelligence in medical imaging', *Phys. Medica Eur. J. Med. Phys.*, vol. 100, pp. 12–17, Aug. 2022, doi: 10.1016/j.ejmp.2022.06.003.

[32] B. khan *et al.*, 'Drawbacks of Artificial Intelligence and Their Potential Solutions in the Healthcare Sector', *Biomed. Mater. Devices N. Y. N*, pp. 1–8, Feb. 2023, doi: 10.1007/s44174-023-00063-2.
[33] L. Jiang *et al.*, 'Opportunities and challenges of artificial intelligence in the medical field: current application, emerging problems, and problem-solving strategies', *J. Int. Med. Res.*, vol. 49, no. 3, p. 03000605211000157, Mar. 2021, doi: 10.1177/03000605211000157.
[34] 'Google DeepMind trial failed to comply with data protection law', Herbert Smith Freehills Kramer | Global law firm. Accessed: Jun. 05, 2025. [Online]. Available: https://www.hsfkramer.com/notes/data/2017-07/google-deepmind-trial-failed-to-comply-with-data-protection-law
[35] C. Véliz, 'Medical privacy and big data: A further reason in favour of public universal healthcare coverage', in *Philosophical Foundations of Medical Law*, T. C. de Campos, J. Herring, and A. M. Phillips, Eds., in Wellcome Trust–Funded Monographs and Book Chapters. , Oxford (UK): Oxford University Press, 2019. Accessed: Jun. 05, 2025. [Online]. Available: http://www.ncbi.nlm.nih.gov/books/NBK550264/
[36] C. S. Kruse, B. Frederick, T. Jacobson, and D. K. Monticone, 'Cybersecurity in healthcare: A systematic review of modern threats and trends', *Technol. Health Care Off. J. Eur. Soc. Eng. Med.*, vol. 25, no. 1, pp. 1–10, 2017, doi: 10.3233/THC-161263.
[37] European Data Protection Board, 'Guidelines 01/2025 on Pseudonymisation', European Data Protection Board, Jan. 2025. [Online]. Available: https://www.edpb.europa.eu/our-work-tools/documents/public-consultations/2025/guidelines-012025-pseudonymisation_en
[38] R. Chevrier, V. Foufi, C. Gaudet-Blavignac, A. Robert, and C. Lovis, 'Use and Understanding of Anonymization and De-Identification in the Biomedical Literature: Scoping Review', *J. Med. Internet Res.*, vol. 21, no. 5, p. e13484, May 2019, doi: 10.2196/13484.
[39] N. Rieke *et al.*, 'The future of digital health with federated learning', *NPJ Digit. Med.*, vol. 3, p. 119, Sep. 2020, doi: 10.1038/s41746-020-00323-1.
[40] M. J. Sheller, G. A. Reina, B. Edwards, J. Martin, and S. Bakas, 'Multi-Institutional Deep Learning Modeling Without Sharing Patient Data: A Feasibility Study on Brain Tumor Segmentation', *Brainlesion Glioma Mult. Scler. Stroke Trauma. Brain Inj. BrainLes Workshop*, vol. 11383, pp. 92–104, 2019, doi: 10.1007/978-3-030-11723-8_9.
[41] '(PDF) Synthetic Data -- what, why and how?', ResearchGate. Accessed: Jun. 05, 2025. [Online]. Available: https://www.researchgate.net/publication/360462365_Synthetic_Data_--_what_why_and_how
[42] M. Giuffrè and D. L. Shung, 'Harnessing the power of synthetic data in healthcare: innovation, application, and privacy', *NPJ Digit. Med.*, vol. 6, p. 186, Oct. 2023, doi: 10.1038/s41746-023-00927-3.
[43] C. Cheng *et al.*, 'A General Primer for Data Harmonization', *Sci. Data*, vol. 11, p. 152, Jan. 2024, doi: 10.1038/s41597-024-02956-3.
[44] A. Zeb, J.-P. Soininen, and N. Sozer, 'Data harmonisation as a key to enable digitalisation of the food sector: A review', *Food Bioprod. Process.*, vol. 127, pp. 360–370, May 2021, doi: 10.1016/j.fbp.2021.02.005.
[45] 'Data Standardization – OHDSI'. Accessed: Jun. 05, 2025. [Online]. Available: https://www.ohdsi.org/data-standardization/
[46] 'What is SNOMED CT', SNOMED International. Accessed: Jun. 05, 2025. [Online]. Available: https://www.snomed.org/what-is-snomed-ct
[47] 'About DICOM- Overview', DICOM. Accessed: Jun. 05, 2025. [Online]. Available: https://www.dicomstandard.org/about

[48] Office of the National Coordinator for Health Information Technology, 'What Is FHIR®?', Office of the National Coordinator for Health Information Technology. [Online]. Available: https://www.healthit.gov/sites/default/files/2019-08/ONCFHIRFSWhatIsFHIR.pdf
[49] J. G. You, T. Hernandez-Boussard, M. A. Pfeffer, A. Landman, and R. G. Mishuris, 'Clinical trials informed framework for real world clinical implementation and deployment of artificial intelligence applications', *Npj Digit. Med.*, vol. 8, no. 1, pp. 1–5, Feb. 2025, doi: 10.1038/s41746-025-01506-4.
[50] H. Alami *et al.*, 'Organizational readiness for artificial intelligence in health care: insights for decision-making and practice', *J. Health Organ. Manag.*, vol. ahead-of-print, no. ahead-of-print, Dec. 2020, doi: 10.1108/JHOM-03-2020-0074.
[51] T. Davenport and R. Kalakota, 'The potential for artificial intelligence in healthcare', *Future Healthc. J.*, vol. 6, no. 2, pp. 94–98, Jun. 2019, doi: 10.7861/futurehosp.6-2-94.
[52] R. Abdullah and B. Fakieh, 'Health Care Employees' Perceptions of the Use of Artificial Intelligence Applications: Survey Study', *J. Med. Internet Res.*, vol. 22, no. 5, p. e17620, May 2020, doi: 10.2196/17620.
[53] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, 'Explainable AI: A Review of Machine Learning Interpretability Methods', *Entropy*, vol. 23, no. 1, p. 18, Dec. 2020, doi: 10.3390/e23010018.
[54] 'Artificial Intelligence in Healthcare: Effective Regulation Can Boost Use of Explainable Innovation in Healthcare', MedTech Europe. Accessed: Jun. 04, 2025. [Online]. Available: https://www.medtecheurope.org/medtech-views/policy-views/artificial-intelligence-in-healthcare-effective-regulation-can-boost-use-of-explainable-innovation-in-healthcare/
[55] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, 'Dissecting racial bias in an algorithm used to manage the health of populations', *Science*, vol. 366, no. 6464, pp. 447–453, Oct. 2019, doi: 10.1126/science.aax2342.
[56] 'Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment | Shaping Europe's digital future'. Accessed: Jun. 05, 2025. [Online]. Available: https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment
[57] D. Atanasovska, 'AI Act: Fundamental Rights Impact Assessments (FRIA) - Who, When, Why, and How to Ensure Ethical AI Deployment', GDPR Local. Accessed: Jun. 02, 2025. [Online]. Available: https://gdprlocal.com/ai-act-fundamental-rights-impact-assessments-fria-who-when-why-and-how-to-ensure-ethical-ai-deployment/
[58] T. Basu, S. Engel-Wolf, and O. Menzer, 'The Ethics of Machine Learning in Medical Sciences: Where Do We Stand Today?', *Indian J. Dermatol.*, vol. 65, no. 5, pp. 358–364, 2020, doi: 10.4103/ijd.IJD_419_20.
[59] 'New Report from Corinium and FICO Finds that Lack of Urgency Around Responsible AI Use is Putting Most Companies at Risk | FICO'. Accessed: Jun. 05, 2025. [Online]. Available: https://www.fico.com/en/newsroom/new-report-corinium-and-fico-finds-lack-urgency-around-responsible-ai-use-putting-most
[60] K. Palaniappan, E. Y. T. Lin, and S. Vogel, 'Global Regulatory Frameworks for the Use of Artificial Intelligence (AI) in the Healthcare Services Sector', *Healthcare*, vol. 12, no. 5, p. 562, Feb. 2024, doi: 10.3390/healthcare12050562.
[61] 'AI Act | Shaping Europe's digital future'. Accessed: Jun. 04, 2025. [Online]. Available: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
[62] 'The landmark AI Act, its ripple effects on the medical technology industry and the challenges yet to be tackled', MedTech Europe. Accessed: Jun. 04, 2025. [Online]. Available: https://www.medtecheurope.org/medtech-views/policy-views/the-landmark-ai-

act-its-ripple-effects-on-the-medical-technology-industry-and-the-challenges-yet-to-be-tackled/
[63] MedTech Europe, 'Medical Technology Industry Perspective on the Final AI Act', MedTech Europe, Mar. 2024. [Online]. Available: https://www.medtecheurope.org/resource-library/medical-technology-industry-perspective-on-the-final-ai-act/
[64] C. for D. and R. Health, 'Artificial Intelligence and Machine Learning in Software as a Medical Device', *FDA*, Mar. 2025, Accessed: Jun. 10, 2025. [Online]. Available: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device
[65] 'Artificial intelligence', OECD. Accessed: Jun. 10, 2025. [Online]. Available: https://www.oecd.org/en/topics/artificial-intelligence.html
[66] 'What is GDPR, the EU's new data protection law?', GDPR.eu. Accessed: Jun. 10, 2025. [Online]. Available: https://gdpr.eu/what-is-gdpr/
[67] 'European Health Data Space Regulation (EHDS) - European Commission'. Accessed: Jun. 04, 2025. [Online]. Available: https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en
[68] D. A. Augustin *et al.*, 'Stanford's Biodesign Innovation program: Teaching opportunities for value-driven innovation in surgery', *Surgery*, vol. 167, no. 3, pp. 535–539, Mar. 2020, doi: 10.1016/j.surg.2019.10.012.
[69] 'Global Initiatives | Stanford Mussallem Center for Biodesign | Stanford Medicine'. Accessed: May 23, 2025. [Online]. Available: https://biodesign.stanford.edu/programs/global-initiatives.html
[70] 'iRhythm | Stanford Mussallem Center for Biodesign | Stanford Medicine'. Accessed: Jun. 09, 2025. [Online]. Available: https://biodesign.stanford.edu/our-impact/technologies/irhythm.html
[71] 'Heart Monitor Patch & Device for Arrhythmia Detection | Zio® ECG Monitors by iRhythm'. Accessed: Jun. 09, 2025. [Online]. Available: https://www.irhythmtech.com/us/en
[72] 'Process', Stanford Mussallem Center for Biodesign. Accessed: May 23, 2025. [Online]. Available: https://biodesign.stanford.edu/about-us/process.html
[73] J. G. Schwartz, U. N. Kumar, D. E. Azagury, T. J. Brinton, and P. G. Yock, 'Needs-Based Innovation in Cardiovascular Medicine: The Stanford Biodesign Process', *JACC Basic Transl. Sci.*, vol. 1, no. 6, pp. 541–547, Oct. 2016, doi: 10.1016/j.jacbts.2016.06.011.
[74] 'Biodesign for Digital Health', Stanford Mussallem Center for Biodesign. Accessed: Jun. 09, 2025. [Online]. Available: https://biodesign.stanford.edu/content/sm/biodesign/programs/stanford-courses/biodesign-for-digital-health.html?utm_source=chatgpt.com
[75] 'Building for Digital Health', Stanford Mussallem Center for Biodesign. Accessed: Jun. 09, 2025. [Online]. Available: https://biodesign.stanford.edu/programs/stanford-courses/building-for-digital-health.html
[76] 'Spezi'. Accessed: Jun. 09, 2025. [Online]. Available: https://spezi.stanford.edu/
[77] K. Lekadir *et al.*, 'FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare', Feb. 2025, doi: 10.1136/bmj-2024-081554.
[78] S. Dykstra *et al.*, 'An institutional framework to support ethical fair and equitable artificial intelligence augmented care', *Npj Digit. Med.*, vol. 8, no. 1, pp. 1–11, Feb. 2025, doi: 10.1038/s41746-025-01490-9.
[79] A. A. Abujaber and A. J. Nashwan, 'Ethical framework for artificial intelligence in healthcare research: A path to integrity', *World J. Methodol.*, vol. 14, no. 3, p. 94071, Sep. 2024, doi: 10.5662/wjm.v14.i3.94071.

[80] A. Callahan *et al.*, ‘Standing on FURM ground -- A framework for evaluating Fair, Useful, and Reliable AI Models in healthcare systems’, 2024, *arXiv*. doi: 10.48550/ARXIV.2403.07911.
[81] Michigan Health & Hospital Association Artificial Intelligence Task Force, ‘Artificial Intelligence Framework for Healthcare’, Michigan Health & Hospital Association, Feb. 2024. [Online]. Available: https://www.mha.org/wp-content/uploads/2024/12/AI-Framework-for-Healthcare_MHA-AI-Taskforce.pdf
[82] K. B. Shiferaw, M. Roloff, I. Balaur, D. Welter, D. Waltemath, and A. A. Zeleke, ‘Guidelines and standard frameworks for artificial intelligence in medicine: a systematic review’, *JAMIA Open*, vol. 8, no. 1, p. ooae155, Jan. 2025, doi: 10.1093/jamiaopen/ooae155.
[83] ‘Building Theories from Case Study Research on JSTOR’. Accessed: May 23, 2025. [Online]. Available: https://www.jstor.org/stable/258557?origin=crossref&seq=1
[84] R. K. Yin, *Qualitative research from start to finish*, Second edition. in Research methods. New York London: The Guilford Press, 2016.
[85] K. M. Eisenhardt, ‘Building Theories from Case Study Research’, *Acad. Manage. Rev.*, vol. 14, no. 4, pp. 532–550, 1989, doi: 10.2307/258557.
[86] W. I. P. Organization, *Global Innovation Index 2024: Unlocking the Promise of Social Entrepreneurship*, 1st ed. Geneva: World Intellectual Property Organization, 2024.
[87] M. Maguire and B. Delahunt, ‘Doing a thematic analysis: A practical, step-by-step guide for learning and teaching scholars.’, *Irel. J. High. Educ.*, vol. 9, no. 3, Art. no. 3, Oct. 2017, doi: 10.62707/aishej.v9i3.335.
[88] ‘Sword Health queria ajudar INEM com IA mas deparou-se com missão “impossível”’, SIC Notícias. Accessed: May 22, 2025. [Online]. Available: https://sicnoticias.pt/especiais/inteligencia-artificial/2025-02-07-video-sword-health-queria-ajudar-inem-com-ia-mas-deparou-se-com-missao-impossivel-0e14a41a?utm_source=chatgpt.com
[89] M. I. Ahmed, B. Spooner, J. Isherwood, M. Lane, E. Orrock, and A. Dennison, ‘A Systematic Review of the Barriers to the Implementation of Artificial Intelligence in Healthcare’, *Cureus*, vol. 15, no. 10, p. e46454, doi: 10.7759/cureus.46454.
[90] A. Gadotti, L. Rocher, F. Houssiau, A.-M. Creţu, and Y.-A. de Montjoye, ‘Anonymization: The imperfect science of using data while preserving privacy’, *Sci. Adv.*, vol. 10, no. 29, p. eadn7053, doi: 10.1126/sciadv.adn7053.
[91] J. Futoma, M. Simons, T. Panch, F. Doshi-Velez, and L. A. Celi, ‘The myth of generalisability in clinical research and machine learning in health care’, *Lancet Digit. Health*, vol. 2, no. 9, pp. e489–e492, Sep. 2020, doi: 10.1016/S2589-7500(20)30186-2.
[92] D. Coldewey, ‘Google medical researchers humbled when AI screening tool falls short in real-life testing’, TechCrunch. Accessed: May 29, 2025. [Online]. Available: https://techcrunch.com/2020/04/27/google-medical-researchers-humbled-when-ai-screening-tool-falls-short-in-real-life-testing/
[93] ‘(PDF) A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy’, in *ResearchGate*, Dec. 2024. doi: 10.1145/3313831.3376718.
[94] A. Wong *et al.*, ‘External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients’, *JAMA Intern. Med.*, vol. 181, no. 8, pp. 1–6, Aug. 2021, doi: 10.1001/jamainternmed.2021.2626.
[95] R. González-Alday, E. García-Cuesta, C. A. Kulikowski, and V. Maojo, ‘A Scoping Review on the Progress, Applicability, and Future of Explainable Artificial Intelligence in Medicine’, *Appl. Sci.*, vol. 13, no. 19, Art. no. 19, Jan. 2023, doi: 10.3390/app131910778.

[96] C. Mennella, U. Maniscalco, G. De Pietro, and M. Esposito, ‘Ethical and regulatory challenges of AI technologies in healthcare: A narrative review’, *Heliyon*, vol. 10, no. 4, p. e26297, Feb. 2024, doi: 10.1016/j.heliyon.2024.e26297.
[97] S. Santra, P. Kukreja, K. Saxena, S. Gandhi, and O. V. Singh, ‘Navigating regulatory and policy challenges for AI enabled combination devices’, *Front. Med. Technol.*, vol. 6, p. 1473350, Nov. 2024, doi: 10.3389/fmedt.2024.1473350.
[98] J. Schmidt *et al.*, ‘Mapping the regulatory landscape for artificial intelligence in health within the European Union’, *NPJ Digit. Med.*, vol. 7, p. 229, Aug. 2024, doi: 10.1038/s41746-024-01221-6.
[99] V. Muralidharan *et al.*, ‘A scoping review of reporting gaps in FDA-approved AI medical devices’, *NPJ Digit. Med.*, vol. 7, p. 273, Oct. 2024, doi: 10.1038/s41746-024-01270-x.
[100] A. Poddar and S. R. Rao, ‘Evolving intellectual property landscape for AI-driven innovations in the biomedical sector: opportunities in stable IP regime for shared success’, *Front. Artif. Intell.*, vol. 7, p. 1372161, 2024, doi: 10.3389/frai.2024.1372161.
[101] S. Jabbour *et al.*, ‘Measuring the Impact of AI in the Diagnosis of Hospitalized Patients’, *JAMA*, vol. 330, no. 23, pp. 2275–2284, Dec. 2023, doi: 10.1001/jama.2023.22295.
[102] R. Al-Halawani, P. H. Charlton, M. Qassem, and P. A. Kyriacou, ‘A review of the effect of skin pigmentation on pulse oximeter accuracy’, *Physiol. Meas.*, vol. 44, no. 5, p. 05TR01, May 2023, doi: 10.1088/1361-6579/acd51a.
[103] C. Radclyffe, M. Ribeiro, and R. H. Wortham, ‘The assessment list for trustworthy artificial intelligence: A review and recommendations’, *Front. Artif. Intell.*, vol. 6, p. 1020592, Mar. 2023, doi: 10.3389/frai.2023.1020592.
[104] ‘Issues of AI and Academic Transparency - Open and Universal Science (OPUS) Project’. Accessed: Jun. 02, 2025. [Online]. Available: https://opusproject.eu/openscience-news/issues-of-ai-and-academic-transparency/
[105] S. H. Bradley *et al.*, ‘Reducing bias and improving transparency in medical research: a critical overview of the problems, progress and suggested next steps’, *J. R. Soc. Med.*, vol. 113, no. 11, pp. 433–443, Nov. 2020, doi: 10.1177/0141076820956799.
[106] P. Saidi, G. Dasarathy, and V. Berisha, ‘Unraveling overoptimism and publication bias in ML-driven science’, *Patterns*, vol. 6, no. 4, p. 101185, Feb. 2025, doi: 10.1016/j.patter.2025.101185.
[107] ‘Breast cancer’. Accessed: Jun. 12, 2025. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/breast-cancer
[108] K. A. Stump-Sutliff, M. V. Durning, and RN, ‘Breast Cancer Survival Rates’, WebMD. Accessed: Jun. 12, 2025. [Online]. Available: https://www.webmd.com/breast-cancer/breast-cancer-survival-rates
[109] A. D. McCrorie, H. Stobart, D. Dodwell, S. A. McIntosh, and S. Potter, ‘Mapping the current landscape of locoregional therapy de-escalation trials in early breast cancer: a systematic review’, *Npj Breast Cancer*, vol. 11, no. 1, pp. 1–11, Mar. 2025, doi: 10.1038/s41523-025-00744-9.
[110] ‘THE 17 GOALS | Sustainable Development’. Accessed: Jun. 16, 2025. [Online]. Available: https://sdgs.un.org/goals

Appendix I

The unstructured interviews with professionals working with AI across various healthcare domains are summarised in Table 11, while the semi-structured interviews with experts possessing extensive experience across all domains are presented in Table 12. Each table includes the contact person, position, institution, country, covered dimensions, and key lessons learned from each interview.

Contact Person **Position, Institution, Country**	**Dimension**	**Key Lessons Learned**
Diogo Martins Researcher on Molecular Modelling, United States	Data Access & AI Model Development	Use of AI in Structural Biology o AlphaFold is an AI-based tool which enables highly accurate protein structure predictions and aids crystallographers in solving previously unresolvable structures. Experimental validation of predictions is still required via lab tests. o This tool was widely adopted in the field, even among initially sceptical researchers, due to outstanding results in CASP (Critical Assessment of protein Structure Prediction). o It was easily integrated into lab workflows, as the models were already mature and easy to use. Role of Data in AI Performance o Protein Data Bank (PDB) is a central open-access repository for protein structures, used by academia and industry. o It provided a vast and robust dataset, which enabled the success of AlphaFold. o For example, AlphaFold 3 has enhanced capabilities, such as predicting not only protein structure but also small molecule interactions, which streamlines previously separate steps in drug design workflows. o However, these models are super reliant on the training set, meaning that the structure of molecules which are not in the training set is poorly predicted. Potential Direction o To overcome the limitations of data-driven approaches, new models which are physics-based should be developed and tested. Pharmaceutical Industry and Data Sharing o Pharmaceutical companies are adopting AI to automate workflows, for example, by training LLMs with programming capabilities. o Despite the potential benefits, companies withhold data (particularly in pharmacokinetics) due to privacy and IP concerns, as well as competitive

		advantage. This limits the comprehensiveness of datasets and reduces the potential of AI-based tools. o Considering that the PDB enabled significant breakthroughs like AlphaFold, similar databases could be built for other fields to accelerate innovation. This, however, brings ethical and strategic discussions on data-sharing vs. proprietary control. Personal Viewpoint o Diogo values manual programming and model building over automated code generation by AI, which highlights the human barrier to AI development and adoption.
Pedro Ferreira Professor and Researcher of Bioinformatics at University of Porto, Portugal	Clinical Implementation & Market Adoption	Use of AI in Biomedical Research o Machine Learning (ML) is essential in genomics, and increasingly applied in histopathology, where models are used to analyse histological images through image and video processing. o The field generates large volumes of data, and the application of ML models emerged naturally, with no major resistance. o AI adoption has shifted biology from a hypothesis-driven discipline to a data-driven one. Current Implementation Status o Unlike in the early 2000s, when tool availability was limited, today there are numerous free and robust AI tools. This facilitates implementation and adoption. o However, hardware and infrastructure can still present challenges, especially for handling large datasets, although modern personal computers already offer strong computational capacity. Technical Challenges o A critical problem in biomedical data is the Curse of Dimensionality, where datasets may contain a high number of attributes (columns) but few samples (rows), leading to poor algorithm performance despite data availability. o There is a shortage of professionals with both biological and computational expertise, which can limit model development and implementation. Data Availability, Ethics, Interpretability and Trust o In the biomedical field, there is no lack of data, with many public and well-organised datasets available. o While ethical issues exist, they are generally well-managed through anonymisation and project-level safeguards. o For researchers, interpretability is not a major concern as they typically understand what they are looking for and how to assess model outputs. Assessing AI's Real-World Value

		o The key question is whether AI models are worth the investment, and how this return on investment (ROI) can be assessed. o An ML model with outstanding accuracy might only yield a modest practical gain (e.g. reducing time by 10% or lowering staffing needs). Moreover, sometimes a less accurate but easier-to-implement model may deliver higher real-world impact. o This would, therefore, require either a custom evaluation of each algorithm's feasibility and relevance in practical settings or the development of a general "blueprint" or checklist to assess real-world applicability and guide implementation. o Evaluating AI tools must move beyond accuracy to practical impact and return on investment. There is a need for more practical examples of case studies to validate models' utility outside of theoretical performance metrics.
Artur Rocha Researcher on secure and privacy-preserving data methods at INESC TEC, Portugal	Data Access & AI Model Development	Data Availability and Privacy o High-quality data is essential for training effective AI models, yet access remains limited due to privacy regulations and institutional constraints. o Hospitals and other data holders are reluctant to share data because they are held accountable for privacy compliance, which is a central barrier to broader AI development. Strategies to Address Data Limitations o Federated Learning: Offers a solution by allowing algorithms to be trained on distributed data without moving it. However, data heterogeneity across institutions can impact model accuracy. Artur notes this is an unexplored area in his work, with no trials yet conducted, but sees it as a promising direction needing validation. o Synthetic Data: Raises ethical and consent concerns, such as whether patients agree to have their data used to generate synthetic datasets. Its usefulness depends heavily on the field of application. For instance, while in immunology, synthetic data may be more reliable than real data, due to the high noise and variability in biological immune responses, in imaging, it may be unreliable or misleading. A key question remains: How much does synthetic data depend on real data? Harmonisation Challenges o The current harmonisation pipeline is partially functional but limited. It requires trust from data providers, depends on anonymisation, and is resource-intensive, currently requiring significant time, cost, and human effort. o Moreover, there are no universal standards, and harmonisation needs vary depending on institutional requirements and the intended use of the data. o Existing standards include the OMOP, a common clinical data standard that facilitates harmonisation across systems by aligning variables but does not cover all use cases and can fall short of user-specific needs; and DICOM, which is specifically used for imaging and annotations. Ethical Priorities

		○ Ethical protection should not be compromised in pursuit of model performance or efficiency. It is crucial to maintain privacy and protect patient data, even though this makes AI development more difficult.
João Matos PhD student at Oxford University on "*AI Ethics & Fairness*", United Kingdom	Ethics & Fairness	Limited Real-World Application of AI in Healthcare ○ Despite the hype, there is limited actual implementation of AI in clinical practice. ○ Challenges are systemic and multifaceted, ranging from regulatory ambiguity to technological misalignment with real needs. Mismatch Between Technology and Real-World Needs ○ Poorly identified market and clinical needs lead to ineffective solutions. ○ There is a lack of multidisciplinary collaboration during development. ○ João stresses the importance of involving patients, clinicians, engineers, pharmacists, and other stakeholders early in the design process to ensure solutions are relevant and feasible within existing clinical workflows. Regulation ○ Complex and unclear regulatory frameworks are a barrier to implementation. ○ There is a need for more adaptive and clear regulatory pathways for AI tools. Ethics, Bias, and Fairness ○ AI models easily replicate societal and systemic biases, especially when training datasets lack representation across ethnicities, genders, socio-economic groups, or geographic locations. ○ There are different reasons behind low representation, such as the absence of certain groups in the population, difficulty in accessing healthcare, and biased measurement tools (e.g. oxygen saturation tools misreading in darker skin tones). ○ Defining fairness remains difficult, and engineers' metrics often don't align with philosophical or ethical standards. ○ Transparency and reporting are critical. Studies have shown that opaque models often outperform transparent ones, raising ethical concerns around model manipulation and publication bias. ○ Reporting guidelines exist and should be strictly followed to ensure accountability. Generalisability and Deployment Challenges ○ Selling AI models as ready-made products is problematic not only because models often fail outside the lab due to environmental shifts, but

		also since models require continuous retraining and validation in local settings. o Generalisation across regions or institutions is largely unrealistic, for example, when Google attempted to deploy models for sepsis prediction, these failed when moved across U.S. states or from the U.S. to Thailand. o AI should move from being a product to a service, encouraging consulting-based partnerships where interdisciplinary teams build, deploy, and continuously monitor models tailored to local datasets and needs. o Continuous deployment and data/model shift are essential considerations for AI model longevity and performance, which highlight the need for ongoing monitoring and adaptation, not static implementation. Systemic and Infrastructural Burden – Healthcare Facility Readiness o AI can exacerbate healthcare system pressures if not properly integrated. In Portugal, when an AI tool was proposed to automate patient distribution in emergency services (INEM), it was quickly realised that this would increase pressure on already overburdened hospitals. o Many hospitals still rely on paper-based systems, which limits AI readiness. To successfully adopt AI-based tools, hospitals need robust digital infrastructure, including digitised patient records, cloud storage, data pipelines and engineering departments. Human-AI Interaction o Studies show that AI + Doctor decision-making can underperform compared to either AI or Doctor alone, due to a lack of trust or misalignment in reasoning. o The Doctor-AI interface needs to be carefully considered, taking into account clinical workflows and user trust. Environmental Impact o Training AI models has a large environmental footprint, involving high energy use, rare materials (e.g. silica), and data centre infrastructure. o This is an often-overlooked issue that should be factored into AI evaluation and policymaking. Explainable AI (XAI) o João views it as conceptually valuable but difficult to implement effectively, given that in practice, explanations may not make sense to clinicians or reinforce pre-established biases. o Trust in AI may rely more on results than transparency, similar to how patients trust medications they don't fully understand. Synthetic Data o João is sceptical about its added value, believing that synthetic data has limited utility, especially if not grounded in high-quality real data. Federated Learning

		o Acknowledges its use in research, particularly for evidence generation, but not yet trusted for building robust, production-grade models. LLMs in Healthcare o LLMs are still emerging, and João sees high potential, especially in processing language data like hospital records. Pilots are already in place in some U.S. healthcare settings.
Martim Quintas e Sousa Data Scientist at MoonLake Immunotherapeutics, Switzerland	Data Access & AI Model Development	Limited Real-world Translation of AI Research o The number of times that research work reaches the real world is minimal. Martim worked on a GAN project to generate rare disease images, but the training was too time-consuming and resource-intensive, and the final output lacked sufficient quality, proving GANs were not reliable enough for rare disease imaging. Promptly Health o Promptly collaborates with hospitals, insurers, and pharmaceutical companies. Its focus is on data harmonisation, primarily via OMOP, a European standard which mostly uses tabular data. While images are usually stored in the same way, tabular data can have several different appearances, highlighting the importance of standardisation o Image data tends to be more standardised and doesn't pose the same challenges. o Harmonisation is a time-consuming task, and using LLMs to automate the process can help save significant time (days to months) while reaching 70–80% accuracy. However, to achieve full accuracy, manual model optimisation is still required, potentially taking years. Data Accessibility: EU vs US o In Europe, data privacy is taken very seriously with highly restrictive data policies resulting in limited access even when data exists, hindering AI model development. o On the contrary, the United States follows a free-market data model, where some companies enable access to data from up to 75% of the population. This data is often anonymised, cleaned and standardised. The US healthcare system has most of its tabular data in a standard format, which facilitates harmonisation. o Insurance data is highly available, while electronic health records (EHRs), which include rich clinical data such as notes and procedure results, are less available. o Pharmaceutical companies use these datasets for sales forecasting, diagnostics, and AI-based optimisation (e.g., Harrison AI). Main Barriers Faced o High-quality data remains the cornerstone of reliable AI models, and creating large, diverse and comprehensive datasets is extremely expensive.

		○ Lack of data access is a fundamental barrier for both research and clinical implementation, which is particularly hard to face in the EU, given the intricacies of European data policies. ○ For AI to be acceptable in medical practice, extremely high precision requirements must be met. If AI models fail to reach near-perfect precision, preference is given to humans (doctors), even when outperformed by AI models. Potential Solutions ○ Research and innovation in efficient and optimised architectures are needed to develop models which are less reliant on high-quality and comprehensive datasets. ○ Datasets should be created in compliance with data-privacy regulations. ○ Given that the regulatory standards for Medical Devices require a transparent understanding of how decisions are made, explainable AI is essential. XAI is useful to increase trust and facilitate the auditing and certification process. ○ Martim is sceptical of synthetic data but believes in the potential of federated learning. AI as a Service vs Product ○ Since the deployment of AI-based solutions requires continuous monitoring and retraining, it should be considered a service, not a one-off product. AI in Clinical Practice ○ One of the challenges in AI deployment is liability: Between doctors and the AI companies, who is responsible for the clinical outcomes? ○ Moreover, trust and user acceptance are also common challenges for AI adoption, particularly in the medical field. To overcome these barriers, AI should be used t augment and not replace human expertise, for instance, by being employed as a support tool for clinicians.
João Fonseca Product Owner and Data Engineer at Promptly Health, Portugal	Data Access & AI Model Development	Promptly ○ Founded in 2017 with the mission of contributing to value-based healthcare by linking treatment cost to patient outcomes. ○ Amongst its various products, it offers Harmonize, a product designed for data harmonisation at the hospital level. This is licensed as a service, so that hospitals retain and process the data locally, while Promptly stays responsible for model updates and maintenance. ○ Promptly uses AI to automate their products and increase efficiency. Main Challenges Identified ○ Access to patient-level data remains one of the biggest barriers. This is particularly true in the EU, due to strict regulation and privacy policies, which limit usable data. ○ Not all collected data is useful or structured properly. Data often lacks quality and standardisation, requiring labour-intensive curation.

		○ Hospitals' operational capacity is another real limitation. Most hospitals lack the necessary technical infrastructure and resources for training AI models, deploying them, or utilising strategies such as federated learning. Federated Learning ○ It's feasible business-wise, but technically challenging. It works by having AI developers share the model training scripts with the hospitals, having the models trained on local data, and getting the results back. ○ Given that the required infrastructure would be owned by the hospitals, these are, in most cases, responsible for its costly set-up. ○ João believes that this strategy has more potential than synthetic data and expects it to grow the most. Synthetic Data ○ João does not believe synthetic data can substitute real-world datasets in most medical use cases. Data Harmonisation ○ Hospital have data in different formats, and harmonisation involves converting formats and types but also matching ontologies and vocabularies. Harmonisation is therefore important, but still immature in practice. ○ It's expensive, and even though AI is being used to automate part of the harmonisation process, it currently still requires human intervention and manual fine-tuning. ○ Different formats exist, such as OMOP for analytics, SNOMED CT for clinical terms and diagnostics, ICD for disease classification, OpenEHR for data management, and FHIR for fast data exchange. Business Model Considerations ○ Business models must adapt to limited hospital resources, and AI solutions must fit within a hospital's operational capabilities. ○ Hospitals should understand and value the benefits to patients, which could justify the investment in infrastructure and model adoption. ○ A shared responsibility model could be established, where AI developers are responsible for technical implementation and support, and healthcare facilities for providing data and enabling integration.
Duarte Rodrigues Director of Innovation at Neuroes, Portugal	Business Strategy	Neuroes ○ Neuroes developed an AI-based ECG helmet designed for mental health monitoring, targeting markets sectors such as sports and corporate wellness. ○ By having the tool classified as a General Wellness product by the FDA, and not a Medical Device, it managed to avoid complex regulatory hurdles, such as clinical trials, and validation was done solely through lab testing. ○ As their business model, Neuroes offers subscription-based services.

		Regulation o Medical Device classification depends on the intended use and invasiveness level. Wellness products are subject to lighter regulation compared to clinical devices, which require Medical Device Certification and clinical trials. o In the EU, following the AI Act is required by regulation. Healthcare Facility Readiness and Adoption o The successful integration of AI-based tools in hospitals is heavily dependent on their infrastructure and operational capacity. o To work with AI, hospitals need to be empowered, requiring digital infrastructures, training, and internal management changes. o Moreover, cultural differences and readiness are also significant barriers to adoption. For example, the higher openness in northern Europe compared to the more conservative systems in countries such as Portugal suggests that certain markets might be easier to reach. o One strategy that can help face these challenges and accelerate adoption is establishing internal contacts within the hospitals. o Another strategy which can facilitate adoption is to introduce AI as a support tool for doctors, not a decision-maker, facilitating regulatory approval and increasing trust by medical professionals and patients. AI Business Models o AI tools require constant monitoring and updates, which is why most companies use AI as a service model. Moreover, this type of business model is subject to recurring revenue, which increases its profit potential. o If an AI-based tool is to be introduced as a product, disclaimers can be an important tool to reduce liability. With this strategy, updates to the AI model can occur via retraining using cloud-shared local data. o Another strategy is for startups to sell validated technologies to larger companies, such as J&J, Medtronic, Bosch, etc. Since these corporations already have distribution networks and hospital integration pipelines, they are better positioned to handle hospital onboarding and regulatory scaling. Main Barrier Identified o Founders develop AI products based on assumptions instead of real user needs, which leads to misalignment between product capabilities and clinical/ hospital requirements. o Given that the number 1 failure point is a lack of early feedback from doctors and patients, a user-centred design from the beginning is essential for long-term adoption.
Inês Campos Application Specialist at ByMe, Portugal	Clinical Implementation & Market Adoption	Healthcare Facility Information Systems o There are multiple systems used in a hospital for data management, such as: PACS (Picture Archiving and Communication System), which stores medical imaging data like MRI and CT results, and different hospitals or units of the same hospital may have separate PACS instance; RIS (Radiology Information System), which manages radiology workflows

		and patient data; VNS, which stores both images and clinical documents; HIS (Hospital Information Systems), the central system that coordinates data and communication between subsystems. o SPMS is the Portuguese public service responsible for overseeing healthcare data access and interoperability. Interoperability Barrier o Hospitals lack interoperability due to having multiple distinct systems, using various software providers with incompatible systems, lacking standardisation, etc. For example, in Portugal, there are still significant paper-based processes, especially outside imaging specialities, such as imageology and gastroenterology, which tend to be more digitally mature. o Overall, resistance to digitalisation remains in certain departments and user groups, often deriving from fear of disruption, usability issues or lack of training. Adoption Strategies o Since interoperability is vital for AI and digital health tools to work, it should be considered early in product development to ease integration and adoption. o Currently, the incompatibility between different software and systems requires doctors to use multiple programs to achieve a single task. Developers, by providing healthcare workers with easy-to-use products, can help increase efficiency and usability, which would accelerate market adoption. o A user-friendly interface (UI) is another way to facilitate adoption, since healthcare workers are more likely to use a tool that's intuitive and fits into their workflow. o To tackle the resistance to new technology, it's important to plan for structured onboarding and training, ongoing support and follow-up, and encouraging habitual use by simplifying daily routines. Key Barriers o Innovation must start from a clearly identified problem and not just a novel idea. o Large volumes of diverse data are required for AI-based innovation. Some strategies could include searching for open datasets, collaborating with hospitals, or data augmentation, which might contribute to bias. o Infrastructure, such as high computational power and server capacity are required for AI model deployment. o For AI-based solutions to be applied in clinical contexts, accuracy and reliability are required to be very high. o End-user training and support are essential for adoption and to achieve a real-world impact.
Vasco Dias Data Protection Officer and a seasoned expert in	Data Access & AI Model Development	Key European Regulations Impacting AI in Healthcare o Medical Device Regulation (MDR): central to AI solutions classified as medical devices.

data protection and intellectual property law at INESC TEC, Portugal		o General Data Protection Regulation (GDPR): a well-established and robust framework which focuses on the fundamental rights of individuals and transparency across the data lifecycle, from collection to processing and disposal. o AI Act: introduces a risk-based classification for AI systems and sets standards, though many are still vaguely defined. It's still in development, subject to constant changes. It includes product safety and risk classification and has some overlap with medical device regulation. o European Health Data Space (EHDS): an upcoming regulation meant to encourage data use across Europe. It introduces conditions and rules for access and usage, aiming to unlock data while maintaining protection. Practical Insights for Project Planning o The process of classifying an AI-based solution is complex and should consider potential health effects. o If classified as a Medical Device, start by assessing the MDR regulation and check the overlap with the AI Act. Anonymisation o It's an irreversible process which often includes the raw data deletion, resulting in reduced data value. Data utility and value is, usually, inversely proportional to anonymisation level. o The process of anonymisation is complex and of high risk for companies, since it's difficult to guarantee full anonymity, and mistakes can lead to heavy legal penalties. This legal uncertainty justifies the rare use of anonymisation. Pseudo-anonymisation o This process uses indirect identifiers (such as coded IDs) to allow data to be traced back only when combined with external information (such as a hospital database), ensuring re-identification is only possible within hospital environments. Therefore, pseudo-anonymisation seems more practical for many healthcare AI use cases. Federated Learning o This approach is gaining traction as a viable alternative to anonymisation due to lower privacy risks and avoidance of complex anonymisation pipelines. It's particularly suitable where regulation enforces that data must remain on-premises. Data Standardisation o Standard data formats should be adopted to ease data exchange, model interoperability, and more efficient AI development and deployment.
Maria Loureiro	Clinical Implementation & Adoption	Loka o Loka uses AI models to simulate laboratory experiments, assisting biotech firms in accelerating pharmaceutical development. The company licenses and sells its models, primarily to partners who have in-house AI

Machine Learning Engineer at Loka, Portugal		teams capable of monitoring and updating the models or can be trained by Loka or associated consultants. Data Access o Data access remains the main challenge for the development of AI-based tools, with different strategies being employed to overcome it. o Federated Learning: not ideal for early model development, and more applicable in later stages when the model architecture is stable. o Data Banks: Since these are often built with data acquired through different methods, the resulting variability can harm model development. o Building Partnerships for data access, where the data is first anonymised – this strategy is used by iLOF. Data Variability and Ethical Barrier o Data variability derives from differences in race, ethnicity, techniques utilised to acquire data, etc, and can lead models to perform differently for different groups. o To face this ethical barrier, transparency is an essential tool, which should include stating model limitations and acknowledging potential biases or gaps in data. Healthcare Facility Readiness o While public hospitals are more likely to lack the digital infrastructure required for AI introduction, market competitiveness and the greater incentive for digitalisation and innovation result in private hospitals typically being more advanced. Cloud Platforms o Cloud suppliers include the AWS, Azure, and Microsoft, who supply platforms that are often used for data storage, processing, and model deployment. o Benefits include data never leaving its proprietary space, ensuring regulatory compliance, and enabling international collaboration, for example, between the EU and the US. o Risks include downtime and cybersecurity threats, which emphasise the need for reliable and secure cloud providers. o However, using the Cloud comes with high operational costs, and strategic partnerships could help reduce expenses. Data Access o Centralisation of data is key to ensure scalability and efficiency, but data access restrictions have to be carefully designed. Healthcare Facility Readiness o A real-world example highlighting the mismatch between model requirements and real-world deployment capabilities was the development of an LLM model for patient distribution by Sword Health, when INEM's system was not ready to adopt such a technology.

		o Early integration with stakeholders and understanding the implementation context are vital for AI-based solutions successful deployment.
Célia Cruz Manager, Founder and Director of Regulatory Affairs of Complear, Portugal	Medical Device Regulation & Reimbursement	Medical Device Classification o A solution qualifies as a medical device if it supports diagnosis, treatment or monitoring of a specific patient. o EHRs are not typically considered, but can be, a medical device. o It's important to distinguish regulations into the Medical Device Regulation (MDR) and the In Vitro Diagnostic Regulation (IVDR) AI Act o This EU regulation highlights the importance of transparency and prohibits certain AI uses, for example, systems considered manipulative or deceptive. Under the AI Act, AI-based solutions are usually either classified as High Risk or must ensure transparency requirements. o If the AI systems fall under medium or high-risk under MDR, which requires assessment by a Notified Body, it is automatically High Risk under the AI Act. o There is a close interplay between MDR, AI Act and GDPR. Practical Guidance for Medical AI Technologies o 1) Assess if the AI-based solution is a medical device under MDR compliance. Rule 11 covers most software tools that support diagnosis or treatment decisions, and classification typically results in Class IIa or higher. Following this rule, classification in Class I is rare, although its acceptance can vary by region, and how the tool is described can heavily impact classification. o 2) Address requirements specific to the AI Act. Many AI Act standards will overlap with MDR if classified as a medical device. o 3) Assess GDPR to ensure data protection, consent, and data lifecycle transparency. o 4) Implement robust AI development frameworks and assess risk. o 5) If not classified as a medical device, follow the AI Act classification and focus on the transparency requirements. o 6) If classified as a medical device, choose and contact a Notified Body. Consider the language, costs, and speed. It's important to note that there is no unified guidance yet, and procedures for software are still vague and inconsistently applied. o 7) It's also important to consider cybersecurity, through the Cyber Resilience Act (CRA), and data interoperability and standardisation through the European Health Data Space Regulation (EHDS). o This process can take around one and half years, although its highly variable and dependent on different factors. Its cost would start around 75.000€ - 100.000€ just for the Notified Body, with added costs related to internal development, testing (cybersecurity, etc.), and possible consultancy. Strategic Advice

		o If possible, to reduce the initial regulatory burden consider launching the device as Class I and expand and upgrade later with additional features. Reimbursement Pathways o In Europe reimbursement for AI tools is currently limited. Strategies to address this barrier include Applying for local reimbursement through national digital health application pathways (for example, DiGA in Germany), negotiating directly with private or public insurers, and participating in national initiatives and competitions. o In the United States, the regulatory system has become increasingly complex and slower, and reimbursement is based on medical codes.

Table 11 Unstructured interviews conducted to explore the field of AI in healthcare and identify key challenges and best practices in AI-based MedTech Innovation.

Contact Person Position, Institution, Country	**Dimension**	**Key Lessons Learned**
Frederico Carpinteiro CEO at Amparo, Portugal	Comprehensive Expertise	Framework Positioning o The framework is useful as an AI-specific complement to Biodesign. It should be applied after identifying clinical and stakeholder requirements, fitting between the Concept Generation and Implementation phases. o Alternatively, the framework should include "Identifying Clinical and Stakeholder requirements" after clinical need identification. New Framework Components o "Requirements Validation" is an important addition to be included before moving into technical development. Question 1 o This is a key question for AI innovation. o Often, the answer to whether usable data exists is "no", due to a lack of data standardisation and poor interoperability across institutions, which are major blockers for AI development and deployment. Question 2 o Replace "hospital" with broader terms like "health institution" to cover varied clinical settings and increase generalisability. o A major challenge is the lack of trained personnel to collect and input data correctly, resulting in poorly entered data, which can harm AI performance. Ensuring access to trained professionals is essential. Question 3

		o The development process should be cyclical: iterate through requirements, validation, and testing until satisfactory results are achieved. After market deployment, a similar cycle begins for future product versions. o "Clinical Integration" could be divided into "Clinical Requirements" (early phase) and "Clinical Deployment" (late phase), to better reflect different concerns at each stage. Question 4 o Stakeholder-specific needs should be revisited after choosing to pursue an AI solution, as different stakeholders may have distinct concerns or expectations related to AI. Question 5 o Data protection principles should be explicitly addressed as a challenge. o Use general terms like "Market Regulatory Requirements," broken down into four dimensions: Cybersecurity, Data Protection, Local Compliance, and Medical Device Development. Then specify according to each key market, like the EU and the US. Question 6 o This question appears to concern requisites and could be incorporated into the initial clinical and stakeholder requirements. o Reimbursement should be considered early, and if codes exist, ensure that their requirements are met. If not, then the process to establish them can take years, and should be factored into business model planning. Question 7 o Usability should be addressed from the start and verified through usability studies before deployment to ensure alignment with user needs. Question 9 o Challenges such as Net Impact and Return on Investment (ROI) are often overlooked, and AI projects proceed without assessing whether the effort is justified by the added value.
Michel Nemery Head of Clinical Development at CEREBRIU Denmark	Comprehensive Expertise	Start with Identifying the Unmet Need o The starting point should always be a clearly defined unmet clinical need. o AI should be designed around users, existing workflows, and concrete use cases. o Solutions must adapt to current workflows, not reinvent them. o A helpful exercise is to ask what a perfect solution would look like, a moonshot idea that guides ambition. o Patients should be involved from the beginning, as they are central to identifying meaningful problems and validating solutions.

		Leadership Commitment o Leadership enthusiasm is essential - doctors, developers, and decision-makers must be excited and aligned. If the proposed value doesn't resonate strongly enough, teams should go back and rework the idea. o Teams should be built around broad perspectives and remain energised and motivated throughout development. o A leadership pitch must be crafted carefully, especially because senior figures often have more to lose from change. o Context also matters rural hospitals tend to be more open to innovation than established city hospitals, where academic status can be a barrier. o Engaging young, ambitious talent is crucial, as they are more likely to drive innovation forward. o Co-creation across stakeholders helps align everyone around a shared vision of change. Organisation (Cross-Silos) o Healthcare institutions are often highly siloed, with slow and complex decision-making, which makes it difficult to introduce and scale innovation. Breaking down silos across departments, by establishing communication channels, is hard but one of the most important enablers of successful innovation. o It's essential to understand how decisions are made within institutions and who holds influence. o Embedding solutions directly in existing equipment (e.g. scanner software) can avoid many organisational dependencies. o The most robust AI solutions are simple, self-contained, and can function within the clinical environment without requiring deep system changes. Question 1 o Data bias is not just a technical issue and often reflects the existing biases of the real world. The idea of completely "debiased" data may be a myth if real-world healthcare itself is biased. o Ongoing access to representative, real-world data is essential for trustworthy AI systems. Question 2 o AI solutions must align with the plug-and-play expectations of clinical users and, wherever possible, be built into existing platforms to avoid disruption. o Teams must understand the entire end-to-end clinical infrastructure to build viable integrations. Question 3 o Model performance must be tested and retrained regularly based on real-world conditions. o It's important to modify models over time by enriching training datasets as new data becomes available. As so, ongoing data curation is required to ensure that the model remains clinically relevant and safe.

		Questions 4 and 5 o Clearly determining the intended use of the AI solution is essential for navigating the regulatory pathway, particularly since regulatory approval becomes significantly more complex if the AI is classified beyond clinical decision support. o Defining the intended use also helps shape go-to-market strategies, especially when partnering with larger companies. Question 6 o The goal of AI should not be limited to speed but meaningfully improving the performance of its clinical users. o Plug-and-play functionality helps accelerate adoption and reduces resistance. o Maximising real-world impact requires clear KPIs tied to actual clinical outcomes. Question 7 o AI could be framed as decision support to preserve, not replace, the authority of clinicians. o Developing AI focused on solving low-complexity, repetitive tasks can lower resistance and demonstrate quick wins. o AI solutions can also target the "normals": common cases that don't require deep clinical expertise. A practical example is stroke triage: 9 out of 10 patients admitted for stroke turn out not to have one, and if this is determined early, especially before weekends, it could reduce unnecessary hospital stays. Since most hospital costs come from extended stays, shortening the length of stay would have a huge economic impact. o Timeliness is often more critical than precision, as faster diagnosis enables earlier treatment, which can be lifesaving and cost-saving. o For AI to be useful, it doesn't have to necessarily outperform the best clinicians, it only has to be more reliable than the poorest alternative, such as an overworked junior doctor on a night shift. Question 8 o Fairness in AI demands careful auditing, but the process must acknowledge real-world imperfections. o Transparency is essential. Patients should not be excluded to achieve ideal datasets, since AI must be developed for unpredictable, real-life conditions. o Local validation is critical as performance can drop dramatically when AI is deployed in new settings. An example: an AI technology with published precision around ~99%, intended to be embedded in a mobile X-ray machine, performed at 14% at Michel's hospital. This was due to training data biased toward healthy young males, a mismatch with the hospital's patient population. o Although population drift is a challenge, technology drift should also be paid attention, since it changes in how data is generated it is equally impactful. Models trained with data acquired with old technologies,

		might perform differently when placed under new technological conditions. o Continuous validation and monitoring are required not only for technical reasons, but also legal and ethical ones under the AI Act. Question 9 o Ultimately, one of the purposes of AI in healthcare is to save lives. o Successful innovation happens when multidisciplinary teams that combine clinical, technical, legal, psychological, and social expertise collaborate toward common goals in real-world settings.
Miguel Amador Chief Innovation Officer at Complear Portugal	Comprehensive Expertise	Role and Limitations of Standards: o The importance of standards in guiding the innovation process was strongly emphasised. Standards help ensure consistency, safety, and interoperability across stages. o Although a growing number of standards specific to AI already exist, they remain underutilised and are not yet widely adopted in practice. o There are also gaps in current standardisation efforts, and in some cases, existing AI-related standards may not be entirely compatible with medical device regulations, potentially leading to conflict or redundancy. Framework Structure and Iterative Nature of Innovation: o Given the inherently cyclical and iterative nature of medical innovation, it was suggested that the framework's questions should not be interpreted as a linear sequence of steps. o Many considerations introduced later in the framework, such as ethical and fairness-related practices (e.g., FRIA analysis, dataset auditing, transparency), must in fact be addressed from the earliest stages. o Among the best practices discussed, there is a mix of regulatory obligations and voluntary recommendations. For example, in the case of AI-based medical technologies intended for general deployment, developers are required to ensure the use of diverse and representative datasets during training. However, local validation post-development is not a regulatory requirement, even though it is widely regarded as a best practice to ensure performance in specific clinical contexts. Regardless, post-deployment monitoring remains mandatory (imposed by the AI Act), particularly to detect performance drift and manage associated risks. o These processes rely heavily on risk identification and mitigation strategies, which should be embedded across the framework, not treated as discrete or isolated steps. Real-World Need Alignment: o The importance of aligning innovation with clearly defined real clinical needs was reaffirmed as foundational to the success of any MedTech project. Concluding Recommendations: o Emphasise the critical role of standards at every stage of the innovation process to ensure consistency and quality. o Clarify that some best practices, particularly those related to ethics, data, and transparency, must be applied from the outset, not deferred until later stages.

Table 12 Semi-structured interviews conducted to validate the proposed VAMI Framework