



CHALMERS
UNIVERSITY OF TECHNOLOGY



Planning a Clinical trial for a Stroke Diagnostic Device in a Resource-Constrained Health System

A Case Study of the StrokeFinder MD100 in Rwanda

Degree project report in Biomedical Engineering and Management and Economics of Innovation

David Fernblad
André de Mots

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

MASTER'S THESIS 2026

Planning a Clinical trial for a Stroke Diagnostic Device in a Resource-Constrained Health System

A Case Study of the StrokeFinder MD100 in Rwanda

David Fernblad
André de Mots



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Planning a Clinical trial for a Stroke Diagnostic Device in a Resource-Constrained Health System

A Case Study of the StrokeFinder MD100 in Rwanda

David Fernblad

André de Mots

© David Fernblad, André de Mots 2026.

Supervisor: Professor Mikael Persson, Head of Division at Signal processing and Biomedical Engineering

Co-supervisor: Dr. Jean-Baptiste Nyandwi, PhD Senior Consultant Pharmacist, University of Rwanda

Examiner: Associate Professor Andreas Fhager, Head of Unit at Signal Processing and Biomedical Engineering

Master's Thesis 2026

Department of Electrical Engineering

Chalmers University of Technology

SE-412 96 Gothenburg

Telephone +46 31 772 1000

Cover: Patient in the StrokeFinder MD100. Newcastle Herald, March 13, 2018, <https://www.newcastleherald.com.au/story/5279734/leading-by-precision/>.

Typeset in L^AT_EX

Printed by Chalmers Reproservice

Gothenburg, Sweden 2026

Planning for Clinical Trials Evaluating the Clinical Value of a Stroke Diagnostic Tool in a Resource-Constrained Healthcare System

A Case Study of the StrokeFinder MD100 in Rwanda

David Fernblad, André de Mots

Department of Electrical Engineering

Chalmers University of Technology

Abstract

Stroke is the leading cause of death in Rwanda, and its incidence has grown sharply over recent decades. Yet confirmatory brain imaging is only available at four tertiary hospitals in the country, and most patients arrive too late to receive the only clot-dissolving treatment available. The Medfield StrokeFinder MD100 is a portable device that uses microwave technology to detect stroke at the bedside, without requiring a CT scanner. It has shown promising results in European and Australian studies, but has never been tested in sub-Saharan Africa, where patients, disease patterns, and care pathways differ considerably. This thesis develops a planning framework for a GCP-compliant clinical validation study of the MD100 in Rwanda. The framework is based on sixteen interviews conducted in Kigali between March and April 2026 with clinicians, hospital researchers, national regulatory bodies, and a trial consultancy. Analysis of the roughly 80,000-word interview corpus identified three key findings: the locations where the device would be most useful are also the hardest places to run a validation study; clinical trial terminology is not consistently understood across the healthcare workforce; and knowledge of AI-specific regulation is concentrated at a single institution. These findings are translated into a four-phase trial-planning framework. The core of the framework is a two-stage trial: A first stage at tertiary hospitals establishes how accurately the device classifies stroke against CT imaging under controlled conditions. A second stage at district hospitals then measures whether deploying the device in routine care actually speeds up the time from first clinical contact to diagnosis.

Keywords: StrokeFinder MD100, clinical trial planning, Rwanda, stroke, medical device validation, resource-constrained healthcare, GCP, pragmatic trial, diagnostic accuracy.

Acknowledgements

We would like to begin by expressing our sincere gratitude to our supervisor, Professor Mikael Persson, for his steadfast support throughout the project, his participation in interviews, and his invaluable insights at every stage. We are grateful to our local co-supervisor, Dr. Jean-Baptiste Nyandwi, for his guidance and for opening doors within the Rwandan healthcare system that made this fieldwork possible. Our thanks also go to Pascal Nyiringago and Peace Ndoli for their support on the ground in Kigali.

To the sixteen professionals who gave their time to be interviewed across the hospitals, government bodies, and institutions of Kigali: this thesis rests entirely on your knowledge and openness. We are genuinely grateful for the enthusiasm and curiosity you brought to every conversation. We would also like to thank the Global Mentorship Programme at Chalmers for providing us with the means to execute this project, and David Ivarsson and Hannes Grubb for their support and companionship along the way.

David Fernblad, André de Mots

Gothenburg, June 2026



List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

AIaMD	Artificial Intelligence as a Medical Device
AUC	Area Under the Receiver-Operating-Characteristic Curve
CBHI	Community-Based Health Insurance
CHUB	Centre Hospitalier Universitaire de Butare
CHUK	Centre Hospitalier Universitaire de Kigali
CIICHIN	Kigali-based Clinical Trial Consultancy
CT	Computed Tomography
EMS	Emergency Medical Services
EU MDR	European Union Medical Device Regulation (2017/745)
EUnetHTA	European Network for Health Technology Assessment
FAST	Face, Arms, Speech, Time
GCP	Good Clinical Practice
GCPR	Guidelines for Good Clinical Practices in Rwanda
GDPR	General Data Protection Regulation
HIC	High-Income Country
HTA	Health Technology Assessment
ICER	Incremental Cost-Effectiveness Ratio
ICH	International Council for Harmonisation
ICP	Intracranial Pressure
IEC	Independent Ethics Committee
INMB	Incremental Net Monetary Benefit
IRB	Institutional Review Board
ISO	International Organization for Standardization
KFH	King Faisal Hospital Kigali
LMIC	Low- and Middle-Income Country
LVO	Large Vessel Occlusion
MD100	StrokeFinder MD100
MDR	Medical Device Regulation
MoH	Ministry of Health
MRI	Magnetic Resonance Imaging
NCD	Non-Communicable Disease
NCST	National Commission for Science and Technology
NGO	Non-Governmental Organisation
NMB	Net Monetary Benefit

PI	Principal Investigator
PRECIS-2	Pragmatic Explanatory Continuum Indicator Summary 2
QALY	Quality-Adjusted Life Year
RBC	Rwanda Biomedical Centre
REDCap	Research Electronic Data Capture
RFDA	Rwanda Food and Drug Authority
RMH	Rwanda Military Hospital
RNEC	Rwanda National Ethics Committee
SaMD	Software as a Medical Device
SAMU	Service d'Aide Médicale Urgente
TBI	Traumatic Brain Injury
TIA	Transient Ischaemic Attack
tPA	Tissue Plasminogen Activator
UAV	Unmanned Aerial Vehicle
WHO	World Health Organization
WTP	Willingness to Pay

Contents

List of Acronyms	x
List of Tables	xix
1 Introduction	1
1.1 Background	1
1.1.1 Stroke	1
1.1.1.1 Ischaemic Stroke	2
1.1.1.2 Haemorrhagic Stroke	2
1.1.2 Stroke Treatments	2
1.1.2.1 Thrombolysis	3
1.1.2.2 Thrombectomy	3
1.1.3 Rwandan Healthcare System	3
1.1.4 Clinical Trials	4
1.1.4.1 ICH Good Clinical Practice	4
1.1.5 StrokeFinder MD100	5
1.1.6 Health Technology Assessment	6
1.1.6.1 Health Economy	6
1.2 Purpose	7
1.2.1 Goals	7
1.2.2 Limitations	7
2 Methods	9
2.1 Study Design	9
2.2 Literature Review	9
2.3 Key Stakeholders	10
2.3.1 Sampling and Recruitment	10
2.3.2 Interview Procedure	11
2.3.3 Documentation and Transcription	11
2.4 Data Analysis	12
2.5 Triangulation with the Prior Interview Corpus	13
2.6 Ethical Considerations	14
2.7 Reflexivity and Methodological Limitations	14
3 Results	17
3.1 Scope and reflexivity.	17
3.2 The current state of stroke care in Rwanda	18

3.2.1	Late presentation as the dominant clinical reality	19
3.2.2	Geographic centralisation of diagnostic imaging	21
3.2.3	The fragmented referral chain	22
3.2.4	Thrombolysis bottleneck	23
3.2.5	Coordination failures within and between hospitals	25
3.2.6	Active innovation: drones, zip-lines, and district-level investment	26
3.3	Conducting clinical trials in Rwanda	26
3.3.1	Regulatory architecture: RNEC, RFDA, and the umbrella IRB model	27
3.3.2	Funding pathways and trial site type by phase	28
3.3.3	Data governance and local data residency	30
3.3.4	Hospital readiness and pre-engagement networks	31
3.3.5	Points of failure: monitoring, follow-up, and adoption	31
3.4	Trial planning for the MD100	32
3.4.1	Trial endpoints	32
3.4.2	Site selection: the validation-versus-impact trade off	33
3.4.3	Trial design: pilot-before-trial, blinding, and crossover risk	35
3.4.4	Patient inclusion, exclusion, and informed consent in the Rwandan context	35
3.4.5	AI-specific regulatory considerations	37
3.4.6	Expanded use cases beyond stroke triage	39
3.5	Cross-cutting findings	40
3.5.1	Clinical need and validation difficulty appear at the same locations.	40
3.5.2	Trial-related terminology was variably shared across the sample.	40
3.5.3	AI-regulation knowledge was absent at clinician level and concentrated at one institutional actor.	40
3.6	The case for local validation	41
3.7	Clinical need and where the device fits	42
3.8	Designing the validation study	43
3.8.1	The performance trial	44
3.8.2	Choosing an endpoint for the value trial	44
3.8.3	The control condition: comparison without contamination	45
3.9	Conducting the trial in the Rwandan context	47
3.9.1	Funding and sponsorship: who can carry the trial	47
3.9.2	Data governance: what actually has to leave Rwanda	48
3.9.3	Securing sites: readiness and pre-engagement	48
3.9.4	What makes it fail	49
3.10	Terminology, literacy, and AI regulation as design constraints	50
3.10.1	Trial vocabulary as a constraint, not a measurement artefact	50
3.10.2	The AI-regulation knowledge gap	51
3.10.3	Two findings, one remedy	51
4	Conclusion	53
5	Future research	55

Bibliography	57
A Interview Guide	I

List of Tables

- 3.1 Overview of interview respondents. Roles and institutional affiliations are reported; respondent names are withheld. Interview formats are recorded as full audio-transcribed interviews (T), written minutes (W), or brief informational meetings (I). Paired interviews are denoted with suffixes a and b for the two respondents. 18

 - 4.1 The trial-planning framework: four phases running from pre-engagement through close-out. Phases 1 and 2 constitute the trial proper; Phase 0 prepares the regulatory, data-governance, and institutional ground before submission, and Phase 3 consolidates and disseminates the evidence. The final column states the gate that must be cleared before the next phase begins. 54
-

1

Introduction

1.1 Background

In the background the pathology and treatment of stroke will be introduced, as well as the current challenges regarding the rise of non-communicable diseases (NCDs), diagnostics and treatment in Rwanda. International guidelines for clinical trials will also be presented followed by a description of the StrokeFinder MD100, hence forth referred to as MD100.

1.1.1 Stroke

Stroke is a major global concern being the second leading cause of death from non-communicable diseases globally and the third leading cause of death overall, killing 7.3 million each year as of 2021[1]. The burden of stroke is however more significant in low-and-middle income countries (LMIC) where the incidence of stroke doubled from 1990 to 2016 compared to high-income countries where it sank by 42% in [2]. Further 70% of all global deaths occur in LMIC as well as 87% of all stroke related disabilities. In Rwanda the incidences of stroke increased by 50% between 1990 and 2020[3]. Stroke has in recent years also become the most common cause of death in Rwanda[4]. Stroke is a neurosurgical disorder that is characterised by an abrupt disruption of blood flow to the brain[2]. Thereby reducing the brain's supply of oxygen, leading to the death of brain cells. During a stroke the brain loses 1.9 million neurons and 14 billion synapses each minute, corresponding to the natural ageing of 3.6 weeks [5]. Highlighting the acute nature of stroke. The symptoms of stroke vary depending on where in the brain the stroke has occurred [6]. Common symptoms of stroke are one sided loss of motor function or paralysis, difficulty of speech and cognitive issues such as thinking and memory. One can also experience severe headache, confusion and impaired vision. The risk factors of stroke can be put in one of two categories, non-modifiable and modifiable[2]. Non-modifiable risk factors include factors such as sex, age, previous transient ischaemic attack (TIA) and genetics. Modifiable risks are lifestyle choices, alcohol and drug abuse, diet and insufficient physical activities. Also part of modifiable risks are medical conditions such as hypertension, diabetes and hyperlipidemia. In sub-Saharan Africa these non-modifiable risk factors have become more prevalent in the population over time[3]. The limited availability of medication to control risk factors is also identi-

fied as a driver for the increased prevalence of stroke in sub-Saharan Africa. This finding is echoed by a recent study performed at the Rwanda Military Hospital that investigated the prevalence of hypertension amongst stroke patients[7]. In the study they found that 82.3% of the patients had hypertension, of those who had hypertension 93.2% reported poor adherence to medication. Beyond stroke, traumatic brain injury (TBI) represents a related burden in Rwanda, where high rates of road traffic accidents produce neurological presentations that can clinically overlap with stroke, complicating triage in settings with limited imaging availability. TBI is defined as an acquired disruption of brain function caused by an external mechanical force, and shares several acute neurological features with stroke, including focal deficits, altered consciousness, and intracranial haemorrhage[8].

1.1.1.1 Ischaemic Stroke

Ischaemic stroke is caused by a blood vessel in the brain becoming blocked and restricting blood flow to a part of the brain [2]. This occurs in two ways, thrombotic and embolic. If it is thrombotic in nature the blockage is made up of plaque that gets built up locally on the walls of the blood vessel. In the embolic case the blockage has travelled from a different part of the body to then get lodged in a vessel within the brain [9].

1.1.1.2 Haemorrhagic Stroke

Haemorrhagic stroke results from the rupturing of blood vessels within the brain, causing bleeding[10]. The resulting haematoma typically expands within 3 to 12 hours, resulting in the increase of intracranial pressure (ICP) which is primary mechanism of brain damage of a haemorrhagic stroke. Haemorrhagic is less common but has a higher mortality rate, making up circa 10% to 15% of all strokes globally with a mortality rate between 30% to 40%. Haemorrhagic stroke results from the rupturing of blood vessels within the brain, causing bleeding[10]. The resulting haematoma typically expands within 3 to 12 hours, resulting in the increase of intracranial pressure (ICP) which is primary mechanism of brain damage of a haemorrhagic stroke.

1.1.2 Stroke Treatments

The treatment of stroke varies greatly between ischaemic and haemorrhagic. The treatments for haemorrhagic are varied but they function to address three main problems: blood pressure, ICP and the bleeding [10]. Blood pressure is regulated via antihypertensive drugs, this serves to reduce further bleeding and haematoma growth. Haemostatic therapy also serves to stop the progression of the haematoma by promoting blood clotting, this especially important for patients that are taking anticoagulants, this achieved via drug administration. There are also several surgical treatment options; aneurysm clipping is a procedure that serves to block off the aneurysm from the blood vessels by placing small clips at the base of the aneurysm[11]. Another is coil embolization in which the aneurysm is blocked off by placing a coil inside of the aneurysm. The coil is guided to the aneurysm via a catheter that is inserted in the artery in the groin area.

1.1.2.1 Thrombolysis

Thrombolysis is used to treat acute ischaemic stroke and is available for patients that are presented for treatment within 4.5 hours of symptom onset [12]. The treatment is performed by injecting tissue plasminogen activator (tPA) that breaks down the blood clots causing the occlusion to dissolve and thereby restoring proper blood flow. The tPA works by catalysing the conversion of plasminogen to plasmin, which degrades the fibrin that makes up the blood clot[13]. Thrombolysis has shown to be effective in reducing the risk of death and disability after an Ischaemic stroke[14]. The treatment does however also compromise the extracellular matrix, increasing the risk of haemorrhage. There are several conditions that could exclude a patient from receiving thrombolysis, such as, recent ICH, previous ischaemic stroke within three months, active bleeding etc[12]. The only hospital in Rwanda currently capable of performing thrombolysis is King Faisal Hospital in Kigali[15].

1.1.2.2 Thrombectomy

For large vessel occlusions (LVO), defined as a major cerebral artery, thrombectomy is a possible treatment for acute ischaemic stroke[9]. Thrombectomy refers to the mechanical removal of the occlusion, the procedure starts with making an entry point in the upper thigh area[16]. In which a catheter is inserted and guided to the occlusion. At which point a retrievable stent is deployed to grip the occlusion so that it can then be extracted. Before 2018, the window for thrombectomy was six hours but studies have shown significant benefits of thrombectomy performed up to 24 hours after onset. However, thrombectomy is not currently available in Rwanda[15].

1.1.3 Rwandan Healthcare System

Rwanda's healthcare system is organised into four tiers: community health workers at the village level, health centres, district hospitals, and tertiary referral hospitals[17]. The country operates a community-based health insurance scheme (CBHI), commonly known as *Mutuelle de Santé*, which covers the majority of the population, particularly in rural areas. While the scheme has significantly improved healthcare access, it imposes a stepwise referral structure: patients must generally progress through each tier before accessing higher-level care. At the top of the referral hierarchy, Rwanda's main tertiary hospitals, CHUK, KFH, RMH, and CHUB, provide neurosurgical and neurological services. However, the access to advanced imaging is unequally distributed. CT scanning is only reliably available at these tertiary facilities and a limited number of secondary hospitals, while many district hospitals lack on-site imaging entirely[18]. Rwanda also operates a national Emergency Medical Services (EMS) system, SAMU, which provides prehospital care and coordinates patient transport in Kigali, though coverage and response times vary considerably in rural areas[17]. This imaging gap affects both stroke and TBI patients, as both conditions may require urgent intracranial assessment that is unavailable below tertiary level. Rwanda has over the past two decades made substantial progress in strengthening its health system, including investments in health infrastructure, community

health worker programmes, and clinical training. The country has demonstrated a commitment to health technology adoption, and the Rwanda Food and Drugs Authority (Rwanda FDA) has established regulatory frameworks for medical device importation and clinical trial oversight, frameworks that are actively evolving to accommodate digital health and AI-based diagnostics[4].

1.1.4 Clinical Trials

A clinical trial is a structured research study designed to evaluate the safety, performance, or clinical effectiveness of a medical intervention, including diagnostic devices, in human subjects[19]. The conduct of clinical trials involving medical devices is governed internationally by ISO 14155, which defines Good Clinical Practice (GCP) requirements for device trials, and by the Declaration of Helsinki, which establishes the ethical principles for medical research involving human participants[19], [20]. In the European regulatory context, clinical trials of medical devices are further governed by the Medical Device Regulation (EU MDR 2017/745), which mandates that devices demonstrate clinical evidence of safety and performance in their intended use population prior to market authorisation[21]. The level of clinical evidence required scales with device risk classification: higher-risk devices require prospective clinical investigation, whereas lower-risk devices may rely on existing clinical data, literature review, or equivalence arguments. For novel or innovative diagnostic devices being introduced in resource-constrained settings, additional planning considerations apply. These include assessment of local regulatory requirements, site infrastructure and readiness, availability of trained investigators, ethical review procedures, and the feasibility of patient recruitment within the target population[19]. In Rwanda, clinical trial approval requires submission to the Rwanda FDA and ethical clearance from the Rwanda National Ethics Committee (RNEC), with GCP training required for all principal investigators[4].

1.1.4.1 ICH Good Clinical Practice

The ICH E6 GCP is an international ethical and scientific standard for clinical trials with human participants[22]. Its main aim is to ensure that rights and safety of participants are respected while ensuring trial results are reliable and trustworthy. The ICH guidelines describe four primary stakeholders for clinical trial that all have different responsibilities to ensure that the guidelines are followed and GCP is maintained for the duration of the trial. These stakeholders are Institutional Review Board/Independent Ethics Committee (IRB/IEC), investigator(also referred to as principal/primary investigator (PI)), sponsor and data governance. **IRB** - The IRB is an independent body that protects the trial participant's rights and safety. This duty is performed by first reviewing the trial protocol, informed consent material, assent material and other relevant material and documents regarding the trial. The IRB must give its approval before any participant may be recruited for the trial. Once the trial has started the IRB shall review any and all updates to the protocol and conduct regular review at intervals appropriate to the risk to the participants. **Investigator** - The investigator should be an individual that has a relevant education and experience to the subject of the trial and for conducting a trial. The

investigator can delegate trial duties to other persons or parties but still retains the final responsibility and needs to maintain oversight on delegated duties in proportion to their trial importance and/or patient risk. The investigator also has the responsibility of managing the investigational product, protocol compliance and to promptly report any adverse events to the sponsor. **Sponsor** - The sponsor can be a company, institution or individual and is responsible for initiating, managing and financing the trial. This includes designing the trial protocol and management system, hiring qualified investigators and overseeing the trial from start to finish. The sponsors have several important obligations, proactive risk management and ensuring patient safety, conducting monitoring. The sponsor is also responsible for handling the safety reporting to regulatory authorities and maintaining the essential trial records. **Data Governance** - Data governance is a shared responsibility of the sponsor and the investigator(s), and covers the full lifecycle of trial data that being: capture, retention and destruction. Data governance must ensure integrity, traceability, confidentiality and reliability of the data.

1.1.5 StrokeFinder MD100

The StrokeFinder MD100 is a portable, non-invasive diagnostic device that was developed by Medfield Diagnostics AB (Gothenburg, Sweden) that uses broadband microwave technology to detect intracranial haemorrhage[23]. The device operates by transmitting low-power microwave signals through the skull and measuring the reflected and transmitted signals. Due to the difference in dielectric properties between normal brain tissue and blood, the presence and approximate location of a haematoma can be inferred from the signal pattern. A machine learning algorithm processes the microwave data to generate a classification output indicating whether a stroke is likely present[23], [24]. The MD100 is designed to be used as a triage support tool, rather than a replacement for CT imaging. Its primary intended use is in settings where CT is unavailable or access is delayed, enabling earlier differentiation between haemorrhagic and ischaemic stroke, and thereby supporting triage and referral decisions. The device produces a result within approximately two minutes and is designed to be operable by trained non-specialist personnel, including emergency nurses and paramedics[23], [25]. The device has also been explored as a potential screening tool for intracranial haemorrhage in TBI, given the clinical and technical overlap between the two indications[25]. These characteristics make it potentially well-suited for deployment at the point of emergency contact in low-resource settings such as district hospitals, ambulances, or primary health facilities. Previous clinical studies of the MD100 have been conducted in Europe and Australia with established CT infrastructure as the reference standard. While these studies have demonstrated encouraging diagnostic performance in differentiating stroke/no-stroke, the device has not yet been clinically validated in sub-Saharan African contexts, where patient demographics, disease profiles, presenting co-morbidities and care pathways differ substantially from those in which the device was originally trained and tested in[24], [26]. Therefore, there is a need for local trials to evaluate whether these contextual differences affect the performance of the MD100.

1.1.6 Health Technology Assessment

Health technology assessment (HTA) is a systematic method that aims to aid the decision making process at the policy level for health care[27]. HTA supports decision making by determining the value of health technology and providing evidence for the technology across several disciplines, social, economical, organisational and ethical to decision makers and other relevant stakeholders[27], [28]. The European Network for Health Technology Assessment (EUnetHTA) has created a model for HTA, the current one as of 2026 being HTA Core Model Version 3.0[29]. It lists all the domains that need to be explored and how to do so in order to perform a proper HTA. Of these the Clinical Effectiveness and Cost and economic evaluation are the most relevant for this project. These two are evaluated via clinical trials and health economy assessments respectively.

1.1.6.1 Health Economy

Health economics is a part of economics that analyses the cost and impact of resources within health care, usually measured in quality adjusted life years (QALYs)[30]. QALYs is a way to combine the quality and quantity of life gained or lost, 0 meaning that a person is dead and 1 means 1 year of perfect health[31]. Incremental cost-effectiveness ratio (ICER) is the economic value of an intervention compared to a comparator measured in cost per QALY[32]. The calculation for ICER is as follows:

C_1 : Cost of the new intervention

C_0 : Cost of the comparator or standard intervention

E_1 : Effectiveness of the new intervention

E_0 : Effectiveness of the comparator or standard intervention

$$\Delta C = C_1 - C_0$$

$$\Delta E = E_1 - E_0$$

Thus,

$$ICER = \frac{\Delta C}{\Delta E}$$

where:

ΔC : Incremental cost

ΔE : Incremental effectiveness

With ICER and the willingness to pay (WTP), a set amount a health care provider is willing to pay to gain one QALY, the incremental Net monetary benefit (INMB) can be calculated as such[33]:

$$NMB = (\lambda \cdot E) - C$$

where:

λ : WTP threshold per QALY

E : effectiveness (QALYs)

C : cost of the intervention

The incremental NMB comparing intervention 1 to intervention 0 is:

$$INMB = NMB_1 - NMB_0$$

- If $INMB > 0$, the intervention is cost-effective at threshold λ
- If $INMB < 0$, the intervention is not cost-effective
- If $\lambda > ICER$, then $INMB > 0$

1.2 Purpose

The aim of this thesis is to plan a Good Clinical Practice (GCP)-compliant pilot clinical validation study for the MD100 in Rwanda. The goal of this trial, which is outside of the scope of this thesis, is to generate the local clinical and operational evidence required to decide whether, and under what conditions, the device could be responsibly introduced as a triage support tool for patients with suspected acute stroke in the Rwandan healthcare system.

1.2.1 Goals

The following goals have been defined to operationalise the purpose of this project:

- To map the current stroke care pathways, diagnostic practices, and referral structures across relevant healthcare facilities in Rwanda.
- To identify regulatory requirements, ethical approval procedures, and site readiness criteria for conducting a clinical trial of the MD100 in Rwanda.
- To assess the feasibility of deploying the MD100 in different care settings, including tertiary hospitals, district hospitals, and prehospital EMS environments.
- To synthesise stakeholder perspectives on the clinical value, implementation barriers, and equity implications of the MD100 in the Rwandan context.
- To produce a structured clinical trial planning framework that can serve as a basis for a future formal clinical investigation of the MD100 in Rwanda.

1.2.2 Limitations

The following limitations apply to this project:

- No clinical testing of the MD100 will be performed as part of this project.
 - No patient data will be collected; the study involves only interviews with healthcare professionals and regulatory stakeholders.
 - There will be no contact with patients.
-

-
- The primary clinical focus of this project is stroke. TBI is discussed as contextual background given the overlap in clinical presentation and device applicability, but the trial planning framework is scoped to stroke as the primary indication.
 - The scope is limited to Rwanda; findings may not be directly transferable to other LMIC settings without contextual adaptation.
 - The project does not include a health economic analysis or formal cost-effectiveness evaluation of the MD100.
-

2

Methods

This chapter describes the research design, data collection approach, interviewee selection, and analytical methods employed in this study. The project is a non-interventional, mixed-methods assessment aimed at planning a GCP-compliant clinical validation study for the MD100 in Rwanda. No patients were involved, no clinical testing of the device was performed, and no patient data were collected.

2.1 Study Design

The study follows a mixed-methods, non-interventional design conducted in two complementary parts. The first part consisted of a literature analysis, comprising a structured review of national policies, clinical guidelines, regulatory frameworks, and existing literature on stroke care and clinical trial conduct in Rwanda. The second part consisted of qualitative key informant interviews with stakeholders involved in the planning, oversight, administration, and conduct of clinical trials in the Rwandan healthcare system.

This design was chosen because the research questions concern contextual feasibility and stakeholder perspectives rather than clinical outcomes. A qualitative approach is well-suited to capturing the institutional knowledge, practical experience, and normative judgements that are necessary for planning a clinical trial in a resource-constrained setting[34]. The literature review component provides the structural and regulatory context within which interview findings are interpreted.

2.2 Literature Review

The literature review involved a review of documents relevant to stroke care, medical device regulation, and clinical trial governance in Rwanda. Sources reviewed included national clinical guidelines for stroke management, Rwanda FDA regulatory documentation on medical device importation and clinical trial approval, published literature on neurosurgical capacity and stroke epidemiology in Rwanda, and international standards governing clinical investigations of medical devices, including ISO 14155 [19], the Declaration of Helsinki [20], and EU MDR 2017/745 [21].

The review was structured around the research questions defined in the planning

phase of the project, covering regulatory approval pathways, ethical oversight requirements, site infrastructure and readiness, and the current state of stroke diagnostic practice across facility types in Rwanda.

2.3 Key Stakeholders

The desired stakeholder perspectives covered three broad categories: clinical, trial, and regulatory. The clinical perspective came from medical experts relevant to stroke care, namely emergency physicians, neurologists, and neurosurgeons. The trial perspective came from those involved in the design, oversight, and conduct of clinical trials, including hospital research-department leadership, an active site principal investigator, and an independent trial consultancy. The regulatory perspective was covered by the Ministry of Health and governmental organisations such as the Rwanda Biomedical Centre (RBC).

Gaps in the sampling frame are discussed in Section 2.7.

2.3.1 Sampling and Recruitment

A purposive sampling strategy was used to identify key informants with relevant professional roles and institutional affiliations. The sampling frame was designed to ensure representation across the main stakeholder categories involved in clinical trial planning, regulatory approval, and stroke care delivery in Rwanda. These categories included clinical staff at tertiary and district hospitals, emergency medicine personnel, hospital research leadership, national policy and programme bodies, and an independent trial consultancy. In total, sixteen informants from ten institutions were interviewed. Tertiary referral hospitals were represented by the University Teaching Hospital of Kigali (CHUK), King Faisal Hospital (KFH), and Rwanda Military Hospital (RMH); the district level by Kibagabaga District Hospital; the primary-care tier by Kinyinya Health Center; national bodies by the Ministry of Health and the Rwanda Biomedical Centre (RBC); and trial consultancy by CI-ICHIN, a Kigali-based consultancy that supports the design and conduct of clinical trials in Rwanda. The full respondent overview, with roles and interview formats, is provided in the Results chapter (Table 3.1).

The initial sampling frame was constructed from the contact list of a prior MSc thesis on the MD100 in Rwanda [15], which interviewed an overlapping set of respondents some twelve months earlier, and was extended through introductions provided by the Chalmers academic supervisor (Prof. Mikael Persson) and by the local co-supervisor at the University of Rwanda (Dr. Jean Baptiste Nyandwi). Once fieldwork had begun, several respondents introduced the research team to additional contacts within their own institutions and across institutions, producing a limited element of snowball sampling within the existing frame. Despite repeated requests during fieldwork, a direct interview with the Rwanda FDA could not be secured within the available time; the regulatory perspective is therefore reconstructed from accounts given by respondents at the Ministry of Health, the Rwanda Biomedical

Centre, and hospital research departments who interface with the Rwanda FDA in their professional capacity. This gap is discussed further in Section 2.7. Prospective respondents were initially contacted by email and, in some cases, telephone. Prior to each interview, all relevant materials were sent in advance, including a short description of the research, a one-page description of the StrokeFinder MD100, the consent form, and documentation of ethical clearance. Respondents were informed of the study purpose, the voluntary nature of participation, their right to withdraw at any time, and the procedures for data storage and confidentiality.

2.3.2 Interview Procedure

Interviews were conducted in person at the interviewees' respective institutions during fieldwork in Kigali, Rwanda. Each interview lasted approximately 45 to 60 minutes. With interviewees' permission, interviews were audio-recorded to support accurate transcription. Where recording was not consented to, detailed contemporaneous notes were taken by the interviewing researcher.

A semi-structured interview guide was used across all interviews, developed in advance based on the research questions of the project (Section 1.2) [35]. The guide was structured around three thematic areas reflecting the logical progression of the study. The first area addressed the current state of stroke care in Rwanda, covering diagnostic practices, care pathways, referral structures, and system-level constraints. The second area addressed clinical trials in general, including regulatory requirements, ethical procedures, site readiness, investigator qualifications, and funding. The third area addressed the application of these considerations specifically to the MD100 and to stroke as the target indication, covering trial design, patient recruitment, workflow integration, and stakeholder perspectives on the device's clinical value. While the guide provided a consistent structure across interviews, the semi-structured format allowed interviewees to raise additional perspectives and for interviewers to follow up on emerging themes. This flexibility is appropriate for qualitative research in complex health system contexts where unanticipated factors may be relevant[35].

2.3.3 Documentation and Transcription

Three documentation formats were used across the sixteen interviews. Most sessions were recorded with the respondent's permission and transcribed using AI-assisted speech-to-text software, with the resulting transcripts reviewed manually against the original audio to verify accuracy and to correct transcription errors. These are recorded as full audio-transcribed interviews (T). Where a respondent declined audio recording, handwritten notes were taken by one of the researchers during the meeting and expanded immediately afterwards into written minutes (W). A small number of sessions took the form of brief informational discussions for which no live notes were taken; these were documented afterwards from memory by the researcher present, and are recorded as informational meetings (I). The format used for each interview is reported in the respondent overview (Table 3.1).

Fieldwork and analytical labour were divided between the two authors. Most inter-

views were conducted jointly, with one author leading the questioning and the other taking contemporaneous notes; a smaller number of interviews were conducted by one author alone due to scheduling constraints. AI-assisted transcription of audio-recorded interviews was the default approach; handwritten minutes and informational notes (formats W and I) were predominantly produced by David Fernblad. Chapter drafts were produced collaboratively, with primary responsibility for individual chapters distributed between the authors and reviewed jointly.

All transcripts and notes were anonymised before analysis by replacing respondent names with the codes R1 to R16, retaining only the professional role and institutional affiliation. Two interviews were conducted with two respondents present jointly, at the request of the institutions concerned; for analytical purposes these are split into separate respondents, denoted with the suffixes *a* and *b* (R12a, R12b; R13a, R13b), since the two respondents in each paired interview held distinct roles and frequently expressed distinct views. The combined corpus comprises approximately 80,000 words of transcript and notes across the sixteen interviews.

2.4 Data Analysis

Interview recordings were transcribed and, together with interview notes, subjected to thematic analysis. Thematic analysis was conducted following the framework described by Braun and Clarke [36], proceeding through the stages of familiarisation with the data, generation of initial codes, identification and review of themes, and interpretation of findings in relation to the research questions. The six phases of Braun and Clarke’s framework were followed iteratively rather than strictly linearly, with movement back to earlier phases where later phases revealed the need for refinement.

The analysis followed a hybrid deductive and inductive coding strategy. The three top-level thematic areas were defined a priori on the basis of the interview guide and correspond to its three thematic sections: the current state of stroke care in Rwanda, the conduct of clinical trials in Rwanda, and trial-planning considerations specific to the MD100. Sub-themes within each top-level area were generated inductively through open coding of the corpus. The first analytical pass produced per-interview summaries under the three thematic areas together with an initial codebook of approximately 117 inductively generated codes, which were subsequently grouped into the sub-themes reported in the Results chapter.

Coding was conducted in parallel by both researchers in most cases, with regular comparison and reconciliation of codes; in a smaller number of cases, an initial pass by one researcher was reviewed by the other. Disagreements were resolved by discussion rather than by formal inter-rater reliability calculation, in line with the reflexive variant of thematic analysis adopted here, which treats coder subjectivity as an analytical resource rather than a source of error. Three cross-cutting findings that traverse the deductive structure were identified during the theme-review phase and are presented as a separate section of the Results chapter.

Findings from individual interviews were synthesised across respondents and triangulated with the literature review to identify areas of convergence and divergence between stakeholder perspectives and the documentary evidence base.

Quantitative figures volunteered by respondents during interviews, for example, estimated proportions of late presentations or estimated numbers of thrombolysis cases, are reported throughout the Results chapter as respondent estimates rather than as measured data, and are not treated as audit-quality observations.

Not every respondent addressed every sub-theme. A respondent's account is reported on a given sub-theme only where they spoke to it from direct professional experience; sub-themes on which a respondent did not speak, or on which they declined to give a substantive answer, are not attributed to them. The weight of any sub-theme in the analysis therefore reflects the number and roles of respondents who could speak to it from their professional vantage, rather than the total number of respondents in the sample.

2.5 Triangulation with the Prior Interview Corpus

A prior MSc thesis on the MD100 in Rwanda, completed approximately twelve months before the present fieldwork, was available to the research team together with the underlying interview transcripts [15]. That thesis was framed by the EUnetHTA Core Model for Health Technology Assessment rather than by a trial-planning framework, and its analytical structure therefore differs from the present study. Its interview guide and respondent selection partially overlap with the present study but do not coincide, since the questions it pursued reflect the HTA Core Model rather than the trial-planning structure adopted here.

Two distinct uses were made of this prior material. First, the written thesis was treated as a secondary source within the literature review described in Section 2.2, on the same footing as published literature on Rwandan stroke care and clinical trial conduct. Findings from the prior thesis informed the framing of the interview guide and the institutional sampling frame of the present study. Second, the prior interview transcripts were available during the analysis phase and were consulted on an ad hoc basis at a small number of points where a finding from the present interviews invited direct comparison. This was an opportunistic rather than a systematic use of the prior corpus: no attempt was made to compare the two datasets point by point, and the comparison was made only where a specific present-study finding happened to have a clear counterpart in the earlier transcripts. Where this selective triangulation illuminates or qualifies a finding, the relevant comparison is recorded in a footnote at the point of the finding in the Results chapter.

A systematic side-by-side re-analysis of the prior interview corpus under the present trial-planning framework was outside the scope of this thesis. Two reasons motivate this scope decision. First, the prior interviews were conducted under a different

interview guide whose questions reflect the HTA Core Model rather than a trial-planning structure, and a systematic re-analysis would therefore require treating the prior corpus as a partially redirected dataset rather than as a directly comparable one. Second, the analytical workload of re-coding the prior corpus from a different framework was not compatible with the project timeline. The present study's relationship to the prior thesis is therefore best characterised as occasional corroboration recorded in footnotes, rather than as triangulation in any systematic sense. The handful of such comparisons reported in the Results chapter should be read as illustrative points of contact between the two corpora, not as a comprehensive reconciliation of them.

2.6 Ethical Considerations

Ethical approval for this study was obtained from the CMHS Institutional Review Board at the University of Rwanda prior to the commencement of fieldwork. All procedures were conducted in accordance with the ethical principles of the Declaration of Helsinki [20] and the requirements set out in the approved IRB protocol. Both researchers held current ICH-GCP certificates compliant with the ICH-GCP R3 update at the time of fieldwork.

Informed consent was obtained verbally from each respondent before the start of every interview. The consent script, derived from the consent form approved as part of the IRB submission, set out the study purpose, the voluntary nature of participation, the right to withdraw at any time without consequence, the procedures for data storage and confidentiality, and the option to consent or decline to audio recording separately from consent to participate. Verbal consent was chosen as appropriate to the professional and non-clinical character of the interviews and to the absence of patient data, in line with accepted practice in qualitative research with professional respondents. Consent was documented in the interview notes, including the respondent's separate decision on audio recording. Where a respondent declined to be recorded, written minutes were taken in place of audio.

Interview data were stored securely and accessed only by the research team. No patient data were collected at any stage of the study.

2.7 Reflexivity and Methodological Limitations

Several methodological limitations apply to this study and are reported here together with the reflexive considerations that bear on their interpretation.

The interviewer team consisted of two MSc students from Chalmers University of Technology, one with a Biomedical Engineering background and one with a Management and Economics of Innovation background. Neither is a clinician, neither is a regulator, and neither is Rwandan. This positionality shaped the conduct of interviews in identifiable ways. Questions in the technical and innovation-management

vocabulary of the interviewers came more easily than questions in the clinical-procedural vocabulary of the respondents, and questions about commercial and procurement dimensions of trial conduct were pursued with more confidence than questions about specific clinical protocols. Where respondents drew distinctions of clinical practice that were unfamiliar to the interviewers, follow-up was sometimes weaker than it could have been; the resulting gaps are visible in the corpus and are flagged in the Results chapter where they bear on a finding.

A more consequential reflexive consideration concerns the iterative clarification of trial-related terminology during fieldwork. The clearest instance is the term *sponsor*, which in the ICH-GCP sense denotes the party that takes responsibility for the initiation, management, and financing of a trial, but which was initially interpreted in a lay financial sense by several respondents. Once this divergence was recognised during the early interviews, the interviewers began clarifying the ICH-GCP meaning during the interview itself. Respondents who first interpreted the term in the lay sense were therefore, after clarification, answering a different question from the one they would have answered without the intervention. As a reflexive consequence, apparent convergence across the sample on procedural definitions reflects in part the interviewer's framing rather than a baseline professional consensus among Rwandan clinicians. This pattern is reported as a substantive finding in the Results chapter rather than treated as a methodological artefact to be controlled out, since it is itself diagnostic of the state of trial-related vocabulary in the sample and therefore of direct relevance to trial planning.

All interviews were conducted in English, which is the working language of higher education and the principal language of clinical practice at the institutions sampled. Respondents in the sample were fluent in English, but English is not the first language of any of the interviewers nor of most of the respondents. This is a source of subtle imprecision that should be acknowledged. The semi-structured format allowed clarification of formulation in real time where this was needed, but no formal back-translation procedure was applied to the transcripts.

Beyond these reflexive considerations, several scope and sampling limitations apply. The sample is a convenience selection within a purposive frame. All sixteen interviews were conducted in Kigali, predominantly with respondents drawn from the major tertiary institutions, one district hospital, one health centre, and the central national bodies in the capital. Health-centre and district-hospital perspectives are represented, but the sample is weighted toward tertiary-care vantages and is not a nationally representative description of stroke care or research practice. Findings throughout the Results chapter are framed accordingly.

The absence of a direct interview with the Rwanda FDA is a substantive limitation for a study oriented toward trial planning, since the Rwanda FDA is one of the two principal regulatory bodies in the trial-approval pathway. The regulatory perspective in this study is reconstructed from accounts given by respondents at the Ministry of Health, the Rwanda Biomedical Centre, and hospital research departments who interface with the Rwanda FDA in their professional capacity, rather

than from the Rwanda FDA directly. The absence of an industry-sponsor respondent with experience of a device trial in Rwanda limits the present study's insights on the practicalities of sponsor-side trial conduct. Partial insight was however gained from the active site PI and trial-consultancy respondent. The absence of a radiologist among respondents leaves a third gap, since radiology is a central reference function for any imaging-comparator trial; the present study reconstructs the radiology workflow from accounts given by clinical and research-leadership respondents who interface with it.

A fourth sampling gap concerns the institutional locus of AI-specific regulatory expertise. R13a identified the Research Innovation Department, Data Science Division at the Rwanda Biomedical Centre as the appropriate point of contact for AI-regulation questions; no respondent from this division was interviewed. The AI-regulation findings reported in Sections 3.4.5 and 3.5 therefore reflect what respondents at clinical and research-leadership level knew about AI regulation, rather than what is in fact regulated.

All four gaps are taken forward into the trial-planning recommendations developed in the Results and Conclusion chapters.

3

Results

This chapter presents the findings from the thematic analysis of sixteen semi-structured interviews conducted in Kigali between 17 March and 21 April 2026 with health-care professionals, hospital research leaders, and regulatory and trial-consultancy stakeholders relevant to the planning of a clinical validation study of the Medfield StrokeFinder MD100 in the Rwandan healthcare system. The analysis followed the six-phase reflexive thematic analysis approach of Braun and Clarke, with a hybrid deductive and inductive coding strategy:

The three top-level thematic areas were defined a priori on the basis of the semi-structured interview guide (Section 2.3.2) and correspond to its three thematic sections: the current state of stroke care in Rwanda, the conduct of clinical trials in Rwanda, and trial-planning considerations specific to the MD100, while sub-themes were generated inductively from open coding of the corpus.

3.1 Scope and reflexivity.

Three methodological caveats apply to the findings that follow and should be borne in mind throughout. First, the sample comprises respondents drawn predominantly from Kigali-based tertiary referral hospitals (CHUK, KFH, RMH) and the central regulatory and research institutions in the capital. Health-centre and district-hospital perspectives are represented, but the corpus is weighted toward tertiary-care vantages and is not a nationally representative description of stroke care or research practice. Statements made about “stroke care”, “clinical trials in Rwanda”, or “the system” below should be read as referring to the institutions and care pathways represented in this sample, rather than as describing Rwandan healthcare in general. Section headings are retained as framing categories, but the empirical claims within each section should be understood within this scope.

Second, the iterative clarification of trial-related terminology by the interviewer during fieldwork (most notably the ICH-GCP sense of the term “sponsor”, as discussed in Section 3.5.2) shaped subsequent responses in several interviews. Respondents who first interpreted a term in a lay sense were, after the interviewer’s clarification, answering a different question from the one they would have answered without that intervention.

Third, the quotations presented throughout this chapter are illustrative rather than exhaustive. The source corpus comprises approximately 80,000 words, of which a small fraction is reproduced verbatim below. Quotes have been selected to convey the substance of a sub-theme as economically as possible, and the absence of a quote from any given respondent under a given sub-theme should not be read as evidence that the respondent agreed or disagreed with the sub-theme’s content.

The chapter is organised in three sections that correspond to the three top-level thematic areas. Section 3.2 presents findings on the current state of stroke care as represented in the sample, Section 3.3 findings on the conduct of clinical trials in Rwanda as described by respondents, and Section 3.4 the trial-planning findings specific to the MD100. A closing section (3.5) synthesises three cross-cutting findings that traverse the deductive structure. Selective comparison against the prior interview corpus [15] is flagged in footnotes where it supports or contrasts with the present findings; a systematic side-by-side re-analysis of that corpus was outside the scope of this thesis, as explained in the Methods chapter.

Respondents are anonymised with codes R1 to R16 and identified by professional role and institution; an overview is provided in Table 3.1. Quotations are presented verbatim with only minor disfluencies removed for readability; substantive content is unchanged.

Table 3.1: Overview of interview respondents. Roles and institutional affiliations are reported; respondent names are withheld. Interview formats are recorded as full audio-transcribed interviews (T), written minutes (W), or brief informational meetings (I). Paired interviews are denoted with suffixes a and b for the two respondents.

ID	Role	Institution	Date (2026)	Format
R1	Technical Advisor	Ministry of Health	17-03	T
R2	Senior ED Clinician	King Faisal Hospital (KFH)	21-03	T
R3	Emergency Physician	Rwanda Military Hospital (RMH)	23-03	T
R4	Emergency Physician	Kibagabaga District Hospital	24-03	T
R5	Emergency Physician (Site PI)	CHUK	26-03	T
R6	Nurse / In-charge	Kinyinya Health Center	30-03	W
R7	Director of Cardiovascular Disease	Rwanda Biomedical Centre (RBC)	02-04	T
R8	Neurosurgeon	CHUK	09-04	T
R9	Director of Quality Assurance and Research	Kibagabaga District Hospital	09-04	T
R10	Head of Research Department	CHUK	10-04	T
R11	Neurologist	King Faisal Hospital	16-04	T
R12a, R12b	ED Clinician (a), Neurosurgical Clinician (b)	CHUK	17-04	T
R13a, R13b	Director of Research (a), Research Officer (b)	King Faisal Hospital	17-04	T
R14	Senior Clinician (Neurology/Epilepsy)	King Faisal Hospital	18-04	W
R15	Director	CIICHIN (trial consultancy)	21-04	W
R16	Director General	Kibagabaga District Hospital	21-04	I

3.2 The current state of stroke care in Rwanda

This section reports respondents’ accounts of stroke care as it is delivered, in their experience, within the institutions represented in the sample. The findings below should be read as describing care pathways at three Kigali tertiary referral hospitals (CHUK, KFH, RMH), one district hospital in the same urban area (Kibagabaga),




Figure 3.1: Thematic map of the analysis. The three deductive thematic areas defined a priori from the interview guide (top row, solid borders) each contain inductively generated sub-themes (white boxes). Three cross-cutting findings (dashed borders) traverse multiple deductive areas and are addressed in Section 3.5.

one health centre (Kinyinya), and the national policy and research institutions (Ministry of Health, RBC). They are not a nationally representative epidemiological description as all sites are clustered geographically and thus the local population has a quicker access to tertiary hospitals than the people living in the more rural parts of Rwanda.

Within that scope, respondents’ accounts can be organised under six inductively generated sub-themes: late presentation as the dominant clinical reality (3.2.1), the geographic centralisation of diagnostic imaging in the four national referral hospitals (3.2.2), the fragmented multi-tier referral chain (3.2.3), the thrombolysis bottleneck at King Faisal Hospital (3.2.4), coordination failures within and between hospitals (3.2.5), and the active national-level innovation initiatives aimed at relaxing these constraints (3.2.6).

3.2.1 Late presentation as the dominant clinical reality

The most consistent observation across the sampled clinicians is that, in their reported clinical experience, the great majority of stroke patients arrive at a treatment facility after the 4.5-hour window within which intravenous thrombolysis is conventionally recommended. The figures volunteered by individual respondents are clinical impressions and self-reported facility estimates rather than measured data, and should be read as such. Within these limits, the estimates traced a clear gradient: at KFH, the only facility in the sample currently capable of administering thrombolysis, approximately half of presenting stroke patients were reported to arrive within window; at a tertiary referral hospital that receives transfers from across

the country, this share fell, in the respondent's estimate, to a few per cent.

R2 (Senior ED Clinician, KFH) estimated that "50/50, 50% came on time, 50% they come late. When I say late it's like above 4 hours," and reported an average arrival time of "between 3 to 7 hours." R3 (Emergency Physician, RMH) gave a higher estimate of late presentations:

R3 estimated that close to 95% of patients presented outside the therapeutic window, attributing this primarily to the time absorbed by road-based transfer — noting that an 80-kilometre journey from a district hospital alone could consume the four-hour threshold. (R3, Emergency Physician, RMH)

R4 (Emergency Physician, Kibagabaga District Hospital) estimated that it is 50% of stroke patients that arrive within the 4 hour window, the same as R4 but at district-level meaning it is earlier in the referral chain. R5 (Emergency Physician and Site PI, CHUK) gave the most categorical estimate of those interviewed, putting late presentations at "almost everyone, like 99%." R7 (Director of Cardiovascular Disease, RBC), drawing on a national policy and programme vantage rather than a single-facility experience, offered the figure of "more than 70%." These estimates are not directly comparable, as they reflect different denominators, different inclusion practices for "suspected stroke," and the different referral positions of the facilities concerned, but they are internally coherent: respondents working further along the referral chain reported higher proportions of late presentations.¹

Among the respondents who addressed the drivers of late arrival, several overlapping causes were named. R3 distinguished an individual-level cause, namely the tendency of patients to wait for symptoms to resolve, from a system-level cause, namely the time absorbed by the multi-tier referral process. R11 (Neurologist, KFH) located the principal cause in awareness, both lay and professional:

The main problem is patients come very late. That is like the main problem. If they come early, you avoid a lot of complication. Hospital stay will be reduced. Even functionality will be much better. (R11, Neurologist, KFH)

R7 offered a similar diagnosis in compact form: "awareness, late consultation due to lack of awareness." R5 explicitly contrasted Rwanda with high-income contexts in which public-awareness campaigns have shifted lay help-seeking behaviour, citing the FAST campaign in the United States as a comparable initiative.

A second-order observation was raised by R3: even when patients arrive well outside the therapeutic window, the clinical value of confirming the diagnosis remained substantial in his account. Confirmation differentiates ischaemic stroke from haemorrhagic stroke, from transient ischaemic attack, and from stroke mimics, and it

¹The prior interview corpus [15], conducted some twelve months before the present fieldwork with partly overlapping respondents, reported similar quantitative estimates for late presentation in the tertiary settings covered; no contradictory estimate has been identified in selective comparison with that corpus.

allows the clinician to plan family education and longer-term management (R3 cited the three-week window for aneurysm surgery in selected haemorrhagic cases). For R3, this defended the routine use of CT imaging in late-presenting patients against the interviewer’s challenge that it was effectively too late to act on the result.

The respondents were almost unanimous that a large majority of stroke patients present after the therapeutic window for thrombolysis. The two outliers were R2 and R4, who reported a 50/50 split at KFH ED and Kibagabaga ED respectively. R2’s account was, however, contradicted by R11, who reported that only 5% of patients received thrombolysis (see Section 3.2.4 for further discussion). This finding is consistent with a study on stroke management and outcomes conducted at KFH and CHUK by Nkusi et al., which reported that no patients presented to the emergency department within the 4-hour window [37]. In contrast to R2, R4’s account reflects the context of a district hospital, where patients must first pass through a referral process before being transferred to tertiary hospitals such as KFH.

The causes identified were primarily delays within the referral process and low public awareness regarding stroke symptoms and the urgency of seeking care. While public awareness lies beyond the scope of the MD100 itself, several respondents suggested that there is potential for the device to reduce diagnostic and referral delays within the health system. Importantly, respondents also argued that the clinical value of stroke confirmation is not exhausted by thrombolysis eligibility. As R3 emphasised, patients presenting outside the therapeutic window still require differentiation between ischaemic stroke, haemorrhagic stroke, transient ischaemic attack, and stroke mimics in order to guide ongoing management and counselling. This suggests that, in the Rwandan context, the value of an early-triage device such as the MD100 may extend beyond enabling thrombolysis alone and include acceleration of diagnosis, referral prioritisation, and post-acute clinical management.

3.2.2 Geographic centralisation of diagnostic imaging

The structural correlate of late presentation, as described by respondents working at multiple tiers of care, is the geographic concentration of confirmatory imaging in the four national referral hospitals (CHUK, CHUB, RMH, KFH), with no CT availability at district level and below. Three of these four tertiary hospitals are located in Kigali, while CHUB is located in Butare in the south of the country. R7 stated the matter directly: “the referral hospitals [have it], but the district hospitals, no, they’re not capable to do all that. But they can clinically suspect.” R3 was more emphatic in his recollection: “I don’t know any district hospital with a CT scan.” MRI was reported as more concentrated still, available exclusively at the four tertiary facilities in Kigali, with R8 (Neurosurgeon, CHUK) noting that CHUB had been “in the process” of acquiring one at the time of the interview.²

²Subsequent interviews, including R12a (CHUK, 17 April 2026), confirmed that CHUK itself had acquired an MRI scanner within the preceding year. This represents a contemporary change from the situation reflected in the prior interview corpus [15], where CHUK was described as lacking on-site MRI.

The consequences of this geography are felt most sharply at the health-centre tier, where R6 (Nurse, Kinyinya Health Center) described the limits of the diagnostic toolkit available to a nurse-only facility: vital signs and clinical inspection, with ECG not permitted at health-centre tier and certain examinations excluded from insurance coverage at health-centre level but covered at district level. The result, in R6's account, was a permanent diagnostic ceiling: any patient with suspected stroke could only be stabilised, briefly assessed, and forwarded.

Within Kigali itself despite having three hospitals with confirmatory imaging, there are still queues to access it. R12a reported that CHUK's ED had, at the time of the interview, "70 patients on a bed capacity of 40 beds," reflecting a level of overcrowding that propagated through triage and into the diagnostic queue.

R7's framing of the diagnostic question was the most general within the sample: in the Rwandan referral hierarchy, district hospitals can "clinically suspect" but cannot "confirm."

From this, it can be concluded that access to confirmatory imaging is unequally distributed across the country. Lower-tier facilities are limited to clinically suspecting stroke, while definitive imaging confirmation is geographically concentrated within the four tertiary hospitals and in only two cities. As a result, stroke diagnosis becomes transfer-dependent, delaying definitive diagnosis and placing additional pressure on tertiary emergency departments and diagnostic queues. In this context, there may be value in devices such as the MD100 that support early triage and facilitate earlier referral escalation.

3.2.3 The fragmented referral chain

A second mechanism behind late presentation reported by respondents is the multi-tier referral structure formalised by the community-based health insurance scheme. Patients are expected to enter the system at the lowest appropriate tier and progress upward as clinical need dictates: home, community health worker or health post, health centre, district hospital and finally tertiary hospital. R3 described the consequences for a stroke patient:

People have to come from a community-based health centre. It's a kind of a health system or a hospital which is nearby in the community, the basic. Then they transfer you, if they find the case is serious; they call ambulance from the district. Then the district hospital transfer you to us. And because people also awareness about stroke or whatever, I don't think it's that much high. Sometimes you find the time is lost during those processes. (R3, Emergency Physician, RMH)

R4 (Emergency Physician, Kibagabaga) gave a structurally similar account from a district-hospital vantage point, and added: the time absorbed by transfer was, in his estimation, less a function of geographic distance than of ambulance availability and the availability of staff to accompany the patient during transfer. The CT scan

to which a Kibagabaga emergency-room patient would typically be transferred was a twenty-minute drive away, in his account, yet the realistic time from suspicion to imaging was reported as “one to two hours, because there are many other factors like finding ambulance, finding the staff to take that person there.”

At the health-centre tier the referral chain has a second issue described in the sample: it is not closed. R6 reported that the Kinyinya Health Center transferred all suspected stroke patients to Kibagabaga District Hospital but received no systematic feedback on whether the suspicion was confirmed. From the referring tier’s vantage this represents a broken quality-improvement loop with educational consequences: the health-centre staff, in R6’s account, could not learn from their own referral pattern because they were never told what had been found.

In Kigali itself, R5 described how fragmentation also occurs at the tertiary level. At CHUK, the pathway to thrombolysis required a formal request to King Faisal Hospital for transfer approval, which itself “can take half an hour or one hour,” in addition to the CT process and the journey time. R5 estimated the cumulative time from symptom onset to thrombolysis administration at approximately eight hours for a patient first presenting at CHUK. This is her clinical impression rather than a measured time series. On her account, it is operationally equivalent to inaccessibility for the great majority of cases.

The referral system described by respondents introduces delays through multiple interacting mechanisms, including sequential tiering, logistical constraints in the availability of ambulances and staff, and formal approval procedures between facilities. These delays are cumulative and can significantly extend time to definitive diagnosis and treatment for acute conditions such as stroke. While much of this fragmentation is structural, the availability of additional clinical information at the first point of contact may influence referral decision-making and escalation speed.

3.2.4 Thrombolysis bottleneck

King Faisal Hospital was identified by all respondents who discussed thrombolysis as the only facility in Rwanda administering intravenous thrombolysis. R5 stated: “King Faisal is the only hospital where thrombolysis can be done.” R3 confirmed that the Rwanda Military Hospital referred eligible patients to KFH for that purpose; R12a confirmed the same pathway from CHUK. The clinical implication described by respondents is severe: a patient who first presents at any tertiary hospital other than KFH must either be transferred (with the time penalty just described) or treated supportively.

R11, the only neurologist in the sample and the clinician most directly responsible for thrombolysis administration, provided insight to the number of thrombolysis reperformed. In her recollection, fewer than twenty patients had received thrombolysis at KFH over approximately two years of practice, against an annual stroke admission rate she estimated at approximately 200 to 250 (a treatment rate she estimated at approximately five per cent). The infrequency of administration was,

in her account, primarily due to late arrival rather than issues of clinical eligibility. These figures are clinical estimates and should not be read as audit-quality data.

A further constraint identified by R3 was the cost of the drug. He quoted weight-adjusted doses of tenecteplase or alteplase at approximately 1.5 million Rwandan francs (RWF), circa 9,500 SEK, per ampoule, with most patients requiring two or three ampoules. One ampoule correspond to roughly twice the average monthly wage (690,000 RWF)[38]. R3 further emphasised this point:

One vial or one ampoule of the treatment is 1.5 million francs. Most people need 2 or 3, so you understand it's a little bit expensive. (R3, Emergency Physician, RMH)

He noted that community-based health insurance covered part of this cost in some cases, but that in his clinical experience the practical financial burden remained a barrier for many patients and a constraint on routine use even when patients arrived within window.

R11 described two active initiatives reported to her as aimed at relaxing this single-centre constraint. The first was a switch from alteplase to tenecteplase, which is administered as a single intravenous bolus rather than as a one-hour infusion and is therefore more practical for delivery at peripheral facilities. The second was the use of Rwanda's medical-cargo zip-line network (originally established in 2016 for blood-product delivery [39]) to dispatch thrombolytic medication from the central pharmacy to peripheral hospitals on demand. The drone network, operated by Zipline in partnership with the Government of Rwanda, has been documented to reduce blood-product delivery times across approximately twenty health facilities outside Kigali [39]. At the time of the interview, in her account, the network had been used once for this purpose, to deliver alteplase to a patient with pulmonary embolism rather than stroke; she mentioned that KFH is presently drafting a national stroke protocol that would extend this practice to ischaemic stroke.

The extremely limited access to thrombolysis, restricted to KFH, further intensifies the bottleneck described earlier in relation to the centralisation of diagnostic imaging within the stroke care pathway. Even patients who receive a timely diagnosis at other tertiary hospitals require an additional transfer to KFH, which further increases the risk of delays that push them beyond the therapeutic window despite being clinically eligible for treatment.

The high cost of thrombolytic agents further limits access for patients who present within the appropriate timeframe. Together, these constraints mean that thrombolysis is very rarely administered, rather than functioning as a routine treatment for ischaemic stroke.

3.2.5 Coordination failures within and between hospitals

A fifth sub-theme, which surfaced most clearly among respondents working in tertiary emergency departments, was the loss of time and clinical certainty caused by recurrent coordination failures. Across these accounts, four recurring mechanisms were identified: failed phone calls, payment-first procedural rules, theatre-availability constraints, and the absence of a stroke-specific code or protocol at receiving facilities. The following accounts illustrate each in turn.

R2 described the multidisciplinary coordination cascade at KFH in the following terms:

When you are managing those multiple calls, I may call neurologist, he's not picking the phone, but I take the risk to take the patient there. And then when the CT is done, maybe the neurologist is not also picking the phone, but at that time it's clear that it's a stroke. (R2, Senior ED Clinician, KFH)

R2 described this as a routine workaround rather than an exception, and identified the absence of a reliable bedside test as the most plausible mitigation in his own setting: "if you have a reliable test, that may be good."

At CHUK, R5 identified a different mechanism of delay: the requirement that imaging be paid for at the point of request introduced a procedural delay between clinical suspicion and CT scan, even when the patient was clinically eligible for thrombolysis. In her account, the hospital's social services unit could intervene for patients unable to pay, and emergency cases could escalate around the queue, but the default rule operated against acute pathways. R12a confirmed a similar pattern, noting that "100 patients, I mean caretakers, waiting to pay" could be present at any given time in his experience, although emergency escalation through the head of clinical division and the head of radiology could compress the path to imaging to under thirty minutes when the case was caught early enough.

R12b added a downstream constraint relevant for haemorrhagic stroke: even when surgery was indicated, operating-theatre availability was itself frequently the next bottleneck in his experience. For ischaemic stroke arriving within window, the cumulative procedural friction (payment, multidisciplinary call, CT queue, transfer approval) meant that, in R12a's account, the very small number of patients who in principle qualified for thrombolysis at non-KFH facilities were in practice rarely treated.

The pattern across these accounts is that the coordination problem, as respondents described it, is not located in any single tier or facility but at the interfaces between them.

3.2.6 Active innovation: drones, zip-lines, and district-level investment

Set against the structural delays described above, several respondents independently described active or planned national-level initiatives aimed at relaxing the identified constraints. R1 (Technical Advisor, MoH) described a Ministry of Health plan to use unmanned aerial vehicles for thrombolytic delivery:

There is a medication for stroke, the thrombolytics, that is not available everywhere. So we want, for patients suspected of cases, to deliver the medication on drones, so that we don't delay. (R1, Technical Advisor, MoH)

R11 described the more developed zip-line initiative referenced in Section 3.2.4, and reported on staff training already in progress in her account: residents at two KFH-affiliated teaching hospitals had received training on thrombolysis administration from visiting nurses, and the planned switch to tenecteplase was intended in part to lower the operational threshold for non-tertiary administration. R2 described a third initiative: the procurement of CT scanners for district hospitals. In his words: "Now they are buying CT scans for like district hospitals to improve that one."

R2 also offered a generalisation about how innovation is adopted by staff in the institutions he has worked in, once initial scepticism is overcome. Drawing on the introduction of bedside i-STAT blood gas analysers and Butterfly handheld ultrasound, he observed that staff tend to be resistant to new tools at first, but become enthusiastic advocates once they see positive results in practice:

When something's new, they may like have easy test, but when it is showing like results, good results, they may be like very active to use. When they see the results, they're the one to request. (R2, Senior ED Clinician, KFH)

A counter-example was offered by R3, who described an extended cautionary case: a vein-localisation device imported from China and adopted in his hospital had, in his account, failed to perform as advertised, and the procurement had been a loss because no local validation had been undertaken before purchase. The lesson R3 drew from this episode was that local validation prior to procurement is a precondition for sustainable adoption. This point recurred independently in the trial-planning section of his interview and has direct implications for the design of any clinical trial of the MD100 (see Section 3.6).

3.3 Conducting clinical trials in Rwanda

This section reports respondents' accounts of clinical trial conduct in Rwanda. The findings reflect the perspectives of one ministry advisor, one national programme director, two hospital research-department leaders (CHUK and KFH), one district-hospital quality-and-research lead, one trial-consultancy director, one active site

principal investigator, and a number of clinicians who have participated in trials in supporting roles. The findings are not, and should not be read as, a comprehensive description of the Rwandan trial-regulatory environment as a whole.

Within that scope, respondents' accounts can be organised under six inductively generated sub-themes: the regulatory architecture and the umbrella role of the Rwanda National Ethics Committee (3.3.1), the typical funding pathways and the boundary between hospital-based and private-institution-based trial phases (3.3.2), data governance and Rwanda's local-residency requirement (3.3.3), institutional readiness and the role of pre-engagement (3.3.4), and the typical points of failure named by respondents (3.3.5). A recurring observation about the operational vocabulary of trial conduct — most visibly the lay-financial reading of the term “sponsor” — first surfaces within the regulatory-architecture sub-theme below and is taken up as a cross-cutting finding in Section 3.5.

3.3.1 Regulatory architecture: RNEC, RFDA, and the umbrella IRB model

All respondents in this sample who addressed the regulatory question identified two central bodies: the Rwanda National Ethics Committee (RNEC), housed within the Ministry of Health, as the ethical clearance authority; and the Rwanda Food and Drug Authority (RFDA), as the authorising body for drug and medical-device trials. In respondents' descriptions, the two operate as a dual gate. R7 stated the requirement as: “You first get ethical clearance from the Rwanda ethics Committee. And then if it is something that includes medical equipment, we also need to see the approval from Rwanda FDA.” R15 (Director, CIICHIN) described the same dual requirement. This is also formalised in the RFDA's "GUIDELINES ON GOOD CLINICAL PRACTICES (GCP) IN RWANDA"(GCPR), that a trial application must first have received the ethical clearance from RNEC before the RFDA can authorize the trial[40].

R13a (Director of Research, KFH) gave the most procedurally detailed account in this sample of how local hospital Institutional Review Boards (IRBs) relate to Rwanda National Ethics Committee:

Every single IRB in hospitals has to be accredited by the National Ethics Committee. So in a way, whatever is approved at that level is also approved where we are because we use the same policies. The RNC is pretty much like an umbrella of IRB. (R13a, Director of Research, KFH)

In R13a's account, the local IRB does not re-approve work that RNEC has already approved; it verifies the documentation, registers the trial locally, and monitors implementation, reporting protocol deviations back to RNEC. R10 (Head of Research, CHUK) added a detail that did not surface in other interviews in this sample: in his recollection, the Rwanda FDA's final approval is conditional on local IRB approval. By his account the sequence is therefore not strictly parallel (RNEC and RFDA both review a submission); it is recursively gated: RFDA waits for the site IRB

before issuing the final letter.

The rationale R10 offered for local IRB involvement was clinical rather than procedural:

If a patient gets an allergy or something else, it is the hospital, the local carrier, which will manage such cases. Not the Ministry of Health. That's why the hospital must be directly involved in the monitoring and approving and monitoring such projects. (R10, Head of Research, CHUK)

Several respondents additionally identified operational requirements relevant for international trial planning. Respondents consistently stated that the principal investigator must either be Rwandan or, in the case of international research, paired with a local co-principal investigator. R4, R9, and R13a further specified that external sponsors require a local institutional collaborator. R13a also noted that trial personnel are expected to hold current Good Clinical Practice certification compliant with the updated ICH-GCP R3 guidance [22].

A recurring issue in respondents' use of clinical-trial terminology concerned the term "sponsor." In formal ICH-GCP usage, the sponsor is the individual, institution, or organisation responsible for the initiation and management of a trial rather than merely its financing. However, seven respondents initially interpreted the term primarily in financial terms before clarification during the interview. This suggests that some international clinical-trial terms are not used in by all clinicians and research administration.

Taken together, respondents' accounts describe a sequential approval chain for clinical trials involving, in order, RNEC, RFDA, and local IRB(s), where RNEC provides ethical clearance, RFDA authorises the trial, and local IRBs verify and monitor implementation at site level. While this general structure is broadly consistent with international clinical-trial governance frameworks, respondents differed in their descriptions of the precise sequencing between these bodies, particularly regarding the relationship between RFDA approval and local IRB review.

3.3.2 Funding pathways and trial site type by phase

Respondents who addressed funding distinguished sharply between two funding ecologies in Rwandan clinical research. The first comprises foreign-funded trials supported by the US National Institutes of Health, the UK government, the US Centers for Disease Control, and the pharmaceutical industry. R13a noted that the most active areas of foreign funding at KFH at the time of interview were dominated by viral infectious disease research. The second comprises locally funded "pragmatic" trials, typified by R13a's example of hand-washing-versus-no-hand-washing studies and surgical-site-infection comparisons: studies that do not require expensive consumables, complex devices, or international supply chains.

R13a was specific about the limits this distinction imposes on medtech trials:

Hardcore device or medications, it's usually external funding. Phase one, phase two is usually taken away from the hospital because the way hospitals here are built, they're built to cover just clinical operations and not research operations. (R13a, Director of Research, KFH)

The hospital-versus-private boundary R13a identified is, in his account, a structural feature of the Rwandan trial system that has direct planning implications for the MD100. Phase 1 and phase 2 trials are conducted at one or two specialised private institutions that maintain the infrastructure for early-phase work (cold-chain reagent storage at -100°C was the example R10 gave); phase 3 and phase 4 trials are conducted at the major hospitals. Because the MD100 has CE-marked status in Europe, its planned validation in Rwanda fits naturally into the phase 3 to phase 4 territory in their description, and is therefore feasible within the existing hospital infrastructure.

Internal funding for trials within the major hospitals was reported as structurally scarce. R10 stated that CHUK had funded only two small clinical trials from its internal small-grants programme to date. The National Commission for Science and Technology (NCST) was identified by R10 as an adjacent funding source for research more generally, but was not described as a typical trial funder.

R10 reported having helped negotiate the removal of a ministerial order that had limited grant inflows to two windows per year: April to June, and December. With the removal of the ministerial order, grants can be received year-round. In R10's view this eases the cash-flow position for financed trials but does not change the scarcity of internal grant capital.

R13a added a category of funder relevant for medtech specifically. While he was not aware of any dedicated medtech-trial funder in Rwanda, R13b (Research Officer, KFH) named the US-originated NGO Team Heart, which has supported cardiac surgery capacity-building since the early 2000s and has invested in follow-up apps and devices as an adjacent activity. Such NGO funders, R13a implied, are among the more plausible non-pharmaceutical sponsors for a medtech validation trial in his experience.

Respondents described clinical-trial funding in Rwanda as strongly dependent on both trial type and trial phase. Foreign funding, primarily from the United States and the United Kingdom, was reported to dominate higher-cost clinical research, particularly in infectious disease and pharmaceutical studies. In contrast, locally funded trials were described as smaller pragmatic studies that require limited infrastructure and consumables.

Respondents further suggested that infrastructure limitations shape where different phases of clinical trials can be conducted. Early-phase device and pharmaceutical trials were described as dependent on specialised private research facilities, while phase 3 and phase 4 studies were viewed as feasible within the existing tertiary-

hospital system. In this context, respondents considered a later-phase validation study of the MD100 to be operationally more feasible than an early-phase device trial within Rwanda.

3.3.3 Data governance and local data residency

Data governance in Rwandan clinical research was described by R13a as subject to a national data-protection law equivalent in scope to the European General Data Protection Regulation. R15 confirmed that, in her practice, a designated Data Protection Officer is required in the trial team. The finding with the clearest operational implications under this sub-theme was the reported enforcement of a local data-residency requirement.

R5, currently a Site PI on a five-hospital trial spanning three African countries, described a concrete instance from her own trial:

Because there is a data regulatory board, we have not yet been able to transfer the data to the global PI; they say data needs to be protected and until it felt safe to transfer the data. (R5, Emergency Physician and Site PI, CHUK)

In R5's trial, the Rwandan portion of the dataset, and in particular the imaging data, could not be exported to the international principal investigator at the time of the interview. Weekly summary reports were transmitted, but the raw data and imaging remained on local servers in her account.

R7 corroborated this from the regulator side and described the operational mechanism. In an instance R7 recalled, a sponsor had proposed to store trial data on a Swiss server; the proposal was rejected on data-residency grounds, and RBC subsequently provided the local server capacity for the data. The use of the REDCap data-capture platform, with per-user login and a clear audit trail, was reported by R5 as a standard practice in her trial.

The local IRBs add a parallel control on hard-copy trial records in R10's description. He reported the procedure at CHUK as follows: trial files are held in locked cabinets to which three separate parties hold keys (the hospital administration, the principal investigator, and a third party external to Rwanda), and the cabinets cannot be opened without the simultaneous presence of all three. He reported that hard-copy files from a clinical trial concluded two years earlier were still held under this protocol.

Taken together, respondents' accounts describe a clinical trial data governance system that is broadly consistent with international norms, including the use of designated data protection roles, audit trails, and institutional review. The most operationally salient feature, however, is the emphasis on local control of data storage and transfer, particularly with respect to raw data export and server location.

In regards to how the trial data is regulated and handled there was likely that there

was a misunderstanding between the respondent and interviewers. Where trial data was understood as equivalent to clinical data by the respondents, which is not the case, trial data would be clearly defined in the protocol and be a subsection of the total clinical data collected from patients during the trial and treatment. Specifically for the MD100 no data such as imaging, that R5 reported to be difficult to transfer out of Rwanda, would be necessary.

3.3.4 Hospital readiness and pre-engagement networks

Two findings under this sub-theme bear directly on the choice of trial sites for the MD100. The first is that, in R10's account, institutional readiness for hospital-based trials in Rwanda has been formally assessed at national level. He reported that the Ministry of Health, in collaboration with the Rwanda Biomedical Centre, had conducted readiness assessments of major hospitals some years earlier and that CHUK, KFH, RMH, and CHUB had been identified as the institutions "almost ready" for hosting clinical trials. These four institutions correspond to the four tertiary hospitals already identified by other respondents in Section 3.2.2 as the centres holding routine confirmatory imaging. This convergence simplifies the operational question of where a phase 3 to phase 4 trial of the MD100 can plausibly be conducted at tertiary level.

The second finding is that, in R10's and R13a's accounts, the operational route into a hospital trial site runs through pre-engagement rather than formal application. Both described the preferred practice as one in which the principal investigator first engages informally with potential collaborating institutions, identifies institutional collaborators, and develops the protocol jointly with them, before submitting the protocol to RNEC. R13a emphasised the consequences of skipping this step:

It's always best to think of the hospitals you're going to work with and think of the people you're going to work with before you head to RNEC. The other way around, [RNEC] will say, you didn't talk before. So you need to be very careful with how you select the hospitals. (R13a, Director of Research, KFH)

Respondents' accounts suggest that readiness for hosting clinical trials is concentrated in the four tertiary hospitals in the country. This is an important consideration when selecting trial sites for the MD100.

In addition, respondents emphasised that early informal engagement with potential trial sites is important for successful trial initiation.

3.3.5 Points of failure: monitoring, follow-up, and adoption

Two distinct failure points for clinical research and for medical-technology adoption were named across the interviews. The first, identified most clearly by R13a, concerns the conduct phase of the trial itself: "lack of monitoring, and lack of hands-on follow-up from the investigators. Many things here die out of lack of follow up." The

implication R13a drew was operational: if the principal investigator is not physically present at the site, the protocol must provide for paid, full-time site staff who report back on a defined cadence. Without that provision, in his account, the trial drifts.

The second point of failure concerns the introduction of a new medical technology after a trial concludes successfully. R10 located the risk in community engagement and adoption rather than in regulatory or technical obstacles: “it is about the level of understanding or mindset of people. Community engagement, I think it is the key, the key factor for success.” R10 contrasted device adoption with drug adoption on this point and observed that he himself could not recall a specific example of unsuccessful device adoption to draw on, because, in his experience, successful introduction had been the norm.

A more concrete point of failure, was brought up by R3 in the form of the vein-localisation device introduced earlier (Section 3.2.6). In R3’s account, the device had been purchased from a foreign supplier and marketed on the basis of evidence from elsewhere. It was used briefly before being abandoned, with the resulting procurement loss attributable to the absence of any pre-purchase local validation.

This anecdote was the strongest argument across the corpus for conducting the MD100 trial before any decision on introduction is taken. As R3 put it: “if this device is taken maybe to my hospital, then I try it. I see it is helping, then I would advocate for it to be purchased, then be introduced.”

3.4 Trial planning for the MD100

This section reports respondents’ suggestions and concerns relating specifically to the design of a clinical validation trial of the MD100 in Rwanda. As in the preceding sections, the findings reflect the perspectives present in the sample (predominantly Kigali tertiary clinicians and hospital research leaders, with one trial-consultancy and two national-institution perspectives), rather than a survey of Rwandan trial expertise at large.

Six inductively generated sub-themes are presented: the spectrum of endpoint preferences respondents identified as relevant for a Rwandan trial (3.4.1); the trade-off between selecting sites that maximise validation quality and those that maximise clinical impact (3.4.2); the trial-design choices respondents recommended (3.4.3); the operational specifics of patient inclusion and informed consent in the Rwandan context (3.4.4); AI-specific regulatory considerations (3.4.5); and the cluster of expanded use cases beyond stroke triage that respondents independently identified (3.4.6).

3.4.1 Trial endpoints

All respondents were clear on the regulatory need for local validation of the MD100 before it can receive market authorisation. Some also expressed that local validation

can serve to support better adoption by healthcare workers if the device has been tested and validated in Rwanda.

Respondents in the sample who then addressed endpoint preferences converged most clearly on diagnostic accuracy, specifically the sensitivity and specificity of the device. Eight respondents (R2, R3, R4, R5, R7, R10, R11, R14) named accuracy first when asked about endpoints. R4 framed the rationale in clinical-decision terms:

Stroke itself has like signs that already has raised the suspicion at a higher level. If you suspect a stroke, you can be like 50% with clinics, without doing any [test]. If someone comes with hemiplegia, aphasia, with the background of diabetes, hypertension, you can be 60% sure that there might be a stroke. So if the device can like raise this suspicion to above that, like making it closer to 100%, that would be a good thing. (R4, Emergency Physician, Kibagabaga)

There was no explicit mentioned desire to evaluate how the MD100 would affect and fit in the current stroke work flows, such as the study by Tsiftsis et. al did in Greek emergency departments [26].

R8 and R11 further contextualised the emphasis on diagnostic accuracy by highlighting perceived differences between the Rwandan and previously studied populations. R8 noted that stroke may present earlier in African populations and may differ in subtype distribution, while also pointing to possible influences of diet, climate, and genetic factors based on clinical experience; R11 similarly raised concerns regarding potential ethnicity-related performance bias in AI-based imaging tools.

Time-to-treatment endpoints were the second most frequently mentioned class. R5 nominated “the efficacy and the impact it would have on how rapid we do things” as a joint primary; R7 specifically named “shorten[ing] the time of diagnosis”; R11 specified two sub-components, namely door-to-imaging time and imaging-to-needle time, and proposed a before-and-after design to measure both:

If you have a clear timeline before utilisation of this [device], then after utilisation, if you see a significant difference, it could be a tool that can be used everywhere in [Rwanda]. (R11, Neurologist, KFH)

3.4.2 Site selection: the validation-versus-impact trade off

Four respondents independently raised a trade-off that shapes the question of trial site selection. The trade-off is between two design priorities. A trial sited at the major Kigali tertiary hospitals (CHUK, KFH, RMH) maximises validation quality, in their accounts, because the comparator CT is available on site, patient flow is high, and trial infrastructure already exists. The same trial sited at district hospitals or below maximises clinical impact, because the diagnostic gap that the device addresses (absence of confirmatory imaging at the point of first medical contact) is largest precisely where validation is hardest.

R13a framed the trade-off as:

Be as representative as you can. The population we have here is fortunate to come to a place where we can diagnose them pretty fast. We can give tPA for ischaemic stroke. We can [provide] neurosurgical services, but it's not the case for everyone in Rwanda. So if you want to come here, you're welcome, but the data you'll have won't be representative. If you really want to provide good impact, go to the most remote places. (R13a, Director of Research, KFH)

R5 endorsed a tertiary-plus-district design from a different vantage point: "definitely all the referral hospitals, and then some of the district hospital with a high workload of patients." R10 added the analytic argument that even at CHUK, where on-site CT is available, the device would have value because the report turnaround on a CT scan can reach an hour. R4, drawing on the district-hospital vantage, advocated district-first siting on the rationale that "it's where the suspicion is being done."

The boundary case that several respondents considered and then rejected was siting the trial in health centres or in the prehospital ambulance environment. R8 and R12a both agreed that ambulance deployment was the most promising long-term use case for the device, but R8 reframed that vision within the trial-versus-applicability distinction:

It's a trial, it makes sense, and the population in the hospital, in a busy hospital like this one, is an ideal place for trial. The ambulance and the health centre will be an extension after a trial. (R8, Neurosurgeon, CHUK)

The resolution that R13a proposed, and which is the most coherent across the sample, was a two-stage approach: a pilot at tertiary level to establish accuracy against a comparator. This would be followed by a formal trial at a representative cross-section of district hospitals (with imaging confirmation arranged by transfer when needed). The pilot-before-trial argument is taken up in the next sub-section.

The responses regarding site selection indicate that trial location is closely tied to the prioritised endpoint. Tertiary hospitals were generally viewed as more appropriate for initial validation, due to higher patient volumes, established research infrastructure, and access to CT imaging as a reference standard. In contrast, district hospitals were seen as more suitable for assessing clinical impact and real-world utility, as they better reflect the settings where diagnostic delays and limited imaging capacity are most pronounced. Overall, respondents framed site selection as a trade-off between methodological rigor in validation and relevance to routine clinical practice.

3.4.3 Trial design: pilot-before-trial, blinding, and crossover risk

R13a provided the most developed account in the sample of how a Rwandan trial of the MD100 should be designed. He offered three specific recommendations. The first was to pilot the device against CT/MRI before initiating a formal trial:

You probably do not have to use it as entirely as a trial. Pilot it. Pilot it first, see how it performs versus imaging. And when it does, you make your point by making a trial. (R13a, Director of Research, KFH)

The pilot phase, in his account, would establish the device's accuracy under Rwandan clinical conditions and supply the local-validation evidence that the trial proper would then build on, with substantially lower regulatory and operational risk.

The second recommendation is that the trial itself must be blinded. R13a's argument was grounded in his perception of the state of trial literacy in Rwanda:

The clinical trial sort of understanding here in Rwanda is not as solid. So what that means is, there could be a little bit of bias in over diagnosing, and even your controls, people would find it moral, OK, to go to that device in order to save them. So you'll have involuntary crossover. (R13a, Director of Research, KFH)

Single-blinding (of patients and of the clinician seeing them), with a third party (the principal investigator or designate) making the device-application decision, was R13a's minimum recommendation. Open-label conduct was framed as posing a serious risk to the validity of the data.

The third recommendation, as an alternative to within-site randomisation, was a case-site/control-site design: two comparable hospitals, one deploying the device and the other not, with comparison drawn across rather than within sites. The case-site/control-site design circumvents the involuntary-crossover problem by removing the device from the control site entirely.

R5 raised a complementary design point: an efficacy endpoint and a time-to-treatment endpoint could be tested in a single trial run, in her view, provided the trial period is long enough ("feasible in one [track], but maybe over a longer period of time"). R3 proposed a four to six month trial period across three or four sites; R4 proposed a six-month time-period sampling design or a sample-size calculation based on local stroke incidence; R5 endorsed a sample-size calculation explicitly. None of the respondents who addressed sample size offered a numeric estimate.

3.4.4 Patient inclusion, exclusion, and informed consent in the Rwandan context

Inclusion. The general direction across respondents who addressed the question is that the inclusion criterion should be permissive. Five respondents (R3, R4, R11,

R12a, R13a) recommended inclusion of any patient with sudden focal neurological deficit or suspicion of acute stroke. Three (R2, R4, R7) added the inclusion of patients presenting with severe headache plus established cardiovascular risk factors, particularly hypertension and diabetes. R12a added altered mental status without other apparent cause. R11 noted, in response to a direct interviewer question, that pregnancy was not a basis for exclusion in her view because the device emits no ionising radiation.

Exclusion. Exclusion was framed by most respondents who addressed the question as topic-dependent. R11 stated explicitly that she “[did]n’t see any reason” to exclude particular subgroups for safety reasons, given the non-invasive character of the device. R5 nominated patients with co-morbidities that would make thrombolysis contraindicated regardless of the device’s reading (for instance atrial fibrillation on anticoagulation) as a potential exclusion class on outcome grounds rather than safety grounds. R5 also identified two Rwandan-specific exclusion categories rooted in the country’s vulnerable-groups classification: prisoners (formally a vulnerable group in Rwanda, on her account) and unaccompanied unidentified patients. R10 reinforced the unaccompanied-patient exclusion: “That person should not be enrolled in a clinical trial,” if no next of kin is available. The age threshold that emerged across the sample varied (R7 proposed ≥ 40 , R4 proposed ≥ 50 , R8 noted strokes occurring even in neonates but accepted the device’s manufacturer-specified ≥ 18 threshold), and no consensus on age cut-off was identified.

Informed consent: typical duration and format. The procedural baseline reported by R2, R4, R5, and R11 is that consent for a non-invasive imaging-related intervention can be obtained in approximately five minutes. R11 reported a duration “not more than 2 minutes” for the analogous thrombolysis consent obtained at KFH. R12a reported, by contrast, a longer duration for trial-specific consent, up to ten minutes in his experience, with some families requesting overnight time to consider. The five-minute baseline reported by the sample therefore applies to the underlying device application but should not, on these accounts, be assumed for the trial-specific consent.

Consent: literacy and cultural adaptations. Two adaptations to standard written-consent protocols are well-established in the Rwandan practice described by respondents. The first is the use of a thumbprint signature in place of a written signature for illiterate participants; R14 and R15 both confirmed this. The second is the witness model described by R10:

There are some scenarios where we need a witness, especially because maybe our context differs from European ones. A patient can be conscious but doesn’t know to read. There will be a witness who will approve that this patient, they read to him what is the content of the consent form, and he understood. (R10, Head of Research, CHUK)

R4 added that the consent form should be available in both English and Kinyarwanda, with verbal reading aloud for participants who cannot read. The use of

thumbprints as signatures and the use of a witness are both in the GCPR as means to receive informed consent from illiterate participants[40].

Consent: unconscious and cognitively impaired patients. Nine respondents (R2, R3, R4, R5, R9, R10, R11, R12a, R15) addressed the question of how consent is obtained from patients who are unconscious or cognitively impaired at the point of clinical contact. The pattern across these accounts was that consent is obtained from the next of kin or a primary caretaker who accompanies the patient; the document is then referred to as an “assent form” (R9, R10) rather than a consent form proper. Where no next of kin is available, on these accounts, the patient may not be enrolled in a research protocol, although clinical care may proceed. In the GCPR it says that if a trial patient is unconscious or mentally disabled a “legally acceptable representative” may sign the informed consent[40].

A culturally specific finding raised by R3 supplements this baseline. In the Rwandan context, in his account, the patient-clinician relationship is structured by a strong default trust in the clinician’s recommendation:

In our cultural context people, they trust, they have hope in a doctor,
“Anything you think can be benefit for the patient, in your knowledge,
please we trust you.” (R3, Emergency Physician, RMH)

R11 referred to a complementary provision in the current American Heart Association guideline: in the absence of a caretaker, and where the benefit of the intervention is clear, the treating team may proceed on the basis of clinical judgement, on her reading of the guideline.

The respondents’ accounts regarding inclusion and exclusion suggest support for relatively broad inclusion criteria in a future trial. Because the MD100 is non-invasive and does not emit ionising radiation, most respondents did not identify strong safety-based reasons for exclusion. Instead, proposed exclusions were primarily related to ethical or methodological concerns, such as inability to obtain proxy consent or co-morbidities that would confound assessment of treatment outcomes.

Taken together, the responses regarding informed consent indicate that respondents generally perceived consent procedures for the MD100 as feasible within the Rwandan clinical context. Participants described existing adaptations for low-literacy populations, including verbal explanation, witness verification, and thumbprint signatures, as already established practices within clinical research. At the same time, the repeated emphasis on next-of-kin assent for unconscious patients and the culturally embedded trust in clinicians suggests that consent in acute stroke settings may rely heavily on family decision-making and confidence in healthcare providers.

3.4.5 AI-specific regulatory considerations

Six respondents (R1, R4, R5, R9, R10, R11) reported, when asked directly, that they did not know whether the AI-driven classification component of the MD100

was subject to regulatory requirements beyond those applied to medical devices in general. This pattern is taken up as a cross-cutting finding in Section 3.5.

R13a was the only respondent in this sample who provided a substantive answer, locating the AI-specific regulatory question within the framework of Rwanda’s data-protection law and referring trial applicants to the Research Innovation Department and Data Science Division at the Rwanda Biomedical Centre as the appropriate institutional contact. He framed engagement with AI-specific regulation as a parallel track running alongside RNEC and RFDA submissions, rather than as part of either. No respondent from this division was interviewed; the AI-regulation findings reported here therefore reflect what clinical and research-leadership respondents knew about AI regulation, rather than the regulatory framework itself (see Section 2.7).

R5 mentioned that all public servants had recently completed an introductory course on AI, which suggests that the knowledge base may be being built at national policy level, but the specifically regulatory application of that knowledge was not established at clinical hospital level within this sample at the time of the interviews.

R9 and R10 each volunteered an interpretive principle that frames how the eventual regulatory position is likely to be specified. R9 stated that, in his view, “AI can’t do anything without your conscience” in clinical medicine, meaning that human clinical judgement must remain in the loop and that the AI output must be verifiable rather than determinative. R10 expressed a similar view:

Artificial intelligence is a tremendous tool to help us, as humans, but it wouldn’t, it shouldn’t replace humans. (R10, Head of Research, CHUK)

R11 added an instructive institutional precedent. KFH was, at the time of the interview, in active discussion to adopt RapidAI, an imaging-based stroke-detection AI tool with wider international currency. The fact that an AI-based decision-support tool was under active consideration for adoption at KFH suggests that the institutional posture toward AI-supported diagnostics in her institution is open rather than restrictive, although the regulatory pathway followed by RapidAI was not documented in the interview corpus.

The aggregate of the accounts regarding AI-specific regulation suggests limited awareness among clinicians and hospital research departments of regulatory requirements specific to AI-based medical devices. Respondents instead tended to frame AI tools within existing medical device and data protection structures, with R13a specifically identifying the Rwanda Biomedical Centre and its Research Innovation Department and Data Science Division as important stakeholders outside the standard RFDA and ethics approval pathways. These findings suggest that early engagement with institutions responsible for data governance and digital health oversight may be important during trial planning for AI-supported diagnostics. Despite the limited regulatory familiarity identified in the sample, respondents generally expressed a positive attitude toward AI-assisted healthcare tools, provided that clinical judgement remains central to decision-making.

3.4.6 Expanded use cases beyond stroke triage

A finding that emerged consistently, and largely independently, from the trial-design discussions concerned the range of clinical use cases that respondents in the sample believed the device could plausibly address beyond its primary stroke-triage indication. The cluster of use cases is presented here not as a recommendation for the present trial's scope, but as a record of the clinical demand that the device elicited from this group of respondents.

The most frequently cited expanded use case, raised independently by three respondents (R11, R8, R12a), was integration with the zip-line thrombolytic delivery architecture (Section 3.2.4). R11 set out the closed-loop logic: a peripheral facility uses the device to identify stroke, the central pharmacy dispatches thrombolytic medication via zip-line, and the local clinician administers it. R8 and R12a endorsed the same architecture from a CHUK vantage, with R8 specifically identifying the device's portability and battery operation as the features that made it operationally suitable for prehospital ambulance use post-trial.

The second cluster of use cases concerned non-stroke clinical conditions that the underlying microwave-imaging technology could in principle detect. R8 identified missed brain tumours, that is, patients who had been treated for epilepsy for years before a delayed CT confirmed an underlying tumour, as a candidate use case for resource-constrained settings, and chronic-headache evaluation in remote areas as another. R3 and R12b raised the possibility of intracranial-bleeding detection in trauma patients who are too unstable to be moved to imaging. R12b framed the use case as:

R12b identified critically unstable patients who could not be transferred to imaging as a particularly compelling use case, arguing that a bedside screening tool could provide diagnostic information that would otherwise remain permanently unavailable for this group.

R8 also raised the use case of transient ischaemic attack detection, but flagged it as uncertain: patients with TIA recover within twenty-four hours, so the time-window for measurement is short, and the validation case would be technically demanding in his view. The detection of silent stroke (clinically asymptomatic small infarcts) was raised by R8 as a further, more speculative possibility.

A third use case that emerged unprompted from R12 was institutional rather than clinical: trial participation as professional development. Both CHUK clinicians (R12a and R12b) explicitly stated that participation in the trial would have value as an opportunity to acquire research skills (data collection, analysis, manuscript preparation), and that the design of the trial should provide for this.

3.5 Cross-cutting findings

This section reports three patterns that traverse the deductive thematic areas and that cannot be placed within any single one of them. Each pattern is described in terms of where in the chapter the supporting observations appear; the implications of these patterns for trial design are taken up in the sections that follow (Sections 3.6–3.10).

3.5.1 Clinical need and validation difficulty appear at the same locations.

Two patterns in the data converge on the same set of geographic locations. The first comprises three sub-themes in Section 3.2: late presentation as the dominant clinical reality (3.2.1), the centralisation of confirmatory imaging in the four national referral hospitals (3.2.2), and the fragmented referral chain between care tiers (3.2.3). Together, these placed the district hospital and below as the point in the Rwandan stroke pathway where the diagnostic gap between clinical suspicion and confirmation is largest, as described by respondents. The second pattern, the validation-versus-impact trade-off reported in Section 3.4.2, placed the same locations as the methodologically hardest to validate a triage device at. The two patterns are reported in separate sections of the chapter, but, on respondents' accounts, point to overlapping geographic locations.

3.5.2 Trial-related terminology was variably shared across the sample.

A second pattern concerns the operational vocabulary of trial conduct. Seven of the sixteen respondents initially interpreted “sponsor” in its lay financial sense and arrived at the ICH-GCP procedural meaning only after the interviewer's clarification (Section 3.3.1), even though the GCPR defines the main trial stakeholder roles in line with the ICH-GCP [40]. The divergence fell more between operational trial leadership and clinical staff than between institutions. It matters for trial planning because the protocol-level meaning of terms such as “sponsor” has to be established jointly with collaborators before submission rather than assumed, which connects directly to the pre-engagement practice reported in Section 3.3.4.

3.5.3 AI-regulation knowledge was absent at clinician level and concentrated at one institutional actor.

A third pattern concerns the distribution of AI-regulation knowledge in the sample. Six respondents drawn from across the institutional landscape (R1 at the Ministry of Health, R4 and R9 at district level, R5 at tertiary level, R10 at hospital research-leadership level, and R11 at tertiary level) reported, when asked directly, that they did not know whether AI-based diagnostic devices were subject to regulatory requirements beyond those applied to medical devices in general (Section 3.4.5). One

respondent (R13a, KFH) provided a substantive answer, which pointed not to clinical respondents but to a specific institutional actor: the Research Innovation Department, Data Science Division, at the Rwanda Biomedical Centre. The pattern in the sample is therefore a contrast between a broad clinician-level knowledge gap and a single concentrated locus of expertise, which has a direct planning consequence: identifying and engaging that institutional actor early is necessary precisely because the knowledge required to do so is not held at the clinical sites where the trial would run.

The sections that follow discuss them with reference to the wider literature on clinical-trial planning in resource-constrained settings and to the prior thesis work of Alshaqouq [15].

3.6 The case for local validation

The MD100 arrives in Rwanda with a CE mark and a body of European and Australian evidence behind it. The first question the findings force is therefore not how to validate the device but whether a local validation is warranted at all, given that prior evidence. Respondents answered consistently, on two independent grounds, and a third follows from the regulatory and adoption pathway.

The first ground is clinical validity. The device’s earlier validation was conducted against populations and care pathways that differ from Rwanda’s on several axes respondents considered material. R8 and R11 noted earlier stroke onset in African populations, a different distribution of stroke subtypes, and possible dietary, climatic, and genetic influences, and R11 specifically flagged the risk of ethnicity-related performance bias in AI-based tools trained elsewhere (Section 3.4.1). For a device whose classifier was trained largely on non-African data, these are not abstract concerns: the signal patterns the algorithm learned to associate with stroke may not hold in a population whose stroke profile differs, and the epidemiology of stroke in Africa is documented to differ on exactly the dimensions respondents raised [3], [24]. Late presentation (Section 3.2.1) compounds the point, since a device evaluated mostly on early-presenting patients cannot be assumed to behave identically on the late-presenting majority that defines the Rwandan caseload.

The second ground is the more persuasive one in the corpus, and it is empirical rather than theoretical. The strongest argument any respondent made for local validation was not an appeal to representativeness but R3’s account of an imported vein-localisation device that was adopted on the strength of foreign evidence and then abandoned when it failed in local use (Sections 3.2.6 and 3.3.5). That episode does in one concrete procurement what the population-difference argument does only in principle: it shows what happens when evidence from elsewhere is treated as sufficient. R3 drew the explicit lesson, that local validation must precede procurement, and we adopt it as the organising rationale for the trial rather than as supporting colour.

These two grounds converge with a third. Local evidence is, on respondents' accounts, a practical precondition for both adoption and authorisation. R3's formulation, that he would advocate for a device's purchase only after seeing it help in his own setting, describes an adoption logic in which clinician confidence is earned through local demonstration rather than transferred from a foreign trial, and the regulator was reported to expect local data before full authorisation. The requirement for local validation is therefore doubly grounded, in clinical validity and in adoption legitimacy, and it is this double grounding, rather than the CE mark, that settles the prior question of whether a trial is warranted at all. The remaining question is how to design it, which the following sections take up.

3.7 Clinical need and where the device fits

Granting that a local trial is warranted, the next question is what clinical problem the device would address and where in the care pathway it would add value. The findings describe a need that is real and specific, but also one the device can only partly meet, and being clear about that boundary matters for setting realistic expectations of what a trial can demonstrate.

The dominant clinical reality is late presentation: the majority of stroke patients arrive after the window in which thrombolysis is conventionally given (Section 3.2.1), a pattern corroborated by the finding of Nkusi et al. that no patients in their Rwandan series reached the emergency department within four hours [37]. The drivers respondents named fall into two kinds, and the distinction is decisive for the device's relevance. The first is low public and professional awareness, a behavioural and systemic problem that no diagnostic device addresses. The second is structural delay, and it is here that a device might intervene: the stepwise referral chain (Section 3.2.3), the concentration of confirmatory imaging in four hospitals, three of them in Kigali (Section 3.2.2), and the confinement of thrombolysis to a single hospital (Section 3.2.4).

It is worth being precise about how these structural features produce delay, because they are easily but wrongly described as a feedback loop. They are not a loop; they are a cumulative chain in which each tier and each transfer adds latency that the next stage cannot recover from. A patient entering at a health centre may pass from health centre to district hospital to tertiary hospital, and, if thrombolysis is indicated, onward again to King Faisal, with R5 estimating roughly eight hours from symptom onset to thrombolysis for a patient first presenting at CHUK. Because the delays compound rather than substitute for one another, diagnosis becomes transfer-dependent, and the point of first contact, where the patient actually is early in the event, is precisely the point with the least diagnostic capability.

This is the gap the MD100 is positioned to address: it offers clinical information at the first point of contact, where confirmatory imaging is absent. Respondents agreed that earlier evidence of stroke at that point could enable the referral chain to be compressed, sending a patient toward tertiary care, or directly to King Faisal,

on the strength of a bedside result rather than after the full stepwise escalation. Crucially, the value of that earlier information is not confined to the minority who could still be treated acutely. As R3 argued (Section 3.2.1), confirming a diagnosis retains substantial value even outside the thrombolysis window, because differentiating ischaemic stroke, haemorrhage, transient ischaemic attack, and stroke mimics governs ongoing management and family counselling regardless of whether thrombolysis follows. The device's value proposition therefore has two components: faster escalation for the few who remain treatable, and earlier correct diagnosis for the many who do not. It is this dual value that makes time-to-diagnosis, rather than time-to-treatment, the endpoint that fits the Rwandan context, an argument developed in Section 3.8.

Two cautions bound the proposition. First, the device cannot resolve the awareness and transfer problems that cause most of the delay; it can reduce only the diagnostic component, and a trial should be scoped to measure that component rather than credited with system-level change it cannot deliver. Second, and more consequentially for the design, the clinical need is largest exactly where the device is hardest to validate: at district level and below, where confirmatory imaging is absent and where, for the same reason, a reference standard for a validation trial is least available (Section 3.5). That tension is not a reason to retreat to tertiary-only validation; it is the central design problem that the next section is built to resolve.

3.8 Designing the validation study

The findings do not point to a single trial but to a sequence of two. The device must answer two distinct questions, and the conditions under which each can be answered well are not the same. The first question is whether the MD100 classifies stroke accurately against a reference standard in a Rwandan population; the second is whether deploying it in routine care changes the speed at which patients reach a diagnosis. We treat these as two trials run in order, performance first and clinical value second, rather than as a single combined study. Three considerations from the findings drive that separation.

The first is the state of trial literacy described in the sample. R13a was explicit that the practical understanding of clinical trials in Rwanda is, in his experience, not yet solid (Section 3.4.3), and the variable familiarity with formal trial vocabulary across the sample points the same way (Section 3.5). This is not a local peculiarity but a recognised feature of the clinical-trial environment across much of sub-Saharan Africa, where device trials in particular remain rare and the supporting research infrastructure is thinly spread [41], [42], [43]. The design response to a setting with limited trial experience is to reduce the number of moving parts that any one trial asks site staff to manage, rather than to load both research questions onto a single complex protocol.

The second consideration is site readiness. As R10 reported, the four tertiary hospitals are the institutions that national readiness assessment identified as prepared

to host trials, and they are also the sites that hold the CT and MRI comparator (Section 3.3.4). A performance trial can therefore be sited where both the comparator and the trial capacity already exist. The clinical-value question cannot, because the gap the device is meant to close is largest precisely where validation is hardest. Running the simpler trial first lets the harder one inherit a trained cohort of sites and investigators rather than starting cold.

The third consideration is logical priority. If the device does not classify stroke accurately under Rwandan conditions, its effect on time-to-diagnosis is not worth measuring. Establishing performance first means the value trial is only mounted on a device that has already earned it, which is also the lower regulatory and operational risk path: the performance stage doubles as the local-validation evidence that R3's cautionary account (Section 3.3.5) identified as the precondition for any procurement decision.

3.8.1 The performance trial

The performance trial is the more straightforward of the two and the findings constrain it tightly. It is sited at the tertiary referral hospitals, where CT, and increasingly MRI, is available on site as the reference standard, where patient flow is high, and where trial infrastructure exists. The endpoint is diagnostic accuracy, reported as sensitivity and specificity against the imaging comparator; this was the endpoint respondents named first and most consistently, and it is the quantity on which a market-authorisation case ultimately rests. Each patient enters and leaves the study at one site, the comparison is internal to that site, and the outcome is binary, so the design asks little inter-site coordination of staff who are new to trial work. Beyond its primary purpose, this stage is the natural point at which to record stroke subtype from the reference scan alongside the device output, so that the resulting dataset can later support evaluation of capabilities the current algorithm does not yet offer, such as ischaemic–haemorrhagic differentiation.

3.8.2 Choosing an endpoint for the value trial

The value trial is harder, and the first difficulty is the endpoint itself. The two endpoints a validation study of this kind would ordinarily reach for, time-to-treatment and clinical outcome, are both compromised in the Rwandan context. Time-to-treatment is not informative where definitive treatment is confined to a single hospital, since for most patients the binding constraint is not the speed of diagnosis but the centralisation of thrombolysis at King Faisal (Section 3.2.4); and a clinical-outcome endpoint would require longitudinal follow-up that the system is not currently configured to support. We therefore take time-to-diagnosis as the most viable primary endpoint for the value trial.

That choice needs a defence it has not yet been given, because the obvious objection is that accelerating diagnosis is pointless if the treatment that diagnosis unlocks is centralised and frequently out of reach. The findings answer the objection in two steps. First, for the minority of patients who do present within the therapeutic

window, earlier diagnosis at or below district level is precisely what would allow the referral chain to be compressed or bypassed, sending a patient toward KFH on the strength of bedside evidence rather than after the full stepwise escalation (Section 3.2.3). Second, and more importantly for the majority who present late, the value of confirming a diagnosis is not exhausted by thrombolysis eligibility. As R3 argued (Section 3.2.1), confirmation still differentiates ischaemic stroke from haemorrhage, transient ischaemic attack, and stroke mimics, and that differentiation governs ongoing management, family counselling, and decisions such as the window for aneurysm surgery in selected haemorrhagic cases. Time-to-diagnosis is therefore a meaningful endpoint in its own right: it measures the point at which clinically actionable information becomes available, and in this setting that information has value independent of whether thrombolysis follows. Notably, none of the respondents framed the trial as an evaluation of how the device would fit into and reshape the existing stroke workflow, in the way the Greek emergency-department feasibility study by Tsiftsis et al. approached the same device [26]; that workflow-integration question is better suited to the value stage than to the performance stage, and the endpoint above captures part of it.

The site question is settled by the same logic that separated the two trials. The validation-versus-impact tension (Section 3.4.2) is not resolved by choosing one horn over the other but by assigning each horn to the trial that can carry it: the tertiary sites carry validation in the performance trial, and the value trial extends down to district hospitals, where the diagnostic gap the device addresses actually occurs. How far below district level to go is left open. Health centres would in principle give the most complete picture of real-world utility, but they are the tier this study had least insight into and the tier least equipped for trial conduct, and we do not recommend including them in the first value trial.

3.8.3 The control condition: comparison without contamination

The remaining design problem is the control condition, and it is the one that most repays careful treatment, because the three candidate solutions each fail in a different way and the choice among them has been made too quickly in the planning literature this study began from.

The natural first option, randomising patients within a site to device or no-device, is ruled out by a contamination risk that R13a identified directly (Section 3.4.3). Where trial literacy is limited, clinicians and families may regard access to a device that could help as something a control patient is morally owed, and the result is involuntary crossover that erodes the contrast the trial is trying to measure. This is a genuine threat and we accept R13a's judgement that open within-site randomisation is not advisable here.

The second option, R13a's own suggestion of a case-site/control-site design in which one hospital deploys the device and a comparable hospital does not, removes crossover by removing the device from the control site entirely. Its weakness is the standard

weakness of any two-cluster comparison: the device effect is confounded with every fixed difference between the two hospitals, and with only one site per arm there is no way to separate the one from the other.

The original planning reasoning is, however, weakest at this point. Faced with the confounders of the case-site/control-site design, that reasoning retreated to a single-arm before-and-after design: time-to-diagnosis is measured first without the device to establish a baseline, the device is then introduced at all sites, and the two periods are compared. The difficulty is that this does not reduce confounding, it changes which confounders apply, and arguably worsens them. A before-and-after comparison with no concurrent control attributes to the device any change that occurs between the two periods, including changes that have nothing to do with it: secular improvement in the referral system, staff learning the diagnostic task over the course of the study, the well-documented tendency of measured performance to improve simply because it is being observed, and seasonal variation in the mix and volume of patients presenting. In a twelve-month study spanning two periods, several of these are not hypothetical. The before-and-after design is operationally the simplest of the three, which is a real virtue given the literacy constraints discussed above, but its simplicity is bought at the cost of the one thing a validation study most needs, a credible counterfactual, and that cost has not so far been stated.

A fourth design avoids the weaknesses of the previous three, and is the one we recommend: a stepped-wedge cluster design. In a stepped wedge, all participating sites begin in the control condition and cross over to deploying the device one after another, in a randomly ordered sequence, until every site has made the transition [44]. The design has three properties that map onto exactly the problems above. Because the crossover happens at the level of the whole site and at a scheduled time, no site is ever running device and control arms simultaneously, so R13a's involuntary-crossover concern does not arise. Because every site contributes time in both conditions, each site serves as its own control and the fixed between-hospital differences that sink the two-cluster comparison are differenced out. And because the analysis adjusts for calendar time, the secular trends and learning effects that contaminate the simple before-and-after are separated from the device effect rather than absorbed into it. The price is coordination: a stepped wedge requires sites to switch on a defined schedule and to report on a common cadence, which is more than the before-and-after asks of them. That price is real, and whether it is affordable depends on how many district sites can be recruited and supported; with only two or three sites a stepped wedge is weak, and the simpler before-and-after, with its limitations stated plainly in the protocol, may be the pragmatic choice. With four or more sites it is the design that gives the device its fairest and most defensible test. Framed in the terms of the PRECIS-2 pragmatic-explanatory continuum, the performance trial sits toward the explanatory end and the stepped-wedge value trial toward the pragmatic end [45], which is the orientation each question calls for.

Two further design parameters were raised by respondents and are left for the protocol-development phase rather than fixed here. Neither the trial duration (R3 proposed four to six months across three or four sites) nor the sample size (en-

dorsed in principle by R4 and R5, but with no numeric estimate offered by any respondent) can be set without the baseline door-to-imaging measurements and the site-level stroke incidence figures that this study did not collect. Both are identified as immediate tasks for the protocol phase.

3.9 Conducting the trial in the Rwandan context

A trial design is only worth specifying if the trial can be funded, sponsored, governed, and sited where it is meant to run. The findings describe an environment that is, on balance, more accommodating to a late-phase device validation than the data-residency and funding accounts first suggest, but the conditions under which it is accommodating shape who can carry the trial and how. This section reads the feasibility findings as a set of structural preconditions on the design of Section 3.8.

3.9.1 Funding and sponsorship: who can carry the trial

The single most consequential feasibility finding is the one that follows from the device’s regulatory status. Because the MD100 already holds CE-marked status, its validation in Rwanda falls into phase-3-to-phase-4 territory rather than early-phase work (Section 3.3.2). This matters operationally because respondents located early-phase device trials at one or two specialised private institutions, and phase-3-to-phase-4 trials at the major hospitals. A validation that can be sited in the existing tertiary-hospital system therefore avoids the narrowest infrastructural bottleneck in the Rwandan trial landscape, and the two-trial structure proposed earlier sits comfortably within it.

The funding route, by contrast, is the binding constraint. Three of the funding channels respondents described do not fit a microwave stroke device. Internal hospital grant capital is scarce, on R10’s account, to the point that CHUK had financed only two small trials from its own programme, and the removal of the two-window-per-year restriction eased cash-flow timing without enlarging that pool. Foreign research funding, which dominates higher-cost work, was reported to be concentrated in infectious disease and pharmaceutical studies, neither of which a stroke-triage device belongs to. And the locally funded “pragmatic” studies respondents described are by definition the low-consumable, low-device trials that the MD100 is not. The plausible sponsor is therefore neither an internal hospital grant nor a conventional pharmaceutical funder. Two candidates remain in the findings: a mission-oriented or NGO funder of the kind R13b named in Team Heart, which has supported device-adjacent work in Rwandan cardiac care, and the device manufacturer itself acting as sponsor. This is the point at which the terminology finding stops being a curiosity and becomes a planning instruction: because “sponsor” was read in its lay financial sense by several respondents (Section 3.5.2), the protocol must fix the term to its ICH-GCP meaning, the party responsible for initiating and managing the trial, so that the division of responsibility between a manufacturer-sponsor and a local institutional collaborator is unambiguous from the outset.

That division is not optional. Respondents were consistent that the principal investigator must be Rwandan or paired with a local co-principal investigator, and that an external sponsor requires a local institutional collaborator (Section 3.3.1). A trial sponsored from Gothenburg cannot, in other words, be run at arm’s length; a local co-PI and a collaborating Rwandan institution are structural requirements written into the approval pathway, not conveniences. This is consistent with the wider picture of medical-device trials in Africa, which remain comparatively rare and are conducted within regulatory systems that are still consolidating, so that local institutional anchoring carries more of the trial’s feasibility than it would in a setting with a denser device-trial track record [41], [42], [46], [47].

3.9.2 Data governance: what actually has to leave Rwanda

The data-residency finding was reported in the Results as the feature with the clearest operational implications, and it appeared at first to be a serious obstacle: R5 described being unable to export the imaging data from her own multi-country trial to the international principal investigator, and R7 corroborated the rejection of an off-shore server proposal on residency grounds (Section 3.3.3). Read literally, this would place a heavy constraint on any internationally sponsored MD100 trial. On closer inspection, however, the constraint that bound R5’s trial does not bind this one in the same way, and the reason is worth making precise, because the Results noted a likely conflation of “trial data” with the full clinical record without resolving it.

The resolution is to specify the trial’s data footprint concretely. For each enrolled patient, the performance trial generates three things: the device’s raw microwave acquisition, the algorithm’s classification output, and a reference label derived from the locally acquired CT or MRI scan. The scientific analysis requires the device output paired with the reference label and the patient’s basic stratification variables; it does not require the CT or MRI images themselves to leave the country. The imaging that R5 could not export is exactly the part of the record that an MD100 validation does not need to move. What has to travel to a sponsor or co-investigator abroad is a de-identified dataset of device classifications, reference labels, and, for the value trial, timing variables, all of which are derived data rather than primary imaging. The operational instruction that follows is therefore narrow and actionable: define the analysable dataset tightly in the protocol, keep all imaging on local servers, and transfer only the derived variables. Handled this way, the local-residency requirement is satisfied without compromising the analysis, which inverts the finding from an apparent obstacle into a manageable design parameter. The standard governance scaffolding respondents described, a designated Data Protection Officer, a REDCap capture platform with a per-user audit trail, and local server provision through the Rwanda Biomedical Centre, applies unchanged.

3.9.3 Securing sites: readiness and pre-engagement

The performance-trial sites are, in effect, pre-identified by the findings. The four tertiary hospitals that R10 reported as having been assessed “almost ready” for hosting

trials are the same four that hold routine confirmatory imaging (Section 3.3.4), so the institutions that can host the trial and the institutions that hold the comparator are the same set. The remaining question is not which sites but how to enter them, and here the findings are unambiguous: the route in is pre-engagement, not application. Both R10 and R13a described the expected practice as engaging potential collaborating hospitals informally, identifying local collaborators, and developing the protocol jointly with them before any submission to RNEC, with R13a noting that skipping this step is actively penalised at the ethics stage. For planning purposes this reclassifies pre-engagement from soft networking into a scheduled task on the critical path: the protocol cannot be finalised until the collaborating institutions have helped shape it. The same logic extends to the one institutional actor that holds AI-regulation expertise, the Data Science Division at the Rwanda Biomedical Centre (Section 3.4.5), which R13a framed as a parallel engagement track to be opened alongside, not after, the RNEC and RFDA submissions. Early, parallel engagement of this kind is a recurrent condition for trial feasibility in sub-Saharan settings, where the relationships that carry a study are typically built before, not during, formal review [43].

3.9.4 What makes it fail

Two failure points named in the findings bear directly on how the trial must be resourced. The first concerns the conduct phase. R13a’s observation that studies “die out of lack of follow up” (Section 3.3.5) is, for a trial whose sponsor and possibly whose PI may be based abroad, the most important single warning in the corpus. It converts into a budget line: if the principal investigator is not physically present, the protocol must fund paid, full-time site staff reporting on a defined cadence. This interacts with the sponsorship question of Section 3.9.1, because an externally sponsored design with an off-site PI is viable only if it is financed to compensate for that distance; an unfunded monitoring provision is not a minor omission but the specific mechanism by which respondents expected such a trial to drift.

The second failure point concerns what happens after a successful trial. R10 located the risk of unsuccessful technology introduction in community engagement and mindset rather than in regulatory or technical obstacles, and R3’s account of the abandoned imported vein-localisation device, discussed in Section 3.8, supplied the corpus’s sharpest illustration of an adoption that failed for want of local validation. Taken together, these reframe the purpose of the value trial: beyond measuring time-to-diagnosis, it is the instrument through which the local credibility that adoption depends on is generated. A trial that produces a clean result but neglects the engagement and follow-through that R10 and R13a emphasised would satisfy the regulator while failing at the thing the device is ultimately meant to achieve.

Across these four considerations the net assessment is that a late-phase MD100 validation is feasible in the Rwandan context, but conditional on three structural commitments that the design must build in from the start: a local co-principal investigator and collaborating institution, a precisely scoped dataset that keeps imaging in country and transfers only derived variables, and a funded monitoring provision

that does not assume the investigator’s physical presence. These are the preconditions on which the framework of Section 3.8 rests.

3.10 Terminology, literacy, and AI regulation as design constraints

Two of the three cross-cutting findings (Section 3.5) are not about stroke care or trial logistics at all, but about the distribution of knowledge: which operational vocabulary is shared across the workforce, and where regulatory knowledge about AI-based devices actually sits. Both are easy to read as incidental observations about the sample. The argument of this section is that they are not incidental. Each one constrains the trial design, and each has already done work in the recommendations made earlier in this chapter.

3.10.1 Trial vocabulary as a constraint, not a measurement artefact

The clearest instance of the vocabulary finding is the term “sponsor,” which seven of the sixteen respondents initially read in its lay financial sense and reinterpreted in its ICH-GCP procedural sense only after the interviewer’s clarification (Section 3.5.2). As the reflexive discussion in the Methods chapter notes (Section 2.7), the convergence that follows clarification is partly an artefact of the interviewer’s framing: respondents who were corrected mid-interview were thereafter answering a slightly different question. The temptation is to treat this as measurement noise to be controlled out. We take the opposite view, and treat it as diagnostic. The fact that a core ICH-GCP role term was not uniformly held, in a sample weighted toward the most research-active tertiary institutions in the country, is itself a finding about the operational readiness of the environment in which the trial would run.

Read this way, the vocabulary finding is not a loose observation but the evidence underwriting a design choice already made. R13a’s judgement that trial literacy in Rwanda is “not solid” was the explicit basis for the blinding and complexity-reduction recommendations of Section 3.8 (Section 3.4.3); the “sponsor” finding is the corpus-level corroboration of that judgement. It also converts into a concrete operational instruction. GCP certification compliant with the ICH E6(R3) update is required of trial personnel (Section 3.3.1), but certification establishes formal qualification, not a shared working vocabulary across a multidisciplinary site team [22]; the protocol and the site-initiation process must therefore define and rehearse the operative roles explicitly rather than assume them. Pre-engagement (Section 3.9.3) is, among other things, where that alignment is built. One boundary on the claim should be stated: the divergence fell between operational trial leadership and clinical staff rather than between institutions, so the constraint is about professional role, not about geography, and the design response is training and definition rather than site exclusion.

3.10.2 The AI-regulation knowledge gap

The second knowledge finding is sharper and the thesis shares in it. Six respondents drawn from across the institutional landscape did not know whether the AI classification component of the MD100 attracts regulatory requirements beyond those applied to medical devices in general, and only R13a could give a substantive answer, which pointed away from the clinical respondents and toward a single institutional actor, the Data Science Division at the Rwanda Biomedical Centre (Section 3.4.5). Because no respondent from that division was interviewed, the honest characterisation of this finding is epistemic rather than substantive: the study establishes where the question should be asked, not what the answer is.

The wider literature makes clear why the question cannot be set aside. The machine-learning classifier that distinguishes the MD100 from a passive measurement instrument is precisely the feature that brings it within the scope of AI- or software-as-a-medical-device regulation, and a CE mark obtained under the European device framework does not settle how that classifier is treated under a national regime [48]. The regulatory landscape for AI in healthcare across Africa is, moreover, still forming and unevenly distributed between countries, so the relevant requirements cannot be inferred from a neighbouring jurisdiction and have to be established locally [47], [49], [50]. The planning consequence has already been drawn in Section 3.9.3: engagement with the Rwanda Biomedical Centre’s Data Science Division is opened as a parallel track alongside the RNEC and RFDA submissions, not deferred until after them.

One element of the findings reduces, without removing, the regulatory uncertainty. The interpretive principle that R9 and R10 articulated, that the AI output must support rather than replace clinical judgement, coincides with the position that AI-as-a-medical-device frameworks generally adopt, and it is also how the MD100 is designed to be used, as a triage-support tool with the clinician retaining the decision. The device’s intended use is therefore congruent with the regulatory posture respondents expected, which lowers the risk of a fundamental mismatch even though the specific approval pathway remains to be confirmed. The institutional signal points the same way: R11 reported that KFH was already in discussion to adopt RapidAI, an imaging-based stroke-detection tool, which indicates an openness toward AI-supported diagnostics rather than a restrictive stance, though, as the Results noted, RapidAI’s regulatory pathway was not documented in the corpus and the precedent is therefore a sign of posture rather than a template to follow.

3.10.3 Two findings, one remedy

The two findings are operationally unrelated, but they share a shape and a remedy. Both are problems of distributed knowledge rather than of infrastructure: in the first, the working vocabulary of trial conduct is unevenly held across roles; in the second, AI-regulation knowledge is held at one institution and absent at the clinical sites where the trial would run. In both cases the remedy is the same, and it is the one this chapter has repeatedly arrived at from different directions: early, role-aware

engagement and the explicit definition of what is otherwise assumed, rather than reliance on a shared baseline that the findings show cannot be taken for granted. This is why the framework leans as heavily as it does on pre-engagement and on reducing the complexity the trial asks of its sites. The design of Section 3.8 and the feasibility conditions of Section 3.9 are, in the end, responses to an environment whose binding constraints are as much about where knowledge sits as about where the CT scanners are.

Taken together, the elements developed across this chapter assemble into a four-phase trial-planning framework, set out in the Conclusion (Table 4.1).

4

Conclusion

From the findings of this thesis there is a clinical need in the current Rwandan stroke care system to reduce the delays patients endure before diagnosis and treatment. These delays stem from the stepwise referral system, the absence of early clinical data at the first point of contact, and the centralisation of diagnostic imaging and treatment. It is a need the MD100 could potentially help to fill, and to evaluate whether it does, we set out a trial-planning framework.

The framework is organised in four phases, summarised in Table 4.1. A preparatory phase (Phase 0) secures the regulatory, data-governance, AI-specific, and institutional commitments before any submission is made. Two trials then run in sequence (Phases 1 and 2), and a close-out phase (Phase 3) consolidates and disseminates the resulting evidence.

The first trial evaluates the performance of the MD100, its sensitivity and specificity, against CT or MRI as the gold-standard comparator. It is hosted at the tertiary hospitals, which are the most prepared for conducting clinical trials and the sites where the comparator is routinely available.

Given a positive outcome of this initial trial, the second trial evaluates the clinical value of the MD100 with the endpoint of time-to-diagnosis, chosen because treatment for ischaemic stroke is confined to King Faisal Hospital. Here the participating sites extend to district hospitals, where the diagnostic gap is largest. Because trial literacy in the sample was limited and within-site randomisation risks involuntary crossover, we recommend a stepped-wedge design: all sites begin without the MD100 and cross over to using it one at a time, in a randomised order, so that each site serves as its own control and the device effect can be separated from background change over time. Where too few district sites can be recruited for a stepped wedge to be viable, a before-and-after comparison remains the pragmatic fallback, provided its weaker control of confounding is stated explicitly in the protocol.

Table 4.1: The trial-planning framework: four phases running from pre-engagement through close-out. Phases 1 and 2 constitute the trial proper; Phase 0 prepares the regulatory, data-governance, and institutional ground before submission, and Phase 3 consolidates and disseminates the evidence. The final column states the gate that must be cleared before the next phase begins.

Phase	Purpose	Setting	Design and end-point	Output / gate to next phase
Phase 0 Pre-engagement	Secure sites, PI, and the regulatory, data-governance, and AI pathway; align trial vocabulary with site staff	Candidate hospitals and regulators (RNEC, RFDA, RBC Data Science Division)	Preparatory; no patient enrolment	Approved protocol; RNEC and conditional RFDA approval; signed site agreements; documented data-governance and AI provisions
Phase 1 Performance trial (explanatory)	Establish diagnostic accuracy against the imaging reference standard	Tertiary referral hospitals (CHUK, KFH, RMH)	Single-arm comparison against CT/MRI; sensitivity and specificity	Accuracy clears the pre-specified sensitivity and specificity thresholds
Phase 2 Value trial (pragmatic)	Measure clinical impact under usual-care conditions	District hospitals, with tertiary imaging confirmation by transfer	Stepped-wedge cluster design; time-to-diagnosis	Evidence of change in time-to-diagnosis under deployment
Phase 3 Close-out and dissemination	Consolidate evidence, complete submissions, disseminate results	—	—	RFDA marketing-authorisation dossier; peer-reviewed publications; RNEC and IRB close-out filings

5

Future research

This thesis plans a trial but does not carry it out. Several clear directions for future work follow from this.

The most immediate next step is to develop the formal trial protocol and run the trial itself. Before the protocol can be finalised, a few things that were outside the scope of this study still need to be established. The baseline door-to-imaging time at the candidate district hospitals has to be measured directly, because the device's impact can only be judged against a real starting point rather than the respondent estimates reported here. A sample-size calculation is needed for both trial stages, based on how many stroke patients each site actually sees. This study also could not interview four stakeholders whose input is important for the protocol: the Rwanda FDA, an industry sponsor with recent experience of a device trial in Rwanda, a radiologist, and the Data Science Division at the Rwanda Biomedical Centre. Engaging these four directly is a priority for the protocol-development phase.

A second direction concerns the device itself. Respondents pointed to two features that the MD100 cannot currently provide but that would make it far more useful in practice: telling haemorrhagic and ischaemic stroke apart, and identifying which side of the brain is affected (Section 3.4.3). Developing the algorithm to support these tasks is work for the manufacturer, but the trial can help. If the first stage of the trial records the stroke subtype from the CT or MRI scan alongside the device's own output, that dataset can later be re-analysed to test how well a future version of the algorithm performs on these tasks, without having to run a new trial from scratch.

A third direction concerns the wider range of uses that respondents saw for the device beyond stroke triage (Section 3.4.6). Several of these are promising enough to deserve their own studies once the main stroke-triage validation is complete. The clearest example is using the device in ambulances and at peripheral facilities, combined with the zip-line system that already delivers medication to remote hospitals, so that a stroke can be identified and treatment dispatched without waiting for transfer. Other possibilities raised by respondents include detecting bleeding in trauma patients who are too unstable to be moved to a scanner, and screening for conditions such as brain tumours in settings where imaging is hard to reach. Each of these would need to be validated separately, but together they suggest that

the device's value in Rwanda may extend well beyond the stroke-triage role studied here.

Bibliography

- [1] V. L. Feigin et al., “Global, regional, and national burden of stroke and its risk factors, 1990–2021: A systematic analysis for the global burden of disease study 2021,” *The Lancet Neurology*, vol. 23, no. 10, pp. 973–1003, 2024.
- [2] D. Kuriakose and Z. Xiao, “Pathophysiology and treatment of stroke: Present status and future perspectives,” *International Journal of Molecular Sciences*, vol. 21, no. 20, p. 7609, 2020. DOI: 10.3390/ijms21207609.
- [3] R. O. Akinyemi et al., “Stroke in africa: Profile, progress, prospects and priorities,” *Nature Reviews Neurology*, vol. 17, no. 10, pp. 634–656, 2021. DOI: 10.1038/s41582-021-00542-9.
- [4] World Health Organization, *Rwanda [country overview]* — *data.who.int*, <https://data.who.int/countries/646>, Accessed: 25 February 2026, 2026.
- [5] J. L. Saver, “Time is brain—quantified,” *Stroke*, vol. 37, no. 1, pp. 263–266, 2006.
- [6] American Stroke Association, *Effects of stroke*, <https://www.stroke.org/en/about-stroke/effects-of-stroke>, Accessed: Apr. 27, 2026.
- [7] A. I. Mugisha, F. Munyaneza, O. Uwishema, and I. Nzayisenga, “Article in bmc neurology,” *BMC Neurology*, 2026. DOI: 10.1186/s12883-026-04645-7. [Online]. Available: <https://link.springer.com/article/10.1186/s12883-026-04645-7>.
- [8] National Institute of Neurological Disorders and Stroke, *Traumatic brain injury (tbi)*, <https://www.ninds.nih.gov/health-information/disorders/traumatic-brain-injury-tbi>, Accessed: 2025-05-23, 2023.
- [9] C. Hui, P. Tadi, M. Z. K. Suheb, and L. Patti, *Ischemic stroke*, StatPearls [Internet], Accessed: Apr. 27, 2026, Treasure Island, FL, 2025. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK499997/>.
- [10] A. K. A. Unnithan and J. M. Das, *Hemorrhagic stroke overview*, StatPearls [Internet], Updated Dec. 13, 2025; Accessed: Apr. 27, 2026, Treasure Island, FL, 2025. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK559173/>.
- [11] National Heart, Lung, and Blood Institute, *Stroke: Treatment*, <https://www.nhlbi.nih.gov/health/stroke/treatment>, Accessed: 2026-05-13, 2023.

-
- [12] M. U. Baig and J. Bodle, *Thrombolytic therapy*, StatPearls [Internet], Updated Aug. 28, 2023; Accessed: Apr. 27, 2026, Treasure Island, FL, 2026. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK557411/>.
- [13] I. M. El-Sherbiny, I. E. Elkholi, and M. H. Yacoub, “Tissue plasminogen activator-based clot busting: Controlled delivery approaches,” *Global Cardiology Science and Practice*, vol. 2014, no. 3, pp. 336–349, 2014. DOI: 10.5339/gcsp.2014.46.
- [14] T. Robinson, Z. Zaheer, and A. K. Mistri, “Thrombolysis in acute ischaemic stroke: An update,” *Therapeutic Advances in Chronic Disease*, vol. 2, no. 2, pp. 119–131, 2011. DOI: 10.1177/2040622310394032.
- [15] H. Alshaqouq, “A health technology assessment of the strokefinder md100 for early detection of stroke and traumatic brain injury in the rwandan healthcare system,” 2025.
- [16] S. Mathews and O. D. Jesus, *Thrombectomy*, StatPearls [Internet], Accessed: Apr. 27, 2026, Treasure Island, FL, 2024. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK562154/>.
- [17] Republic of Rwanda Ministry of Health, *Health sector policy*, https://www.moh.gov.rw/fileadmin/user_upload/Health_Sector_Policy__19th_January_2015.pdf, Accessed: 2026-05-08, Jan. 2015.
- [18] F. Ntirenganya et al., “Surgical capacity and neurosurgical access in rwanda: Current gaps and strategies,” *World Neurosurgery*, vol. 150, e367–e376, 2021.
- [19] International Organization for Standardization, *ISO 14155:2020 – clinical investigation of medical devices for human subjects – good clinical practice*, 2020.
- [20] World Medical Association, *WMA declaration of Helsinki – ethical principles for medical research involving human subjects*, <https://www.wma.net/policies-post/wma-declaration-of-helsinki/>, Adopted 1964; last amended by the 64th WMA General Assembly, Fortaleza, Brazil, October 2013, 2013.
- [21] European Parliament and Council of the European Union, “Regulation (EU) 2017/745 of the European Parliament and of the Council on medical devices,” Official Journal of the European Union, Tech. Rep. L 117, 2017, Amended by Regulation (EU) 2020/561.
- [22] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use, “ICH E6(R3): Guideline for good clinical practice,” ICH, Step 4 (final) document, Jan. 2025. [Online]. Available: <https://www.ich.org/page/efficacy-guidelines>.
- [23] Medfield Diagnostics AB, *Product description: Strokefinder MD100*, Document no. MD2016-127, revision A, Gothenburg, Sweden, Feb. 2016.
- [24] M. Persson et al., “Microwave-based stroke diagnosis making global prehospital thrombolytic treatment possible,” *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 11, pp. 2806–2817, 2014. DOI: 10.1109/TBME.2014.2330554.
-

-
- [25] A. Fhager, S. Candefjord, M. Elam, and M. Persson, "Microwave diagnostics ahead: Saving time and the lives of trauma and stroke patients," *IEEE Microwave Magazine*, vol. 19, no. 3, pp. 78–90, 2018. DOI: 10.1109/MMM.2018.2801646.
- [26] D. Tsiftsis, E. A. Manioti, G. Touris, E. Kyriakakis, N. Tsamopoulos, and M. Gamvroudi, "Detecting stroke at the emergency department by a point of care device: A multicenter feasibility study," *Medical Devices*, vol. 17, pp. 107–112, 2024. DOI: 10.2147/MDER.S445075.
- [27] World Health Organization, *Health technology assessment*, <https://www.who.int/teams/health-product-policy-and-standards/assistive-and-medical-technology/medical-devices/assessment>, Accessed: Apr. 27, 2026, n.d.
- [28] World Health Organization, *Health technology assessment*, <https://www.who.int/health-topics/health-technology-assessment>, Accessed: Apr. 27, 2026, n.d.
- [29] EUnetHTA Joint Action 2, Work Package 8, *HTA Core Model® Version 3.0*, <http://www.htacoremodel.info/BrowseModel.aspx>, Jan. 25, 2016; Accessed: Apr. 27, 2026, 2016.
- [30] The Public Health Agency of Sweden, *Health economics*, <https://www.folkhalsomyndigheten.se/the-public-health-agency-of-sweden/communicable-disease-control/health-economics/>, Accessed: 2026-05-07, Mar. 2023.
- [31] L. Prieto and J. A. Sacristán, "Problems and solutions in calculating quality-adjusted life years (qalys)," *Health and Quality of Life Outcomes*, vol. 1, no. 80, Dec. 2003. DOI: 10.1186/1477-7525-1-80. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC317370/>.
- [32] York Health Economics Consortium, *Incremental cost-effectiveness ratio (icer)*, <https://www.yhec.co.uk/glossary-term/incremental-cost-effectiveness-ratio-icer/>, Accessed: 2026-05-07, 2016.
- [33] York Health Economics Consortium, *Net monetary benefit*, <https://www.yhec.co.uk/glossary-term/net-monetary-benefit/>, Accessed: 2026-05-07, Oct. 2016.
- [34] J. W. Creswell and J. D. Creswell, *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*, 5th ed. Thousand Oaks, CA: SAGE Publications, 2018.
- [35] S. Kvale and S. Brinkmann, *InterViews: Learning the Craft of Qualitative Research Interviewing*, 2nd ed. Thousand Oaks, CA: SAGE Publications, 2009.
- [36] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006. DOI: 10.1191/1478088706qp063oa.
- [37] A. E. Nkusi et al., "Stroke burden in rwanda: A multicenter study of stroke management and outcome," *World Neurosurgery*, vol. 106, pp. 462–469, 2017. DOI: 10.1016/j.wneu.2017.07.014.
-

-
- [38] Lloyds Bank Trade, *Rwanda: Work conditions*, <https://www.lloydsbanktrade.com/en/market-potential/rwanda/work-conditions>, International Trade Portal, Lloyds Bank, n.d. Accessed: May 25, 2026.
- [39] M. P. Nisingizwe et al., “Effect of unmanned aerial vehicle (drone) delivery on blood product delivery time and wastage in rwanda: A retrospective, cross-sectional study and time series analysis,” *The Lancet Global Health*, 2022. DOI: 10.1016/S2214-109X(22)00048-1.
- [40] Rwanda Food and Drugs Authority, *Guidelines on good clinical practices (gcp) in rwanda*, Doc. No. FDISM/PVSM/GDL/010, Rev. 01, Apr. 2023. Accessed: May 25, 2026. [Online]. Available: [https://rwandafda.gov.rw/wp-content/uploads/2023/04/Guidelines%20on%20Good%20Clinical%20Practices%20\(GCP\)%20in%20Rwanda.pdf](https://rwandafda.gov.rw/wp-content/uploads/2023/04/Guidelines%20on%20Good%20Clinical%20Practices%20(GCP)%20in%20Rwanda.pdf).
- [41] B. Matovu et al., “A compound analysis of medical device clinical trials registered in Africa on clinicaltrials.gov,” *Trials*, vol. 25, no. 1, p. 658, 2024. DOI: 10.1186/s13063-024-08427-9.
- [42] B. Matovu et al., “Review of investigational medical devices’ clinical trials and regulations in Africa as a benchmark for new innovations,” *Frontiers in Medical Technology*, vol. 4, p. 952 767, 2022. DOI: 10.3389/fmedt.2022.952767.
- [43] E. I. Nebie et al., “Opportunities and challenges for decentralised clinical trials in sub-Saharan Africa: A qualitative study,” *BMJ Open*, vol. 13, no. 9, e075903, 2023. DOI: 10.1136/bmjopen-2023-075903.
- [44] K. Hemming, T. P. Haines, P. J. Chilton, A. J. Girling, and R. J. Lilford, “The stepped wedge cluster randomised trial: Rationale, design, analysis, and reporting,” *BMJ*, vol. 350, h391, 2015. DOI: 10.1136/bmj.h391.
- [45] K. Loudon, S. Treweek, F. Sullivan, P. Donnan, K. E. Thorpe, and M. Zwarenstein, “The PRECIS-2 tool: Designing trials that are fit for purpose,” *BMJ*, vol. 350, h2147, 2015. DOI: 10.1136/bmj.h2147.
- [46] S. Hubner et al., “The evolving landscape of medical device regulation in East, Central, and Southern Africa,” *Global Health: Science and Practice*, vol. 9, no. 1, pp. 136–148, 2021. DOI: 10.9745/GHSP-D-20-00578.
- [47] N. Nasir, S. Molyneux, F. Were, A. Aderoba, and S. S. Fuller, “Medical device regulation and oversight in African countries: A scoping review of literature and development of a conceptual framework,” *BMJ Global Health*, vol. 8, no. 8, e012308, 2023. DOI: 10.1136/bmjgh-2023-012308.
- [48] B. A. Mateen, J. Ordish, and M. J. A. Reid, “Your AI tool is subject to medical device regulation,” *Nature Health*, vol. 1, no. 1, pp. 9–10, 2026. DOI: 10.1038/s44360-025-00003-y.
- [49] B. A. Townsend, I. Sihlahla, M. Naidoo, S. Naidoo, D.-L. Donnelly, and D. W. Thaldar, “Mapping the regulatory landscape of AI in healthcare in Africa,” *Frontiers in Pharmacology*, vol. 14, p. 1 214 422, 2023. DOI: 10.3389/fphar.2023.1214422.
-

-
- [50] World Health Organization, African Health Observatory, “Analytical fact sheet: Medical device regulation in African countries,” WHO Regional Office for Africa, Tech. Rep., Jun. 2024. [Online]. Available: https://files.who.int/afahobckpcontainer/production/files/Regional_Factsheet_on_Medical_device_regulation_EN.pdf.
-

A

Interview Guide

The following semi-structured interview guide was used across all sixteen interviews, as described in Section 2.3.2. The guide provided a consistent thematic structure; the semi-structured format allowed follow-up questions and the pursuit of emerging themes, so not every question was posed verbatim in every interview, and the order was adapted to the flow of each conversation. Bracketed items below denote prompts and sub-questions used to probe further where relevant.

Current stroke care

- How often do you encounter stroke patients at this facility?
- What are the current emergency-department standard operating procedures (SOPs) or workflow for suspected strokes?
- How does a typical suspected stroke case play out at present?
- What is currently the typical outcome for a stroke patient?
- What are the current gaps or limitations in stroke care from your perspective?
- How long does it typically take for a patient to receive a CT scan?

Clinical trials: general

- How does the hospital work with clinical trials?
 - What are the key actors when conducting clinical trials in Rwanda?
 - Ministry of Health (MoH)
 - Rwanda FDA
 - Rwanda National Ethics Committee (RNEC)
 - Trial team
-

-
- Is there any differentiation between academic and commercial clinical trials?
 - What regulatory approvals are required for running clinical trials?
 - How are clinical trials of medical devices typically financed in Rwanda?
 - What are the requirements on the primary stakeholders?
 - Principal Investigator
 - Institutional Review Board (IRB) / Independent Ethics Committee
 - Sponsor
 - Data governance
 - How is trial personnel recruited, and whose responsibility is this?

Clinical trials: MD100 specific

- What would you like the clinical trial to investigate?
 - What is the endpoint?
 - Where would you like to see the clinical trial conducted?
 - Multi-centre?
 - What would the requirements be of the sites?
 - In what setting would you like the trial to be conducted?
 - Patient recruitment:
 - How are patients chosen?
 - Inclusion criteria?
 - Exclusion criteria?
 - How many?
 - Control group?
 - How are ethical concerns handled?
 - Acute patients
 - Possible cognitive impairment
-

- Unconscious patients
 - Design of new SOPs for the trial:
 - Who will be mainly responsible? the PI?
 - How can it be ensured that the trial does not interfere with current SOPs?
 - What support is needed?
 - Are there any extra regulations given that the MD100 is AI-powered?
 - Do you think the data protection and privacy office would have to be involved in a trial with the MD100?
 - Who is responsible for technical support?
 - Medfield only?
 - Is it possible to train the hospitals' own technical departments?
-

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se



CHALMERS
UNIVERSITY OF TECHNOLOGY