THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Virtuality – From Idea to Practical Use: A Simulation Environment for Interoperable and Artificial Intelligence Enabled Healthcare Processes

Mattias Seth

Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden, 2026

Virtuality – From Idea to Practical Use:
A Simulation Environment for Interoperable and Artificial Intelligence Enabled Healthcare Processes

Mattias Seth
ISBN 978-91-8103-481-3

Acknowledgements, dedications, and similar personal statements in this thesis, reflect the author's own views.

© Mattias Seth 2026 except where otherwise stated.

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 5938
ISSN 0346-718X
DOI https://doi.org/10.63959/chalmers.dt/5938

Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg, Sweden
Phone: +46 (0)31 772 1000

Cover:
Architecture overview of the Virtuality Digital Health Sandbox environment and its interoperable components. Created with ChatGPT 5.5, 2026

Printed by Chalmers Digital Printing
Gothenburg, Sweden, October 2026

Abstract

Healthcare is increasingly dependent on digital innovations, yet translating these innovations into clinical practice remains challenging. This thesis investigates how a Digital Health Sandbox can be designed and implemented to support the transition from promising ideas towards clinical implementation. Following a Design Science Research-inspired approach, the work combines literature studies, triple-helix collaborations, and iterative implementation work.

The research began with a protocol and subsequent scoping review of Internet of Medical Things (IoMT) platforms for home and prehospital care (Papers A–B), revealing a heterogeneous landscape of technologies and standards, with limited support for higher levels of interoperability. A study of health information system procurement (Paper C) identified six recurring challenges in specifying interoperability requirements. Together with challenges encountered during practical implementation and AI integration work (Papers E–F), and emerging regulatory requirements, these findings revealed the need for a controlled testing environment. Together, these findings informed the development of a testing environment (Paper D and G).

Building on this work, the central contribution of this thesis is the design and implementation of Virtuality – Virtual Care Process Simulator, a process-oriented simulation environment built around a modular and interoperable architecture. The architecture combines a Scenario Engine for configuring simulations with reusable Building Blocks representing data providers, services, and user-facing applications (consumers). Through standardized interfaces, explicit data contracts, and validation mechanisms, these components can be connected into executable simulation pipelines that represent complex healthcare processes. Virtuality's separation of build-time and run-time environments supports both the creation and execution of simulation scenarios.

The architecture's applicability was explored through two healthcare use cases involving AI-based trauma triage and stroke assessment. Together, these contributions establish a foundation for testing digital health innovations before clinical deployment, with the potential to support safer AI integration, greater end-user involvement, and ultimately better patient care.

Keywords: Digital Health Sandbox, Healthcare Interoperability, Clinical Decision Support, Artificial Intelligence, HL7 FHIR, Simulation Environment, Medical Informatics.

To Cecilia, for her endless support throughout this journey.

And in loving memory of my mom and dad.

List of Publications

This thesis is based on the work contained in the following publications, with author contributions shown in Table 1.

[A] **Mattias Seth**, Hoor Jalo, Åsa Högstedt, Otto Medin, Ulrica Björner, Bengt Arne Sjöqvist, Stefan Candefjord , " Technologies for Interoperable Internet of Medical Things Platforms to Manage Medical Emergencies in Home and Prehospital Care: Protocol for a Scoping Review ". Published in JMIR Research Protocols, vol. 11, no. 9, article e40243, Sep. 20, 2022. DOI: 10.2196/40243.

[B] **Mattias Seth**, Hoor Jalo, Åsa Högstedt, Otto Medin, Bengt Arne Sjöqvist, Stefan Candefjord , " Technologies for Interoperable Internet of Medical Things Platforms to Manage Medical Emergencies in Home and Prehospital Care: Scoping Review ". Published in Journal of Medical Internet Research, vol. 27, article e54470, Jan. 23, 2025. DOI: 10.2196/54470.

[C] **Mattias Seth**, Hoor Jalo, Eunji Lee, Anna Bakidou, Otto Medin, Ulrica Björner, Bengt Arne Sjöqvist, Stefan Candefjord , " Reviewing Challenges in Specifying Interoperability Requirement in Procurement of Health Information Systems ". Published in Studies in Health Technology and Informatics, vol. 310, pp. 8–12, Jan. 25, 2024. DOI: 10.3233/SHTI230917.

[D] Fanny Apelgren, **Mattias Seth**, Hoor Jalo, Bengt Arne Sjöqvist, Stefan Candefjord , " Introducing Virtuality – Virtual Care Process Simulator: A Concept Utilizing Synthetic Data and a Digital Health Sandbox for Care Process Simulations ". Published in BIOSTEC 2025 Proceedings, vol. 1, pp. 1127–1135, 2025. DOI: 10.5220/0013401300003911.

[E] **Mattias Seth**, Fanny Apelgren, Anna Bakidou, Peter Lundgren, Hans Blomberg, Bengt Arne Sjöqvist, Stefan Candefjord , " FHIR Data Mapping Quality Assessment: A Scoping Review and Suggested Requirements for Future Pre-Mapping Frameworks ". Accepted for publication in Computational and Structural Biotechnology Journal. Manuscript ID: CSBJ-D-26-00542.

[F] **Mattias Seth**, Fanny Apelgren, Anna Bakidou, Peter Lundgren, Hans Blomberg, Bengt Arne Sjöqvist, Stefan Candefjord , " Introducing FHIR

Mapping Assessment and Performance Score (F-MAPS) – A Pre-Mapping Quality Assessment Tool ". Manuscript in preparation.
s

[G] **Mattias Seth**, Fanny Apelgren, Anna Bakidou, Peter Lundgren, Hans Blomberg, Bengt Arne Sjöqvist, Stefan Candefjord , "From Concept to Implementation: Design and Architecture of Virtuality — A Virtual Care Process Simulation Environment". To be submitted.

Other publications by the author, not included in this thesis, are:

[H] Jalo H, **Seth M**, Pikkarainen M, Häggström I, Jood K, Bakidou A, Sjöqvist BA, Candefjord S, "Early identification and characterisation of stroke to support prehospital decision-making using artificial intelligence: a scoping review protocol". BMJ Open. 2023 May 22;13(5):e069660. DOI: `https://doi.org/10.1136/bmjopen-2022-069660`. PMID: 37217266; PMCID: PMC10230929.

[I] Jalo H, Lee E, **Seth M**, Bakidou A, Pikkarainen M, Jood K, Sjöqvist BA, Candefjord S, "Stroke Prehospital Decision Support Systems Based on Artificial Intelligence: Grey Literature Scoping Review". In Proceedings of the 17th International Joint Conference on Biomedical Engineering Systems and Technologies – HEALTHINF. SciTePress, 2024, pp. 458–465. DOI: `https://doi.org/10.5220/0012380400003657`.

[J] Jalo H, **Seth M**, Pikkarainen M, Häggström I, Jood K, Bakidou A, Sjöqvist BA, Candefjord S, "Early Identification and Characterization of Stroke to Support Prehospital Decision Making Using Artificial Intelligence". Submitted to BMJ Open.

[K] Bakidou A, Andersson Hagiwara M, Lee E, Caragounis EC, Sjöqvist BA, **Seth M**, Jonsson A, Candefjord S, "Human-AI collaboration for prehospital trauma triage: Designing the On Scene Injury Severity Prediction (OSISP) model as a clinical decision support system". Digit Health. 2025 Dec 12;11:20552076251403207. PMID: 41393842; PMCID: PMC12701220 DOI:`https://doi.org/10.1177/20552076251403207`.

[L] Bakidou A, **Seth M**, Hagiwara MA, Jalo H, Caragounis EC, Sjöqvist BA, Jonsson A, Candefjord A, "Usability of a Prehospital Clinical Decision

Support System Based on Artificial Intelligence: A case study of EMS personnel's experience of using On-Scene Injury Severity Prediction (OSISP)". *Manuscript in writing.*

Table 1: Author contributions to Papers A–G. – No contribution, ● Contributed, ● Contributed to high extent, ■ Main responsible.

Part	A	B	C	D	E	F	G
Conceptualization	■	■	■	●	■	■	■
Methodology	■	■	■	●	■	■	■
Planning	■	■	■	●	■	■	■
Data curation	■	■	■	●	■	■	■
Writing - Original draft	■	■	■	●	■	■	■
Writing - Review & editing	■	■	■	●	■	■	■
Visualization	■	■	■	●	■	■	■
Project administration	●	●	●	●	●	●	●
Funding acquisition	–	–	–	–	–	●	●

Supervised projects

Supervised projects by the author, not included in this thesis, are:

D. Erikmats and R. Celiker, "Interoperable Medical Device Analytics in Microservice Architectures: Development of a Hybrid Electrochemical and Gaussian Process Regression Framework for Lithium Carbon Monoflouride Battery Prognostics". Master's thesis, Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2026. Accessed: Sep. 16, 2026. [Online]. Available: `https://hdl.handle.net/20.500.12380/312481` Role: Co-supervisor.

A. Cider, N. Ronnby, "Designing a Web-Based Graphical User Interface for Managing Digital Twin Simulations in Healthcare". MSc thesis, Department of

Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2025. [Online]. Available: https://hdl.handle.net/20.500.12380/309698.
Role: Main supervisor.

W. Bergman, S. On Huang, "Designing a Sandbox for Smarter Health Solutions: Creating a Usability-Centered GUI Sandbox for Innovators Within Healthcare to Simulate Traffic Events and Validate Digital Health Algorithms". MSc thesis, Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2025. [Online]. Available: https://hdl.handle.net/20.500.12380/310624.
Role: Main supervisor.

A. Ragnevad, J. Petersson, S. Albinsson, "System solution with data fusion for optimal handling of fall accidents in home healthcare". BSc thesis, Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2024.
Role: Main supervisor.

A. Alami Alamdari, S. Oguz, V. Olafsson, G. Rydner, "Blockkedja för säkrare hälsodata och ökad patientsäkerhet: En litteraturstudie om blockkedjeteknikens funktion inom sjukvården". BSc thesis, Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2023. [Online]. Available: https://hdl.handle.net/20.500.12380/306616.
Role: Main supervisor.

P. Amanzadeh, M. Bryntheson, T. Kvarnström, I. Stake, M. Vikla, "Fallolyckor – Data fusion och UI/UX för en tryggare tillvaro för anhöriga till äldre inom hemsjukvården". BSc thesis, Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2023. [Online]. Available: https://hdl.handle.net/20.500.12380/306598.
Role: Main supervisor.

Contents

Abstract i

List of Publications v

Acknowledgements xiii

Acronyms xiv

I Overview 1

1 Thesis outline 7

2 Introduction 9

2.1 Background 9
2.2 Problem Statement 10
2.3 From innovation to safe implementation 12
2.4 Research Trajectory 12
2.5 Why Digital Health Sandboxes Matter 14
2.6 Thesis Aim and Research Questions 16

3 Interoperability and Digital Health Foundations 19
3.1 Healthcare as a Complex Sociotechnical Environment 20
3.2 Interoperability in Healthcare 20
3.2.1 Technical, Syntactic, and Semantic Interoperability . . . 21
3.2.2 Beyond Technical Data Exchange 22
3.3 Digital Health Infrastructures and the Internet of Medical Things 23
3.4 Standards and Frameworks for Healthcare Interoperability . . . 24
3.4.1 HL7 FHIR 25
3.4.2 Other Interoperability Approaches and Implementation Frameworks 27
3.5 Interoperability Challenges in Practice 27
3.6 Testing and Validation of Healthcare Interoperability 28
3.7 How to Bridge the Gap 30

4 Frameworks, concepts and Project Outcomes Guiding the Research 31
4.1 Building Upon and Extending Existing Knowledge 32
4.1.1 Verified Innovation Process for Healthcare Solutions . . 33
4.1.2 Acute Support Assessment and Prioritizing (ASAP) Concept 34
4.1.3 Model for AI Utilization (MAIN) 35

5 Development of Virtuality 37
5.1 Design Rationale and Objectives 37
5.2 The "How" 38
5.2.1 Virtuality - The first prototype 38
5.2.2 System Requirements 40
5.3 Virtuality - The Second Prototype 43
5.3.1 Virtuality Starting to Take Form 45
5.3.2 Development Begins 46
5.4 Architectural Synthesis of Virtuality 48
5.4.1 Architectural Overview 49
5.4.2 Validation and Execution Mechanisms 52
5.5 From Building Blocks to Executable Simulations 54

6 Example Use Cases of Virtuality 57
6.1 Use Case 1: Trauma Triage 58
6.2 Use Case 2: Stroke Neglect Study 62

6.3 Other Use Cases: Traffic and Defense Medicine 65

7 Overview and Motivation of Included Studies 67
7.1 From Interoperability Challenges to Practical Needs 68
7.2 Conceptualizing a Simulation Environment 69
7.3 Addressing Information Exchange and Data Representation Challenges 70
7.4 From Concept to Practical Implementation 71

8 Summary of included papers 73
8.1 Paper A (Scoping review protocol) 73
8.2 Paper B (Scoping review) 76
8.3 Paper C (Full Conference paper) 78
8.4 Paper D (Full conference paper) 79
8.5 Paper E (Scoping review) 80
8.6 Paper F (Framework paper - early draft) 81
8.7 Paper G (Architecture and design paper) 82

9 Conclusions and Final Remarks 85
9.1 Answering the Research Questions 86
9.2 Final Remarks and Future Outlook 88

References 91

II Papers 103

A IoMT Interoperability: Protocol A1

B IoMT Interoperability: Scoping Review B1

C Interoperability Requirements in HIS Procurement C1

D Introducing Virtuality D1

E FHIR mapping quality assessment E1

F F-MAPS F1

G Virtuality design and architecture **G1**

Acknowledgments

Looking back at these years, I am most of all grateful for how much I have learned and for all the people I have met along the way. First, I would like to thank my supervisors, Stefan Candefjord and Bengt Arne Sjöqvist, for your guidance, support, feedback, and many interesting discussions throughout my PhD. Thank you for challenging my ideas, sharing your knowledge and experience, and giving me the freedom to explore this journey in my own way. I have learned a lot from both of you, and for that I am truly grateful. Most of all, thank you for believing in me and in my ideas, and for giving me the opportunity and tools to pursue them. I would also like to thank my co-supervisors Hans Blomberg and Peter Lundgren for your valuable feedback, insights, and interesting discussions.

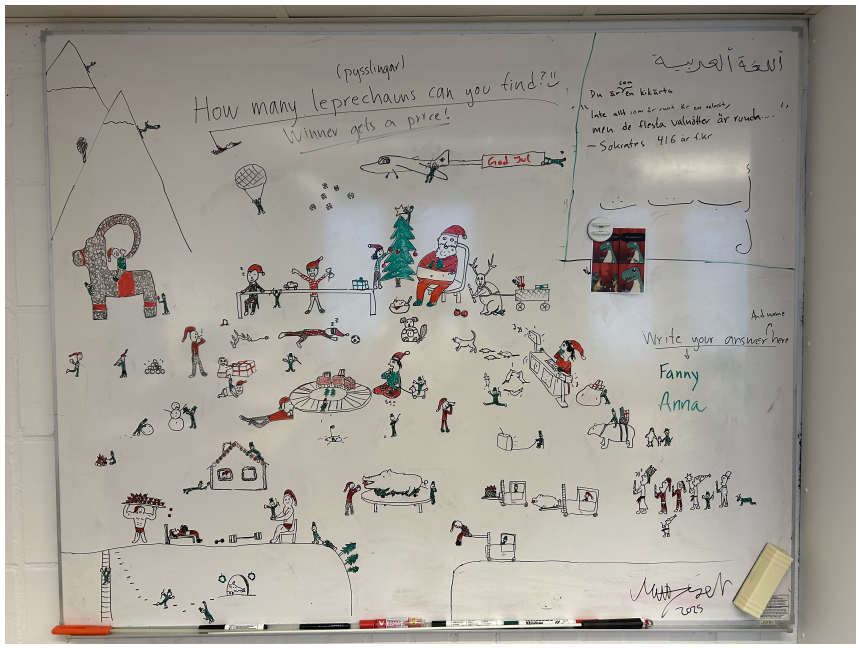
A special thanks goes to my fellow PhD colleagues and friends in the Care@Distance research group, Anna, Hoor, and Fanny. Thank you for making these years so much more enjoyable. From conferences and hotel breakfasts to pre- and after-works, board-games, research discussions, and everything in between, you have been a constant source of support, laughter, and motivation. Thank you for sharing both the good and frustrating parts of life as a PhD student, and for always motivating me to keep the whiteboard busy with my questionable art. A special thank you also goes to Victor and Sofie for supporting my more artistic side. You have made many ordinary days at work a lot more enjoyable. I have learned a lot from all of you, but most importantly, I am very grateful that I got to share this journey with you.

A big thank you also goes to all collaborators, co-authors, healthcare professionals, researchers, engineers, students, and project partners who have contributed to the work presented in this thesis. Much of this research has been shaped by practical projects and collaboration across disciplines, and I am grateful to everyone who has shared their time, expertise, and ideas.

To my friends and family, thank you for your support and encouragement.

Finally, and most importantly, to Cecilia. Thank you for always supporting me, encouraging me, and making sure I stay on track. You have been there through both the good and the stressful parts of these years and have always reminded me that there is more to life than work. This thesis would simply not have been possible without you.

Mattias Seth,
Hällingsjö, September 2026



Not all PhD work resulted in publications.

Statement on the use of artificial intelligence

Generative AI tools (GPT-5.6 Sol) were used to assist with language editing, improving readability, brainstorming, and feedback on the structure of this thesis. All scientific content, analyses, interpretations, and conclusions are the author's own. The author takes full responsibility for the accuracy and content of the thesis.

Acronyms

AI:	Artificial Intelligence
API:	Application Programming Interface
ASAP:	Acute Support Assessment and Prioritizing

CDSS:	Clinical Decision Support System
DSR:	Design Science Research
EEHRxF:	European Electronic Health Record Exchange Format
EHDS:	European Health Data Space
EHR:	Electronic Health Record
EIF:	European Interoperability Framework
EMS:	Emergency Medical Services
FAIR:	Findable, Accessible, Interoperable and Reusable
FHIR:	Fast Healthcare Interoperability Resources
F-MAPS:	FHIR Mapping Assessment and Performance Score
GUI:	Graphical User Interface
HIS:	Health Information System
HL7:	Health Level Seven
ICD:	International Classification of Diseases
IoMT:	Internet of Medical Things
IoT:	Internet of Things
IP:	Intellectual Property
IT:	Information Technology
JSON:	JavaScript Object Notation
LOINC:	Logical Observation Identifiers Names and Codes
MAIN:	Model for AI Utilization
MDR:	Medical Device Regulation
OECD:	Organisation for Economic Co-operation and Development

OSISP:	On-Scene Injury Severity Prediction
OWL:	Web Ontology Language
RPM:	Remote Patient Monitoring
SHACL:	Shapes Constraint Language
SNOMED CT:	Systematized Nomenclature of Medicine – Clinical Terms
TCP/IP:	Transmission Control Protocol / Internet Protocol
UI:	User Interface
VIPHS:	Verified Innovation Process for Healthcare Solutions
XML:	eXtensible Markup Language

Part I

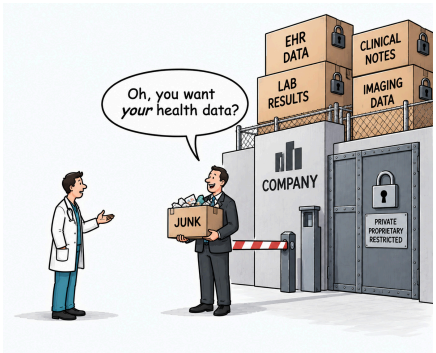
Overview

Preface

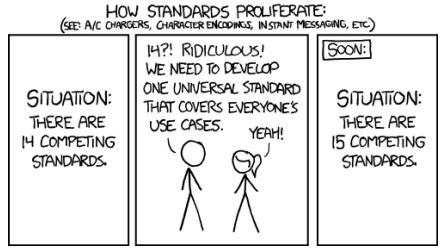
During my first year as a PhD student, I came across several memes about interoperability that I found both funny and unexpectedly insightful (Figure 1). It was not until later that I realized they had captured something that would become central throughout my research: although interoperability is widely recognized as important, it is often difficult to define precisely and even harder to achieve in practice. These memes were later used in my teaching in digital health at Chalmers.

To better understand interoperability and what it means in practice, I like to think of the well-known telephone game. In the game, a message is passed (often whispered) from one person to another. Due to various reasons, such as mishearing, misinterpretation, or simply poor memory, the message often changes along the way. By the time it reaches the final person, the initial meaning of the message may have been distorted, altered, or even completely lost. From a purely technical perspective, communication was successful. However, from a practical perspective, one could probably question whether the communication really was successful or not.

Healthcare information systems (HIS) often behave in much the same way as people in the telephone game. Data can be exchanged through various mechanisms, while differences in systems' internal data models, terminologies, assumptions, and workflows still cause meaning to shift along the way. In other words, successfully exchanging information does not necessarily mean that it



(a) The healthcare interoperability paradox: More digital health data and systems do not necessarily lead to better information exchange or integrated care. Illustration created by Mattias Seth, 2026 using ChatGPT.



(b) The introduction of additional standards to address interoperability challenges can itself contribute to further fragmentation. Source: Randall Munroe, xkcd. Licensed under CC BY-NC 2.5.

Figure 1: Illustrations of two recurring interoperability challenges in digital health.

has been correctly understood or can be meaningfully used.

This brings us to (one of many definitions of) interoperability, often defined as the ability of different systems to exchange, interpret, and use information in a standardized manner. Regardless of definition, interoperability is a fundamental prerequisite for the effective use of healthcare data, enabling information to flow across systems and organizational boundaries while supporting coordinated care, clinical decision-making, and digital innovation.

Due to its importance, researchers, healthcare personnel, engineers, and entrepreneurs continue to work towards ensuring that new digital health innovations are interoperable. Unfortunately, despite these efforts, many HIS implementation initiatives continue to fail.

This dilemma shaped my PhD work: if interoperability is both important for the successful use of digital technologies and, at the same time, highly complex, how can we ensure that digital health innovations and research ideas reach patients and clinicians, and that they provide the intended value?

Throughout my PhD, it became clear that existing tools could support parts of the problem, but there seemed to be a missing component. At the same time, regulatory bodies had been legislating and recognizing the need

for testing environments, so-called *sandboxes*, in which new innovations could be studied in a more systematic and realistic way, without exposing any risks to patients. This eventually led me to Virtuality - Virtual Care Process Simulator. Rather than being developed from the start, Virtuality grew out of the needs, insights, and practical hurdles I struggled with in research projects and practical implementation work. Through my studies and projects, it became clear that I'm not alone in thinking this. A virtual, digital, isolated testing environment was requested by many others. In this thesis, I will guide you through this journey: starting from trying to better understand interoperability in order to improve decision precision for falls among the elderly, to exploring how a Digital Health Sandbox might help make digital health innovations more testable, visible, and practically manageable. Hopefully, this work can guide future sandbox development initiatives.

CHAPTER 1

Thesis outline

This thesis is divided into two parts. Part I provides an overview of the research and describes how the work evolved from early studies of healthcare interoperability, project outcomes, expertise within my research group at Chalmers University of Technology, Care@Distance, triple-helix collaborations and regulatory drivers, towards the design and implementation of a Digital Health Sandbox.

Chapter 2 introduces the background, problem statement, research trajectory, and the motivation for Digital Health Sandboxes. Thereafter, the aim of the thesis and the research questions are presented.

Chapter 3 presents the conceptual and technical foundations of the thesis, including healthcare interoperability, digital health infrastructures, interoperability standards, and current approaches to testing and validation.

Chapter 4 describes the frameworks, concepts, research projects, and practical experiences that guided the research and gradually shaped the development of Virtuality.

Chapter 5 presents the design and development of Virtuality, from the initial prototype and system requirements to the architectural synthesis of the platform and its validation mechanisms.

Chapter 6 demonstrates how Virtuality can be applied through two representative healthcare use cases.

Chapter 7 explains how the included studies relate to one another and how they collectively contributed to the evolution of the research and the development of Virtuality.

Chapter 8 summarizes the scientific papers included in the thesis and highlights their individual contributions.

Finally, Chapter 9 synthesizes the findings of the thesis by answering the research questions, discussing the overall contributions, limitations, and implications of the work, and concluding with final remarks and future research directions.

Part II contains the seven scientific papers upon which this thesis is based.

CHAPTER 2

Introduction

"While it's relatively straightforward to edit a software component or replace one element with another, it's nowhere near as easy to change an organization."

— George Westerman

2.1 Background

Healthcare systems are under growing pressure to deliver high-quality care under increasingly difficult conditions [1]. Aging populations, multimorbidity, constrained resources, and rising expectations on accessibility and efficiency are driving the need for digital solutions that can support care delivery across all healthcare domains [2]. In parallel, advances in areas such as remote patient monitoring (RPM), sensor technologies, artificial intelligence (AI), and clinical decision support systems (CDSSs) have expanded the range of available tools, with the potential to improve care delivery and support more efficient use of

limited resources [3].

However, despite technical progress and a growing number of scientific contributions, transition from promising prototype to a clinically ready product remains difficult. Many digital health innovations never progress beyond research and pilot stage, and those that do are frequently associated with major problems [4], [5]. Solutions that perform well in controlled and isolated environments, may fail when introduced into real healthcare settings, not necessarily because the technology is ineffective, but because successful implementation requires more than technical functionality [6]. Factors including interoperability, healthcare workflows, regulatory compliance, security, trust and clinical acceptance are equally important [7].

Some of these challenges have been highlighted through several failed interoperability and implementation projects surrounding large-scale health information systems (HIS) [8]. Although standards, frameworks, and interoperability initiatives have expanded the options available to support these processes, similar projects continue to fail [9]. Questions therefore remain: how can innovations be tested, validated, and refined before they are introduced in healthcare?

2.2 Problem Statement

Over the past decade, healthcare has seen rapid growth in digital innovations aimed at improving care delivery [10]. More recently, this development has been further accelerated by the emergence of AI-powered development environments, cloud infrastructure, and various digital services, which together make it possible to translate ideas into functional prototypes faster than ever before [11], [12], even for users without advanced technical expertise. In this sense, some of the constraints that historically shaped software development, such as slower development and release cycles, rigid system boundaries, expensive and limited processing power have become less dominant. In principle, this creates new opportunities for multidisciplinary teams to explore and develop clinically relevant digital health solutions. However, while these tools accelerate innovation, they do not remove the need to ensure that resulting solutions can be integrated, validated, and safely used in complex healthcare environments. Nor do they guarantee that innovations, such as AI-based CDSS [13], become practically usable, clinically meaningful, and available be-

yond the research or prototype stage [7].

However, achieving this is not trivial. The path from an initial idea to a clinically validated product is long and involves a wide range of steps. In many cases, these steps are becoming increasingly difficult to navigate, placing a growing burden on researchers, developers, and entrepreneurs. One reason to this is due to the evolving European regulatory and policy landscape, including the European Health Data Space (EHDS), the European Union AI Act, and the Medical Device Regulation (MDR). Together, these frameworks reinforce the importance of trustworthy data use, interoperability, accountability, clinical safety, and evidence for digital health technologies, particularly when software and AI-based functionality are intended to support clinical decision-making [14], [15]. Although the intent to protect patients is well intentioned, these requirements intensify the need for systematic approaches to test, validate, and evaluate digital innovations before they are introduced into routine care.

This tension raises a practical and methodological question for digital health research and innovation: how can cross-disciplinary teams take advantage of the increasingly rapid pace of digital transformation and the opportunities it creates, while simultaneously receiving the support needed to test and validate innovations with the aim of bringing promising ideas closer to clinical practice?

Existing interoperability efforts and testing approaches do not fully address this challenge. Traditional tools, such as conformance testing, standard validation tools, and interface-level verification, remain important for assessing certain aspects of technical adherence. However, such approaches do not necessarily provide a means to evaluate how multiple components interact within a realistic healthcare workflow, how data are transformed and interpreted across heterogeneous systems, or how digital innovations (and its users) behave when introduced into a broader clinical process. As a result, there remains a gap between the rapid creation of digital health prototypes and the ability to systematically evaluate whether those solutions actually bring any real clinical value.

This thesis addresses that gap by investigating how a Digital Health Sandbox can be designed to bridge this gap. More specifically, the thesis contributes the development of Virtuality- Virtual Care Process Simulator, a web-based simulation environment intended to enable controlled yet realistic testing and validation of digital healthcare solutions.

2.3 From innovation to safe implementation

The problem outlined above is not only how to develop digital innovations, but how to progress them toward clinical implementation. In digital health research, technical feasibility and scientific publication are important milestones. They do not, however, answer whether an innovation can create meaningful value in clinical practice. Between a functioning prototype and a clinically usable solution lies a translational phase with several challenges, including access to high-quality data, dependencies on external systems and vendors, interoperability with surrounding legacy Information Technology (IT) infrastructures, workflow integration, long-term maintenance, intellectual property (IP), clinical utility, and regulatory readiness.

Within the Care@Distance research group at Chalmers University of Technology, these challenges repeatedly emerged across several triple-helix projects, involving academia, industry and the public sector. It became clear that a successful prototype or validated AI algorithm was only one part of the journey (or rather the very beginning of that journey). A central issue was not only whether an innovation worked, but whether it could be matured into something meaningful and sustainable in a real healthcare setting?

These observations contributed to the development of several frameworks within the research group, including VIPHS (Verified Innovation Process for Healthcare Solutions) [16], MAIN (Model for AI Utilization), and ASAP (Acute Support Assessment and Prioritizing). Although these frameworks are developed in different project contexts, they share a common concern involving how digital health innovations can progress beyond isolated technical results through staged validation, interdisciplinary collaboration, and attention to implementation conditions. Together, they form an important conceptual foundation for this thesis. The frameworks that informed my research are described in more detail in Chapter 4.

2.4 Research Trajectory

Although this thesis is mainly based on Virtuality, my research did not begin with the ambition of designing a Digital Health Sandbox (Figure 2.1). Instead, my initial work focused on interoperable Internet of Medical Things (IoMT) solutions for home and prehospital care, together with practical im-

plementation projects within the Care@Distance research group.

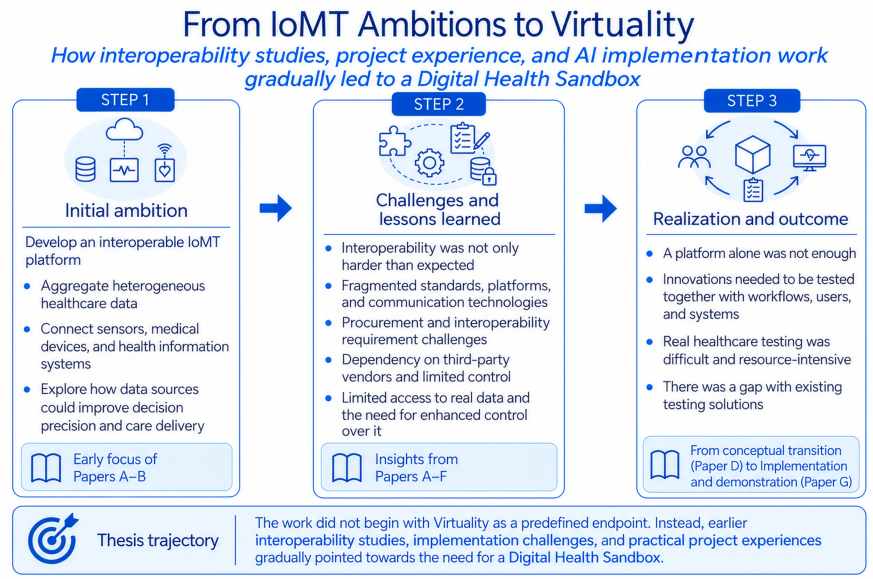


Figure 2.1: The trajectory of this project began with an attempt to better understand interoperability in order to develop an Internet of Medical Things (IoMT) platform. However, as the research progressed, new insights and practical implementation challenges emerged, gradually shifting the focus of the project. (Created with ChatGPT, 2026)

During these first projects, several recurring challenges became apparent. While interoperability between systems proved more difficult than expected, additional obstacles also emerged, including limited access to realistic healthcare data, dependence on third-party platforms, devices and applications, and the lack of environments where digital innovations could be tested before deployment. These observations gradually shifted the focus of the research from developing individual digital health applications towards investigating how such innovations could be developed, integrated, tested, and evaluated more systematically.

Figure 2.1 illustrates this evolution. The figure should not be interpreted as a predefined research plan, but rather as a reflection of how the direction of the thesis was shaped.

2.5 Why Digital Health Sandboxes Matter

Innovation in many sectors depends on environments in which new technologies, workflows, and practices can be explored before they are introduced into routine use. In healthcare, this principle is already well established through simulation-based training, where controlled environments are used to support learning, repeated practice, and experimentation without exposing patients to unnecessary risk [17]. More broadly, such environments serve an important translational role by allowing innovations to be developed, tested, refined, and assessed before they are implemented in real-world care settings.

These environments can take different forms depending on their purpose, level of realism, and position in the innovation process [18]. Common examples include laboratory environments, simulation environments, test beds, living labs, and sandboxes (Figure 2.2).

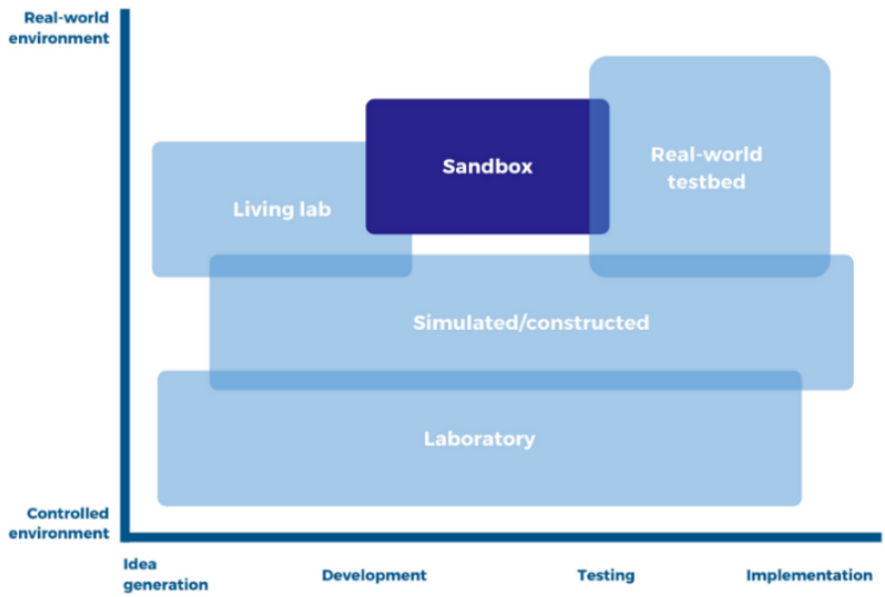


Figure 2.2: Different types of testing environments. Sandboxes are often positioned later in the innovation process than early laboratory work and are intended to provide a bounded space in which technologies, services, or workflows can be explored and evaluated before wider implementation. Adapted from Leckenby et al. [18].

Although these concepts overlap to some extent, they differ in terms of what functionality they are designed to provide, the degree of realism they provide, the extent of end-user involvement, and the type of evidence or learning they are intended to generate [18]. Some primarily support technical aspects under tightly controlled conditions (e.g., laboratory environments and certain test beds), whereas others emphasize workflow realism (e.g., simulation environments), organizational learning and co-creation (e.g., living labs), or experimentation in near-real-world settings (e.g., sandboxes).

Among these different forms of testing environments, the concept of the *sandbox* has become increasingly prominent, not least in relation to the governance of digital and AI-based innovations. This growing relevance is reflected in the regulatory landscape, where the European Union, through the AI Act, mandates all Member States to ensure that at least one AI regulatory sandbox is established at national level and operational by 2 August 2026 [14]. Although this deadline has now passed, it remains unclear how many Member States have fulfilled this requirement.

In the regulatory context, a sandbox is typically understood as a controlled environment in which innovative products, services, or business models can be tested under defined conditions and regulatory oversight [19]. Importantly, the sandbox concept does not refer only to the act of testing itself. It also encompasses the surrounding structures needed to make such testing meaningful and responsible, including planning, stakeholder engagement, risk management, data collection, monitoring, evaluation, and learning from the experiment [19]. In other words, a sandbox should not only be seen as a technical setup, but rather as an organized framework for experimentation and knowledge generation.

However, the sandbox concept used in this thesis differs somewhat from its dominant regulatory interpretation. While the characteristics described by the Organisation for Economic Co-operation and Development (OECD) [19] remain highly relevant, the challenge addressed in this thesis is not primarily one of regulatory flexibility or supervised market preparation. Rather, it concerns how digital health innovations can be tested and validated in a way that reflects the complexity of healthcare systems, with the ambition to move ideas and innovations as close as possible to successful clinical implementation, while keeping the end user involved throughout the process.

Accordingly, this thesis adopts a broader and more technical interpretation

of the sandbox concept. In this sense, the sandbox is seen as a healthcare-oriented testing environment that supports technical integration, simulation, and coordinated evaluation of heterogeneous components, including applications, AI models, medical devices, data sources, scenario definitions, and user-facing interfaces. It also needs to accommodate practical concerns that are central to healthcare innovation, such as data management, confidentiality, traceability, and intellectual property.

However, despite growing interest in sandbox-oriented approaches, there is still limited knowledge about how such environments should be designed and implemented in practice [20]. This thesis therefore addresses not only the need for such an environment, but also how it can be structured and developed. This is where the *Virtuality - Virtual Care Process Simulator* comes in.

Virtuality can be understood as a response to the challenges and gaps outlined in the preceding sections. Building upon research outcomes, ideas, expertise, and practical knowledge within the Care@Distance research team, Virtuality has been developed to address these challenges and gaps.

A central ambition is to reduce the risk that promising innovations and ideas remain as prototypes, pilot studies, or publications once a research project ends, and instead create better conditions for them to progress toward clinical use.

2.6 Thesis Aim and Research Questions

Given this background, the aim of this thesis is to investigate how digital health innovations, particularly interoperable and AI-enabled CDSS, can be systematically tested and validated prior to clinical deployment, and to contribute the conceptual, architectural, and technical foundations of a Digital Health Sandbox for such pre-deployment simulation and validation.

To address the aim of this thesis, the thesis is guided by one overarching research question (RQ) and three secondary research questions (SQ1–SQ3). Together, these questions form a sequential research process in which each question builds upon the findings of the previous one and contributes towards answering the overall research objective.

Throughout this thesis, Virtuality, Virtuality – Virtual Care Process Simulator, and Digital Health Simulation Environment refer to the same platform. The term Digital Health Sandbox (or simply sandbox) refers to the sandbox

environment provided by Virtuality. Unless otherwise stated, the platform is simply referred to as Virtuality. The following terminology is use throughout this thesis:

- **Virtuality**, **Virtuality – Virtual Care Process Simulator**, and **Digital Health Simulation Environment** refer to the same platform. The term **Digital Health Simulation Environment** is used because the research showed that supporting digital health innovation requires functionality beyond a traditional sandbox.
- **Digital Health Sandbox** (or simply **sandbox**) refers to one of Virtuality's central capabilities, providing a controlled environment for developing, integrating, simulating, testing, and evaluating digital health innovations.
- **Simulation** refers to the imitation of real-world healthcare processes and events within a controlled environment. By reproducing relevant clinical workflows, events, and interactions between digital systems, simulations enable digital health innovations to be tested and evaluated under defined and realistic conditions before clinical deployment.

To summarize, the *sandbox* provides the controlled environment for experimentation, while a *simulation* provides a means of representing and executing healthcare scenarios within that environment. *Virtuality* combines these capabilities in a single platform.

Main Research Question

> **RQ:** How can a Digital Health Sandbox be designed and implemented to support the pre-deployment simulation, testing, and validation of interoperable digital health innovations in healthcare?

To answer this overarching question, the thesis addresses three secondary research questions:

SQ1. Understanding the problem space: What characterizes the current landscape of digital health innovations and the interoperability-related challenges that arise across their development, procurement, integration,

and use in healthcare settings?
Addressed primarily in: (Paper A), Paper B, Paper C

SQ2. Deriving sandbox requirements and architecture: What design requirements, architectural principles, and core components are needed for a Digital Health Sandbox in which digital health innovations can be integrated, simulated, tested, and evaluated prior to clinical deployment?
Addressed primarily in: Paper D, Paper E, Paper G
Contributing to: Paper C, Paper F

SQ3. Instantiation and evaluation: How can the identified requirements, architectural principles, and interoperability mechanisms be instantiated in a working Digital Health Sandbox, and how can such an environment be demonstrated through realistic healthcare use cases?
Addressed primarily in: Paper G

Taken together, these questions reflect a sequential research process in which the problem space is first characterized, then translated into architectural and interoperability-related requirements, and finally instantiated and evaluated through the development of *Virtuality*.

CHAPTER 3

Interoperability and Digital Health Foundations

"Standards are good because you have so many to choose from."

— Andrew Tanenbaum

This chapter introduces the conceptual and technical background needed to understand the work presented in this thesis. The chapter does not aim to provide a complete overview of all aspects of healthcare digitalization or healthcare interoperability. Rather, it focuses on the concepts, challenges, and technological foundations most relevant to the research problem addressed in this thesis.

The chapter begins by outlining the complexity of contemporary healthcare information environments and the role of healthcare interoperability. It then introduces key standards and architectural approaches relevant to digital health innovation, and implementation-oriented interoperability efforts. Finally, the chapter reviews current approaches to interoperability testing and validation, and discusses why these are often insufficient.

3.1 Healthcare as a Complex Sociotechnical Environment

Healthcare is a highly complex and regulated domain in which digital technologies must operate across organizational, technical, and clinical boundaries. The Care delivery depends on the complex and coordinated interaction of professionals, information systems, medical devices, administrative processes, legal frameworks, and patient-facing services. Such complex environments, where people, tasks, tools, technologies, the physical environment, and organizational structures interact, are referred to as socio-technical systems [21]. These socio-technical characteristics must be considered when developing digital innovations for healthcare.

A central challenge is that many of today's healthcare IT systems were developed for local or organization-specific purposes, and not designed to interoperate with a broader digital ecosystem. Legacy electronic health records (EHRs), laboratory systems, imaging systems, administrative systems, and device platforms were often developed according to older design paradigms, meaning that their internal architectures differs from the data models, integration mechanisms, terminology systems, and interfaces that are frequently used today [22].

These discrepancies create challenges for the successful implementation of digital health innovations. This is because successful implementation cannot be reduced to whether an application or algorithm performs well in isolation. It also depends on whether the innovation can exchange information with surrounding systems, fit into existing clinical workflows, and be managed within relevant regulatory, organizational, and technical constraints. For this reason, interoperability is not merely a desirable technical feature in healthcare, but a central prerequisite for digital innovation to function in practice [7].

3.2 Interoperability in Healthcare

Interoperability is commonly understood as the ability of various systems, devices, applications, or organizations to exchange, interpret and use information in a standardized way [23]. However, as discussed in previous sections, the term is often regarded as a complex, multilayered concept, involving more than technical aspects [24].

As a result, several interoperability models have been proposed to describe its different dimensions. Although the terminology varies across these models, the overall principle is similar: successful interoperability requires more than simply transmitting data from one system to another.

3.2.1 Technical, Syntactic, and Semantic Interoperability

Considering the rapid growth of health data over the years, interoperability is not only essential for data exchange, but also a prerequisite for systems and applications to also search, find, (re-)use and analyze data (FAIR Data Principles) [25], [26]. In the development of AI applications, where access to high-quality data is essential, the FAIR Data Principles are particularly important.

Achieving these objectives, however, requires interoperability at several different levels. To illustrate these levels, various interoperability models are used to divide the concept into different layers [24] (Figure 3.1). At the most fundamental interoperability layer (or level), sometimes referred to as *technical* or *foundational* layer, interoperability concerns the basic ability of systems to connect and exchange data [23]. In modern healthcare systems, this often involves Internet-based communication using protocols (rules for how data should be sent and received) such as TCP/IP (Transmission Control Protocol/Internet Protocol).

However, even if two systems can communicate technically, they must still agree on how data should be organized in order for it to be processed. This level of interoperability is often called *syntactic or structural interoperability* and include message formats, such as XML (eXtensible Markup Language) or JSON (JavaScript Object Notation). Together with technical interoperability, fulfilling these two layers enables receiving systems to parse and handle incoming data in a predictable way [27].

To prevent systems from misinterpreting or misusing exchanged data, a third layer of interoperability is required: *semantic interoperability.* This layer addresses whether the receiving system interprets exchanged information in the same way as the sending system intended. To achieve semantic interoperability, systems need to adhere to shared information models, coding systems, and terminology standards such as SNOMED CT (Systematized Nomenclature of Medicine – Clinical Terms), LOINC (Logical Observation Identifiers Names and Codes) and ICD (International Classification of Diseases).

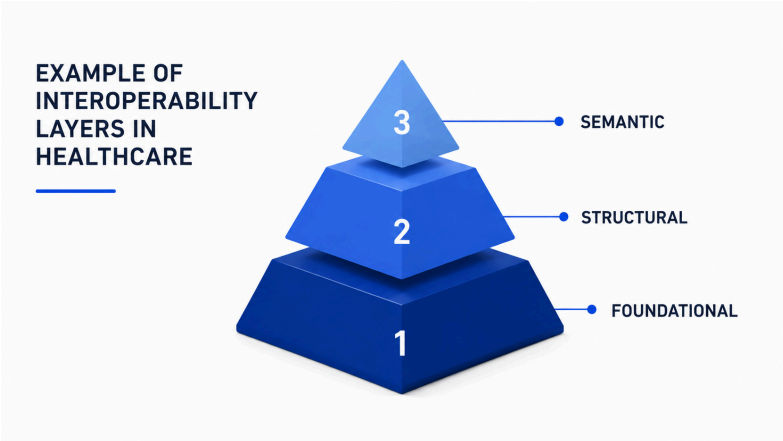


Figure 3.1: Example of a hierarchical interoperability model comprising three layers: foundational, structural, and semantic interoperability. Each higher layer depends on the successful achievement of the preceding lower layers (Created with GPT-5.6 Sol, 2026).

3.2.2 Beyond Technical Data Exchange

In healthcare, technical, syntactic, and semantic interoperability are necessary, but they are rarely sufficient [23], [28], [29]. Digital systems are embedded within care processes, professional routines, organizational structures, and legal frameworks that often differ across healthcare domains [30]. Because patients are likely to receive care from multiple institutions and across several domains throughout their lives, it is essential that systems operating within each domain are interoperable. For that reason, interoperability must also be understood as an organizational, legal and workflow related concept.

This is often addressed through organizational [29] or legal [31] interoperability layers, which concern how responsibilities, routines, coordination mechanisms, privacy legislation, and security requirements are aligned across actors and institutions [28], [29], [31]. In some contexts, these broader dimensions of interoperability are discussed in terms of *cross-domain* - or *cross-border interoperability* [32], [33]. This broader view of interoperability is also increasingly reflected in European initiatives aimed at enabling health data exchange across national boundaries. Examples include the ongoing work on a Euro-

pean Electronic Health Record Exchange Format (EEHRxF), which provides common specifications for interoperable exchange of EHRs and has become a key component of the European Health Data Space (EHDS) to support cross-border exchange of health information [34], [35], [36]. These initiatives are not only intended to improve data availability for healthcare delivery and innovation, but also to strengthen the role of patients by making it easier for individuals to access, share, and exercise greater control over their own health data. Related efforts seen in the European Interoperability Framework (EIF) further illustrate that interoperability is not only a technical concept [31].

Together, these perspectives are important in the present thesis, since the challenge addressed here is not simply whether one system can send data to another, or if the FAIR Data Principles are followed, but rather how these insights can be used in practice to support development, testing and validation of digital health innovations to support safe and meaningful clinical use.

3.3 Digital Health Infrastructures and the Internet of Medical Things

With the outbreak of the COVID-19 pandemic in March 2020 [37], the need for early detection, remote monitoring, quarantine support, social distancing measures, and faster coordination of care became increasingly apparent [38]. At the same time, long-term demographic trends such as population ageing, increasing multimorbidity, and a growing number of frail older adults living with chronic conditions, have further highlighted the potential of sensor technologies and Internet of Medical Things (IoMT) solutions in healthcare [39].

The IoMT is commonly described as the healthcare-oriented extension of the Internet of Things (IoT), comprising interconnected medical devices, sensors, software, and communication infrastructure that operate in highly regulated clinical environments [40].

From an interoperability perspective, IoMT environments are particularly challenging (and in my opinion, interesting) because they involve heterogeneous data sources, varying communication mechanisms, different device capabilities, proprietary solutions, and multiple organizational actors. This creates both technical and organizational complexity, especially when sensitive data must be aggregated, transformed, and reused in near real-time, across multiple applications and care settings [41], [42]. To highlight these challenges,

interoperability within this context is often defined using domain-specific interoperability models 3.2.

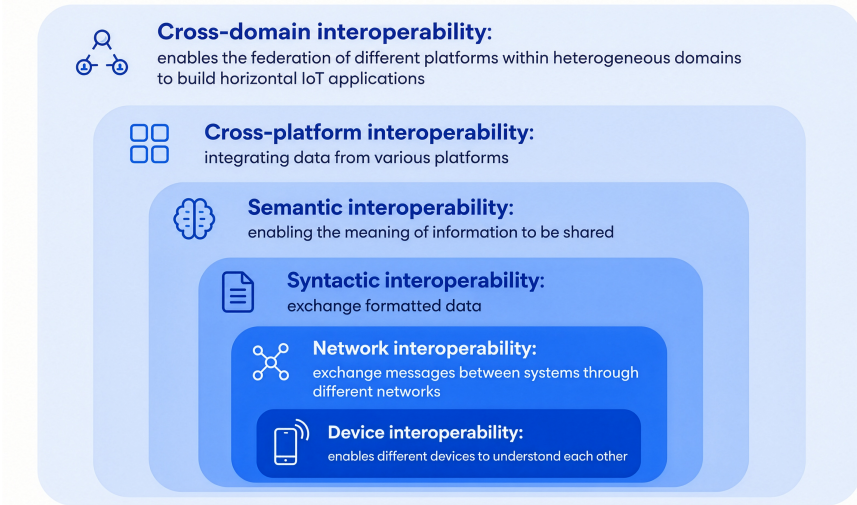


Figure 3.2: Example of an interoperability model adopted to the IoMT context. Adapted from Seth et al. [43], licensed under CC BY 4.0.

The studies included in this thesis emerged partly from this context. In particular, the work on interoperable IoMT platforms in home care and prehospital care in Paper A and B, highlighted the practical difficulty of selecting, combining, and implementing standards and technologies across distributed healthcare environments. These observations are an important part of the motivation for Virtuality as a Digital Health Sandbox, which further informed which technologies, standards and tools were relevant for the platform to support and use.

3.4 Standards and Frameworks for Healthcare Interoperability

A wide range of standards, frameworks, and implementation initiatives have been developed to improve healthcare interoperability, including IoMT. Although these efforts address different aspects of interoperability, they do not

eliminate the practical challenges of designing interoperable digital health solutions.

3.4.1 HL7 FHIR

Health Level Seven (HL7) Fast Healthcare Interoperability Resources (FHIR®) has become one of the most widely adopted standards for health data exchange, with recent international surveys indicating growing adoption across healthcare systems worldwide [44]. The development of FHIR began in 2011 and was first published as a Draft Standard for Trial Use in 2014 [45]. FHIR was developed to support more flexible and web-oriented exchange of healthcare data by combining healthcare-specific information modeling with technologies commonly used in modern software development, such as REST APIs (Representational State Transfer Application Programming Interface), JSON (JavaScript Object Notation), and XML (eXtensible Markup Language) [46].

Due to its primary focus on health-related information exchange between software platforms, FHIR is sometimes regarded as a platform specification [47]. To support this focus, the standard structures information into modular resources (FHIR Resources) representing clinical and administrative concepts such as patient demographics (Patient Resource), clinical observations (Observation Resource), medical conditions (Conditions Resource) and medications (Medications Resource). As of the latest FHIR version R5, there are currently 157 FHIR Resources that can be exchanged individually or assembled into larger structures such as Documents and Bundles [48]. FHIR also supports the use of external terminologies (such as SNOMED-CT) and profiles, allowing information to be tailored to particular use cases while still building on a common base standard.

CDS Hooks and SMART on FHIR

While FHIR provides a standardized way to structure and exchange healthcare data, it does not by itself specify how external CDSS or applications should be integrated into clinical workflows. To address the needs associated with advancements in AI and to facilitate integration of AI-based application and CDSS, complementary specifications based on HL7 FHIR have emerged. These include the CDS Hooks and SMART on FHIR specifications, where the challenge is not only to exchange data, but also to ensure that model outputs

can be triggered, accessed, and presented in a clinically meaningful way [49], [50], [51].

The HL7 CDS Hooks specification enables a clinical system, such as an EHR, to invoke (or call) an external CDSS at predefined points in a clinical workflow. This is achieved through so called *hooks*, which represent an event-driven integration mechanism that allows external CDS services to be triggered at specific moments, such as when a patient overview page is opened, a medication is prescribed, or an order is initiated. The external service receives contextual information about the current workflow state and can return suggestions, warnings, recommendations, or links to additional actions, based on the contextual information and how the CDSS has been implemented [49].

SMART on FHIR addresses a related but complementary need. Rather than focusing on event-driven hooks, SMART on FHIR provides a framework for securely launching third-party applications within or alongside HITs [50]. This makes it possible to build interoperable applications that can securely integrate into existing systems. This approach offer more advanced user interaction than is typically possible through lightweight decision-support messages alone.

Together, CDS Hooks and SMART on FHIR are highly relevant to recent developments in digital health and AI-based CDSS [52]. The two specifications should not be seen as competing alternatives, but as complementary mechanisms for integrating interoperable digital services into healthcare systems.

In the context of this thesis, HL7 FHIR alongside CDS Hooks and SMART on FHIR Specification is particularly relevant for two reasons. First, it is a well-established standard supported by a strong international community, including major technology companies [46], [53], [54], healthcare organizations, and national initiatives across Europe [15], [55]. Second, it plays a central role in enabling interoperability in healthcare and exemplify the type of standards-based integration mechanisms that a Digital Health Sandbox may need to accommodate when simulating, testing, and evaluating digital health innovations [56].

3.4.2 Other Interoperability Approaches and Implementation Frameworks

HL7 FHIR represents only one part of the broader interoperability landscape. Given the complexity of healthcare interoperability, no single technology or standard is sufficient on its own. Instead, multiple tools, standards, and approaches need to be combined. This is where Integrating the Healthcare Enterprise (IHE) becomes particularly relevant. IHE is an initiative by healthcare professionals and industry, aiming to improve healthcare interoperability [57]. Through its integration profiles and testing initiatives, such as Connectathons, IHE focuses on how existing standards can be combined and operationalized.

3.5 Interoperability Challenges in Practice

Although several standards, frameworks and interoperability initiatives, interoperability challenges in healthcare persist at multiple levels.

One recurring challenge concerns data mapping and transformation, where data originating from one system often need to be translated before they can be reused in another context [58]. Returning to the telephone game analogy introduced earlier, this can be compared to two people speaking different languages. For them to understand each other, mechanisms are needed to translate between the two languages while preserving the meaning of the message. Similarly, healthcare systems may require mappings between different data structures, formats, and terminologies to correctly exchange and interpret information (Figure 3.3).

However, data mapping is not a trivial task. If performed incorrectly, there is a risk that semantic meaning is misinterpreted or misaligned, which may in turn affect downstream decision support, simulation, analytics, and ultimately patient safety [59]. Although several tools exists to facilitate this process [60], [61], data mapping is not the only problem. Other recurring interoperability challenges include vendor lock-in, legacy systems that are difficult to integrate [62], contextual misalignment between systems and workflows, and governance, privacy, and access-related constraints [63]. These challenges may arise at different stages of the innovation and implementation process and are not always easy to anticipate. One way to structure and analyze them is to relate them to different interoperability levels. Doing so helps clarify where in

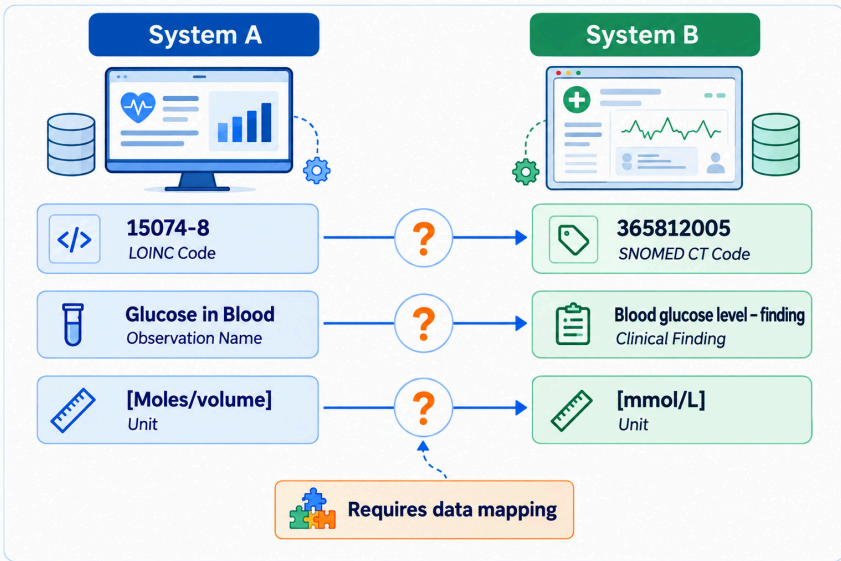


Figure 3.3: Example of how two systems need data mapping or data translation before they can understand each other. Even after mapping, there is a risk for semantic loss (Created with GPT-5.6 Sol, 2026).

the interoperability chain (and thus where in the software development lifecycle) a particular challenge emerges, what type of misalignment it represents, and which kinds of mitigation strategies may be required (Figure 3.4).

3.6 Testing and Validation of Healthcare Interoperability

Because interoperability has remained a challenge in healthcare for several decades, with problems arising at different levels and stages of development, a range of testing and validation approaches have been developed [64]. These include various sandboxes, conformance and integration tests, profile validation tools, and more recent collaborative interoperability initiatives such as Connectathons. Tools such as the HL7 FHIR Validator [65], open-access testing tools such as Touchstone and Crucible [66], test kits provided by Inferno

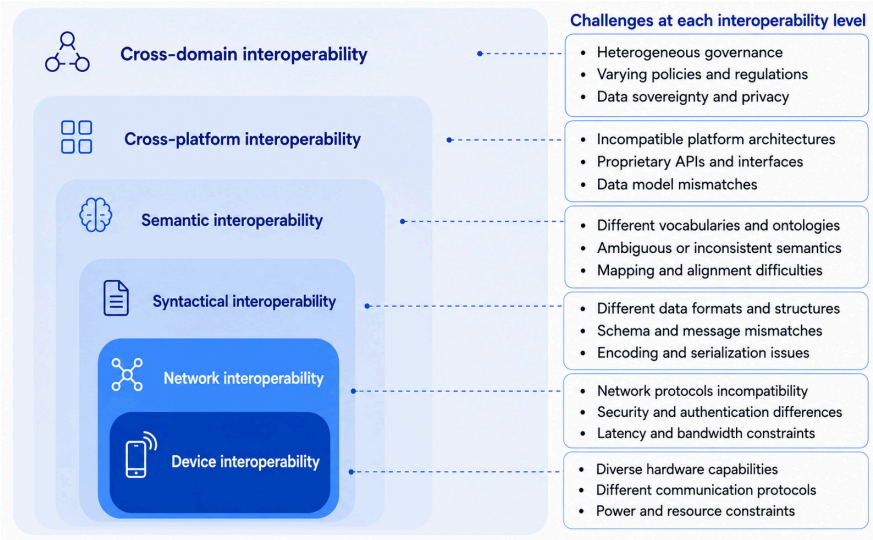


Figure 3.4: Example of various interoperability issues appear at different interoperability levels. Each interoperability level can be related to various stages the software development cycle, such as design, development, test and maintenance.

[67], and the open-source testbed Gazelle developed by IHE-Europe [68], [69], all support various aspects of interoperability testing. Furthermore, large-scale testing initiatives such as IHE Connectathons and Argonaut project [70] have also contributed significantly to improving technical interoperability between systems by enabling vendors and organizations to come together and test specific interactions under controlled conditions.

Although these approaches are valuable, they also have limitations [66]. They primarily assess whether particular messages (FHIR Resources), APIs, or transactions conform to predefined specifications. Typical aspects addressed by these tools are: "Does the server support the right FHIR REST interactions?", "Does it return the expected status codes/content types?" or "Does it conform to the required FHIR profile(s) in an implementation guide?". Although these approaches are valuable, they do not necessarily reveal how a digital innovation behaves when embedded within a broader clinical workflow. Nor do they, in an accessible manner, support the transition from early

idea and prototyping to evidence generation, validation, and eventual clinical implementation.

Taken together, the preceding sections illustrate that the barriers preventing successful implementation of digital health innovations rarely arise from a single technical issue. Existing tools typically address individual parts of this challenge, but rarely provide an environment in which these aspects can be investigated together.

3.7 How to Bridge the Gap

Imagine that you have developed an AI-based CDSS. You have collected data, trained and validated the model on external datasets, documented the methodology, and published the results in a well-recognized scientific journal.

Now what?

For many digital health innovations, this is where the journey ends. Although idea may be scientifically sound, as we have seen in preceding chapters, considerable work often remains before it can be used in clinical practice.

When these steps are not addressed, promising innovations risk remaining as prototypes, publications, or isolated research projects. This represents not only a loss of time and financial resources, but also of valuable knowledge, expertise, and innovation potential that could otherwise have contributed to improving healthcare.

This creates a central tension in contemporary digital health innovation. Although a wide range of interoperability standards, tools, frameworks, and testing approaches are available, it remains far less clear how they should be combined and applied in practice when a digital health innovation is to move beyond early prototyping.

This creates a need for environments that do more than verify isolated interfaces or technical conformance. There is a need for environments that can support users from different professional backgrounds in moving from early ideas and prototyping to structured validation, evidence generation, and validation, while also accommodating interoperability, workflow realism, and the documentation demands associated with healthcare innovation. It is within this context that the work presented in this thesis is positioned. More specifically, it motivates the development of *Virtuality - Care Process Simulator*, a simulation environment to bridge this gap.

CHAPTER 4

Frameworks, concepts and Project Outcomes Guiding the Research

> *"The best way to have a good idea is to have lots of ideas, and throw away the bad ones."*
>
> — Linus Pauling

My research is somewhat different from what is traditionally expected in many scientific disciplines. In a conventional setting, one typically starts with an idea, formulates a hypothesis, and then designs experiments, tests, or calculations to evaluate whether that hypothesis holds. Empirical data are collected, analyzed, and used to either support or reject the initial assumption.

However, within the field of applied digital health and software architecture, this approach is often not sufficient—or even feasible.

The systems being studied are not isolated phenomena that can easily be controlled or reproduced in a laboratory setting.

As a result, research in this domain often shifts from purely hypothesis-driven experimentation towards the design, development, and evaluation of

solutions in realistic—but controlled—environments. Rather than asking only "is this hypothesis true?", the question becomes: how can we design systems that work in practice, under real-world conditions, without introducing unacceptable risk?

These type of questions reflect ongoing discussions within the Care@Distance research team at Chalmers Univeristy of Technology.

4.1 Building Upon and Extending Existing Knowledge

The work presented in this thesis builds on a broader research trajectory within the Care@Distance research group. Over more than a decade, the group has investigated digital health solutions for remote and prehospital care, including simulation-based environments [71], prehospital clinical decision support [72], and the effects of information technology on clinical work and care processes [73]. Together, these studies contributed practical and conceptual knowledge about the challenges of developing digital tools that are not only technically feasible, but also clinically useful.

A recurring question in this work has been how promising research outputs can be moved beyond the stage of prototypes, algorithms, and scientific publications. For example, AI-based models such as the On Scene Injury Severity Prediction (OSISP) trauma algorithm, developed within the Care@Distance research group, raised questions that extended far beyond model performance alone. How should such models be packaged and governed? How should intellectual property, data access, and responsibility be handled? How can the model be integrated with existing clinical systems and workflows, while meeting regulatory requirements? And how can it be tested in a way that supports trust, safety, and eventual clinical use?

Similar questions motivated the need for more structured approaches to innovation and implementation. Within the research group, this led to the development of frameworks and concepts such as VIPHS, MAIN, and ASAP, which all address different aspects of how digital health innovations can be matured, evaluated, and moved closer to clinical use. Although these frameworks were developed in different project contexts, they share a common concern: the need for staged validation, cross-disciplinary collaboration, and early attention to implementation conditions.

Building on these experiences, this thesis takes the next step by investigating how such ideas can be operationalized in a technical environment. Before introducing Virtuality in detail, the following sections therefore describe the key models, frameworks, and concepts that have shaped its development.

Three central frameworks are considered:

- **VIPHS (Verified Innovation Process for Healthcare Solutions)** — explains *why* a structured and stepwise validation process is necessary.
- **ASAP (Acute Support Assessment and Prioritizing)** — describes *how* systems should be designed, emphasizing principles such as data fusion, AI, interoperability, and user interaction.
- **MAIN (Model for AI Utilization)** — defines *what* AI utilization in healthcare should look like at a conceptual and strategic level, empathizing the need for a practical bridge from research to clinical implementation.

Together, these frameworks provide a foundation that have informed the development of Virtuality.

4.1.1 Verified Innovation Process for Healthcare Solutions (VIPHS)

The challenges described in the previous sections—where digital health initiatives fail to move beyond pilot stages or struggle during clinical implementation—are not unique to this thesis. Similar observations have been reported in previous research, and learning outcomes from several failed transformation and implementation initiatives point in the same direction. To overcome common barriers, the Care@Distance research team has developed the Verified Innovation Process for Healthcare Solutions (VIPHS), a stepwise model developed to support digital health innovation from idea to clinical use [16] (Figure 4.1).

VIPHS was developed through case studies in prehospital care and is based on service design principles (iterative, collaborative, and human-centred). The model addresses a key gap in digital health: the lack of structured guidance for how to test, validate, and implement new solutions when both technology and clinical workflows are affected.

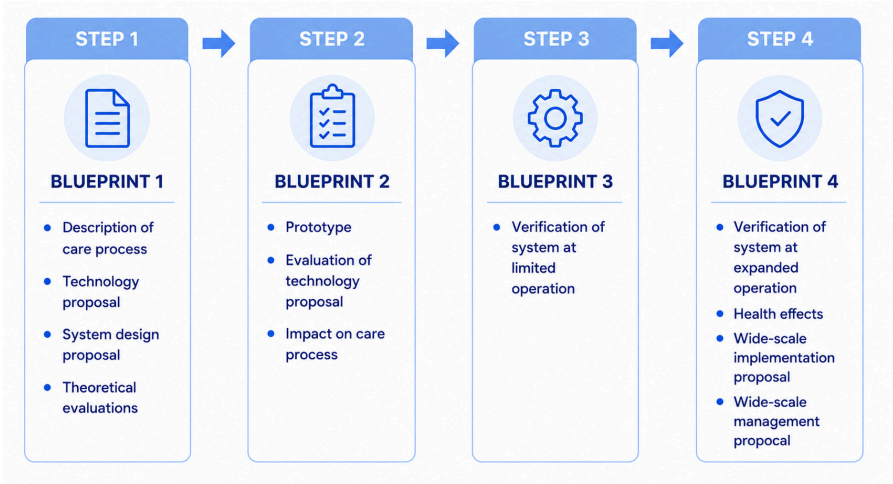


Figure 4.1: The VIPHS model describes a step-wise approach, where each step and associated outcomes create blueprints on how to proceed [16]. Used with permission.

At its core, VIPHS describes a step-wise process where innovation progresses through increasingly realistic environments—from early analysis and laboratory testing, to simulation, pilot studies, and eventually broader clinical deployment. Each step produces a "blueprint" that documents findings and reduces uncertainty before moving forward to the next step. This reflects an important idea: validation in healthcare is not a single event, but a continuous and iterative process. In this sense, VIPHS helps explain *why* a structured, stepwise approach is needed when it comes to maximizing clinical utilization of digital healthcare innovations.

4.1.2 Acute Support Assessment and Prioritizing (ASAP) Concept

If VIPHS establishes the *why*, the Acute Support Assessment and Prioritizing (ASAP) Concept—developed within the Care@Distance research initiative in 2022—provides important guidance on *how* Virtuality should be designed and implemented. ASAP is a generic concept that promotes open-source software, modularity and standardized tools to support data fusion (aggregat-

ing data from various sources), development, and implementation of CDSS, AI, telemedicine, and innovative user interaction in healthcare applications (Figure 4.2). Guided by the ASAP concept, Virtuality was developed using appropriate tools to support these desired capabilities.

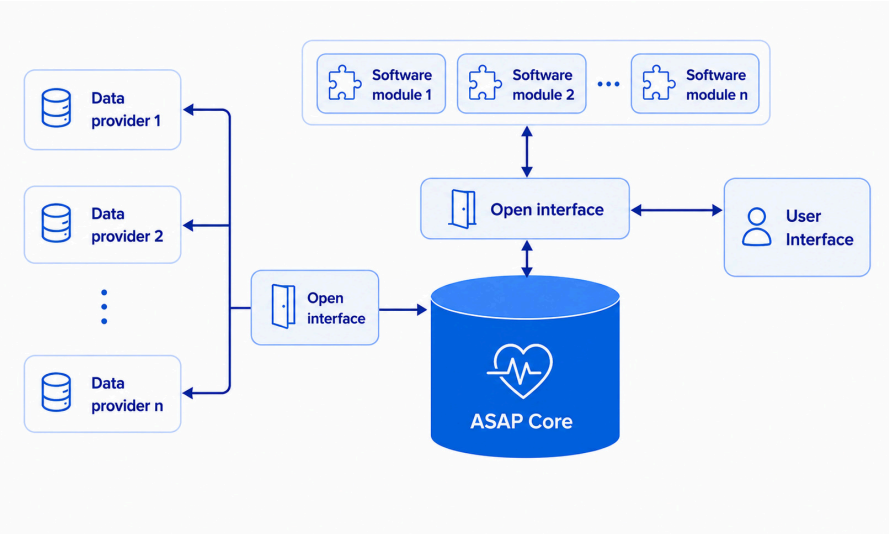


Figure 4.2: The Acute Support Assessment and Prioritizing (ASAP) Concept, illustrating how data aggregation can be supported by open-source software, modularity and standardized interfaces.

4.1.3 Model for AI Utilization (MAIN)

In 2022, Bengt Arne Sjöqvist, Professor Emeritus at Chalmers University of Technology, introduced and started to develop the Model for AI Utilization (MAIN) concept and framework.

MAIN is a conceptual framework that outlines how AI-based solutions can be effectively integrated into healthcare systems by aligning technological capabilities with clinical needs, organizational structures, and user contexts.

To achieve this, MAIN adopts a triple-helix perspective, emphasizing the importance of interaction between key stakeholders—such as healthcare providers, technology developers, and academic research— to increase the likelihood that

academically driven medical technology research, and in particular AI-based CDSS, will achieve utilization and upscaling in industry and healthcare. In this sense, MAIN can be seen as an organizational and collaboration model which highlights that successful AI utilization requires more than accurate algorithms. It requires innovations to be packaged and adapted in a way that provides opportunities to transfer research results into regulatory-approved products and solutions.

However, MAIN does not prescribe a specific implementation process. Rather, it identifies critical dimensions such as data availability, system integration, intellectual properties and CE marking, thereby framing the conditions and criteria necessary for achieving clinical impact. In relation to Virtuality, MAIN provides a broader perspective on where digital health innovations should ultimately lead and what is required to get there.

CHAPTER 5

Development of Virtuality

"Everything in software architecture is a trade-off."
— Mark Richards & Neal Ford

Up to this chapter, we have discussed various factors influencing the need for a dedicated testing environment. In this chapter, I will describe more in detail how Virtuality was developed.

5.1 Design Rationale and Objectives

Imagine a healthcare environment in which new ideas could emerge directly from clinical practice. An environment in which solutions are rapidly shaped by needs identified within daily work and in which digital innovations can be iteratively refined, tested, and validated by cross-disciplinary teams without compromising patient safety, data integrity, or regulatory requirements.

While iterative development and early user involvement are common in many other industries, the healthcare landscape makes such approaches con-

siderably more difficult. Stricter regulatory requirements, concerns around patient safety and data protection, and an increasing reliance on large-scale, monolithic systems have all reduced the capacity for local experimentation and rapid innovation. At the same time, healthcare professionals rarely have the time, organizational support, or technical resources required to actively participate in the iterative development and evaluation of new digital solutions [74].

However, recent advances in AI have fundamentally changed the conditions for innovation [75], making it possible to transform ideas into functional prototypes faster and with fewer technical barriers than before. One of the main objectives with Virtuality is to support healthcare in taking advantage of this leap, by ensuring that promising innovations and research outcomes do not stall at the prototype stage.

5.2 The "How"

The development of Virtuality did not begin with a specific technology stack or software platform. Instead, the process was inspired by a Design Science Research (DSR) approach [76], starting with a collection of concepts, frameworks, and lessons learned from previous research projects, which together helped shape the objectives and requirements of the platform. Later phases involved iterative design and implementation of the platform, followed by artifact (Virtuality) validation.

Paper A and B started this process [43], [77] by influencing where to look and gave a solid understanding of which technologies can be used and how they could be combined.

Rather than selecting a single framework or architectural paradigm, Virtuality was conceptualized as a synthesis of these experiences. The ambition was to create an environment capable of supporting the complete innovation lifecycle, from concept exploration and prototyping to simulation, validation, and pre-deployment validation.

5.2.1 Virtuality - The first prototype

The development of Virtuality began in 2022, when the Care@Distance research group had only a preliminary vision of what such an environment might

eventually become. At that stage, Virtuality was not yet regarded as a comprehensive Digital Health Sandbox, but rather as an initial attempt to create a digital environment in which healthcare innovations, based on the ASAP Concept, could be explored and evaluated before clinical deployment.

The first prototype was developed in the context of the Autumn Leaves project, a triple-helix collaboration between academia, healthcare, and industry aimed at creating an interoperable IoMT platform for the management of falls among older adults. The project involved multiple actors, data sources, and technical systems, and quickly exposed practical challenges related to healthcare interoperability. Accessing, sharing, and integrating data across organizational and technical boundaries proved difficult, and questions arose regarding which standards, technologies, and architectural approaches were most appropriate for enabling interoperability in practice. These challenges were not unique to the project, but have also been widely reported in the digital health and interoperability literature, including concerns about sharing valuable data, limited awareness of how data will be used, and a lack of understanding of the benefits of interoperable data sharing [78]. Together, these challenges and insights made it clear that there was a need for a simulation environment, with easier access to synthetic data and enhanced control over applications and technologies. Without such an environment, we saw a clear risk that future innovation work would stall. These early experiences became an important driver behind the development of Virtuality.

Technically, the first prototype was implemented as a web-based application using NuxtJS for the frontend, FastAPI for the backend, and PostgreSQL for data storage. Compared with later versions of Virtuality, the prototype was relatively simple. Its development was guided by six patient personas defined within the Autumn Leaves project, each representing a typical patient group encountered in prehospital or home care settings. The central idea was that data, regardless of origin (including data from sensors and third-party service providers), should be able to enter a common environment where they could be processed, combined, and evaluated in relation to a clinical use case. In this way, the prototype provided an initial basis for exploring how heterogeneous data sources and digital functionality could be brought together within a shared environment.

Although the prototype emerged from the Autumn Leaves project, it was later also used in early work related to the OSISP trauma model and its

evaluation. In that context, it provided an initial environment in which an AI-based model could be presented and assessed in a more realistic digital setting. Compared to running the model in isolation on a standalone machine, this use case highlighted additional requirements beyond data integration alone, including the need to package AI models as reusable services, present outputs through a user-facing application, and evaluate such functionality in relation to clinical workflows and user interaction. It was also during this work that I began to look more closely at standards and specifications such as HL7 FHIR, CDS Hooks, and SMART on FHIR.

At the same time, the work with the first prototype made clear that the initial implementation was not sufficient. It was closely tied to a limited number of use cases, had not been developed from a formalized set of design requirements, and provided only limited support for orchestration and reuse of components. The environment also lacked several of the mechanisms later considered necessary for a reusable simulation platform, including more systematic support for modular integration, scenario configuration, persistence, and traceability. In addition, the project exposed a number of practical barriers that slowed the innovation process, including dependencies on third-party vendors, limited control over access to external data sources, and the need to develop supporting applications and interfaces largely from scratch. Together, these challenges highlighted the need for a more structured and reusable environment in which healthcare innovations could be integrated, configured, and evaluated with greater technical and organizational independence.

Despite these limitations, the first prototype played an important role in the development of Virtuality. It helped communicate the concept to collaborators, supported early stakeholder feedback, and, perhaps most importantly, revealed practical and architectural challenges that would shape a set of requirements for the second generation of Virtuality.

5.2.2 System Requirements

The preceding chapters, including the first Virtuality prototype, identified recurring challenges related to healthcare interoperability, translational implementation, and the evaluation of AI-enabled digital health innovations. These challenges formed the starting point for the design requirements of Virtuality, which later informed the development of the second prototype. However, the requirements were not derived from the literature alone. They were

gradually refined throughout the research process and informed by several complementary sources, including the frameworks presented in the preceding chapter, learning outcomes from earlier research projects, limitations of available testing environments, discussions and practical experience from the Care@Distance research group.

Several of the requirements were also shaped by design-oriented work carried out in parallel with the development of Virtuality. In particular, two master's thesis projects supervised within the work of this thesis contributed important insights into how a simulation environment should be structured, visualized, and made usable for different stakeholders [79], [80]. Through interviews, field studies, and iterative design work, these projects provided input on what users perceived as meaningful simulation functionality and how such functionality could be represented in a graphical user interface (GUI).

In addition, practical implementation work related to the OSISP trauma model and its integration into an interoperable solution intended for use in research involving ambulance personnel [81]. That work highlighted challenges related to packaging AI models, integrating them with surrounding systems, preparing user-facing applications, and making such solutions usable in realistic healthcare contexts. These experiences provided concrete insights into the kinds of functionality, flexibility, and traceability that Virtuality would need to support.

At the same time, in 2024, while the second Virtuality prototype was starting to emerge, the AI Act, with the aim to foster trustworthy AI in Europe, was published in the Official Journal of the European Union [14], marking a significant milestone in the regulation of AI within the European Union.

Taken together, these insights, experiences and emerging regulatory directions shaped the design and development of the second prototype.

Table 5.1: Functional and non-functional requirements for Virtuality. Each requirement has been informed by one or several sources.

ID	Requirement	Description	Informed by
Simulation and usability			
R1	**Process-driven simulation**	The primary unit of simulation shall be a complete healthcare process rather than an isolated application, API, or algorithm.	VIPHS [16]; scientific literature; master's thesis projects [79], [80]

Continued on next page

ID	Requirement	Description	Informed by
R2	**Real-world fidelity and flexibility**	The environment shall support realistic representations of clinical workflows, stakeholder interactions, system dependencies, and operational constraints.	VIPHS [16]; Care@Distance research project outcomes
R3	**Iterative experimentation**	The environment shall support rapid modification, repetition, and refinement of simulations.	VIPHS [16]; master's thesis projects [79], [80], Care@Distance research project outcomes
R4	**Rapid setup and use**	Users shall be able to configure and execute simulations with minimal technical effort and within a short time frame.	master's thesis projects [79], [80]; Care@Distance research project outcomes
R5	**Multi-stakeholder accessibility**	The environment shall support collaboration between clinicians, researchers, developers, managers, auditors, and policymakers.	MAIN; VIPHS [16]; EU AI Act; master's thesis projects
R6	**Location independence**	The environment shall not depend on dedicated physical facilities or co-located participants.	MAIN, ASAP
Interoperability and platform architecture			
R7	**Standards-first interoperability**	All inter-component communication and data exchange shall support established healthcare interoperability standards.	MAIN, VIPHS, ASAP; master's thesis projects [79], [80], research within Care@Distance [43], [77]
R8	**Modularity, reuse, and extensibility**	Simulation components shall be reusable across scenarios, projects, and organizations, and new components shall be addable without modifying the core platform.	ASAP; master's thesis projects [79], [80]; Care@Distance research project outcomes
R9	**Synthetic data support**	The environment shall support the generation, management, and use of synthetic data, while allowing integration of real-world data when appropriate.	ASAP; research within Care@Distance [82]; master's thesis projects [79], [80], Care@Distance research projects
Governance, validation, and deployment			
R10	**Regulatory alignment**	The environment shall support evaluation activities aligned with relevant regulatory frameworks and standards.	MAIN; EU AI Act; VIPHS [16]
R11	**Persistence and reproducibility**	Simulation configurations, inputs, outputs, and execution states shall be persistently stored.	Simulation practice; reproducibility principles; AI evaluation needs

Continued on next page

ID	Requirement	Description	Informed by
R12	**Auditability and traceability**	The environment shall maintain a complete record of simulation events, data flows, user actions, and system decisions.	EU AI Act; MAIN
R13	**Deployability and environment flexibility**	The environment shall support deployment across different technical and organizational settings.	
R14	**Security and access control**	The environment shall provide role-based access control, authentication, authorization, and data protection mechanisms appropriate for healthcare and regulatory use.	GDPR; healthcare security requirements; AI governance expectations

These requirements guided the design and implementation of Virtuality and are reflected in the architectural components described in the following sections.

5.3 Virtuality - The Second Prototype

While the first prototype helped establish the technical foundations of Virtuality, several fundamental questions remained unanswered. How should users interact with the platform? What does a healthcare simulation actually look like in practice? How can complex healthcare processes be represented in a way that is understandable for both clinicians and technical stakeholders? Most importantly, *what is a simulation*, and which steps should a user go through to create one?

To explore these questions, two master's thesis projects were initiated in 2025 under my supervision. Rather than focusing on technical implementation, the projects focused on the user experience and visual design of Virtuality. The objective was to better understand how simulations should be represented, configured, and executed within the platform. Also, their insights were further used to iterate the requirements in Table 5.1.

The projects approached the challenge from two different perspectives. One group focused on falls among elderly patients receiving home care [79], while the second group focused on traffic accident scenarios [80]. These scenarios were deliberately chosen because they represent two different types of healthcare processes with different characteristics. The fall scenario emphasized long-term monitoring, home-based care, and preventive interventions, whereas the traffic accident scenario focused on acute care, emergency response, and time-critical decision-making. Together, the projects were intended to com-

plement one another and provide broader insight into how Virtuality could be designed to support diverse healthcare contexts.

Design questions explored in the master's thesis projects

- What should constitute a simulation within Virtuality?
- What information should be required before a simulation can be executed?
- How should users connect data sources, applications, and AI services?
- How can complex healthcare processes be visualized in an intuitive manner?
- How much technical detail should be exposed to different user groups?

Through workshops, iterative prototyping, and stakeholder discussions, the students explored how users could construct simulations without requiring programming knowledge. One of the most important outcomes was the realization that users naturally viewed simulations as compositions of reusable components rather than traditional software systems. Clinicians often described healthcare processes as a sequence of actors, decisions, and information exchanges, while technical stakeholders tended to think in terms of applications, services, and data flows [79], [80]. Interestingly, both perspectives could be represented using the same visual metaphor: connecting modular components together to form a complete workflow. They further helped identify the user journey required to move from an empty workspace to a completed simulation, laying the groundwork for the Scenario Engine and user interface implemented in the second prototype and described more in detail in Chapter 6.

The integration work surrounding the OSISP trauma model also contributed significantly to the design of Virtuality. During development, several practical questions emerged. What if a clinical scenario needed to be adjusted? Or what if we wanted to simulate a scenario in which certain clinical data became available only later in the care process?? How would these changes influence application behavior, user interactions, and the overall execution of the clinical

workflow?

Addressing these questions highlighted the need for a flexible and modular architecture that allowed individual components to be added, replaced, or reconfigured without redesigning the entire system. These insights led to the development of the Building Block concept, a plug-in-based architecture inspired by the process of assembling reusable LEGO® bricks onto a common baseplate. This concept later became one of the central architectural principles of Virtuality.

5.3.1 Virtuality Starting to Take Form

While first prototype established the technical foundations of Virtuality, the master's thesis projects provided the first concrete visions of how users could interact with the platform. Together, these efforts transformed Virtuality from a conceptual architecture into a more tangible and user-centered design.

Informed by the concepts developed through the two master's thesis projects, the second prototype of Virtuality was initiated. Several important insights emerged from these projects.

Key design insights from the master's thesis projects

- **Simulation was understood as a care process rather than a technical system.** Participants rarely described simulations in terms of databases, APIs, or interoperability standards. Instead, they understood them as workflows composed of actors, events, technologies, and outcomes.
- **Different stakeholder groups required different views of the same simulation.** Clinical users preferred patient-centred and process-oriented views, whereas technical users focused more on data flows, integrations, and system behaviour.
- **A single representation was not sufficient.** The findings suggested that the same simulation needed to be represented through multiple synchronized views, each tailored to the needs of a particular stakeholder group.
- **Interoperability needed to be supported without expos-**

ing unnecessary technical complexity to all users. The projects highlighted the need to make data exchange, integration, and system orchestration visible to technical users while keeping the interaction model understandable for clinical and non-technical stakeholders.

- **Simulation environments must support both acute and longitudinal care scenarios.** The contrast between home-care fall scenarios and traffic accident scenarios showed that Virtuality needed to accommodate both long-term monitoring and preventive workflows, as well as time-critical emergency care processes.

5.3.2 Development Begins

When implementation of the second prototype began, the emergence of AI-powered no-code and low-code development platforms significantly accelerated the development process. Ideas that previously would have required substantial software engineering effort could now be visualized, prototyped, and iteratively refined much faster. This made it possible to focus more attention on architecture, usability, workflow design, and integration logic rather than on low-level implementation details.

For Virtuality, this marked an important shift. The first prototype had mainly functioned as an exploratory integration platform and proof of concept. With the support of Large Language Model (LLM)-based development tools, the second prototype could evolve more rapidly into a broader virtual testing environment with richer functionality and a more ambitious architectural scope.

At the same time, AI-assisted software development introduces several risks. Recent studies have highlighted problems such as technical debt accumulation, reduced code ownership, hallucinated or inconsistent code, and uncertainty regarding the long-term maintainability of AI-generated software [83], [84], [85]. Although a growing body of work discusses prompt engineering, documentation, manual review, and testing as possible mitigation strategies, there is still limited evidence on how to build sustainable and maintainable systems as AI-assisted codebases grow over time [83], [84], [85].

These concerns were particularly relevant in the development of Virtual-

ity. Unlike a generic prototype or research demonstrator, Virtuality was intended to support healthcare-oriented simulation, interoperability, AI integration, and eventually regulatory-aligned evaluation. In such a setting, architectural drift, hidden technical debt, weak traceability, or unpredictable code generation are not minor inconveniences. They risk undermining the platform as a whole. For that reason, several countermeasures were taken. Some of these were inspired by the emerging literature on AI-assisted development, while others grew out of iterative prompting, practical development work, and established software architecture principles [86], [87].

Together, these measures aimed to preserve deterministic behaviour, architectural consistency, and traceability while still benefiting from the speed of AI-assisted implementation.

Architectural safeguards applied during AI-assisted development of Virtuality

- **Deterministic, schema-driven core.** The system was built around explicit and validated structures, including FHIR R4/R5 bundles, ontologies, and schema-validated contracts. LLMs were used as supporting tools, for example for mapping assistance and synthetic data generation, but were not allowed to control validation logic or core execution paths.
- **Explicit architectural constitution.** Architectural decisions, core contracts, and stable interfaces were documented explicitly and treated as governing artefacts during development. This was intended to reduce drift across development sessions and maintain continuity when AI-assisted changes were introduced incrementally.
- **Frozen architectural surfaces.** Certain critical components and interfaces were treated as stable boundaries that could not be modified without explicit review. These included core data exchange interfaces, authentication boundaries, persistence mechanisms, and other components central to interoperability and system integrity.
- **Automated architectural fitness functions.** Architectural constraints were encoded as executable checks and run continu-

ously during development. These checks enforced rules related to layering, data access boundaries, validated inputs, access control, and separation of responsibilities across the platform.

- **Layered enforcement of compliance.** Architectural compliance was not left to manual discipline alone, but was enforced through several layers, including local checks during development, pre-commit validation, and continuous integration. This reduced the risk that inconsistent or unsafe changes would silently enter the codebase.
- **Deterministic validation and reporting.** Compliance-related artefacts, such as regulatory classifications, simulation pipeline journals, and structured evaluation outputs, were generated through deterministic template- and metadata-driven logic. This was important for reproducibility and auditability.

Taken together, these measures meant that the development of Virtuality was of an *AI-assisted* nature, rather than an *AI-governed* one. LLMs were used to accelerate implementation, but the overall architecture remained constrained by explicit contracts, deterministic validation, and documented governance mechanisms. In the context of this thesis, this was important not only from a software engineering perspective, but also because a simulation environment intended to support healthcare innovation must itself be designed in a way that preserves traceability, reproducibility, and trust.

5.4 Architectural Synthesis of Virtuality

The preceding sections described how Virtuality emerged from a combination of factors. Rather than presenting every technical detail of the resulting implementation, this section summarizes the core architectural ideas that were shaped through the work reported in this thesis and in the included papers.

A more detailed and formal description of Virtuality's architecture, design principles, and evaluation is provided in Paper G.

5.4.1 Architectural Overview

At a high level, Virtuality can be described as a process-oriented simulation environment, in which healthcare simulations are composed from reusable software components and executed within a controlled, standards-based environment (Figure 5.1).

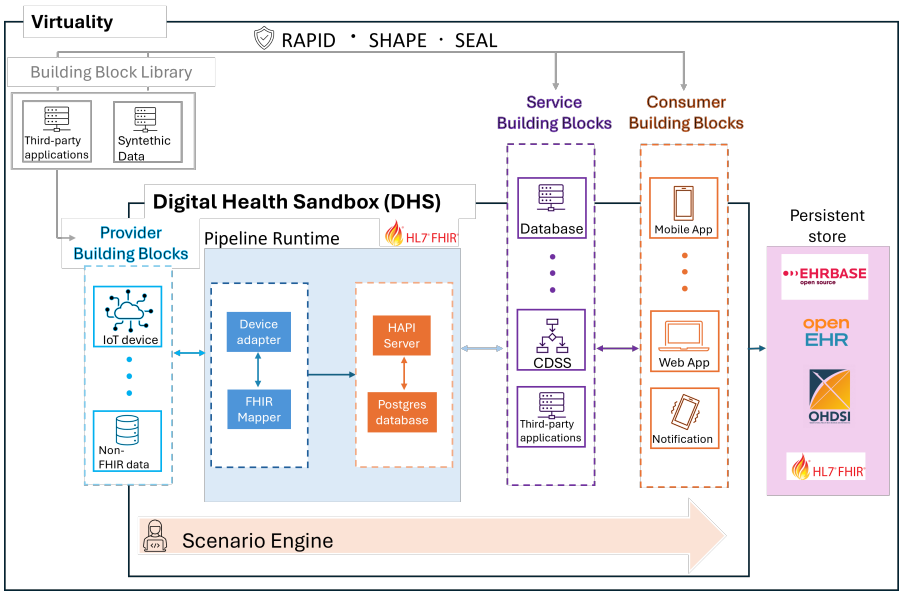


Figure 5.1: Conceptual architecture view of Virtuality showing the main platform components.

The hybrid architecture, based on the ASAP concept, consists of a core platform surrounded by modular and isolated Building Blocks that can be added, configured, and connected depending on the purpose of the simulation (Figure 5.3). These Building Blocks are divided into three main categories: *Providers*, *Services*, and *Consumers* that communicate over an event-driven FHIR-based communication layer. Providers represent sources of data, such as sensors, synthetic data, databases, and APIs. Services represent processing functionality that acts on incoming data, for example AI-based CDSS, transformation logic, orchestration services, or routing components, and Consumers represent user-facing applications through which simulation participants in-

teract with the simulated workflow, such as dashboard applications, or other interfaces used by clinicians, patients, or other stakeholders (Figure 5.2).

Similar to building with LEGO, a user of Virtuality can combine multiple Building Blocks to expand a simulation in different ways. A simulation can be extended *vertically* by adding additional Building Blocks of the same type, or *horizontally* by combining Building Blocks of different types to represent more complex healthcare workflows (as illustrated in Figure 5.2).

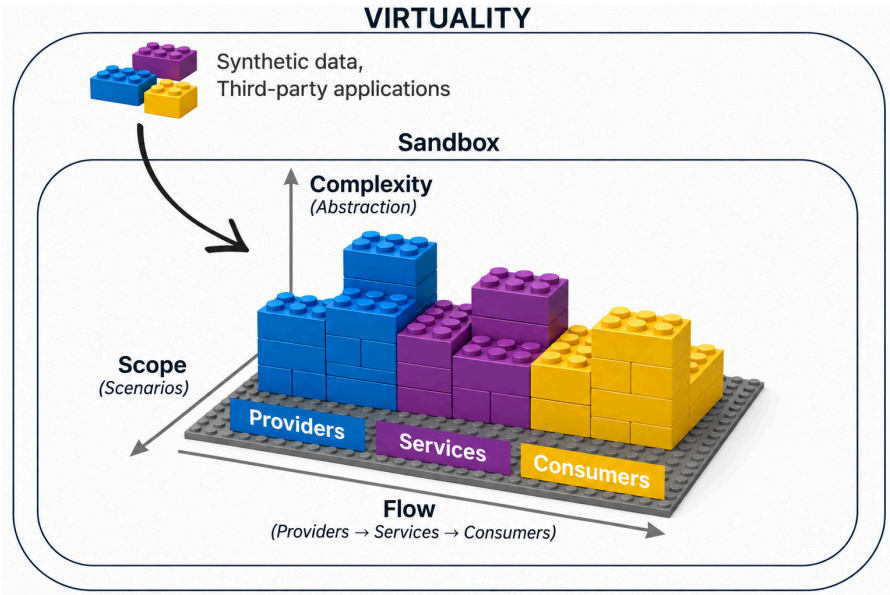


Figure 5.2: Illustration of the Building Block concept in Virtuality. Simulations are composed from modular Providers, Services, and Consumers that can be connected to represent a healthcare process.

Users assemble simulations through the Scenario Engine, a graphical user interface with a visual drag-and-drop canvas, by selecting Building Blocks from a shared library, connecting them into a simulation pipeline, and defining the context in which the simulation should operate. This context may include environmental conditions, organizational constraints, security settings, or assumptions about the healthcare process being simulated. To enable such composition, each Building Block adheres to well-defined data contracts and

interfaces based on the HL7 FHIR R4/R5 standard.

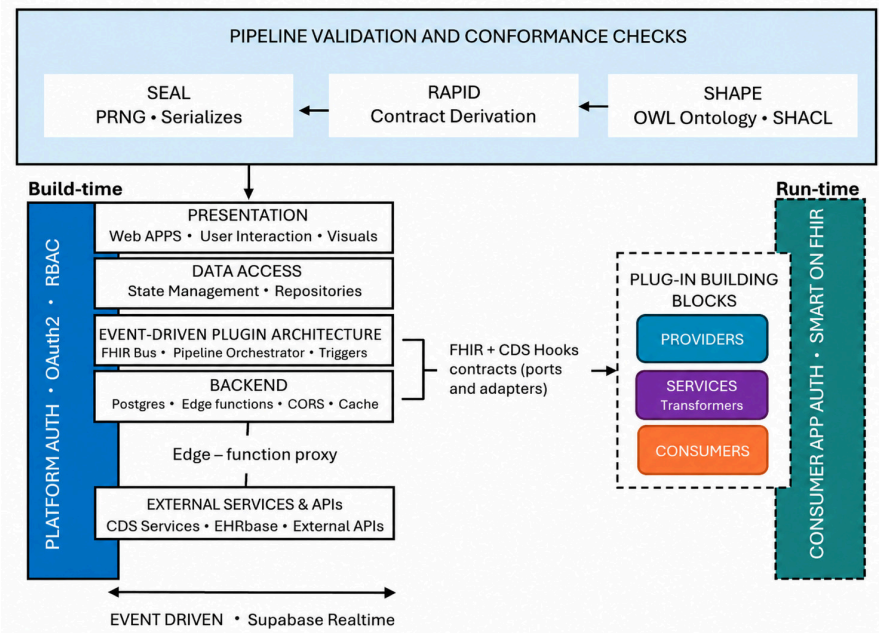


Figure 5.3: High-level architectural view of Virtuality illustrating the hybrid architecture of Virtuality.

Inspired by the stepwise model in the VIPHS framework, Virtuality aims to to maximize clinical utility and utilization by supporting mechanisms for interoperability validation, documentation, and controlled execution. This is achieved through a pipeline validation and conformance checks layer, where the platform validates that selected building blocks are structurally compatible, that expected data exchanges can be supported, and that the simulation configuration can be reproduced (Figure 5.3).

This overall architecture reflects the dual role of Virtuality. On the one hand, it functions as a design-time environment in which simulations can be created, configured, and validated. On the other hand, it functions as a run-time environment in which participants interact with the resulting applications (Consumer building blocks) and workflows under conditions intended to resemble realistic healthcare settings.

5.4.2 Validation and Execution Mechanisms

One of the most important architectural ideas in Virtuality is the use of *Building Blocks* as the primary abstraction for representing simulation components. This concept emerged gradually during the development of the platform and was influenced by both practical implementation work and the master's thesis projects described earlier in the chapter. Across these activities, it became increasingly clear that clinical users tended to view a simulation as a series of patient-centric events. The Building Block model provided a way to bridge this perspective by offering a common abstraction that could be understood both conceptually and technically.

However, the ability to visually connect Building Blocks is not sufficient on its own. Building Blocks may appear compatible at a conceptual level while still failing when connected in practice. Data may be represented differently, required fields may be missing, assumptions about timing or workflow may be inconsistent, and systems may support different subsets of a standard. For Virtuality to function as a meaningful simulation environment, the platform must also determine whether the simulation is structurally coherent, and whether its execution can be reproduced.

To address this, Virtuality incorporates a set of complementary validation and execution mechanisms that operate between simulation design and runtime execution on the top layer in Figure 5.3. These mechanisms are referred to as *SHACL-based Architecture and Pipeline Enforcement (SHAPE)*, *Reproducible Architecture for Pre-deployment Interoperability Determination (RAPID)*, and *Serialized Execution Artifact Lock (SEAL)*. These mechanisms are not new technologies in themselves. Instead, they combine existing standards and tools in a new way to support validation and conformance checks for Building Blocks and data pipelines in Virtuality. In particular, Virtuality relies on semantic web technologies, including ontologies expressed in OWL (Web Ontology Language), which support representation of domain knowledge and inference over incomplete data, and SHACL (Shape Constraint Language), which provides a constraint language for validating data against predefined structural and semantic rules[88].

SHACL-based Architecture and Pipeline Enforcement (SHAPE)

The first mechanism, *SHAPE*, focuses on structural and semantic consistency. Its purpose is to determine whether the simulation has been assembled in a way that is meaningful according to the ontology and architectural rules of the platform. In practice, this includes checking whether the selected Building Blocks are assigned appropriate roles, whether required relationships between components are present, and whether the overall simulation structure is coherent before execution. After passing SHAPE - validation, the validation continues with the next step, RAPID.

Reproducible Architecture for Pre-deployment Interoperability Determination (RAPID)

The second mechanism, *RAPID*, focuses on interoperability validation. Its role is to assess whether the selected Building Blocks can actually exchange the information required by the simulation pipeline. Rather than assuming interoperability only because two components both claim FHIR compliance, RAPID evaluates the expected inputs and outputs contracts of the connected components and checks whether the necessary resources, fields, structures, and terminologies are available across the pipeline. Once RAPID - validation has been confirmed, the process continues with SEAL.

Serialized Execution Artifact Lock (SEAL)

The third and last mechanism, *SEAL*, is concerned with controlled execution and reproducibility. Once a simulation has been validated by SHAPE and RAPID, SEAL creates a frozen execution representation containing the relevant configuration, selected components, mappings, and execution settings required to run the simulation in a reproducible way. This is important because one of the ambitions of Virtuality is not only to support experimentation, but also to allow simulations to be revisited, compared, and audited afterwards.

Taken together, SHAPE, RAPID, and SEAL combine and operationalize existing technologies in a novel way to support validation of process-based healthcare simulations (Figure 5.4). Their role is not limited to checking isolated technical artifacts, such as profile conformance, data structures, or API integrations. Instead, they evaluate the simulation as a whole, including the complete Building Block pipeline, the relationships between its components,

and the environmental and execution configurations under which the simulated healthcare process is intended to operate.

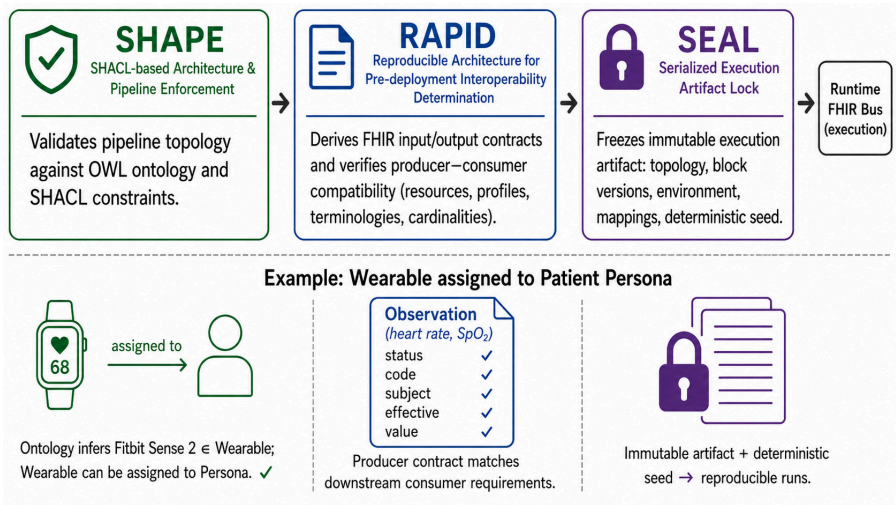


Figure 5.4: How SHAPE, RAPID and SEAL relate and work (Created with GPT-5.6 Sol, 2026).

5.5 From Building Blocks to Executable Simulations

Before a simulation can be executed, the user first selects the type of simulation to create (e.g., trauma, stroke, home healthcare). This selection determines which Environment Setup and Building Blocks that become available within the Scenario Engine and the Building Block Library. In this way, users are presented only with components that are relevant to the selected simulation context.

If a required Building Block is not already available, it can be created using the Building Block Builder, a guided configuration process that assists the user in creating interoperable and reusable Building Blocks needed for Virtuality to understand how the new component should behave.

During this process, the user specifies:

the Building Block type (e.g., Provider, Service, or Consumer), the semantic class or role (e.g., Persona, Wearable, Medical Device, Laboratory System), supported terminology systems, the required input and output data, and the corresponding FHIR resources and mandatory fields that define the Building Block's data contract.

The semantic role is particularly important because it determines how the Building Block is classified within the ontology. This makes it possible for the SHAPE validation mechanism to infer relationships between Building Blocks, and to identify incompatible connections and missing data fields.

When the user initiates the simulation (e.g., by clicking on "Build Simulation"), the SHAPE mechanism validates that each Building Block satisfies its semantic constraints and that all required relationships are consistent with the ontology. The RAPID mechanism subsequently verifies that all declared FHIR data contracts can be satisfied throughout the scenario, ensuring that downstream Building Blocks receive the information they require before execution begins. Only when these validation steps have been successfully completed is the scenario considered executable.

CHAPTER 6

Example Use Cases of Virtuality

"Just because something doesn't do what you planned it to do doesn't mean it's useless."

— Thomas Edison

As discussed by Richards and Ford in *Fundamentals of Software Architecture* [89], software architecture is fundamentally shaped by trade-offs. This reflects an important conclusion throughout this thesis: there is rarely a single correct way to design or implement a digital health platform. Instead, architectural decisions involve balancing different qualities, constraints, and practical needs.

This also raises the question of how Virtuality should be evaluated. Within the scope of this PhD, Virtuality as a platform was not formally evaluated through a usability study. Although some aspects of the GUI were informed by earlier Master's thesis projects, the technical foundation of the platform has not been evaluated in isolation. Instead, its suitability and applicability have been explored through several example use cases within the Care@Distance group.

In this thesis I present two example use cases to demonstrate how Virtuality can be used in different contexts to support the development, simulation, and evaluation of digital health applications.

The two use cases highlight different aspects of the platform. The first focused on the evaluation of the On-Scene Injury Severity Prediction (OSISP) model, an AI-based CDSS intended to support prehospital trauma triage [90]. The second focused on the development of a web-based voting application that enabled neurologists to assess synthetic stroke images and videos, with the goal of generating labeled data for future AI development. Together, the examples illustrate Virtuality's applicability, flexibility and modularity.

6.1 Use Case 1: Trauma Triage

The first use case was focused on evaluating the OSISP model in a realistic prehospital setting. OSISP is an AI-based CDSS designed to support ambulance personnel during the assessment of adult trauma patients by predicting whether a patient is severely injured and may require transport to a trauma center [90].

To evaluate the usability and clinical applicability of the OSISP model, a user study involving 34 Emergency Medical Services (EMS) personnel was conducted. The original simulation study was initially set up and executed using the first prototype of Virtuality. Afterwards, the same simulation was recreated in the second prototype of Virtuality, where a simulation pipeline consisting of five interoperable Building Blocks was configured (Figure 6.1). I refer to this recreated setup as the *OSISP replica simulation*.

The purpose of the OSISP replica simulation was not to repeat the original user study, and no real users were involved in this second setup. Instead, the aim was to examine whether the second prototype of Virtuality could support the same simulation scenario in a more structured, reusable, and configurable way through the Building Block architecture. In other words, the OSISP replica simulation was used to explore whether the architectural ideas of the second prototype could successfully recreate and support a previously developed simulation scenario.

In the OSISP replica simulation, three Provider Building Blocks were configured to represent three different patient personas. Each patient profile contained unique demographic and clinical characteristics, including age, gender,

injury mechanism, and physiological parameters. The personas were designed to represent different trauma scenarios and severity levels.

A Service Building Block was configured to represent the OSISP model and exposed standardized interfaces for receiving patient information and returning predictions. To create this Service Building Block, the developer guide provided within the Virtuality platform was followed. This allowed the model to be wrapped in a RESTful API adhering to the CDS Hooks specification.

The Consumer Building Block consisted of a web-based trauma triage application, accessed through a tablet, enabling potential participants to review patient information, perform assessments, and interact with the OSISP predictions (Figure 6.3). The graphical user interface (GUI) of the application was informed by the Master thesis conducted with the Care@Distance research group [91].

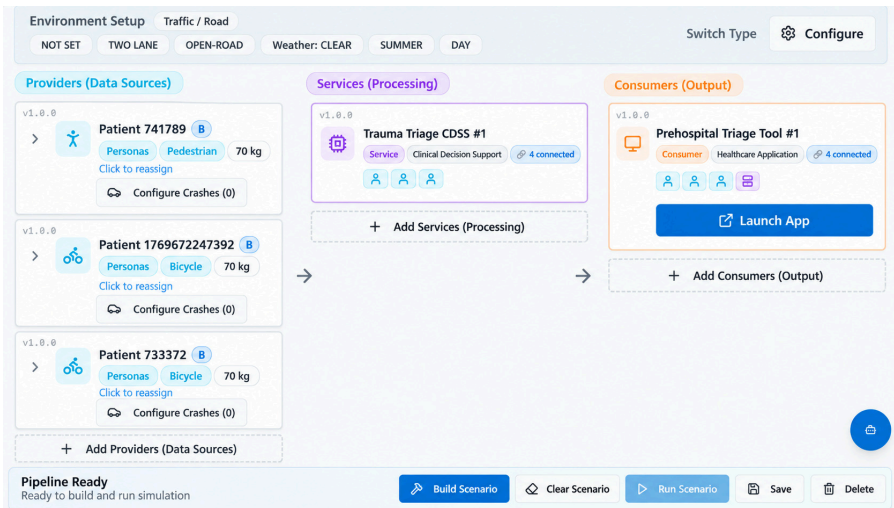


Figure 6.1: Virtuality simulation pipeline, accessed through the Scenario Engine, used to recreate the OSISP evalutaion study. The pipeline consists of three Provider Building Blocks representing patient data, one Service Building Block representing the OSISP decision support logic, and one Consumer Building Block providing the user-facing trauma triage application. When the simulation was executed, the Consumer Building Block was exposed as a web application through a pipeline-specific URL.

The idea with this simulation setup was for users to selected one of the available patient personas (represented by one of the three Patient Persona Provider Building Blocks) and performed a trauma assessment using the Consumer application. The selected patient data were routed through the simulation pipeline to the OSISP service, which generated a prediction that was subsequently displayed in the user interface.

Figure 6.1 illustrates the simulation pipeline, while Figure 6.2 presents the data flow between the different Building Blocks. Figure 6.3 shows the trauma triage application (Consumer Building Block).

While Figure 6.1 emphasizes the structural composition of the simulation through the *Card view*, it does not fully show how information moves between the components during execution. For that reason, Virtuality also provides a complementary *Data flow view*, illustrated in Figure 6.2.

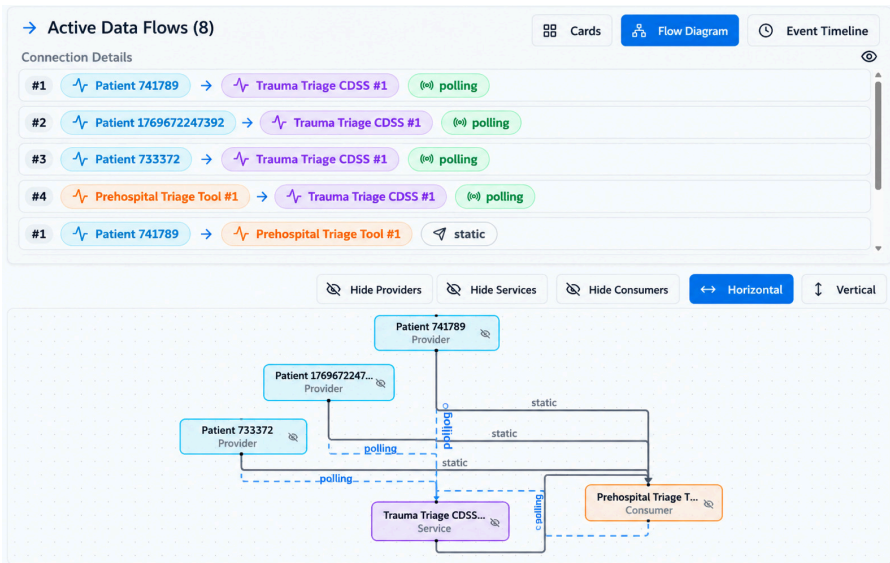


Figure 6.2: Data flow between Building Blocks during the trauma triage simulation.

The view in Figure 6.2 presents the same simulation from a different perspective and allows the user to inspect how data routing and interactions between Building Blocks have been configured within the simulation pipeline.

The final perspective is the one seen by the participants of the simulation study. Rather than interacting with the pipeline view, participants access the simulation through the Consumer Building Block, which is exposed as a web-based trauma triage application on a tablet (Figure 6.3).

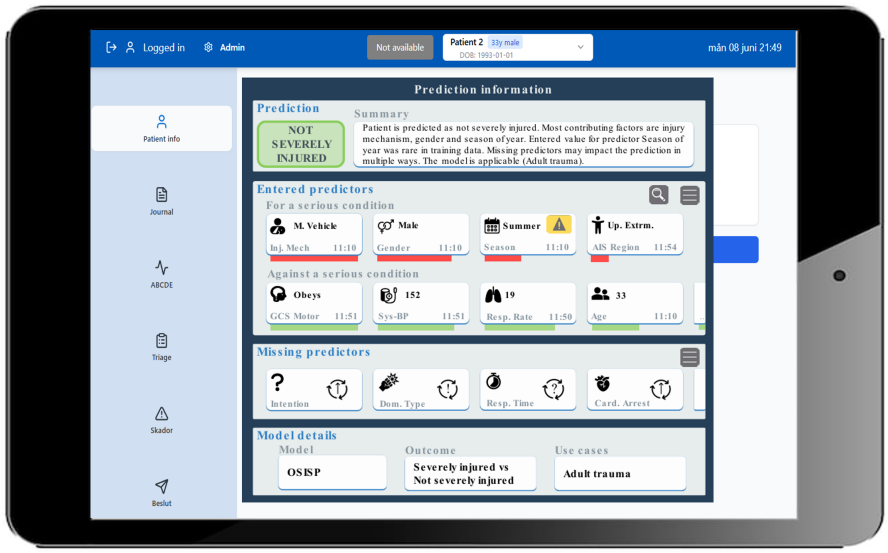


Figure 6.3: Web-based trauma triage application representing the Consumer Building Block. The selected patient persona is displayed in the header (Patient 2). The pop-up window emulates the call to the external OSISP CDSS and presents the model prediction to the participant.

Figure 6.3 shows this participant-facing interface, where users can review the patient case, perform their assessment, and receive the corresponding OSISP output. Behind the scenes, both the data flow and the decision logic execute according to the configuration defined in the Scenario Engine and the underlying simulation pipeline. However, these internal mechanisms are never exposed directly to the simulation participants themselves. Instead, participants interacted only with the Consumer Building Blocks, represented by a trauma triage application in this case (shown in Figure 6.3).

6.2 Use Case 2: Stroke Neglect Study

To further demonstrate the flexibility of Virtuality, a second use case was conducted in collaboration with researchers investigating stroke neglect (simulation pipeline in (Figure 6.4). Unlike the previous use case, this study involved real users and aimed to generate labeled data that could later be used to develop an AI-based CDSS for stroke.

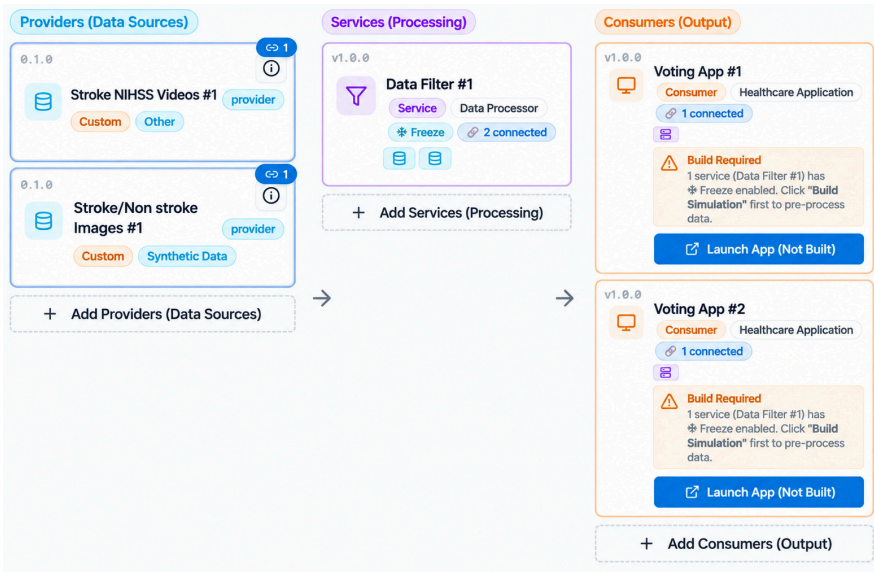


Figure 6.4: Virtuality simulation pipeline used in the stroke neglect study.

The objective was to develop a web-based voting application (Consumer Building Blocks represented as Voting Apps in Figure 6.4) that enabled four neurologists to independently assess synthetic stroke images and videos (Provider Building Blocks represented by Stroke NIHSS Videos and Stroke/Non stroke images in Figure 6.4). For each case, the neurologists were asked to determine whether the patient exhibited signs of stroke and whether the observed symptoms indicated left-sided neglect, right-sided neglect, or no neglect.

Using Virtuality, separate Building Blocks were created following the developer guidelines in Virtuality (with copy-paste templates). After creation, they appeared in the Building Block Library and were subsequently added to

the pipeline (Figure 6.4).

Two Provider Building Blocks were used to store the synthetic image (PNG-files) and video datasets (MP4). A Service Building Block was configured to retrieve cases from the available datasets, randomize their presentation order, and divide the resulting cases into two independent assessment batches (Data Filter Service Building Block in Figure 6.4).

Two Consumer Building Blocks were then configured as web-based voting applications (Figure 6.6). Although the applications provided identical functionality, they received different randomized batches generated by the Service Building Block (Figure 6.5). This enabled the neurologists to evaluate different subsets of cases while maintaining a consistent assessment workflow across all participants.

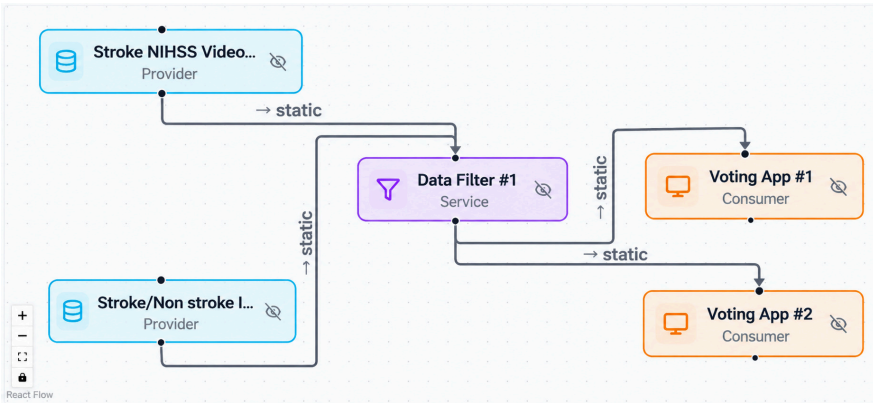


Figure 6.5: Data flow between Building Blocks in the Stroke Neglect study.

When a neurologist logged in to the application (Figure 6.6), the Service Building Block selected content from the available datasets and delivered the corresponding image or video to one of the voting applications. The neurologist's assessment was subsequently recorded and stored for later analysis through the admin-login page, which was managed by my college.

This use case demonstrates how Virtuality can be used not only for evaluating AI-based systems but also for supporting data collection workflows and collaborative annotation studies. Furthermore, it illustrates how Service Building Blocks can encapsulate study-specific logic, such as randomization and batch generation, while Consumer Building Blocks provide independent

user interfaces for interacting with the resulting datasets. By reusing existing Building Blocks and configuring new simulation pipelines, I was able to rapidly iterate a specialized solution without modifying the underlying Virtuality architecture.

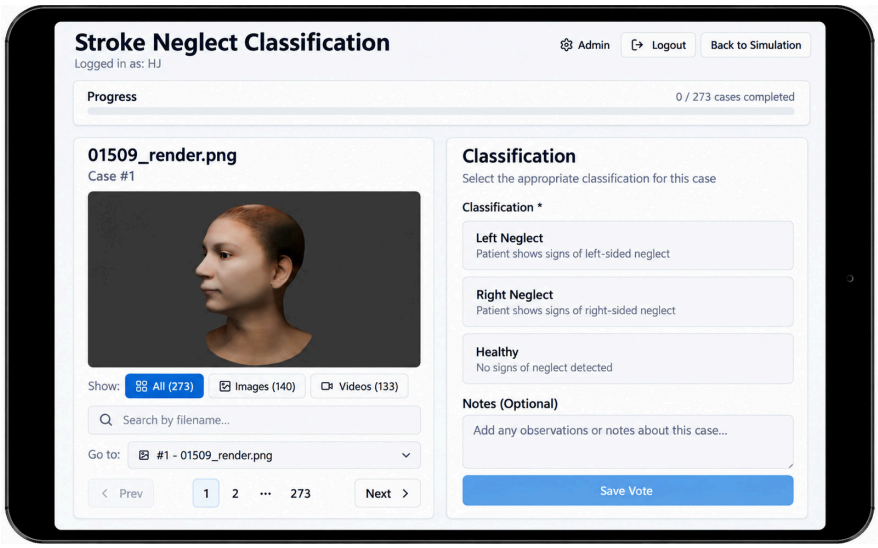


Figure 6.6: The resulting voting application represented by the Voting App Consumer Building Block in Figure 6.4

The two use cases presented in this chapter illustrate different ways in which Virtuality can be instantiated in practice. The two use cases illustrate an important architectural characteristic of the platform: although the first implementation of a new simulation may require the creation and configuration of several Building Blocks from the ground up, these components become reusable assets once added to the shared library.

This means that subsequent simulations can potentially be assembled more efficiently by reusing and adapting existing Building Blocks rather than starting from scratch. In the stroke neglect case, for example, the voting application, the filtering service, and the data provider components can in principle be reused for future studies involving other image or video datasets, alternative annotation tasks, or modified voting and labeling procedures. More broadly, the examples show that Building Blocks are not intended as fixed,

one-off components, but as configurable modules designed to be adjusted and recombined across different simulation contexts. In this way, each new use case contributes not only a specific application, but also additional reusable capabilities that strengthen the Virtuality platform for future research and innovation work.

6.3 Other Use Cases: Traffic and Defense Medicine

Although the use cases presented in this chapter focus on healthcare, the Scenario Engine also allows users to navigate between different domains. This reflects the ambition to make Virtuality's modular architecture applicable beyond healthcare, for example, to traffic safety and Mass Casualty Incident (MCI) simulations (Figure 6.7).

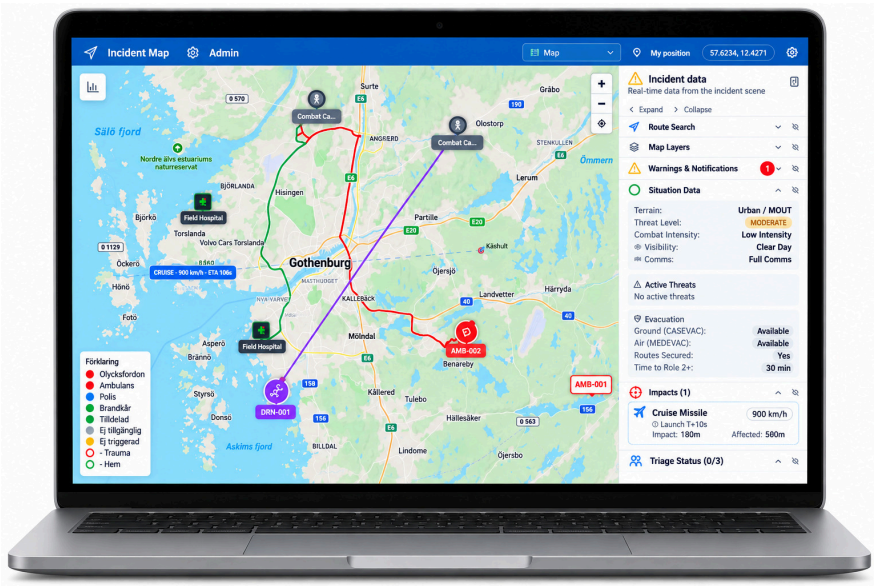


Figure 6.7: A consumer application illustrating a simulated mass-casualty scenario with emergency resources, evacuation routes, and real-time incident information.

CHAPTER 7

Overview and Motivation of Included Studies

"I may not have gone where I intended to go, but I think I have ended up where I needed to be."

— Douglas Adams

As described earlier, the work presented in this thesis did not emerge from a single predefined hypothesis or a linear sequence of events. Instead, it has been shaped by iterative work addressing real, practical challenges observed in previous research projects and triple helix collaborations. The work has also been guided by learning outcomes experienced from cross-disciplinary collaborations, AI implementation efforts, regulatory drivers, research, and of course, from the frameworks discussion in Section 4.

Consequently, each study included in this thesis has contributed a piece of understanding that gradually shaped the development of Virtuality. Together, they represent a progression from problem identification to conceptualization, design, and implementation.

7.1 From Interoperability Challenges to Practical Needs

In my first project as a PhD student, the Autumn Leaves project, the ambition was to aggregate and standardize data from multiple heterogeneous sources in order to improve decision precision and enhance care delivery in the context of falls among elderly. Early in the project, it became clear that this was not straightforward. I already knew that interoperability is a complex topic and that many different standards, tools, and architectural approaches can be used to achieve it. However, it was less clear how others had approached similar challenges in practice, particularly in the context of IoMT solutions in home care and prehospital care. I therefore wanted to better understand how interoperability had been addressed in the literature, which technologies had been used, and how these had been combined to achieve various levels of interoperability. This motivated the first two papers in the thesis, Paper A and Paper B.

Paper A presented the protocol for a scoping review, and Paper B completed the review and synthesized existing research on interoperable IoMT platforms for home care and prehospital settings. Paper B showed that IoMT interoperability depends on a broad and fragmented ecosystem technologies, tools and standards. Many of the reviewed studies focused on lower levels of interoperability, such as connectivity and data exchange in smaller networks, typically Local Area Networks (LAN) and Personal Area Networks (PAN). Fewer addressed how these technologies could support higher levels of interoperability. The review also made clear that achieving higher levels of interoperability often requires combining multiple technologies and standards across different layers rather than relying on any single solution [77].

Through Paper A and Paper B, I gained a much clearer picture of the fragmented interoperability landscape surrounding IoMT and digital health platforms. At the same time, the findings raised a broader concern. If the interoperability landscape is difficult to navigate even for researchers and technical experts, how can healthcare organizations make informed decisions about procurement and implementation initiatives, in contexts where such decisions may affect patient safety and care delivery?

This question motivated Paper C, which shifted the focus from technical solutions to the organizational processes through which interoperability am-

bitions are translated into practice. Paper C examined challenges in specifying interoperability requirements during procurement of health information systems (HIS). The paper showed that organizations often struggle to formulate concrete and measurable interoperability requirements, particularly when needs are cross-disciplinary and technical expertise is unevenly distributed across stakeholders. In other words, even when interoperability is recognized as important, it can still be difficult to express what should actually be procured, implemented, or evaluated [92].

Taken together, Papers A, B, and C revealed a gap between the technical possibilities of interoperability and the practical realities of using it in healthcare. On the one hand, these papers provided valuable insights into the kinds of standards, tools, and interoperability mechanisms that a platform such as Virtuality should support. On the other hand, they also highlighted that clinical users, researchers, and healthcare organizations need practical ways to explore, test, and evaluate such solutions without being entirely dependent on external vendors or highly specialized technical expertise. This became an important motivation for Virtuality: not only to support interoperability at a technical level, but also to provide an environment in which new solutions, requirements, and integration strategies could be explored in a way that is more accessible and practically useful for healthcare stakeholders.

7.2 Conceptualizing a Simulation Environment

Building upon the insights from Papers A, B, and C, the next phase of the research shifted from identifying interoperability challenges toward exploring how such challenges might be addressed in practice. If interoperability is difficult to achieve, difficult to specify in procurement, and difficult to evaluate before implementation, then there is also a need for environments in which new solutions can be tested, configured, and discussed under more realistic but still controlled conditions.

This motivated Paper D, which was the first paper to introduce the concept of *Virtuality - Virtual Care Process Simulator*. The paper proposed a Digital Health Sandbox approach in which synthetic data, interoperable applications, and configurable simulation scenarios could be combined to recreate realistic healthcare situations.

Paper D was therefore an important conceptual shift in the thesis. Rather

than focusing primarily on how to achieve interoperability, the paper focused on what kind of environment would be needed to work with interoperability in practice. This phase also marked the first more systematic effort to review and position existing solutions relevant to Virtuality, including digital sandboxes, synthetic data environments, and simulation-based approaches in healthcare. In this sense, Paper D marked the transition from understanding the interoperability problem space, towards beginning to formulate a practical response to it.

7.3 Addressing Information Exchange and Data Representation Challenges

As Virtuality evolved from concept toward design and implementation, practical challenges related to information exchange and data representation emerged. One of the first major challenges was related to integration work of the OSISP model. Since the OSISP model had been developed before Virtuality, the AI-model expected the input data to be provided in a specific format. Its syntactic and semantic expectations were therefore not aligned with HL7 FHIR, which was used as the main communication layer in Virtuality. This meant that the model could not simply be integrated into the platform without first transforming its expected inputs and outputs into FHIR-compatible representations.

This work involved translating between the data structures expected by the OSISP model and the FHIR resources used within Virtuality, that is, *data mapping*. Working with these mappings quickly made it clear that this is not a trivial task. It raised questions about how model inputs should best be represented in FHIR, how semantic meaning could be preserved across systems, and how the quality of a mapping could be assessed and validated.

These experiences motivated Paper E, which investigated challenges associated with the quality of FHIR data mapping. The review showed that data mapping often relies heavily on manual work and subjective interpretation, and that validation methods are inconsistent and often insufficiently described [93]. It also highlighted the limited methodological support available for assessing mapping quality in a systematic way. These findings aligned closely with the challenges encountered during the integration of the OSISP model, and were further supported by four independent FHIR experts consulted as

part of the review.

The need for a more systematic and practically usable approach to data mapping was then pursued further in Paper F. Whereas Paper E focused on the quality challenges associated with FHIR mapping, Paper F examined how mapping processes could be facilitated in practice and how implementers, clinicians, and researchers can be supported in this work. This became particularly important in the context of Virtuality, where the ambition was not only to support technically correct transformations, but also to make mapping more accessible for users who may not have extensive prior FHIR expertise. In the context of emerging AI-based CDSS, the challenge is not limited to mapping clinical observations or administrative data, but may also involve representing information about AI models themselves, including training, validation, bias, intended use, and human oversight. This is particularly relevant in light of emerging regulatory frameworks such as the AI Act, which place increasing emphasis on transparent and traceable AI systems. However, many of these concepts do not map easily to existing core FHIR resources. For that reason, ongoing work is being carried out to formalize a dedicated implementation guide intended to support this process, although at the time of writing it remains in an early draft stage.

Taken together, Papers E and F contributed to the development of functionality in Virtuality aimed at making transformation to FHIR more accessible, including graphical support for configuring mappings in a way that is more transparent and usable for implementers, clinicians, and researchers.

7.4 From Concept to Practical Implementation

The final phase of the research focused on turning the insights gained in the earlier papers into something more concrete and practically useful. As described previously, this has been an iterative process in which functionality has been developed, revised, changed, and refined over time. In that sense, Virtuality will probably never be "finished". New functionality will continue to be added as healthcare needs change, as new technologies emerge, and as regulatory requirements evolve.

Paper G presents the design and architecture of Virtuality as an interoperable simulation platform where systems, applications, devices, synthetic data, and AI models can be connected and evaluated within realistic healthcare

workflows [94]. In many ways, this paper represents the point at which the different strands of the thesis come together. The earlier papers had identified interoperability challenges, procurement barriers, mapping difficulties, and the need for a process-oriented simulation environment. Paper G brings these insights together in the form of a concrete platform description. In other words, Paper G acts as a bridge between the conceptual work and the practical implementation, showing how the ideas developed throughout the thesis were translated into Virtuality.

CHAPTER 8

Summary of included papers

"Seven papers, countless iterations, and one question: What is the result?"

— Mattias Seth

This chapter provides a summary of the included papers that has shaped this thesis, and the development of Virtuality.

8.1 Paper A (Scoping review protocol)

Seth M, Jalo H, Högstedt Å, Medin O, Björner U, Sjöqvist BA, Candefjord S Technologies for Interoperable Internet of Medical Things Platforms to Manage Medical Emergencies in Home and Prehospital Care: Protocol for a Scoping Review JMIR Res Protoc 2022;11(9):e40243 doi: 10.2196/40243 PMID: 36125863 PMCID: 9533201

Summary

Paper A established the conceptual foundation for the thesis by defining the interoperability within Internet of Medical Things (IoMT). In the paper, IoMT was further defined as an infrastructure of interconnected medical devices, sensors, applications, and supporting systems that communicate over the Internet to enable remote monitoring, data exchange, and clinical decision support across distributed care settings (Figure 8.1 and Figure 8.3).

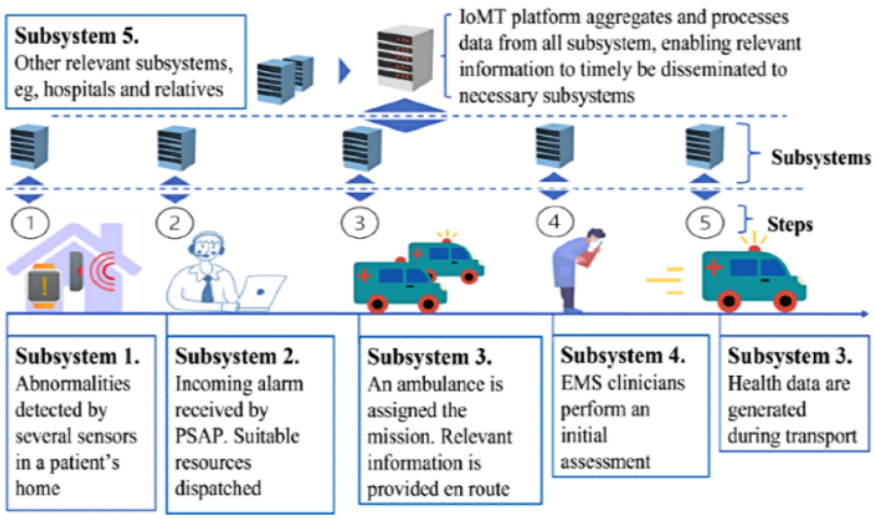


Figure 8.1: Example of events and data being transferred through systems in home care and prehospital care settings. CC BY 4.0.

Within this context, interoperability was framed as a prerequisite for safe and effective care processes, referring not only to the ability of systems to exchange information, but also to do so in a predictable, standardized, and meaningful way. To structure this complexity, Paper A conceptualized interoperability through a six-level model, highlighting that interoperability challenges can arise at multiple levels (Figure 8.2).

In parallel, the paper introduced an IoMT reference model to describe the main building blocks of IoMT platforms and to support later mapping of technologies, standards, and protocols to different layers of interoperability (Figure 8.3). In this way, Paper A did not merely outline a review protocol,

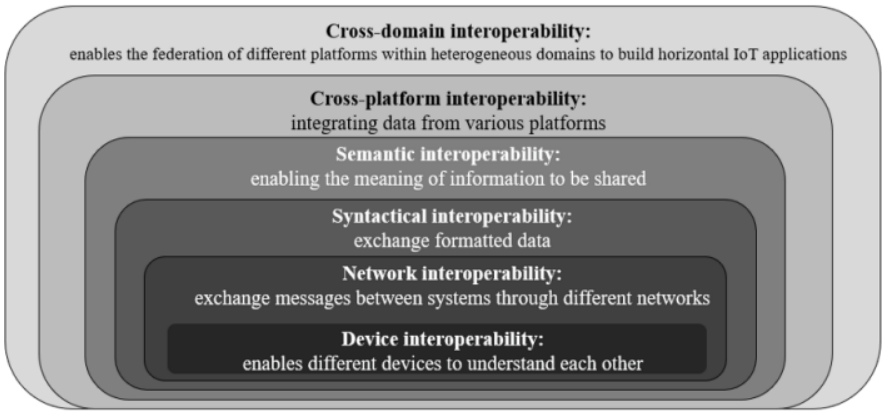


Figure 8.2: The IoMT interoperability model defined in the protocol, consisting of Device, Network, Syntactic, Semantic and Cross-domain Interoperability. CC BY 4.0.

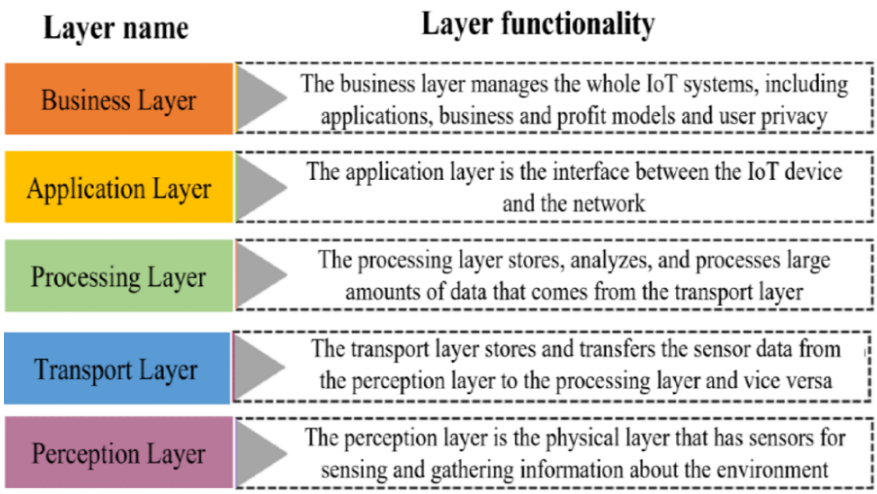


Figure 8.3: IoMT was defined using a five-level protocol stack. CC BY 4.0.

but laid the conceptual groundwork for the remainder of the thesis.

Author contributions

I had the main responsibility for designing the study and drafting the manuscript. All co-authors contributed to the refinement of the review protocol and revised the manuscript critically.

8.2 Paper B (Scoping review)

Seth M, Jalo H, Högstedt Å, Medin O, Sjöqvist BA, Candefjord S Technologies for Interoperable Internet of Medical Things Platforms to Manage Medical Emergencies in Home and Prehospital Care: Scoping Review J Med Internet Res 2025;27:e54470 doi: 10.2196/54470 PMID: 39847768 PMCID: 11803335

Summary

Paper B completed the scoping review outlined in Paper A and provided a structured overview of technologies, standards, and tools used to support interoperability in Internet of Medical Things (IoMT) platforms for home-based care and prehospital care. Building on the six-level interoperability model and the five-level IoMT reference model defined in the protocol, the paper mapped how previous studies addressed interoperability challenges across device, network, syntactic, semantic, cross-platform, and cross-domain layers.

A total of 2949 records were identified across the database searches, and after screening and snowballing, 158 studies were included in the final synthesis. Of these, 30 studies describing concrete IoMT platform development efforts. These studies were examined in greater detail to identify the enabling technologies and standards used in practice. The review showed that interoperability in IoMT depends on a broad and fragmented ecosystem of communication technologies, standards, middleware, and tools. Among the 30 platform-development studies, the most frequently reported technologies included cloud computing (19/30, 63%), Representational State Transfer Application Programming Interfaces (REST APIs; 17/30, 57%), Wi-Fi (17/30, 57%), gateways (15/30, 50%), and JavaScript Object Notation (JSON; 14/30, 47%). Technologies supporting real-time communication, such as Message

Queuing Telemetry Transport (MQTT; 7/30, 23%) and WebSocket (7/30, 23%), were commonly used to enable emergency alerts, time-sensitive messaging, and streaming of patient data (Figure 8.4). By contrast, technologies more closely associated with higher levels of interoperability were reported less frequently, including Substitutable Medical Applications and Reusable Technologies on Fast Healthcare Interoperability Resources (SMART on FHIR) (2/30, 7%), Blockchain (2/30, 7%), openEHR (2/30, 7%), and Kubernetes (3/30, 10%).

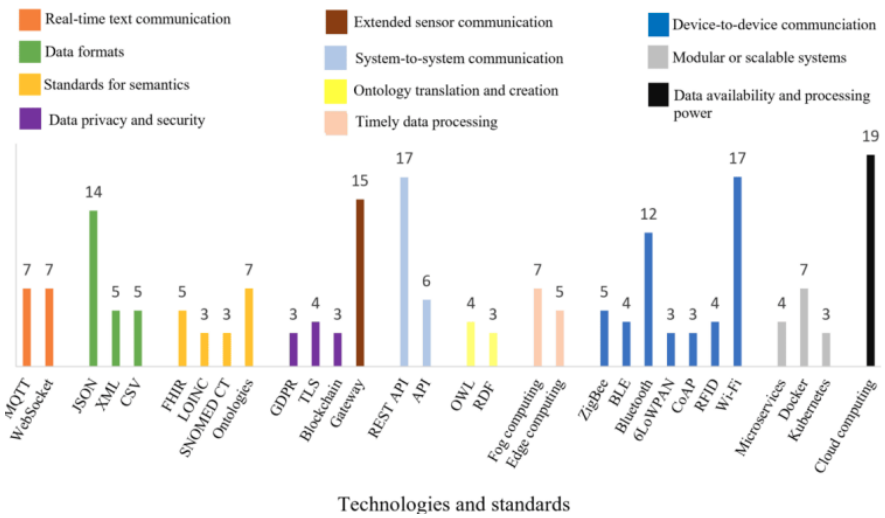


Figure 8.4: Occurrences of common technologies addressing interoperability in Internet of Medical Things (IoMT) platforms for prehospital care and home-based care, grouped by application area (each color). CC BY 4.0.

The review also highlighted that most identified solutions focused primarily on lower interoperability layers, such as connectivity, messaging, and data exchange. In the material, 17 of the 30 examined platform-development studies (57%) were judged to primarily emphasize lower levels of interoperability. In addition, among the 25 studies from which interoperability challenges were extracted, the most frequently reported challenges related to the volume and complexity of data (13/57 challenge occurrences, 23%), latency (12/57, 21%), privacy and security (12/57, 21%), proprietary protocols (8/57, 14%), seman-

tic heterogeneity (8/57, 14%), and poor connectivity (4/57, 7%).

To distinguish between platforms and research primarily focused on lower levels of interoperability and those addressing higher levels, the paper motivated the complementary term *Internet of Medical Systems* (IoMS), with IoMT understood as a subset of this ecosystem.

Author contributions

I had the main responsibility for the study design, literature review process, data extraction, synthesis of the results, and drafting of the manuscript. Author HJ contributed to the screening. All co-authors contributed to the interpretation of the findings and revised the manuscript critically.

8.3 Paper C (Full Conference paper)

Seth M, Jalo H, Lee E, Bakidou A, Medin O, Björner U, Sjöqvist BA, Candefjord S. Reviewing Challenges in Specifying Interoperability Requirement in Procurement of Health Information Systems. Stud Health Technol Inform. 2024 Jan 25;310:8-12. doi: 10.3233/SHTI230917. PMID: 38269755.

Summary

Paper C examined challenges in specifying interoperability requirements during procurement of health information systems (HIS). Through thematic analysis of 152 included peer-reviewed articles and public documents, the paper identified six recurring challenges: conflicting needs, knowledge gaps, difficulties in formulating concrete requirements, changing or incomplete requirements, premature specification of requirements, and insufficient stakeholder communication. The findings showed that lack of technical expertise and difficulties in clearly describing needs and requirements were among the most frequently shared challenges, highlighting that interoperability in procurement is not only a technical issue but also a challenge of communication, governance, and interdisciplinary competence.

Author contributions

I was primarily responsible for the design of the study, the literature review and analysis, and the drafting of the manuscript. All co-authors contributed to the interpretation of the findings and revised the manuscript critically.

8.4 Paper D (Full conference paper)

Apelgren F, **Seth M**, Jalo H, Sjöqvist BA, Candefjord S
Introducing Virtuality – Virtual Care Process Simulator: A Concept Utilizing Synthetic Data and a Digital Health Sandbox for Care Process Simulations
Published in Proceedings of the 18th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2025), Volume 1
pp. 1127–1135, 2025
DOI: `https://doi.org/10.5220/0013401300003911`

Summary

Paper D introduced *Virtuality*, a conceptual framework for realistic simulation of healthcare processes in a digital environment by combining synthetic health data, third-party applications, and a Digital Health Sandbox (DHS). A small review of existing sandbox and simulation solutions suggested a gap in healthcare simulation environments. The paper therefore addressed the need for safe and realistic environments in which digital health innovations can be verified, validated, and refined before clinical trials and real-world implementation, while reducing dependence on sensitive patient data. A central contribution was the conceptual design of the DHS around four interoperable building blocks—Providers, Services, Consumers, and a Scenario Engine—that together enable configurable simulations of healthcare workflows. The concept was illustrated through two representative use cases, namely fall detection in home care and stroke assessment in prehospital care.

Author contributions

I had the main responsibility for the conceptualization of the DHS and architecture description, including the Concept section. I also revised the manuscript.

8.5 Paper E (Scoping review)

Seth M, Apelgren F, Bakidou A, Lundgren P, Blomberg H, Sjöqvist BA, Candefjord S
FHIR Data Mapping Quality Assessment: A Scoping Review and Suggested Requirements for Future Pre-Mapping Frameworks
Accepted for publication in Computational and Structural Biotechnology Journal. Manuscript number: CSBJ-D-26-00542

Summary

Paper E examined how the quality of Fast Healthcare Interoperability Resources (FHIR) data mappings is currently assessed and what gaps hinder reliable and reproducible mapping practices. Through a scoping review of 39 included studies, screened from 742 unique records, the paper showed that mapping quality assessment remains fragmented and largely dependent on manual post-mapping review. Manual review was the dominant validation approach (31/39, 80%), while automated validators were reported less frequently and were mainly used to assess structural conformance of the generated FHIR resources rather than the correctness of the source-to-target mapping itself. Coverage was the most commonly reported quality dimension (28/39, 72%), followed by structural conformance (20/39, 51%) and semantic correctness (16/39, 41%). Terminology accuracy (13/39, 33%) and completeness (11/39, 28%) were identified less frequently, while inter-rater reliability (4/39, 10%) and retrievability (3/39, 8%) were rarely addressed. Based on the identified gaps and challenges, the paper proposed seven conceptual mapping quality dimensions and six requirements for future pre-mapping quality assessment frameworks, which were subsequently refined through consultation with four FHIR experts.

Author contributions

I was responsible for study design, selection, conducting the questionnaire and collecting FHIR expert feedback, synthesizing the result, and writing the manuscript. Manuscript revised with help from all co-authors.

8.6 Paper F (Framework paper - early draft)

Seth M, Apelgren F, Bakidou A, Lundgren P, Blomberg H, Sjöqvist BA, Candefjord S.
Introducing FHIR Mapping Assessment and Performance Score (F-MAPS) – A Pre-Mapping Quality Assessment Tool

Summary

Paper F builds directly on the challenges identified in Paper E and aims to move the work one step further by focusing on how FHIR data mapping can be facilitated in practice. Whereas Paper E examined mapping quality as a research problem, Paper F aims to focus on the development of a more systematic approach for supporting FHIR mapping processes.

In the paper, I explore how mapping workflows can be structured, documented, and assessed in a way that reduces ambiguity and supports more transparent transformation of heterogeneous data into FHIR-compatible representations. One ambition is for mapping work to not remain a hidden technical activity performed solely through ad hoc solutions or manual interpretation, but should instead be supported through clearer processes, visual guidance, and explicit mechanisms for validation and quality assessment. In the context of Virtuality, this work will hopefully contribute to the development of functionality intended to help platform users to configure mappings more systematically and to make the transformation process easier to understand and reproduce. The article is still an early draft, but the concept and ambition is clearly defined.

Author contributions

I am responsible for the conceptualization and design of the framework, the implementation-oriented analysis, and the drafting of the manuscript. All co-authors contributed to the refinement of the framework and revised the manuscript critically.

8.7 Paper G (Architecture and design paper)

Seth M, Apelgren F, Bakidou A, Lundgren P, Blomberg H, Sjöqvist BA, Candefjord S From Concept to Implementation: Design and Architecture of Virtuality — A Virtual Care Process Simulation Environment

Summary

Paper G presents the design and architecture of *Virtuality* as an interoperable simulation environment for digital health innovation.

A central contribution of the paper is the architectural formalization of Virtuality as a process-oriented simulation environment rather than a traditional isolated sandbox. In Virtuality, synthetic data, third-party applications and a Digital Health Sandbox are combined to allow Virtuality function as a build-time environment (users design and constructs simulations) and a run-time environment (participant interact with the configured simulation). This dual focus allows users to construct and run simulations, by combining reusable *Building Blocks* representing data providers, services, and user-facing applications. These can be configured into executable healthcare workflows through a Graphical User Interface (GUI), referred to as the *Scenario Engine*. The hybrid system architecture, inspired by layered - and plug-in architectures supports not only the integration of heterogeneous systems, but also simulation design, scenario configuration, execution, persistence, and traceability.

The paper also describes several of the mechanisms introduced to support interoperability, reuse, and validation within the platform. These include standards-based communication through HL7 FHIR, architectural support for modular and extensible Building Blocks, and platform mechanisms such as SHAPE, RAPID, and SEAL, which together support validation, orchestration, traceability, and reproducibility of simulations. By combining existing technologies—including semantic web technologies, interoperability standards, synthetic data, and modular software components—in a new way, Virtuality extends beyond isolated interface testing and instead enables evaluation of complete healthcare processes consisting of multiple interacting components and contextual conditions.

In the context of the thesis, Paper G represents the practical synthesis of the research journey. It formalizes how the challenges identified in earlier papers were translated into architectural principles, design requirements, and

implemented functionality.

Author contributions

I had the main responsibility for the architectural design of Virtuality, the synthesis of the design requirements, and the drafting of the manuscript. All co-authors contributed to the refinement of the platform concept and revised the manuscript critically for important content.

CHAPTER 9

Conclusions and Final Remarks

"The best way to predict the future is to invent it."
— Alan Kay

This thesis argues that the main challenge in translating digital health innovations into clinical practice is not technical functionality or interoperability in isolation, but the absence of practical environments in which interoperability, workflow integration, data transformation, and evaluation can be addressed together. Through a progression of studies on IoMT interoperability, procurement processes, FHIR data mapping, and healthcare simulation, the thesis has shown that digital health innovations must be understood as components embedded within larger care processes, organizational conditions, and data ecosystems.

From this perspective, the central contribution of the thesis is the design and implementation of *Virtuality*, a process-oriented simulation environment, consisting of synthetic data generation mechanisms, integration of third-party application and a Digital Health Sandbox, that together intend to support

simulation, testing, and evaluation of interoperable digital health innovations.

Additionally, an important finding of this research is that the concept of a Digital Health Sandbox often captures only part of the functionality required to support digital health innovation. Consequently, Virtuality evolved into a broader Digital Health Simulation Environment, of which the sandbox represents one central capability.

In conclusion, the thesis contributes both a conceptual reframing of what a Digital Health Sandbox for healthcare should support, and a concrete architectural instantiation of how such an environment can be realized.

Based on these findings, the research questions can be answered as follows.

9.1 Answering the Research Questions

RQ1. What characterises the current landscape of digital health innovations and the interoperability-related challenges that arise across their development, procurement, integration, and use in healthcare settings?

The thesis shows that one of the main challenges in digital health innovation is the difficulty in navigating available tools for interoperability and operationalizing interoperability requirements.

These challenges arise not only from technical complexity, but also from fragmented responsibilities, uneven expertise across stakeholder groups, and the challenge of translating clinical and organizational needs into system requirements. One important implication of this thesis is therefore that digital health innovations cannot be meaningfully evaluated in isolation from the wider care processes, stakeholder contexts, and implementation conditions in which they are intended to operate.

Another important finding is that existing interoperability tools, validation mechanisms, and healthcare sandboxes primarily address isolated parts of the innovation process. Many focus on message validation, standards compliance, application testing, or individual interoperability layers, while providing limited support for evaluating digital innovations as complete healthcare processes. At the same time, the thesis shows that expectations on Digital Health Sandboxes and testing environments are increasing, while there appears to be a lack of concrete examples demonstrating how these expectations can be translated into practical implementation

RQ2. What design requirements, architectural principles, and core components are needed for a Digital Health Sandbox in which digital health innovations can be integrated, simulated, tested, and evaluated prior to clinical deploy- ment?

A central conclusion of the thesis is that a Digital Health Sandbox should not be viewed solely as an environment for executing simulations. Rather, it should provide an ecosystem surrounding the broader innovation and testing process, supporting the preparation, design, integration, execution, validation, and evaluation of digital health innovations. Simulation studies represent one important use of such an environment, but the same underlying capabilities can also support other activities. The thesis has shown that one way to achieve this, is to use modular and process-oriented technologies and processes to separate the build-time environment from the runtime.

To be practically useful in healthcare, the Digital Health Sandbox must also be user-friendly and accessible for anyone, while simultaneously supporting different stakeholder perspectives, potentially through no-code or low-code interfaces. This can enable clinical users to rapidly set up, adapt, and iterate on simulations without relying entirely on technical developers. The sandbox should also support synthetic data, integration of third-party applications, and functionality for investigating interoperability across several levels, from technical connectivity and data exchange to semantics, workflows, governance, security, and regulatory requirements.

To implement these capabilities, the Digital Health Sandbox can benefit from a hybrid architecture that combines strengths from several architectural paradigms, including layered architecture, modular design, and plug-in capabilities. Within such an architecture, quality attributes such as reusability, configurability, transparency, auditability, accessibility, and reproducibility become especially important in a healthcare context. These qualities should not only guide the design of the platform, but also be reinforced through validation mechanisms that check the full simulation setup.

Finally, the thesis indicates that a Digital Health Sandbox should include concrete functionality to support practical healthcare innovation work, such as data mapping capabilities, synthetic data generation for mimicking IoMT, and automated generation of documentation, reports, and user feedback needed for evaluation, refinement, and potentially regulatory compliance.

RQ3. How can the identified requirements, architectural principles, and interoperability mechanisms be instantiated in a working Digital Health Sandbox, and how can such an environment be demonstrated through realistic healthcare use cases?

The thesis demonstrates that the identified principles can be instantiated in practice through the development of *Virtuality*, a web-based Digital Health Sandbox for simulation and evaluation of interoperable digital health innovations. Its web-based design supports cross-platform accessibility, allowing users to access the environment through a standard web browser without installing dedicated software. This also lowers the technical barrier for involving users from different organizations and technical environments.

Virtuality was implemented as a modular platform in which reusable Building Blocks can be assembled into simulation pipelines through a Scenario Engine, while a common interoperability layer supports data exchange, validation, orchestration, persistence, and controlled execution.

The practical feasibility of this architecture was demonstrated through multiple use cases. In one use case, Virtuality was used to evaluate the workflow integration of the OSISP trauma triage model in a realistic prehospital simulation involving EMS personnel. In another, the platform was used to rapidly configure a voting application for neurologists to assess synthetic stroke videos and images in order to generate labeled data for future AI development. Together, these examples show that Virtuality can be used both to evaluate AI-enabled decision support systems and to support the development of new digital health applications and research workflows. Although the thesis does not claim that Virtuality is a finished or fully validated platform, it does show that the proposed architectural principles can be operationalized in a concrete environment capable of supporting realistic and configurable simulation of digital health innovations prior to clinical deployment.

9.2 Final Remarks and Future Outlook

Although Virtuality is still evolving, the work presented in this thesis suggests that the platform, in its current form, addresses a real and important need in digital health innovation. Across the included studies, as well as in the practical use cases and surrounding project activities, it became increasingly clear that researchers, developers, clinicians, and healthcare organizations need better ways to explore, test, and refine digital health innovations before they are

introduced into clinical practice. In that sense, the relevance of Virtuality lies not only in its technical architecture, but in the type of work it makes possible.

The practical use cases presented in this thesis, together with its use in ongoing projects, indicate Virtuality is applicable beyond a single application or study context. The platform has already been used to support various use cases, and future simulations are already being prepared (such as traffic safety simulations). At the same time, several actors have recognized a broader potential in Virtuality, not only as a platform for testing and evaluation, but also for onboarding, training, and requirements gathering, interestingly, not only in healthcare contexts.

Virtuality should also not be understood as a replacement for living labs, test beds, or real-world evaluation. Rather, it has the potential to complement such environments by providing a digital and configurable space in which innovations can be explored before moving into more resource-intensive and clinically sensitive settings. In that sense, Virtuality could eventually be deployed locally within hospitals or innovation environments, allowing clinicians, researchers, and innovators to collaboratively test and refine new digital solutions in a more iterative and accessible way.

It should also be acknowledged that the general idea behind Virtuality is not entirely new. Related environments, sandboxes, and simulation tools have been proposed in healthcare and in adjacent domains [95]. Furthermore, requirements for sandbox-based environments are also beginning to emerge more explicitly, not only from technical and research perspectives, but increasingly from regulatory and innovation-oriented initiatives [19], [95]. However, what distinguishes Virtuality is not a single isolated feature, but rather the way in which several needs are brought together in one environment.

At the same time, Virtuality is far from finished. Several parts of the platform remain to be expanded, refined, and formally evaluated. Future work should therefore focus not only on extending the technical capabilities of the platform, but also on studying how it is used in practice by different stakeholder groups. This includes more systematic evaluation of usability, usefulness, and adoption, broader support for different types of healthcare workflows.

To summarize, I believe that Virtuality has the potential to become increasingly relevant in the highly regulated and complex healthcare landscape. At

the same time, I also hope that Virtuality can serve as an illustrative example of how future testing environments may be designed, and provide ideas, guidance, and practical inspiration for how tomorrow's healthcare innovations can be brought one step closer to clinical use and patient benefit.

References

[1] B. Page, D. Irving, R. Amalberti, and C. Vincent, "Health services under pressure: A scoping review and development of a taxonomy of adaptive strategies," *BMJ Quality & Safety*, vol. 33, no. 11, pp. 738–747, 2024. DOI: 10.1136/bmjqs-2023-016686.

[2] Z. Fu, E. Zitkus, S. Cook, and D. W. Maidment, "Barriers of digital health self-management in older adults with multimorbidity using a novel user-focused approach: A systematic review," *Universal Access in the Information Society*, vol. 25, p. 78, 2026. DOI: 10.1007/s10209-026-01341-8.

[3] A. Karnwal et al., "Translational applications of artificial intelligence and digital health technologies in real-world healthcare: Opportunities and challenges," *Informatics in Medicine Unlocked*, vol. 64, p. 101 769, 2026, ISSN: 2352-9148. DOI: 10.1016/j.imu.2026.101769.

[4] J. Persson and C. Rydenfält, "Why are digital health care systems still poorly designed, and why is health care practice not asking for more? three paths toward a sustainable digital work environment," *Journal of Medical Internet Research*, vol. 23, no. 6, e26694, 2021. DOI: 10.2196/26694.

[5] B. D. Naughton, M. J. Gill, E. Niarchou, M. Field, and S. Shaw, "An overlooked reason why digital innovation fails: The stigmatisation of digital health innovators," *BMJ Innovations*, vol. 11, pp. 207–211, 2025. DOI: 10.1136/bmjinnov-2025-001423.

[6] L. Petersson et al., "Challenges to implementing artificial intelligence in healthcare: A qualitative interview study with healthcare leaders in sweden," *BMC Health Services Research*, vol. 22, p. 850, 2022. DOI: 10.1186/s12913-022-08215-8.

[7] R. T. J. Loo et al., "Recommendations for successful development and implementation of digital health technology tools," *Journal of Medical Internet Research*, vol. 27, e56747, 2025. DOI: 10.2196/56747.

[8] M. Hertzum, G. Ellingsen, and Å. Cajander, "Implementing large-scale electronic health records: Experiences from implementations of epic in denmark and finland," *International Journal of Medical Informatics*, vol. 167, p. 104 868, 2022. DOI: 10.1016/j.ijmedinf.2022.104868.

[9] KPMG AB, "Granskning: Införande av millennium," KPMG AB, Sweden, Independent review report, Mar. 2025, Rapport för Västra Götalandsregionen, 2025-03-18.

[10] A. I. Stoumpos, F. Kitsios, and M. A. Talias, "Digital transformation in healthcare: Technology acceptance and its applications," *International Journal of Environmental Research and Public Health*, vol. 20, no. 4, p. 3407, 2023. DOI: 10.3390/ijerph20043407.

[11] L. Banh, F. Holldack, and G. Strobel, "Copiloting the future: How generative ai transforms software engineering," *Information and Software Technology*, vol. 183, p. 107 751, 2025, ISSN: 0950-5849. DOI: 10.1016/j.infsof.2025.107751.

[12] S. Abrahão, J. Grundy, M. Pezzè, M.-A. Storey, and D. A. Tamburri, "Software engineering by and for humans in an ai era," *ACM Transactions on Software Engineering and Methodology*, vol. 34, no. 5, p. 129, 2025. DOI: 10.1145/3715111.

[13] N. Newton, A. Bamgboje-Ayodele, R. Forsyth, A. Tariq, and M. T. Baysari, "A systematic review of clinicians' acceptance and use of clinical decision support systems over time," *npj Digital Medicine*, vol. 8, p. 309, 2025. DOI: 10.1038/s41746-025-01662-7.

[14] European Parliament and Council of the European Union, *Regulation (eu) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence and amending regulations (ec) no 300/2008, (eu) no 167/2013, (eu) no*

168/2013, (eu) 2018/858, (eu) 2018/1139 and (eu) 2019/2144 and directives 2014/90/eu, (eu) 2016/797 and (eu) 2020/1828 (artificial intelligence act), Official Journal of the European Union, OJ L 2024/1689, 12 July 2024, Regulation (EU) 2024/1689, 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj.

[15] S. Beyer et al., "Preparing for the european health data space: An open-source compiler for fast, transparent, and portable health data transformations," *Frontiers in Medicine*, vol. 12, p. 1 661 091, 2025. DOI: 10.3389/fmed.2025.1661091.

[16] E. Lee, B. A. Sjöqvist, M. Andersson Hagiwara, H. Maurin Söderholm, and S. Candefjord, "Development of verified innovation process for healthcare solutions (viphs): A stepwise model for digital health," in *Caring is Sharing – Exploiting the Value in Data for Health and Innovation: Proceedings of MIE 2023*, vol. 302, 2023, pp. 736–740. DOI: 10.3233/SHTI230250.

[17] M. Saleem and Z. Khan, "Healthcare simulation: An effective way of learning in health care," *Pakistan Journal of Medical Sciences*, vol. 39, no. 4, pp. 1185–1190, 2023. DOI: 10.12669/pjms.39.4.7145.

[18] E. Leckenby, D. Dawoud, J. C. Bouvy, and P. Jónsson, "The sandbox approach and its potential for use in health technology assessment: A literature review," *Applied Health Economics and Health Policy*, vol. 19, no. 6, pp. 857–869, 2021. DOI: 10.1007/s40258-021-00665-1.

[19] OECD, *Regulatory sandbox toolkit: A comprehensive guide for regulators to establish and manage regulatory sandboxes effectively*. Paris: OECD Publishing, 2025. DOI: 10.1787/de36fa62-en. [Online]. Available: https://doi.org/10.1787/de36fa62-en.

[20] A. Raz, L. B. Lifshitz-Milwidsky, and Y. Inbar, ""Sailing the ship while building it": Sandbox regulation and the political economy of Israeli healthcare AI innovation," *SSM - Qualitative Research in Health*, vol. 8, p. 100 642, 2025, ISSN: 2667-3215. DOI: 10.1016/j.ssmqr.2025.100642.

[21] M. E. Salwei and P. Carayon, "A sociotechnical systems framework for the application of artificial intelligence in health care delivery," *Journal of Cognitive Engineering and Decision Making*, vol. 16, no. 4, pp. 194–206, 2022. DOI: 10.1177/15553434221097357.

[22] V. Petri, T. Antoine, B. Emma, H. Hervé, F. Emma, and V. Hana, "Clinical data integration and processing challenges in healthcare caused by contemporary software design," *Digital Health*, vol. 11, p. 20 552 076 251 374 233, 2025. DOI: 10.1177/20552076251374233.

[23] M. Lehne, J. Sass, A. Essenwanger, J. Schepers, and S. Thun, "Why digital medicine depends on interoperability," *npj Digital Medicine*, vol. 2, p. 79, 2019. DOI: 10.1038/s41746-019-0158-1.

[24] C. A. Marino and C. Diaz Paz, "Enhancing interoperability for a sustainable, patient-centric health care value chain: Systematic review for taxonomy development," *Journal of Medical Internet Research*, vol. 27, e69465, 2025. DOI: 10.2196/69465.

[25] C. N. Vorisek et al., "Towards an interoperability landscape for a national research data infrastructure for personal health data," *Scientific Data*, vol. 11, p. 772, 2024. DOI: 10.1038/s41597-024-03615-3.

[26] M. D. Wilkinson et al., "The fair guiding principles for scientific data management and stewardship," *Scientific Data*, vol. 3, p. 160 018, 2016. DOI: 10.1038/sdata.2016.18.

[27] E. E. Umberfield, C. J. Staes, T. P. Morgan, R. W. Grout, B. W. Mamlin, and B. E. Dixon, "Syntactic interoperability and the role of syntactic standards in health information exchange," in *Health Information Exchange*, B. E. Dixon, Ed., 2nd ed., Cambridge, MA: Academic Press, 2023, pp. 217–236, ISBN: 9780323908023. DOI: 10.1016/B978-0-323-90802-3.00004-6.

[28] C. Ntafi, S. Spyrou, P. Bamidis, and M. Theodorou, "The legal aspect of interoperability of cross border electronic health services: A study of the european and national legal framework," *Health Informatics Journal*, vol. 28, no. 3, p. 14 604 582 221 128 722, 2022. DOI: 10.1177/14604582221128722.

[29] Healthcare Information and Management Systems Society (HIMSS), *Interoperability in healthcare*, https://www.himss.org/resources/interoperability-healthcare, Accessed: 2026-06-28, 2026.

[30] K. Adegoke et al., "Interoperability as a catalyst for digital health and therapeutics: A scoping review of emerging technologies and standards (2015–2025)," *International Journal of Environmental Research and Public Health*, vol. 22, no. 10, p. 1535, 2025. DOI: 10.3390/ijerph22101535.

[31] European Commission, "New european interoperability framework: Promoting seamless services and data flows for european public administrations," Publications Office of the European Union, Luxembourg, Tech. Rep., 2017. DOI: 10.2799/78681.

[32] L. Tangi, M. Combetto, J. Martin Bosch, and A. P. Rodriguez Müller, "Artificial intelligence for interoperability in the european public sector: An exploratory study," Publications Office of the European Union, Luxembourg, Tech. Rep. JRC134713, 2023. DOI: 10.2760/633646.

[33] European Parliament and Council of the European Union, *Regulation (eu) 2024/903 laying down measures for a high level of public sector interoperability across the union (interoperable europe act)*, Official Journal of the European Union, OJ L 2024/903, 22 March 2024, 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/903/oj.

[34] EHRxF. "European electronic health record exchange format (ehrxf)." Accessed June 24, 2026, Accessed: Jun. 24, 2026. [Online]. Available: https://ehr-exchange-format.eu/.

[35] N. da Silva Carvalho et al., "Accelerating european electronic health record exchange format (eehrxf) implementation across europe: A policy perspective," *Frontiers in Medicine*, vol. 12, p. 1 648 170, 2025. DOI: 10.3389/fmed.2025.1648170.

[36] European Commission, *Commission recommendation of 6 february 2019 on a european electronic health record exchange format*, 2019. Accessed: Jun. 28, 2026. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019H0243.

[37] World Health Organization, *WHO Director-General's opening remarks at the media briefing on COVID-19 – 21 August 2020*, https://www.who.int/news-room/speeches/item/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---21-august-2020, Accessed: 2026-06-22, Aug. 2020.

[38] A. Castiglione, M. Umer, S. Sadiq, M. S. Obaidat, and P. Vijayakumar, "The role of internet of things to control the outbreak of covid-19 pandemic," *IEEE Internet of Things Journal*, vol. 8, no. 21, pp. 16 072–16 082, 2021. DOI: 10.1109/JIOT.2021.3070306.

[39] D. Alshehri, N. Noman, R. Chiong, S. J. Miah, A. L. Sverdlov, and D. T. M. Ngo, "Factors influencing the adoption of internet of medical things for remote patient monitoring: A systematic literature review," *Computers in Biology and Medicine*, vol. 192, no. Part B, p. 110 142, 2025, ISSN: 0010-4825. DOI: 10.1016/j.compbiomed.2025.110142.

[40] IEEE Standards Association, *What is the internet of medical things (iomt) & what does it mean for medical device connectivity?* https://standards.ieee.org/beyond-standards/the-internet-of-medical-things-iomt/, Accessed: 2026-06-22, Jun. 2026.

[41] B. A. Mohammed et al., "Security challenges and solutions in internet of medical things (iomt) communication: A review," *Journal of King Saud University – Computer and Information Sciences*, vol. 38, p. 254, 2026. DOI: 10.1007/s42979-026-03980-8.

[42] S. Jayousi et al., "From innovation to integration: Bridging the gap between iomt technologies and real-world health management systems," *Sensors*, vol. 25, no. 21, p. 6660, 2025. DOI: 10.3390/s25216660.

[43] M. Seth et al., "Technologies for interoperable internet of medical things platforms to manage medical emergencies in home and prehospital care: Protocol for a scoping review," *JMIR Research Protocols*, vol. 11, no. 9, e40243, 2022. DOI: 10.2196/40243.

[44] HL7 International and Firely, *2026 state of fhir survey results*, Global survey report based on responses from 101 participants across 63 countries, May 2026. [Online]. Available: https://www.hl7.org/fhir/.

[45] HL7 International, *Fhir version history*, https://build.fhir.org/history.html, Accessed: 2026-06-24, 2026.

[46] M. Ayaz, M. F. Pasha, M. Y. Alzahrani, R. Budiarto, and D. Stiawan, "The fast health interoperability resources (fhir) standard: Systematic literature review of implementations, applications, challenges and opportunities," *JMIR Medical Informatics*, vol. 9, no. 7, e21929, 2021. DOI: 10.2196/21929.

[47] HL7 International, *Getting started with fhir*, https://build.fhir.org/overview-dev.html, FHIR CI-Build, accessed 2026-06-24, 2026.

[48] HL7 International, *FHIR Specification (R5): Resource Index*, https://hl7.org/fhir/resource.html, FHIR Release 5, Version 5.0.0; accessed 2026-06-22, 2023.

[49] HL7 CDS Hooks, *Cds hooks 2.0.1 specification*, Accessed 2026-06-22, 2025. [Online]. Available: https://cds-hooks.hl7.org/.

[50] HL7 FHIR, *Smart app launch implementation guide*, Accessed 2026-06-22, 2023. [Online]. Available: https://hl7.org/fhir/smart-app-launch/.

[51] HL7 FHIR, *Smart app launch: Launch and authorization*, Accessed 2026-06-22, 2024. [Online]. Available: https://www.hl7.org/fhir/smart-app-launch/app-launch.html.

[52] S. Makhni et al., "Meeting the challenges of electronic health record (EHR) optimization," *npj Digital Medicine*, vol. 9, p. 8, 2026. DOI: 10.1038/s41746-025-02178-w.

[53] Apple Inc., *Accessing Health Records*, https://developer.apple.com/documentation/healthkit/accessing-health-records, Apple Developer Documentation, HealthKit; accessed 2026-06-22, 2026.

[54] Microsoft Research. "Health data interoperability: Standards and policy," Microsoft Research, Accessed: Jun. 22, 2026. [Online]. Available: https://www.microsoft.com/en-us/research/project/health-data-interoperability-standards-and-policy/.

[55] A. Marfoglia, F. Nardini, V. A. Arcobelli, S. Moscato, S. Mellone, and A. Carbonaro, "Towards real-world clinical data standardization: A modular fhir-driven transformation pipeline to enhance semantic interoperability in healthcare," *Computers in Biology and Medicine*, vol. 187, p. 109 745, Mar. 2025. DOI: 10.1016/j.compbiomed.2025.109745.

[56] T. K. H. Le, W. Kongsiriwattana, A. Meny, R. Kaewsaeng-On, C. Nimyungdee, and K. Namee, "Enhancing healthcare interoperability with fhir: A systematic approach to online data management," in *Proceedings of the 2024 9th International Conference on Information Systems Engineering*, ser. ICISE '24, New York, NY, USA: Association for Computing Machinery, 2025, pp. 35–40, ISBN: 9798400717369. DOI: 10.1145/3711954.3711958.

[57] IHE International, *Integrating the healthcare enterprise (ihe)*, `https://www.ihe.net/`, Accessed: 2026-06-24, 2026.

[58] P. Tabari, G. Costagliola, M. De Rosa, and M. Boeker, "State-of-the-art fast healthcare interoperability resources (fhir)-based data model and structure implementations: Systematic scoping review," *JMIR Medical Informatics*, vol. 12, e58445, 2024. DOI: `10.2196/58445`. [Online]. Available: `https://medinform.jmir.org/2024/1/e58445`.

[59] A. Marfoglia, V. A. Arcobelli, S. Moscato, A. A. La Mattina, S. Mellone, and A. Carbonaro, "Challenges of health data standard adoption and usage: A systematic review," *Journal of Biomedical Informatics*, vol. 178, p. 105 022, 2026, ISSN: 1532-0464. DOI: `10.1016/j.jbi.2026.105022`.

[60] S.-C. Haw, T.-E. Tai, W.-E. Kong, K.-W. Ng, and P. Naveen, "Modernizing legacy healthcare systems utilizing semantic-based automated mapping," in *2025 IEEE International Conference on Smart Internet of Things (SmartIoT)*, Institute of Electrical and Electronics Engineers (IEEE), Dec. 2025, pp. 271–277. DOI: `10.1109/SmartIoT66867.2025.00045`.

[61] J. Kruse, J. Wiedekopf, A.-K. Kock-Schoppenhauer, A. Essenwanger, J. Ingenerf, and H. Ulrich, "A generic transformation approach for complex laboratory data using the fast healthcare interoperability resources mapping language: Method development and implementation," *JMIR Medical Informatics*, vol. 12, e57569, 2024. DOI: `10.2196/57569`.

[62] T. O. Ebo, A. C. David-Olawade, D. M. Ebo, E. Egbon, and D. B. Olawade, "Transforming healthcare delivery: A comprehensive review of digital integration, challenges, and best practices in integrated care systems," *Digital Engineering*, vol. 6, p. 100 056, 2025, ISSN: 2950-550X. DOI: `10.1016/j.dte.2025.100056`.

[63] O. Olakotan, R. Samuriwo, H. Ismaila, and S. Atiku, "Usability challenges in electronic health records: Impact on documentation burden and clinical workflow: A scoping review," *Journal of Evaluation in Clinical Practice*, vol. 31, no. 4, e70189, 2025. DOI: `10.1111/jep.70189`.

[64] D. E. Vega, "Automated interoperability testing of healthcare information systems," in *Advances in Computers*, vol. 85, Elsevier, 2012, pp. 213–276. DOI: `10.1016/B978-0-12-396526-4.00005-9`.

[65] HL7 International, *Using the FHIR Validator*, https://confluence.hl7.org/display/FHIR/Using+the+FHIR+Validator, FHIR Confluence documentation, last updated 19 March 2026; accessed 2026-06-12, 2026.

[66] J. Walonoski, R. Scanlon, C. Dowling, M. Hyland, R. Ettema, and S. Posnack, "Validation and testing of fast healthcare interoperability resources standards compliance: Data analysis," *JMIR Medical Informatics*, vol. 6, no. 4, e10870, 2018. DOI: 10.2196/10870.

[67] Inferno Framework, *FHIR Validation*, https://inferno-framework.github.io/docs/writing-tests/fhir-validation/, Inferno Framework documentation; accessed 2026-05-22, 2026.

[68] A.-G. Bergé, *Gazelle testing environment*, Presentation, IHE Europe / Kereval, Slide deck describing Gazelle validation tools, simulators, proxy, and test management platform, n.d.

[69] IHE Europe, *New testing tools on gazelle testing platform*, https://www.ihe-europe.net/new_testing_tools_on_gazelle, Accessed: 2026-06-10, Oct. 2020.

[70] HL7 International, *Accelerator program overview*, https://confluence.hl7.org/spaces/AP/overview, HL7 Confluence; accessed 2026-06-22, 2026.

[71] H. Maurin Söderholm et al., "Research challenges in prehospital care: The need for a simulation-based prehospital research laboratory," *Advances in Simulation*, vol. 4, no. 1, p. 3, Feb. 2019. DOI: 10.1186/s41077-019-0090-0.

[72] M. A. Hagiwara, B. A. Sjöqvist, L. Lundberg, B.-O. Suserud, M. Henricson, and A. Jonsson, "Decision support system in prehospital care: A randomized controlled simulation study," *The American Journal of Emergency Medicine*, vol. 31, no. 1, pp. 145–153, Jan. 2013. DOI: 10.1016/j.ajem.2012.06.030.

[73] M. Andersson Hagiwara, L. Lundberg, B. A. Sjöqvist, and H. Maurin Söderholm, "The effects of integrated it support on the prehospital stroke process: Results from a realistic experiment," *Journal of Healthcare Informatics Research*, vol. 3, no. 3, pp. 300–328, May 2019. DOI: 10.1007/s41666-019-00053-4.

[74] N. Alotaibi, C. Brown Wilson, and M. Traynor, "Enhancing digital readiness and capability in healthcare: A systematic review of interventions, barriers, and facilitators," *BMC Health Services Research*, vol. 25, no. 1, p. 500, 2025. DOI: 10.1186/s12913-025-12663-3.

[75] Y. Zhang, X. Wang, S. Tarba, and A. Neely, "Digital transformation through innovation: The human-ai decision spectrum," *Technovation*, vol. 153, p. 103 524, 2026, ISSN: 0166-4972. DOI: 10.1016/j.technovation.2026.103524.

[76] J. Venable, J. Pries-Heje, and R. Baskerville, "A comprehensive framework for evaluation in design science research," in *Design Science Research in Information Systems. Advances in Theory and Practice*, K. Peffers, M. Rothenberger, and B. Kuechler, Eds., ser. Lecture Notes in Computer Science, vol. 7286, Berlin, Heidelberg: Springer, 2012, pp. 423–438. DOI: 10.1007/978-3-642-29863-9_31.

[77] M. Seth, H. Jalo, Å. Högstedt, O. Medin, B. A. Sjöqvist, and S. Candefjord, "Technologies for interoperable internet of medical things platforms to manage medical emergencies in home and prehospital care: Scoping review," *Journal of Medical Internet Research*, vol. 27, e54470, 2025. DOI: 10.2196/54470.

[78] I. Jussen, F. Möller, J. Schweihoff, A. Gieß, G. Giussani, and B. Otto, "Issues in inter-organizational data sharing: Findings from practice and research challenges," *Data & Knowledge Engineering*, vol. 150, p. 102 280, 2024, ISSN: 0169-023X. DOI: 10.1016/j.datak.2024.102280.

[79] A. Cider and N. Ronnby, "Designing a web-based graphical user interface for managing digital twin simulations in healthcare," Master's Thesis, Chalmers University of Technology, 2025. [Online]. Available: https://hdl.handle.net/20.500.12380/309698.

[80] W. Bergman and S. On Huang, "Designing a sandbox for smarter health solutions," Master's thesis, Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, 2025. [Online]. Available: https://hdl.handle.net/20.500.12380/310624.

[81] A. Bakidou et al., "Usability of a prehospital clinical decision support system based on artificial intelligence: A case study of ems personnel's experience of using on-scene injury severity prediction (osisp)," Manuscript in writing, 2026.

[82] F. Apelgren, M. Seth, H. Jalo, B. A. Sjöqvist, and S. Candefjord, "Introducing virtuality – virtual care process simulator: A concept utilizing synthetic data and a digital health sandbox for care process simulations," in *Proceedings of the 18th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2025) - Volume 1*, 2025, pp. 1127–1135, ISBN: 978-989-758-731-3. DOI: 10.5220/0013401300003911.

[83] A. Aljohani and H. Do, "Promptdebt: A comprehensive study of technical debt across llm projects," in *Proceedings of the 29th International Conference on Evaluation and Assessment in Software Engineering (EASE '25)*, Association for Computing Machinery, 2025, pp. 371–382. DOI: 10.1145/3756681.3756976.

[84] A. Martin-Lopez et al., "More code, less understanding? on the impact of ai assistants on developers' productivity and code ownership," *IEEE Transactions on Software Engineering*, 2026. DOI: 10.1109/TSE.2026.3679627.

[85] X. Chen, C. Gao, C. Chen, G. Zhang, and Y. Liu, "An empirical study on challenges for llm application developers," *ACM Transactions on Software Engineering and Methodology*, vol. 34, no. 7, Sep. 2025, ISSN: 1049-331X. DOI: 10.1145/3715007. [Online]. Available: https://doi.org/10.1145/3715007.

[86] *Iso/iec/ieee 42010:2022 software, systems and enterprise — architecture description*, International standard, Geneva, Switzerland: ISO, IEC, and IEEE, 2022.

[87] *Iso/iec 25010:2023 systems and software engineering — systems and software quality requirements and evaluation (square) — product quality model*, International standard, International Organization for Standardization and International Electrotechnical Commission, 2023.

[88] A. Oudshoorn, M. Ortiz, and M. Šimkus, "Shacl validation in the presence of ontologies: Semantics and rewriting techniques," *Artificial Intelligence*, vol. 352, p. 104 483, 2026, ISSN: 0004-3702. DOI: 10.1016/j.artint.2026.104483.

[89] M. Richards and N. Ford, *Fundamentals of Software Architecture*, 2nd ed. O'Reilly Media, Inc., Mar. 2025, ISBN: 9781098175504.

[90] A. Bakidou, E.-C. Caragounis, M. Andersson Hagiwara, A. Jonsson, B. A. Sjöqvist, and S. Candefjord, "On scene injury severity prediction (osisp) model for trauma developed using the swedish trauma registry," *BMC Medical Informatics and Decision Making*, vol. 23, no. 1, p. 206, 2023. DOI: 10.1186/s12911-023-02290-5.

[91] D. Wallstén, G. Axton, E. Lee, A. Bakidou, B. A. Sjöqvist, and S. Candefjord, "Design for integrating explainable ai for dynamic risk prediction in prehospital it systems," in *Artificial Intelligence, Social Computing and Wearable Technologies*, vol. 113, 2023, pp. 268–278. DOI: 10.54941/ahfe1004199.

[92] M. Seth et al., "Reviewing challenges in specifying interoperability requirement in procurement of health information systems," *Studies in Health Technology and Informatics*, vol. 310, pp. 8–12, 2024. DOI: 10.3233/SHTI230917.

[93] M. Seth et al., "Fhir data mapping quality assessment: A scoping review and suggested requirements for future pre-mapping frameworks," Accepted for publication in *Computational and Structural Biotechnology Journal*. Manuscript ID: CSBJ-D-26-00542. Accepted, 2026.

[94] M. Seth et al., "From concept to implementation: Design and architecture of virtuality — a virtual care process simulation environment," To be submitted, 2026.

[95] Innovative Health Initiative. "New project bridge to bring regulatory sandboxes into health," Accessed: Jun. 25, 2026. [Online]. Available: https://www.ihi.europa.eu/news-events/newsroom/new-project-bridge-bring-regulatory-sandboxes-health.