INSTITUTO SUPERIOR DE ENGENHARIA DO PORTO
MESTRADO EM ENGENHARIA INFORMÁTICA

O Futuro da Monitorização de Idosos: Integração da Realidade Aumentada para Auxílio na Prestação de Cuidados

RAFAEL GONÇALVES MARTINS
Setembro de 2025

POLITÉCNICO
DO PORTO

ISEP INSTITUTO SUPERIOR DE ENGENHARIA DO PORTO

The Future of Elderly Monitoring: Integrating Augmented Reality for Improved Care

Rafael Gonçalves Martins

A dissertation submitted in partial fulfilment of the requirements for the degree of Master in Informatics Engineering, Specialization Area of Games, Graphics and Interactive Systems

Advisor: Luís Conceição

Co-advisor: Goreti Marreiros

Porto, September 2025

Statement of Integrity

I hereby declare that I have conducted this academic work with integrity.

I have not plagiarized or applied any form of undue use of information or falsification of results along the process leading to its elaboration.

Therefore, the work presented in this document is original and authored by me, having not previously been used for any other end. The exceptions are explicitly recognized in the section "Ethical considerations" of the first chapter. This section also states how AI tools were used and for what purpose.

I further declare that I have fully acknowledged the Code of Ethical Conduct of P.PORTO.

ISEP, Porto, September 27, 2025

Dedicatory

To my family and colleagues

Abstract

Remote patient monitoring systems are a viable approach to address the evolving needs of healthcare facilities. However, while conventional systems can reduce constant in-person supervision, they require caregivers to remain near a computer, presenting a challenge since caregiving is inherently a hands-on role and requires freedom of movement. This dissertation project designed and implemented an augmented reality window-based interface system, AleRa, that overlays patient information onto the caregiver's field of vision, fetched from a FHIR compliant server that is part of the RM4Health ITEA project's architecture. AleRa aimed to explore if augmented reality could enable continuous monitoring while allowing caregivers to remain actively engaged in direct patient care, improving response time to changes in health conditions. Sixteen participants tested the system in a case study, praising its intuitiveness, visual quality, and strong potential as a helping tool for caregivers. Improving responsiveness of certain interface elements, and lowering the dependency on the companion mobile app, were the main constructive feedback, with other additional features also being suggested for future work. Overall, the dissertation project achieved its objectives, made relevant scientific contributions to augmented reality and health monitoring through implementation and publications, and contributed to the dissemination of related international projects.

Keywords: Augmented Reality, Interactive Systems, Remote Monitoring, Artificial Intelligence

Resumo

Os sistemas de monitorização remota de pacientes têm vindo a ganhar relevância como forma de resposta às crescentes necessidades de cuidadores e instituições de saúde. Múltiplos aspetos da vida diária podem beneficiar de monitorização remota, nomeadamente a monitorização continua de sinais vitais para deteção precoce de problemas de saúde, controlo de acessos e de localização, e gestão de prescrições e tomas de medicação. Estes sistemas têm vindo a ser explorados na literatura e observa-se que permitem reduzir a necessidade de supervisão presencial constante. No entanto, nas suas versões mais tradicionais, estes sistemas obrigam os cuidadores a permanecer próximos de um computador para aceder à informação disponibilizada. Tal entra em conflito com a natureza prática do trabalho de prestação de cuidados, que exige mobilidade, liberdade de movimentos e contacto direto com os utentes.

Este projeto de dissertação, denominado AleRa, propôs-se a explorar o potencial da realidade aumentada (AR) como solução para este desafio. O sistema desenvolvido consiste principalmente numa aplicação para óculos de realidade aumentada que sobrepõe interfaces tridimensionais ao campo de visão do cuidador, complementada por uma aplicação móvel de apoio a certas funcionalidades. O projeto AleRa derivou inicialmente do caso de uso português do projeto ITEA multinacional RM4Health, e é capaz de integrar com este sistema, mais especificamente com o seu servidor FHIR de armazenamento de dados e o seu servidor de autenticação, para obter os dados médicos a apresentar nas suas interfaces AR. Visto ser um projeto multidisciplinar, envolvendo realidade aumentada, saúde e inteligência artificial, especial atenção foi dada para garantir conformidade com regulamentos de ética e de proteção de dados.

Após uma revisão da literatura, identificação de requisitos, e desenho da solução, a implementação do sistema recorreu ao motor de jogo Unity, utilizando o pacote OpenXR, ao invés de pacotes de fornecedores específicos (ex.: Meta ou Microsoft), para maximizar compatibilidade com múltiplos dispositivos de AR. Em termos de interfaces, foi desenvolvido um modelo modular de janelas virtuais, em que cada janela controla o seu próprio conteúdo, as janelas são independentes do seu posicionamento. Para o sistema de posicionamento, foram implementadas duas grelhas que gerem as posições possíveis de janelas no mundo real, permitindo diferenciar entre janelas “próximas” (interativas) e “distantes” (principalmente visuais). Estas grelhas seguem os movimentos do cuidador para garantir visualização contínua de informação mesmo em movimento. Em termos de janelas específicas, foram desenvolvidos componentes para monitorização em tempo real de sinais vitais, listagem de alertas, visualização de tendências históricas dos sinais vitais, e acompanhamento da toma de medicação semanal.

O sistema foi avaliado quanto à sua usabilidade e relevância de informação através de um estudo de caso com 16 participantes, divididos em três grupos. Os participantes consideraram o sistema intuitivo e de boa qualidade visual, e demonstraram forte interesse na aplicação prática deste tipo de sistemas não só em contextos profissionais de saúde, mas também nas suas próprias áreas de atividade profissional. Entre os aspetos menos positivos, identificou-se a falta de responsividade de alguns elementos da interface, a dependência do sistema na aplicação móvel, e o desconforto causado após uso prolongado devido ao peso do HMD utilizado. Os participantes sugeriram também várias novas funcionalidades, que poderão ser abordadas em trabalho futuro, como a personalização das posições das janelas, a integração de um sistema de localização de pacientes dentro de instituições (ex.: hospitais), e o enriquecimento da gestão de medicação com informações adicionais (ex.: janelas de folhetos e instruções de administração).

O projeto AleRa alcançou os objetivos definidos, e contribuiu para a área de Jogos, Gráficos e Sistemas Interativos, ao propor e validar um sistema de interfaces de AR orientada para contextos de saúde, desenhado também para potencial reutilização noutros contextos profissionais, como a Indústria 4.0. Este trabalho integrou-se nas atividades do centro investigação GECAD, relacionando-se com os projetos internacionais Inno4Health e RM4Health, o que resultou também na submissão e publicação de artigos científicos em conferências e revistas indexadas.

Palavras-Chave: Realidade Aumentada, Sistemas Interativos, Monitorização Remota, Inteligência Artificial

Acknowledgement

First and foremost, I would like to thank all of my family and friends, for always being there to listen, and for motivating me every step of the way. It is their unconditional support that made it possible for me to dedicate myself fully to the development of this master's dissertation project.

I am deeply grateful to my advisors, Luís Conceição and Goreti Marreiros, for their guidance and support throughout this work, and for trusting me with the equipment and resources needed to research and develop such an uncommon system as this augmented reality interface.

To my dear friend Márcia Carvalho, a very special thanks for all the encouragement over the years, and for sharing with me both the joys and the challenges of completing this master's course and dissertation project.

To Hugo and Rita, colleagues at GECAD, a big thank you for all the work carried out on the RM4Health project, from which the idea for this dissertation originated.

I would also like to thank all my other GECAD colleagues, whose contributions and valuable input, whether direct or indirect, helped shaping this dissertation.

My sincere thanks go as well to the participants in the case study conducted as part of this dissertation, who gave their time to travel to the GECAD laboratories, showed great curiosity in exploring the developed system, and provided invaluable feedback for future iterations.

Additionally, I would like to thank the RM4Health project (ITEA-2021-21022-RM4Health), funded by the European Regional Development Fund (ERDF) within the project number COMPETE2030-FEDER-00391100. The work carried out in RM4Health, particularly within its Portuguese use case, directly inspired and informed this dissertation, with the developed interfaces integrating with RM4Health servers to obtain the data necessary for visualization.

Finally, I would like to thank ISEP, its teachers, staff, and students, for providing me with a place to belong, learn, and grow, from the first year of my bachelor's in 2018, to the completion of my master's in 2025. Thank you for these seven years of unforgettable memories and experiences.

Contents

List of Figures

List of Tables

List of Acronyms

3D	Three-dimensional
AI	Artificial Intelligence
API	Application Programming Interface
APK	Android Package Kit
AR	Augmented Reality
C#	C Sharp (programming language)
CRUD	Create, Read, Update, and Delete (the four fundamental operations of data management)
DOI	Digital Object Identifier
EU	European Union
EWS	Early Warning Score
FHIR	Fast Healthcare Interoperability Resources
GECAD	Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development
GPS	Global Positioning System
GRASP	General Responsibility Assignment Software Patterns
HL7	Health Level Seven
HMD	Head-Mounted Displays
HMI	Human Machine Interface
HTML	HyperText Markup Language
ICU	Intensive Care Unit
IMCU	Intermediate Care Unit
IoT	Internet of Things
IP	Internet Protocol
ISBN	International Standard Book Number
ISEP	Instituto Superior de Engenharia do Porto
IPP	Instituto Politécnico do Porto
ITEA	Information Technology for European Advancement
JSON	JavaScript Object Notation
JWT	JSON Web Tokens
LCD	Liquid Crystal Display
LERP	Linear Interpolation
LLM	Large Language Model
MEWS	Modified Early Warning Score
ML	Machine Learning
MR	Mixed Reality
MTCNN	Multi-task Cascaded Convolutional Network
NASA	National Aeronautics and Space Administration
NASA-TLX	NASA Task Load Index
NEWS	National Early Warning Score
ONNX	Open Neural Network Exchange
PICOC	Population, Intervention, Comparation, Outcome, and Context

PIN	Personal Identification Number
PREPD	Dissertation Preparation (course unit at ISEP)
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
QR	Quick Response (as in “QR Code”)
RAM	Random-Access Memory
Raw-TLX	Raw version of NASA-TLX (simplified scoring variant)
REST	Representational State Transfer
RFID	Radio-Frequency Identification
RM	Remote Monitoring
SDK	Software Development Kit
SLR	Systematic Literature Review
SOLID	Single Responsibility, Open/Closed, Liskov Substitution, Interface Segregation, and Dependency Inversion (object-oriented design principles)
SUS	System Usability Scale
UDP	User Datagram Protocol
UI	User Interface
URL	Uniform Resource Locator
URP	Universal Render Pipeline (from Unity engine)
VR	Virtual Reality
XR	Extended Reality

1 Introduction

This dissertation was developed in partial fulfilment of the requirements for the degree of Master in Informatics Engineering – Specialization Area of Games, Graphics and Interactive Systems. This dissertation was carried out at the Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD), localized in Instituto Superior de Engenharia do Porto (ISEP), originally deriving from, and capable of integrating with, the Portuguese use case of the Information Technology for European Advancement (ITEA) 4 international project RM4Health [1], [2].

This chapter introduces this dissertation project, how it derived from and can integrate with RM4Health's system architecture, a list of ethical considerations, and finally this dissertation document's structure.

1.1 Contextualization

Advances in sensor technology and the Internet of Things (IoT) have made remote monitoring (RM) of patients an increasingly viable approach to address the evolving needs of healthcare facilities. In particular, remote monitoring of the vital signs and daily activities has emerged as a prominent research area that aims to provide faster and more effective care while promoting independence. Multiple aspects of daily life can benefit from RM, including continuous tracking of vital signs for early detection of health anomalies, activity monitoring, location tracking and access control, and medication intake management.

Building on these possibilities, RM4Health is an international ITEA 4 project committed to the development of wearable sensors and open technology platforms for continuous vital sign monitoring in both healthcare and sports [2]. RM4Health covers six use cases, bringing together a consortium of 18 partners from four countries: Finland, Portugal, Spain, and the Netherlands. In particular, the Portuguese use case focuses on the creation of an intelligent health monitoring ecosystem, depicted in Figure 1, for seniors who live in residential homes or in domiciliary care, that allows comprehensive monitoring of various aspects of a senior's daily life, and facilitates the tasks of the assigned caregivers [3].

RM4Health utilizes, for the storage of all health-related data, a Fast Healthcare Interoperability Resources (FHIR) compliant server and database. FHIR, initially presented in a work by Bender and Sartipi [4], is a specification developed by the Health Level 7 (HL7) organization, which has been increasingly adopted by the healthcare industry for exchanging health-related data [5]. It specifies a set of resource formats for data storing (e.g.: a Patient resource must allow to store a name, address, and date of birth) and a Representational State Transfer (REST) Application Programming Interface (API) to access and manage said resources.

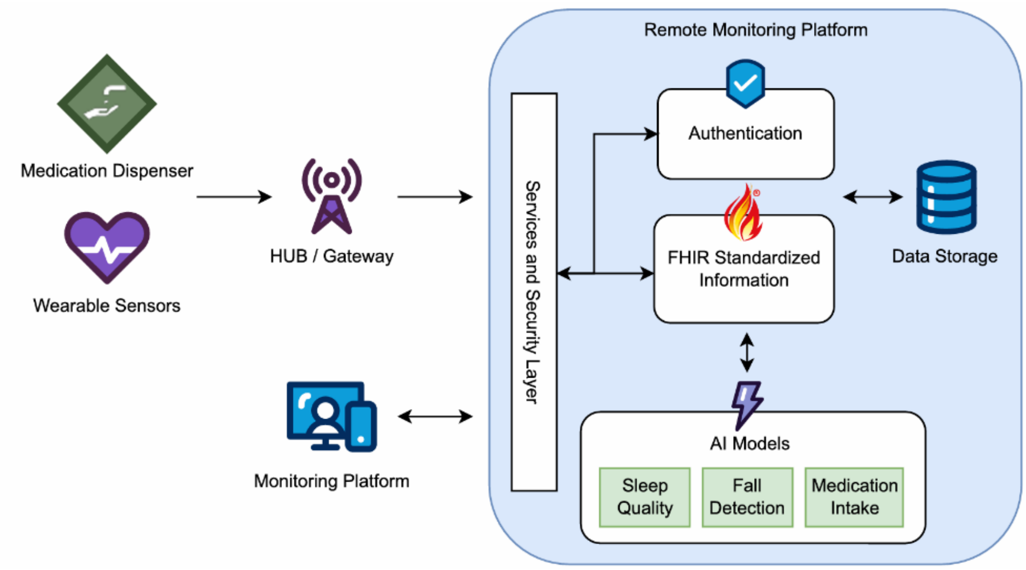

Figure 1 – General representation of the system architecture for the Portuguese RM4Health use case.[1]

However, while such systems can reduce the need for constant in-person supervision, they often require caregivers to remain near a computer to track patient data. This presents a challenge, as caregiving is inherently a hands-on role and caregivers require the freedom to roam around their facility as needed.

Augmented reality (AR) and mixed reality (MR) technologies present a promising solution to this challenge. Currently, AR is an actively explored technology in literature, with many studies evaluating or confirming its potential application in different fields. For instance, in Industry 4.0, Canito et al. [6] present an architecture for an AR solution capable of combining historical and contextual information about the user to best assist them in pending tasks. Meanwhile, Marreiros et al. [7] mention how AR solutions could be used as an adaptive human-machine interface (HMI) to visualize data gathered from physical machines in smart factories, presenting

---

[1] Source: A publication work by Martins et al. [3] detailing the system architecture of the Portuguese use case of the RM4Health project.

relevant information according to the user and context. In the health field, Martinho et al. [8] review how AR is being used as a gamification technique to help users exercise, or as a physical or cognitive rehabilitation tool.

As such, this dissertation project proposed and aimed to verify that by providing context-sensitive overlays of patient information directly in a caregiver's field of vision, AR/MR systems have the potential to enable continuous unobtrusive monitoring while allowing caregivers to remain actively engaged in direct patient care, improving response time to emergencies and changes in health status.

Its resulting system was named AleRa, a hybrid term derived from the Portuguese word "Alerta" (meaning "alert" in English), reflecting the dissertation project's focus on health monitoring and awareness, while also being intentionally stylized with the letters A and R capitalized, referencing augmented reality (AR), the system's technological foundation.

The resulting AleRa system is an AR application for head-mounted displays (HMD) that translates familiar healthcare interfaces to the real-world environment, including features such as real-time vital sign charts, patient health history visualizations, alerts for abnormal conditions, and medication intake adherence. It is designed to utilize, as its main source of health information, any FHIR-compliant data storage server, allowing for easy integrating with the RM4Health solution, or similar architectures.

A key focus of the dissertation was to design and development an intuitive informational AR interface that could serve as a foundation for similar future applications in healthcare or other professional contexts. In addition to the system development, this dissertation project also involved carrying out a case study to validate the usability and intuitiveness of the developed interfaces, and the usefulness of the medical information displayed.

1.2 Objectives

As described previously, the aim of this dissertation project was to explore the potential of AR and MR technology in continuous health monitoring, develop an AR interface to caregivers in patient monitoring, and integrate it with the RM4Health project architecture.

From this, a set of core objectives were able to be identified:

- Research the state of the art in remote health monitoring and the use of AR / MR technology in healthcare, particularly in health monitoring.
- Develop an AR interface that allows caregivers to monitor various health parameters and receive alerts for one or multiple patients simultaneously. If necessary, the interface should also enable interaction with the system.
- Integrate the AR interface with RM4Health's system architecture, testing the visualization of medical information through AR overlays.

Since this dissertation was carried out at GECAD’s research center, alongside the RM4Health project, an additional objective was that the developed work was expected to contribute to the submission of one or more scientific papers to international conferences and journals indexed in Scopus and Web of Science, related to the research areas of this dissertation.

1.3 Ethical Considerations

Given the interdisciplinary nature of this dissertation project, encompassing engineering, healthcare, and user interactions, it was of the utmost importance to guarantee adherence to all relevant ethical guidelines and codes of conduct. The following sub sections outline the ethical commitments that guided the dissertation project's execution, as well as other relevant disclosures such as previously developed work utilized, and the usage of artificial intelligence (AI) tools during this dissertation.

1.3.1 Compliance with Institutional Academic Integrity Policies

As a student of the Instituto Politécnico do Porto (IPP), this dissertation project adhered to the institution's academic integrity policies. This included, but was not limited to:

- **Proper Attribution**: Ensuring all ideas, sentences, paragraphs, or texts from external sources were correctly cited and referenced to avoid plagiarism, in order to respect intellectual property.
- **Original Work**: Committing to producing original research, findings, and system designs and implementations, to contribute authentically to the academic community.

1.3.2 Adherence to Professional Engineering Ethics

In alignment with the ethical standards of the engineering profession, this dissertation project also upheld principles such as:

- **Honesty and Integrity**: Maintaining transparency and truthfulness in all aspects of the project, including data collection, analysis, and reporting.
- **Accountability**: Taking responsibility for the project's outcomes and ensuring that all engineering practices met the highest professional standards.

1.3.3 Compliance with Data Protection Regulations

Given the dissertation project's involvement with medical applications and user data, strict compliance with data protection regulations was essential:

- **Data Anonymization**: Implementing measures to anonymize user data, protecting individual identities and ensuring privacy.
- **Data Security**: Utilizing encryption and secure data storage solutions to safeguard sensitive information against unauthorized access.
- **Data Minimization and Deletion**: Collecting only the data necessary for the project's purposes and ensuring that data was deleted upon request or when it was no longer needed, in accordance with relevant data protection laws.

1.3.4 Observance of Case Study Research Ethics

Engaging in a case study, with voluntary participants, involved additional ethical responsibilities, such as:

- **Informed Consent**: Ensuring that all participants were informed about the nature of the dissertation project and had provided their voluntary consent to participate.
- **Participant Welfare**: Prioritizing the safety, well-being, and rights of all participants throughout the case study.

1.3.5 Statement on Previously Developed Work

As affirmed in the Statement of Integrity, the work presented in this document is original and authored by me, having not previously been used for any other end. However, an exception must be acknowledged: due to the formative structure of the master’s degree where this dissertation project is inserted, some sections were initially developed during a previous course unit in the first semester of the second year, namely “Dissertation Preparation” (PREPD). Specifically, the Introduction, State of the Art, and Requirements sections were originally written in a previous document in that context and have since been expanded and refined for inclusion in this final dissertation. The Introduction section also assimilated work written in an originally separate section “Project Planning” of this previous document.

1.3.6 Statement on the Use of Artificial Intelligence Tools

AI tools, and in particular Large Language Models (LLMs), have proven to be capable of helping software engineers in terms of code generation and software development. However, much of the content generated by LLMs contains errors and misinformation, and an overreliance on these tools also compromises integrity policies such as the commitment to producing original work, be it research findings or system implementations. As such, in this dissertation project, both the number of AI tools and their respective use were kept to a minimum necessary.

One of the uses of an AI tool during this dissertation, namely ChatGPT, was to obtain some preliminary knowledge about the AR development process in Unity, especially regarding initial system configuration, as some issues, especially those requiring adjustments in the Unity Editor menus rather than in code, were very specific and not well documented. Still, this information was always compared with available discussions in software forums such as “Stack Overflow”, or with official Unity AR courses, or tutorials made by independent creators. ChatGPT was also used to generate a base version of the icon part of the AleRa system’s logo, which was then manually edited to alter its colors, add transparency, and add the system name. Additionally, ChatGPT was used to generate simulated images of patients, to be used during the system’s and in the dissertation, to avoid ethical and privacy concerns related to using real patient data.

The other AI tool utilized during the implementation of the system was an online text-to-speech model, provided by ElevenLabs. This model was used to generate audio narrations for the authentication and tutorial steps, providing the users better usability and immersion.

1.4 Document Structure

This dissertation document was structured in eight main chapters, following the traditional process of software development, namely: Introduction, State of the Art, Requirements, Design, Solution, Case Study, Conclusions, and Limitations and Future Work.

This first chapter, Introduction, contextualizes this dissertation project, including how it derives from the Portuguese use case of the RM4Health ITEA project, and states the dissertation objectives. Additionally, it describes the domains of ethical concerns taken into account during the development of this dissertation.

The second chapter, State of the Art, explores the current state of existing work in this field from two different perspectives: on one hand, through a systematic literature review that gathers current literature work and analysis how AR / MR is being utilized in remote health monitoring; and on the other hand, through assessment of existing technologies that could be potentially employed in this dissertation project to developed the its proposed solution.

In the third chapter, Design, system requirements are identified based on the dissertation objectives and anticipated caregiver needs, followed by the detailing of the overall system design and all aspects taken into consideration during the AleRa system’s planning.

The fourth chapter, Solution, presents the resulting system of this dissertation project, and details the most important implementation aspects and decisions taken during development.

In the fifth chapter, Case Study, an empirical study done with a population of 16 participants is reported, where the developed system was tested in terms of usability and perceived workload (both physical and cognitive), and to obtain quantitative and qualitative feedback on how this system could be enhanced in the future.

Finally the sixth chapter, Conclusions, consolidates the conclusions derived from the development of this dissertation project, states the main achievements reached, including a list of scientific work contributed to literature, and reports on identified limitations during development and on possible future work to extend this system or similar ones, in the themes of AR and health monitoring.

2 State of the Art

A thorough exploration of the State of the Art was, as is the case for the vast majority of research and engineering work, one of the first and most fundamental stages of the development of this dissertation project. It was by reviewing what exists and what is being done in current literature that an analysis was able to be done to understand potential studies or functionalities that may be missing or underexplored. Then, by also surveying currently available and future trends in relevant hardware and software, a feasible project design was conceived. As such, the following sections present, respectively, a systematic literature review on the topic of augmented/mixed reality applied to health monitoring, and a survey of existing enabling technologies in this field, both made at the start of development of this dissertation project.

2.1 Systematic Literature Review

Remote health monitoring has become a widely researched topic in recent years, with increasing attention on the development of innovative, interactive, continuous, and non-intrusive solutions. Beyond the popularity of web and mobile applications, advancements in virtual world technology have led to a significant body of work integrating Virtual Reality (VR) into these healthcare settings, for example to provide patients with immersive, therapeutic experiences.

Despite this growing interest in VR, the integration of AR or MR into healthcare monitoring settings remains, to the best of my knowledge, a relatively underexplored area. Notably, there seems to be a shortage of solutions that utilize these technologies to assist practitioners or caregivers with patient monitoring functions, rather than targeting patients directly. This section presents a focused Systematic Literature Review (SLR) on this subject, consolidating existing work, that established a support base for the project conducted in this dissertation.

As defined by Kitchenham [9], an SLR, being a form of secondary study, is a means of consolidating and evaluating all available research related to a particular research question, topic area, or phenomenon of interest. Following Kitchenham's proposed methodology, a review protocol was developed and utilized by defining the research questions and search strategy, reviewing the articles, and then discussing the results obtained in relation to the research questions defined.

2.1.1 Research Questions

For this SLR, the main research question was defined as: **"What is the current state-of-the-art, perceived benefits, and identified limitations of using mixed or augmented reality to assist health professionals and caregivers in patient remote health monitoring?"**. This main question encompasses the global theme of the SLR and was later used to derive research query terms and inclusion and exclusion criteria for article search and screening.

From this main research question, six research sub questions were also derived, corresponding to the main information topics to be extracted and analyzed from the selected articles. These sub questions are presented in Table 1, and subsequently described.

Table 1 – Sub research questions

Identifier	Research question
RQ1	What are the medical areas where these augmented reality-based remote monitoring systems have been tested in?
RQ2	What methods are used to collect the data shown to the caregivers?
RQ3	What are the types of data shown to caregivers in these systems?
RQ4	Do these systems include active caregiver-interface interactions, or is the focus only on passive information display?
RQ5	What is the feedback of caregivers on the usability and feasibility of these systems?
RQ6	Has artificial intelligence ever been integrated with these systems?

The first question sought to map the application landscape of AR-based remote monitoring systems in healthcare, to try to understand and categorize their scope, effectiveness, and relevance in diverse healthcare scenarios. The second question explored the patient data collection methods employed in each study, such as wearable sensors or questionnaires, to identify how patient health information can be captured and processed. The third question focused on categorizing the information provided to health professionals or caregivers, and how they are presented in an AR setting, to evaluate the systems' comprehensiveness and alignment with practitioners' needs. The fourth question aimed to investigate whether these designed systems allow caregivers and practitioners to interact actively with the AR interface, or whether they consist solely of passive information displaying, which helps in assessing the level of engagement required and the implications for system functionality and usability. The fifth question gathered and evaluates user feedback on the usability and feasibility of AR-based systems, their satisfaction, perceived effectiveness, and challenges faced, providing critical insights into the practical adoption of AR in remote health monitoring. Finally, the sixth question tried to examine if, and how, AI has been integrated within AR-based remote monitoring solutions, to investigate how it can be used to enhance their capabilities or if this is an underexplored area in literature.

Answering these sub questions resulted in the extraction of relevant information to build a comprehensive review of the current literature work in the topic of using augmented and mixed reality in patient remote health monitoring.

2.1.2 Search Strategy

According to Kitchenham's methodology [9], the search strategy was divided into the definition of three key components: search sources, search terms (research query) and inclusion and exclusion criteria. The following sections describe each step of this process.

Selection of Search Sources

The first step of the search strategy was to define which search sources would be utilized to collect articles related to this SLR. For this work, searches were performed in four electronic databases, presented in Table 2 listing each one's respective Uniform Resource Locator (URL).

Table 2 – Search Sources.

Identifier	Source	URL
ED1	ACM Digital Library	https://dl.acm.org/
ED2	PubMed	https://pubmed.ncbi.nlm.nih.gov
ED3	Web of Science	https://login.webofknowledge.com/
ED4	IEEE Xplore	https://ieeexplore.ieee.org/Xplore/home.jsp

Definition of Search Terms

The definition of the search terms that build the SLR's research query is one, if not the most crucial step in a systematic review, since it materialized the research question. Due to the importance of this step, the "Population, Intervention, Comparation, Outcome, and Context" (PICOC) method [10], [11] was used to define search terms and build the complete research query. PICOC is a popular method used to describe a searchable question within the five elements that constitute its acronym. Table 3 presents, for each component, the research question part that corresponds to it, and the respective search terms derived. The final research query was, then, the joining of the search terms defined with the conjunction keyword "AND".

Table 3 – PICOC-derived search terms and consequent research query

Component	Derived from Research Question	Search terms
Population	Recent studies in patient remote health monitoring	("Remote Monitoring" OR "Digital Health" OR eHealth OR mHealth OR "mobile health" OR Telehealth) **AND**
Intervention	Using mixed / augmented reality to assist in patient health monitoring	("mixed reality" OR "augmented reality" OR ar OR xr) **AND**
Comparation[2]	-	-
Outcome	Perceived benefits, limitations, and viability of such systems	(benefit OR benefits OR limitation OR limitations OR experience OR experiences OR impact OR impacts OR viability OR feasibility) **AND**
Context	Target users are health professionals, other practitioners, formal or informal caregivers	(caregiver OR caregivers OR practitioner OR practitioners OR "health professional" OR "health professionals")

Definition of Inclusion and Exclusion Criteria

After the articles were gathered from the selected sources, they needed to be filtered using appropriate inclusion and exclusion criteria, to remove undesirable results (regarding the main research question defined initially). For this review, following convention, only studies that met all three inclusion criteria and did not meet any exclusion criteria were accepted during the screening phases. The inclusion and exclusion criteria defined for this SLR can be consulted in Table 4 and Table 5, respectively.

Table 4 – Inclusion criteria

Identifier	Inclusion criterion
IC1	The study explores how artificial / mixed reality systems can be integrated into patient health monitoring settings
IC2	The study reports results / feedback of tests performed with real users
IC3	The main target of the systems designed in the study are health professionals, other practitioners, or caregivers

---

[2] As this dissertation did not follow a comparative study design, the “Comparation” component was not utilized in this review.

Table 5 – Exclusion Criteria

Identifier	Exclusion Criterion
EC1	Sources not written in English
EC2	Sources published before 2020
EC3	Sources besides journal articles and conference proceedings
EC4	Sources that consist of systematic or other literature reviews, or otherwise not original research studies

2.1.3 Study Selection and Data Extraction

With both the research questions and search strategy defined, the four electronic databases defined as sources were queried using the defined search terms, and the results gathered, selected, and reviewed. To systematize this process, this SLR made use of the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) methodology [12], which according to the recent officially provided flow diagrams [13], divides the study selection process in three phases: Identification, Screening, and Included.

As can be seen in Figure 2, in the Identification phase a total of 877 results were retrieved from the electronic databases, with 2 of them being instead from manual reference tracking. Then, the results were compared by its Digital Object Identifier (DOI), in the case of conference proceedings and journal articles, or by its International Standard Book Number (ISBN), in case of books and book sections, and 17 duplicates were removed, leaving a total of 860 articles to be screened.

With the most recent flow diagrams, the PRISMA screening phase encompasses two steps that are often considered separate: Abstract Screening, and Full Text Eligibility Assessment (or simply, “Eligibility”). An important aspect to note is that, according to the PRISMA screening updated guidelines [14], the screening of articles in both the Abstract Screening and the Eligibility steps should be done by at least two independent reviewers. However, due to the nature of this master dissertation work, these screenings had to be carried out by a single reviewer, which may have introduced higher bias on study selection. To try to mitigate this, the inclusion and exclusion criteria were always applied in their strictest interpretation.

In the Abstract Screening step, Microsoft Excel was first used to automatically filter out results that met exclusion criteria 2 and 3, that is, filter out results published before 2020 and results besides journal articles and conference proceedings (namely, books and book sections). This automatic process removed 552 articles. The remaining 308 articles were then manually screened by applying the inclusion and exclusion criteria to each title and abstract. From this manual process, an additional 279 articles were screened out, totalling 831 results removed.

The remaining 29 articles were sought for retrieval to assess their eligibility, however 3 of them could not be retrieved due to being closed off behind a paywall. The 26 articles successfully retrieved were then screened again, this time assessing eligibility through a full body reading.

Of these, 10 articles were rejected due to not being related to the theme of exploring AR / MR in health monitoring settings, 4 were rejected by not reporting tests or results of the designed solutions, another 4 were excluded due to their systems' targets not being health professionals or caregivers, and 2 were rejected by their nature of being themselves literature review articles, which according to the Cochrane Handbook [15] should not be included since a systematic review should focus on primary studies to avoid duplicating evidence synthesis.

Finally, the 6 remaining studies were deemed eligible, and were as such included for data extraction and analysis. This entire process flow was, again, synthetized in a PRISMA 2020 flow chart, shown in Figure 2.

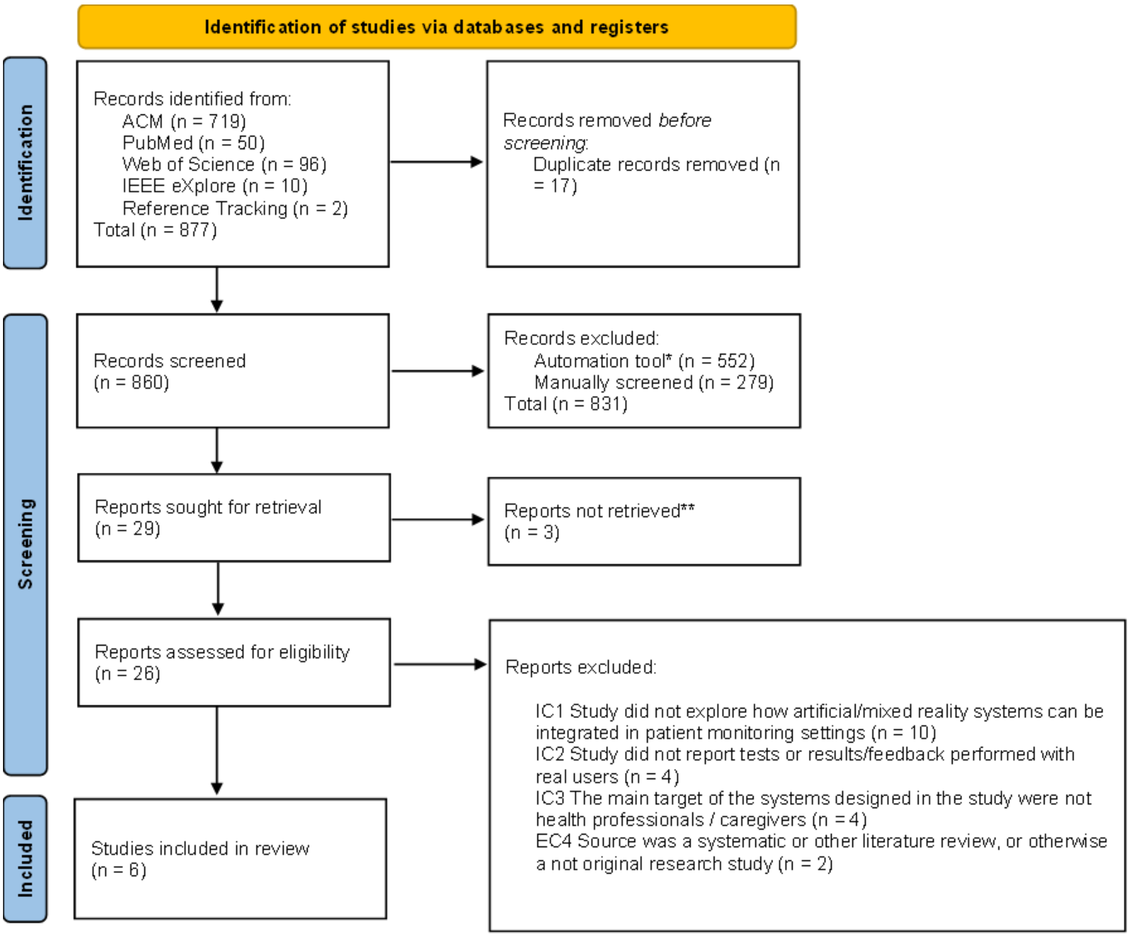

*Microsoft Excel table filters were used as an automation tool, to exclude records according to Exclusion Criteria 2 and 3

**Reports not retrieved consist of articles which have a public abstract, accepted in the previous review stage, but do not present an accessible full article body, and as such cannot be assessed for eligibility.

Figure 2 – PRISMA 2020 flow chart, applied to this SLR.

2.1.4 Results and Discussion

In this section, the information extracted from the 6 reviewed papers, present in Table 6, was consolidated and used to answer the previously defined sub research questions. For each question, a first paragraph reports objective results observed, related to that question, and the second paragraph consists of a brief subjective commentary on those results.

Table 6 – Final papers included in the review

Identifier	Authors	Year	Title
P1	Kimmel et al.	2021	opticARe - Augmented Reality Mobile Patient Monitoring in Intensive Care Units [16]
P2	Vock et al.	2021	IDIAR: Augmented Reality Dashboards to Supervise Mobile Intervention Studies [17]
P3	Lorenz et al.	2023	Creating Routes for Landmark-Training with the CompanionApp: A Pilot User Study [18]
P4	Albright et al.	2024	Opportunities and Challenges for Augmented Reality in Family Caregiving: Qualitative Video Elicitation Study [19]
P5	Tanbeer e Sykes	2024	MiVitals-Mixed Reality Interface for Vitals Monitoring: A HoloLens based prototype for healthcare practices [20]
P6	Kim et al.	2023	The Multidomain Metaverse Cancer Care Digital Platform: Development and Usability Study [21]

RQ1 What are the medical areas where these augmented reality-based remote monitoring systems have been tested in?

The reviewed studies show that these AR-based remote monitoring systems are starting to be experimented and employed in various different medical areas. For example, the study by Albright et al. [19] implements an AR tool to assist formal and informal caregivers to monitor and care for cancer patients, the study by Vock et al. [17] utilizes AR to allow domain experts such as health psychologists to supervise patient physical activity, and the study by Kimmel et al. [16] employs AR to help nurses monitor patients in critical care settings, namely in Intensive Care Units (ICUs).

This high variety of medical areas where AR-based remote monitoring has been employed shows the potential that AR brings to helping caregivers and health professionals in patient health monitoring. However, since few studies yet exist, and they are applied to different medical areas, it is difficult to extract more exact conclusions about the effectiveness of AR in each one. As such, there is space in literature for more AR-based health monitoring systems to emerge.

RQ2 What methods are used to collect the data shown to the caregivers?

Of the 6 studies reviewed, 2 articles, by Tanbeer e Sykes [20] and Vock et al. [17] utilize wearable sensors as the main method of collecting patient data, while the one by Kimmel et al. [16] utilizes the traditional sensors of the ICUs. The other studies report using images and medical records as their main source of patient health information.

This variety of data collection methods is no doubt related to the fact the studies have been applied to different medical areas, as seen in RQ1, so no exact conclusions can be made on the effectiveness of each method. Nonetheless, this variety shows that AR-based health monitoring systems are capable of integrating with a large number of systems, giving more opportunities for exploring new applications in literature.

RQ3 What are the types of data shown to caregivers in these systems?

Of the 6 studies reviewed, 2 articles, by Tanbeer e Sykes [20] and Kimmel et al. [16], although using different methods, both focus on collecting and showing patient vital signs data, such as heart rate, respiratory rate, temperature, oxygen saturation, and blood pressure, to the caregivers, with the latter also providing alerts for problematic values. On the other hand, the study by Vock et al. [17] instead collects and displays physical activity data such as patient step count and activity level, while the study by Kim et al. [21] focuses on creating an environment where multiple caregivers can be shown a patient's various medical records and health information simultaneously.

Once again, this variety of types of data collected and shown are a reflection of the different medical areas where these studies have employed AR-based remote monitoring systems. Still, again, this variety of data shows the potential for larger, more complex systems that integrate all types of data, for example in a similar way to the experience provided by the system developed by Kim et al. [21] but with integrated real time sensor data.

RQ4 Do these systems include active caregiver-interface interactions, or is the focus only on passive information display?

4 of the 6 articles reviewed not only display patient health information passively, as an overlay screen only, but also support active interaction with said overlays. For example, the study by Tanbeer e Sykes [20] allows clinicians to interact with holographic displays to zoom, rotate, and focus on specific data, while the study by Vock et al. [17] incorporates multimodal interaction through three options: head-mounted displays, smartphones, and voice commands. On the other hand, the studies by Kimmel et al. [16] and Lorenz et al. [18] do not support direct interaction with the AR overlays, focusing primarily on passive health information visualization.

The majority of studies integrating active caregiver-interface interactions shows how AR systems have the potential to be integrated into more complex systems, as more than just an alternative display, while not diminishing their usability as passive health information visualization devices.

RQ5 What is the feedback of caregivers on the usability and feasibility of these systems?

In general, all 6 studies reported having obtained positive feedback on the usability of AR systems in health monitoring. Namely, the participants of the study by Tanbeer e Sykes [20] rated an average of 84 on the System Usability Scale (SUS), indicating high usability among clinicians, and the domain experts in the study by Vock et al. appreciated the multimodal interaction approach and found it effective for monitoring studies [17]. However, 2 studies in particular also revealed negative or important aspects to consider when developing these AR systems. When the study by Kimmel et al. [16] compared the use of three sizes of AR overlays, large, small, and context-dependent, the nurse participants largely preferred the context-dependent one, and indicated that the large one was the worst option, since it hindered their vision and therefore their work capabilities. Meanwhile, the article by Albright et al. [19] studied how, although beneficial for caregivers, the use of AR may also instead aggravate their burdens, since the caregivers may require time and training to learn how to use the AR systems, and if the system gives incorrect information, their tasks will worsen.

As such, the results to this research question highlight a crucial factor to keep in mind when developing AR systems for health monitoring, which is that ultimately the user experience has to be the priority, and great care must be used during development so that the systems do not overwhelm the end-users. As such, this indicates that there is space in literature for deeper studies, following Albright et al.'s [19] model, on how these AR systems may unintentionally disrupt caregiver activities instead of helping, and perhaps the creation of AR medical interface standards, so that health professionals would not need to learn new technologies every time.

RQ6 Has artificial intelligence ever been integrated with these systems?

Of the 6 studies reviewed, only the study by Kim et al. [21] reports having artificial intelligence integrated into their AR multipurpose system, being used to for data analysis and supporting decisions in cancer care. The third elicitation study in the article by Albright et al. [19] describes that the idea is to use AR graphic overlays of the patient's body position, achievable by body tracking artificial intelligence models, but that in their study, for simplicity, they used video editing instead to create the overlays.

This shows that, currently, AR systems are still only being used in simpler use cases, where no artificial intelligence is employed, which highlights an opportunity in literature to explore the integration of these two knowledge areas.

2.2 Existing Technology Assessment

A literature review is widely recognized as an effective way of researching and extracting information from current literature related to a topic, in this case the use of augmented / mixed reality for remote health monitoring. However, as an engineering project, more direct insights into the currently available hardware and software in the market were required, to allow designing and developing an operational and feasible system.

As such, this section presents the assessment done on existing technologies that could be helpful for the development of this dissertation project. At this stage this research focused on the possible hardware devices and primary development tools to be used for the development of an HMD AR system, as more specific technologies such as libraries or dependencies could only be explored after project requirements were identified and a system design was established.

To ground the system in realistic development options, the first step was to survey currently available AR HMDs, and identify the software development kits (SDKs) that each one provided or supported. The results of this assessment are summarized in Table 7, which lists major HMD devices, their operating systems, and their associated SDKs, together with universal tool compatibility such as OpenXR or WebXR[3].

Table 7 – AR / MR compatible headset hardware

Headset	**Operating System**	**Primary SDK/Platform**	**Universal Tools**	**References**
Meta Quest 3	Android-based	Meta Spatial SDK	WebXR; OpenXR	[22]
HTC Vive	Windows	Vive Wave SDK	OpenXR	[23]
Microsoft HoloLens	Windows	Visual Studio with Windows 10 SDK	WebXR; OpenXR	[24]–[26]
PICO 4	Android-based	PICO Integration SDKs	WebXR; OpenXR	[27]
Magic Leap 2	Android-based	Unity SDK Unreal Engine SDK	WebXR; OpenXR	[28]–[30]

Next, a search for AR development engines and environments was carried out to determine existing possibilities, and which SDKs they natively support. This analysis, shown in Table 8, includes established VR/AR game engines such as Unity and Unreal, lighter game engines such as Godot, as well as web-based environments and Android Studio.

---

[3] The acronym present on these tools' names stands for Extended Reality (XR), an umbrella term for VR, AR and MR technologies.

Table 8 – Augmented / Mixed Reality Development Platforms

Platform / Tool	Type	Supported SDKs	References
Unity	Game Engine	OpenXR; OpenXR: Meta	[31]
Unreal Engine	Game Engine	OpenXR; Oculus; SteamVR	[32]
Godot	Game Engine	OpenXR	[33]
Web Frameworks	Web Frameworks (e.g., Angular)	WebXR	[34]
Android Studio	Development IDE	Jetpack XR SDK; OpenXR; WebXR	[35]

Finally, the findings from both tables were combined into a compatibility matrix, presented in Table 9. This cross-mapping illustrates which platforms can be used with which HMDs, and through which SDKs. The resulting overview made it possible to identify the feasible development pathways for the AleRa system, balancing portability, functionality, and maintainability.

Table 9 – Compatibility table between AR HMDs and development platforms.

	Unity	Unreal Engine	Godot	Web Frameworks	Android Studio
Meta Quest 3	OpenXR; OpenXR: Meta	OpenXR	OpenXR	WebXR	OpenXR; WebXR
HTC Vive	OpenXR,	OpenXR	OpenXR	-	OpenXR
Microsoft HoloLens	OpenXR	OpenXR	OpenXR	WebXR	OpenXR; WebXR
PICO 4	OpenXR; PICO Unity Integration SDK	OpenXR; PICO Unreal Integration SDK	OpenXR	WebXR	OpenXR; WebXR
Magic Leap	OpenXR; Unity SDK	OpenXR; Unreal SDK	OpenXR	WebXR	OpenXR; WebXR

From this assessment it was possible to observe that OpenXR is a well-established SDK, supported by all examined AR HMDs and development platforms. OpenXR is an open standard that provides a common set of APIs for developing AR / VR applications across a wide range of devices [36]. This made it the main candidate for the SDK to use during the system's implementation.

Regarding the development platform, while most support OpenXR, Unity and Unreal Engine have the additional advantage of also supporting some vendor specific SDKs. An objective of this dissertation project was to develop a system compatible with various devices, so OpenXR was already predicted to be the chosen SDK at this stage. However, these vendor-specific SDKs often provide additional development features, which could boost performance or user immersion. As such, it was wise to opt for a development platform that supported these additional SDKs so that they would be available, if necessary, in future work extending the system developed in this dissertation. Following this reasoning, Unity was slightly favored at this stage due to supporting an additional SDK for the Meta Quest 3.

As for the AR HMD, the ones examined offer similar features, all being compatible with OpenXR, so the decision on which to use during development and testing only depended on the available equipment at GECAD, where the work for this dissertation was conducted.

2.3 Summary of Findings

From the reviewed literature, it was observed that there have been previous studies into the use of AR in different patient monitoring contexts, from oncology to intensive care. These systems demonstrate the versatility of AR for displaying a variety of information, such as vital signs, medical records, or physical activity, and they often include some level of interactive functionality. Caregivers generally responded positively to AR systems, but also raised important concerns about potential cognitive overload, training requirements, and the risks of inaccurate information. Moreover, integration of AI into these systems is still rare, highlighting a research opportunity for future AR solutions.

Meanwhile, the technology assessment showed that the AR ecosystem is currently diverse, both in terms of existing HMD hardware and development software, but that it converges through intercompatibility. All major headsets and development platforms were found to support the OpenXR standard, which, as such, emerges as the most sustainable option for cross-device development. Unity and Unreal Engine stood out as development platforms for their maturity and breadth of support, with Unity being slightly favored due to additional compatibility. Since the devices themselves offer similar core capabilities, the final hardware choice was determined according to availability within GECAD’s laboratories.

Together, these findings provided a dual perspective: the literature highlighted the opportunities and pitfalls of applying AR to health monitoring, while the technology assessment clarified the existing development options. This state of the art synthesis then informed design requirements and decisions for the AleRa system during the next development stage, namely the adoption of Unity with OpenXR to ensure interoperability, the focus on usability and intuitive interface design, and the possibility of prototyping AI functionalities, while trying to ensure extensibility for future work.

3 Design

With the main state of the art review completed, the development process of the AleRa system could begin. As is common in software engineering, the first stage involved a careful analysis of the project's objectives and current needs of the target fields, namely augmented reality and remote health monitoring, to define concrete system requirements that should be reached to guarantee a successful implementation. These requirements formed the initial step of the design process, serving as the foundation for the system's architecture and operation logic. This design process was then essential not only to explore alternatives, but also to establish the core decisions that would guide implementation.

To try to ensure maintainability and extensibility, well-known software engineering principles were considered during this stage, including the General Responsibility Assignment Software Patterns (GRASP), the "Single Responsibility, Open/Closed, Liskov Substitution, Interface Segregation, and Dependency Inversion" (SOLID) design principles, and broader concepts such as modularity.

The following sections present the requirements defined for the AleRa system, the overall system architecture, and the key design considerations taken into account prior to entering the implementation stage.

3.1 Requirements

The core objectives of this dissertation project were to develop an AR interface system that allows health professionals or caregivers to monitor health parameters and receive alerts for patients under their care, and to integrate this system with the RM4Health project's architecture (and other similar architectures) for information feeding, which translated into obtaining data from an existing HL7 FHIR compliant server. Thus, some essential requirements identified were the displaying of patient health parameters (e.g.: vital signs data, medication intake history), the emission of alerts tied to dangerous detected conditions, and the fetching of data from a remote backend FHIR server.

In addition to the ones directly derived from the dissertation objectives, other requirements were identified as fundamental or useful for this system.

One such requirement is user authentication, since appropriate authorization is required for the system to access, and display to caregivers, sensible medical data from a FHIR-compliant server. The most common approach to authentication is the use of a username and password combination, and it is the method used by RM4Health’s authentication server. However, virtual keyboard-based text input, with a traditional QWERTY format, is reportedly inaccurate and time-consuming in head-mounted AR environments [37]. For this dissertation project, given the need to protect sensitive medical records while also reduce, rather than increase, caregiver workload, a login solution that was both secure and easy to use had to be implemented.

Another useful functionality detected was patient context awareness, that is, knowing when to present each patient’s information to the caregiver. While a completely passive system was feasible, it would force caregivers to navigate through menus to select and visualize a specific patient’s data. Since one objective of this AR system is to ease caregivers’ burden when monitoring patients, it was deemed useful for the system to recognize which patient is being interacted with and show the appropriate information, in other words, a quick patient selection feature.

Finally, again given the novelty of using an AR interface, and considering that the system could eventually be deployed in real-world scenarios, a requirement for an interactive tutorial was also established. This would allow users to be trained directly within the AR environment, without needing to rely on an assisting team.

These requirements were synthesized in Table 10, organized according to the order in which users would first encounter or use them, and their completion status was tracked throughout the dissertation project’s development to ensure that the final product aligned with the dissertation’s objectives.

Table 10 – System requirements identified and respective completion status.

ID	Requirement	Status
REQ1	The system should allow caregivers to authenticate in a secure and easy to use way.	Completed
REQ2	The system should provide an interactive tutorial to teach caregivers about its functionalities.	Completed
REQ3	The system should allow to quickly change which patient is selected and adapt to show their respective information.	Completed
REQ4	The system should be capable of reading patient data from a FHIR-compliant backend server.	Completed
REQ5	The system should be able to display patient vital sign data.	Completed
REQ6	The system should emit alerts when dangerous conditions are detected in the patient's vital signs.	Completed
REQ7	The system should be able to display medication intake data.	Completed

3.2 Development Environment

An important first step was to decide which development environment and tools to use since the AR development tools available depend strongly on both the selected engine and target HMD. As this system aims to assist a variety of caregivers in patient monitoring, it was important to develop it in a device-independent way. As noted in the assessment of existing technologies (see section 2.2), HMDs usually provide their own proprietary SDKs, but many also support open standards such as OpenXR and WebXR. As such, to maximize portability across platforms, it was decided that development would rely exclusively on OpenXR tools. The trade-off is that OpenXR, due to its universal nature, only provides basic, generic functionality, requiring developers to assemble more complex features manually. Vendor-specific SDKs often provide enhanced tools, but their use would restrict deployment to a single device ecosystem.

From the engines that support OpenXR, two stand out: Unity and Unreal Engine. Both are mature and widely adopted platforms for AR/VR development, and both allow vendor-specific SDK features, if needed in future extensions of this system. Since their functionality is largely comparable, Unity was chosen for this dissertation project due to existing prior experience with the platform, as well as due to the high availability of tutorials, community support, and structured AR courses.

3.3 Companion Mobile App

Another very important system design choice arose from limitations imposed by most HMD vendors. Although these devices use cameras to detect surrounding structures, obstacles, and, in some cases, enable AR passthrough functionality, in most devices developers are denied direct access to raw camera footage, for privacy reasons. Given that one of the identified requirements was for the system to include some context awareness capability, it was determined that a companion mobile application would be developed to complement the main HMD application, providing access to camera-based features that the HMD cannot support directly, and that the two applications would be capable of communicating with each other to ensuring real-time synchronization of information as required. This idea of a companion mobile app that communicates with the HMD was also inspired by the work done by Vock et al. [17], analysed during the literature review stage.

3.4 General Architecture

The development logic in Unity is different to the full-stack structure of a typical frontend–backend system. In a Unity app, all user interactions happen within one or more Scenes, which are somewhat analogous to views in conventional applications, and within each scene, objects represent interactive elements, which in the case of the AleRa system are mostly user interface (UI) elements. However, each object can carry its own scripts, which define their behavior, and can control other objects both from within its own script or by calling the other objects script functions. It is also common to include dedicated objects that manage global aspects of the scene. This creates a very flexible development environment, which while powerful, can easily lead to architectural inconsistencies.

As such, during this design stage, principles such as Single Responsibility, Modularity, and Information Expert were followed to try to ensure maintainability and clarity. Based on these principles, a general architecture was outlined, approximating a Component Diagram (third level of the C4 Model), as presented in Figure 3. For clarity, Scenes and global scene managers were omitted from the diagram, as their implementation is context-dependent and not fixed during the design stage. Nevertheless, certain core scene manager components were anticipated from the start, including an Audio Component, an AR Environment Manager, and a Tutorial Component.

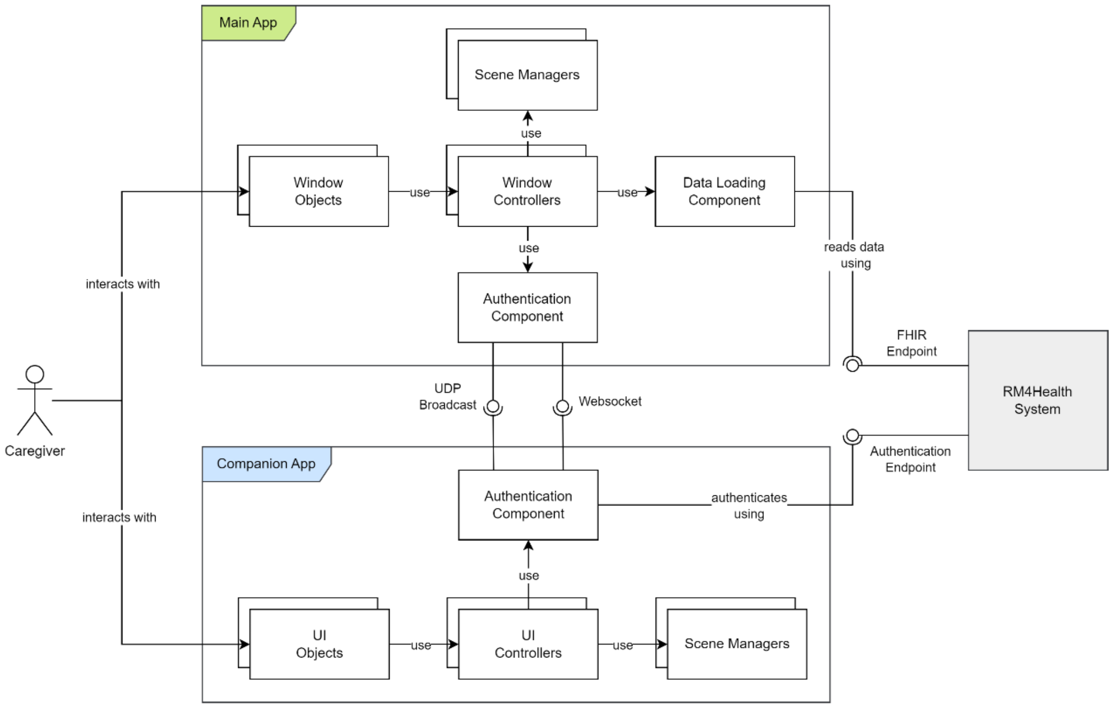

Figure 3 – General architecture of the AleRa system.

3.5 Interface Considerations

In terms of interface design, it was anticipated that the novelty of using an AR system would place users outside their comfort zone and require a learning process. To minimize this burden on caregivers and avoid further alienation, it was established that the system's interfaces should be designed to feel familiar, drawing inspiration from technologies caregivers already use. As such, a window-based interface was chosen to ease the transition into the AR environment. This decision was further solidified by the previous successes with window-based interfaces observed in the literature review stage, namely in the works by Kimmel et al. [16], Vock et al. [17], Albright et al. [19] and by Tanbeer and Sykes [38], who all employed them for different functions, from charts to image viewing

Taking inspiration from SOLID's Single Responsibility principle, a guiding principle when designing the AleRa system was modularity. All features and windows were developed as independent components, allowing them to serve as building blocks for constructing tailored interfaces for different scenarios as needed.

Leveraging the three-dimensional (3D) nature of AR environments, two types of windows were envisioned to organize information and intuitively convey possible interactions to the user. In a near plane, within arm's reach, "near" windows were planned for detailed content and direct interaction, such as scrolling, tab navigation, or opening sub windows (e.g., for medication intake information). On a further plane, "far" windows were planned for information requiring prolonged monitoring, such as real-time vital signs. The envisioned window positioning design was represented in Figure 4, from the user's point of view. Finally, while "anchoring", that is, fixing virtual objects to real-world coordinates, is a common AR feature, it was considered unsuitable for the AleRa system. Since caregivers are expected to move freely while monitoring patients, fixed windows would likely disrupt usability, even if they could be manually repositioned. Instead, the planned solution was for windows to follow the user's view automatically.

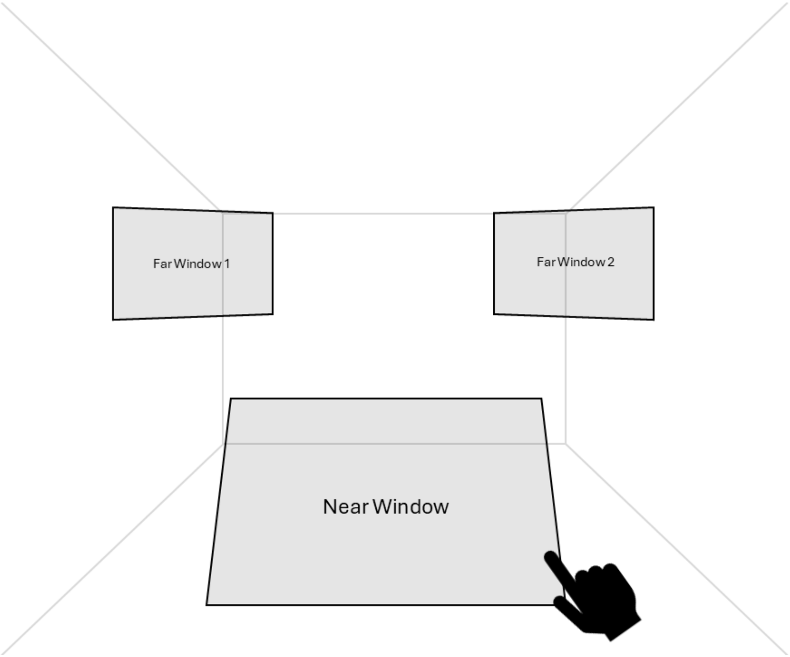

Figure 4 – Envisioned window-based interfaces, from the point-of-view of the user.

3.6 Authentication and Device Pairing

As noted during the requirements definition (see section 3.1), a username and password combination is the most common approach for authentication, and it is the method used by RM4Health's authentication server. However, text input using a traditional keyboard layout is rather challenging in HMDs, a limitation widely recognized in the literature [37], [39], [40]. Various alternative methods have been explored, from gesture-based to voice-driven input, but no clear standard has emerged. For instance, Derby et al. [37] reviewed how users have been found to type up to 32.5 words per minute (wpm) in a traditional QWERTY cell phone keyboard, but only 6.5 wpm in the default virtual keyboard used in HoloLens, a well-established AR HMD, and only 6 wpm with their alternative solution.

Rather than risking alienating users with an unfamiliar method, it was decided that the AleRa system would delegate text input, and consequently user authentication, to the companion mobile app. This decision was also inspired by the work carried out by Vock et al. [17], analysed during the literature review stage, who employed a synchronized smartphone app that adapted its UI based on selections made in their main HMD app, allow different types of user input methods such as text, buttons, and voice.

In RM4Health's architecture, the FHIR-compliant data storage server delivers medical data in JavaScript Object Notation (JSON) format through a REST API. To access this data, the main HMD application must present a valid JSON Web Token (JWT) for authorization, generated by an authentication server when the user logs in. However, according to the previous decision, the mobile app became the one responsible for authentication. As such, a sequence flow had to be designed to securely connect both applications and allow the transferring of this token to the HMD application. Since the mechanism of synchronization chosen could strongly affect the design of this functionality, technology alternatives were explored at this stage, and Websockets were chosen due to their security and low latency.

Since WebSocket communication requires knowledge of the peer device's internet protocol (IP) address, a discovery mechanism was designed: the HMD broadcasts its name and IP address over the User Datagram Protocol (UDP), which the mobile app listens for. When the caregiver authenticates against the existing authorization server using the mobile app, this app receives a JWT bearer token. The app then presents to the caregiver a list of discovered HMDs, allowing the user to select the target device and initiate a WebSocket connection.

Furthermore, for added security, an additional mechanism was planned: After the Websocket connection between devices is established, the HMD app generates a confirmation personal identification number (PIN) that is displayed to the user. The mobile app prompts the user to enter this PIN, and the attempt is verified by the HMD via the WebSocket channel. Only if the PINs match does the HMD application confirm pairing, after which the mobile app securely transmits the JWT token. At this point, both applications can proceed with their respective workflows, and the HMD is now in possession of the necessary JWT bearer token, being authorized to access the FHIR server. The overall logic designed for this process was documented through a sequence diagram, presented in Figure 5.

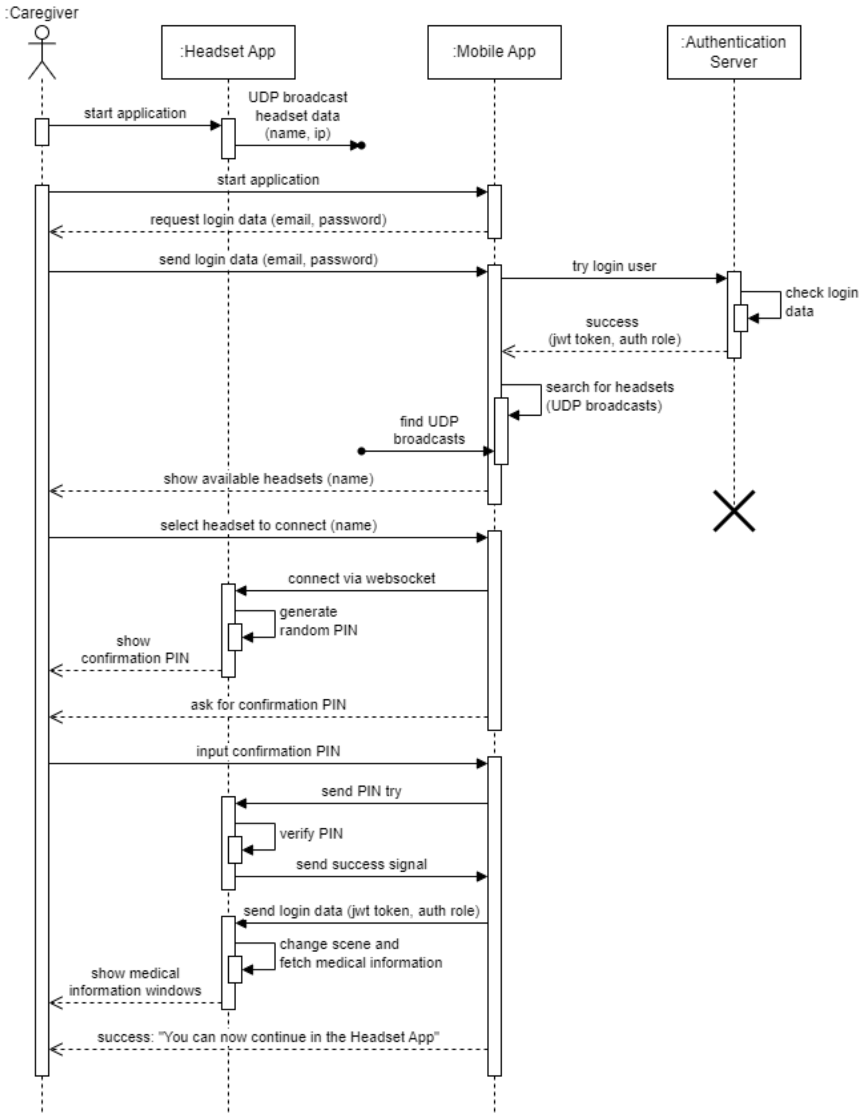

Figure 5 – Sequence diagram designed for the Authentication and Device Pairing functionalities.

Text input alternatives considered

The adopted authentication and device pairing solution is admittedly complex, and stems from the difficulties involving text input in HMDs. Before committing to this approach, several alternatives text input solutions were explored, but were ultimately discarded prior to implementation due to identified limitations. Table 11 consolidates these alternative solutions, ruled out prior to implementation because their design was anticipated to either be impractical for users, difficult to implement effectively, or currently infeasible with available technology. Although these attempts did not advance beyond the design stage, they were documented here as future research may overcome the present limitations and render these solutions more viable.

Table 11 – Text input alternatives considered.

Solution	Description	Advantage(s)	Limiting Factor(s)
Speech-To-Text input	Leverage the use of voice recognition and speech-to-text technology to input spoken words as text	(i) No dependence on a mobile app; (ii) Handless text input; (iii) Quickest text input capability	(i) Accuracy of voice recognition AI models, affected by noisy environments; (ii) Only dictionary words can be recognized (limits random letter passwords); (iii) Sensitive data (such as passwords) have to be said out loud
Virtual Writing Input	Allow users to input text by writing cursive text on the air	(i) No dependence on a mobile app; (ii) Intuitive text input method, mimicking real life handwriting	(i) Accuracy of cursive character recognition AI models, which are underexplored compared to machine writing models; (ii) Text recognition models usually approximate to the closest dictionary word, which limits random letter passwords
Phone companion app (Bluetooth connection)	Utilize Bluetooth for the connection and data transfer with the phone companion app, instead of a Websocket	Higher connection security since Bluetooth operates only on nearby devices	Dependence on a third-party library to use Bluetooth in Unity, which could compromise data security

3.7 Patient Selection (Context Awareness)

According to the requirements identified, the system should allow the caregiver to quickly change which patient is selected and adapt to show their respective information, bringing context awareness to the main HMD application.

During this design stage, several alternative mechanisms for this feature were considered, namely Facial Recognition, Quick Response (QR) code reading, and Radio-Frequency Identification (RFID) tag scanning. Facial recognition using the HMD cameras would provide the most seamless experience. However, these approaches have several limitations:

1. Due to privacy concerns, direct access to HMD cameras is restricted by most vendors. For example, Meta only recently allowed developer access to its devices cameras through the Passthrough Camera API [41], a change that occurred after the development of the present system had already started. This hindered the potential use of facial recognition and QR code reading, which would have to be delegated to the companion mobile app.
2. Because of these privacy and ethical concerns, regulations such as the European Union (EU) AI Act strongly scrutinize and restrict the use of AI models. Article 2(6) of Regulation (EU) 2024/1689 [42] explicitly exempts the use of AI models and its outputs for the sole purpose of scientific research and development activities from strict regulation, which made facial recognition prototyping possible (see section 4.2.9), but Article 2(8) clarifies that "testing in real world conditions shall not be covered by that exclusion", so testing it in a case study was not allowed.
3. RFID tags remove the limitations derived from utilizing cameras, and most modern phones can read them, they require the closest proximity to the patient of all alternatives, which may inconvenience the caregiver, and it may not yet be universally supported across devices.

Taking into consideration these restrictions, the final design for this feature consisted in using QR code reading through the mobile app's phone cameras, leveraging the previously designed feature that allows both applications to communicate via WebSockets.

4 Solution

With the dissertation project requirements identified and the AleRa system's base design established, the next step, and the longest in duration, was the development of the planned solution, comprising a main HMD AR application, and a companion mobile application. This section describes the various software and tools used during the full development process, and presents the resulting system components, highlighting their final appearance and detailing the main decisions and elements used to create each one.

4.1 Software and Tools

As designed, the main AleRa application was developed using the Unity game engine (Unity Technologies, version 2022.3.38f1) [43], where the main scripting tool is the C Sharp (C#) programming language. The implementation of the AR-related features relying exclusively on OpenXR packages to try to ensure compatibility with multiple HMDs (e.g., Meta Quest, Microsoft HoloLens, Apple Vision Pro). By prioritizing OpenXR over vendor-specific SDKs, AleRa was designed to maximize portability and reach a broader user base.

In addition to the main HMD application, a companion mobile app was also created within the same Unity environment. As decided during the design stage, this secondary app provides functionalities that were either difficult or impractical to handle directly through the HMD, namely user authentication and QR code scanning for patient context switching.

For development and testing of the main HMD app, the GECAD research group provided access to a Meta Quest 3. The Quest 3 weighs 515 grams, is powered by a “Qualcomm Snapdragon XR2 Gen 2” chipset, with 8 GB of Random Access Memory (RAM). If features dual color liquid crystal display (LCD) panels with a resolution of 2064 × 2208 pixels per eye, supports refresh rates up to 120 Hz, and has a field of view of 110° horizontal and 96° vertical [44]. It also includes front-facing color cameras that it utilizes to provide passthrough AR functionality, that is, to display the real world with virtual elements superimposed.

To setup development for the HMD application, the first steps of Unity’s “VR Development Pathway” course were followed. Since the Quest 3 runs on Android, the “Android Build Support” module was installed, enabling Unity to compile Android Package Kit (APK) files. Developer mode was enabled on the device, and it was connected to Windows via the Meta Quest Link application, allowing Unity to build and deploy applications directly to the HMD through a link cable.

For the application’s Unity project itself, the Universal 3D template, that uses the Universal Render Pipeline (URP) was selected. URP was chosen because it offers optimized rendering and flexibility across devices, making it useful for AR development. Then, the necessary packages were installed using Unity’s Packet Manager, namely AR Foundation, OpenXR, and XR Interaction Toolkit. The latter provides a ready-made XR Rig object, that maps real-world head and body movement to the virtual scene, along with providing hand-tracking interaction methods integrated through the Unity’s Input System.

From the interaction methods provided by the XR Interaction Toolkit, two were adopted for the AleRa system:

- Direct touch, where the user selects UI elements by tapping directly with the index finger.
- Ray-based pointing, where a ray is projected between the thumb and index finger toward the desired window, with selection confirmed via a pinch gesture.

With these essential capabilities provided by Unity and its relevant packages, this dissertation’s development could focus on AleRa’s interfaces and caregiver-centered functionalities, rather than low-level AR feature implementation.

To maximize compatibility, the companion mobile app was developed using the same Unity version and again starting from the Universal 3D template with URP. Since no AR features were required for this application, such packages were not included in this secondary app. During this dissertation project, only an Android build was produced and tested, as iOS builds in Unity require a MacOS environment [45]. Nonetheless, the project structure should be compatible with future iOS builds.

4.2 System Components

With the base Unity project configured, development proceeded to the implementation of the individual system components. Each component underwent several refinement iterations, exploring different information presentation approaches and interaction mechanisms. These feature implementations relied on functionality already provided by Unity’s built-in UI system (e.g.: buttons, scrollbars, and text), base AR and interaction mechanisms provided by OpenXR and XR Interaction Toolkit, and a few other selected external packages for additional functionality.

The following subsections describe in detail the technical design and implementation of individual components of both the HMD and mobile applications. Most of these features were later tested through a case study, but this section also documents a few prototype components that were not finalized for testing due to different reasons, such as ethical concerns or time constraints, but still represent valuable feature exploration which could be further refined in future work.

As a note, since all expected participants in the case study were native Portuguese speakers, the interfaces were developed in Portuguese. This is reflected in the screenshot figures presented throughout this section.

4.2.1 Window-based Interfaces

As designed, to maintain familiarity for caregivers, the AleRa system was developed with a window-based UI, composed of three-dimensionally placed panels within the augmented reality environment. Each window was implemented as an independent object, managing its own content, following the information expert and single responsibility principles, contributing to a modular and extensible architecture suitable for future enhancements.

Windows

The windows themselves were created using Unity’s base UI components, each being primarily composed of a Canvas in “World Space” render mode, to behave as a 3D placed object in the real world, with various child elements such as buttons, text, and layout defining components. Transparency was applied to avoid obstructing the user’s view of the real world, and a black background color was selected to maintain sufficient contrast and readability under different lighting settings. Furthermore, by joining the capabilities of the Canvas Group component, which allows to control the transparency of all its child objects, with the external library LeanTween, which simplifies the linear interpolation of this property’s value, a window state system was implemented, allowing each window to alternate between multiple sets of content with a fading animation. An example of the general appearance of an AleRa window, placed in AR space, is shown in Figure 6.

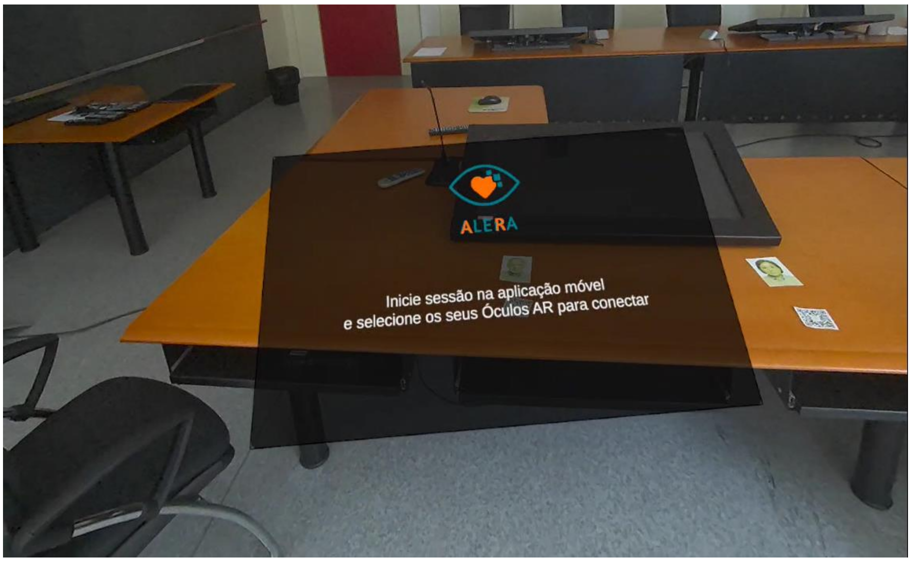

Figure 6 – General appearance of a window in the AleRa system.

Window Positioning

Once a window base was defined, a mechanism to position them relative to the user was required. Taking inspiration from the Quest 3's operating system UI and from the XRSocketInteractor functionality provided by XR Interaction Toolkit, both of which allow to snap different objects to the same predefined place, named a socket by the latter, a flexible positioning method for the AleRa's windows were developed. Instead of binding specific windows to fixed coordinates directly, a grid of socket positions was implemented, to which any window could be dynamically parented or unparented at runtime through scripts (similar to manipulating elements in a HyperText Markup Language (HTML) document via JavaScript). This separation of concerns made it possible to manage layout logic independently from window content in the Unity object hierarchy (see Figure 7.a).

Since some windows were intended to show constant information to the user, it was necessary for them to accompany the caregiver as they moved around. During early prototyping, two approaches were explored:

- Spatially fixed windows, anchored in world space, only manually repositionable by the user. In Unity this translated to the socket grid being a different object in the scene, not parented to the XR Rig camera.
- Head-locked windows, directly following the user's movements without delay. This was achieved by parenting the socket grid object to the XR camera in the scene hierarchy, which made the grid automatically copy all camera rotations and translations, which in turn mirrored the caregivers head movements.

The first approach proved cumbersome for prolonged use, particularly when moving between different locations, while the second resulted in an unnatural and distracting experience. As a better solution, a smooth follow mechanism was implemented: the socket grid was not parented to the XR camera, but instead controlled by a script (see Figure 7.b) that linearly interpolated (LERP) between the window's current position and the camera position, creating a delayed following effect. Since during the design stage a distinction was planned between windows near the user for active interaction and windows further away for more passive visualization, this socket grid logic was split into two instances, each with its own "smooth follow" script, which resulted in a subtle parallax-effect where the distant windows moved slower than the near ones, contributing to a more natural experience.

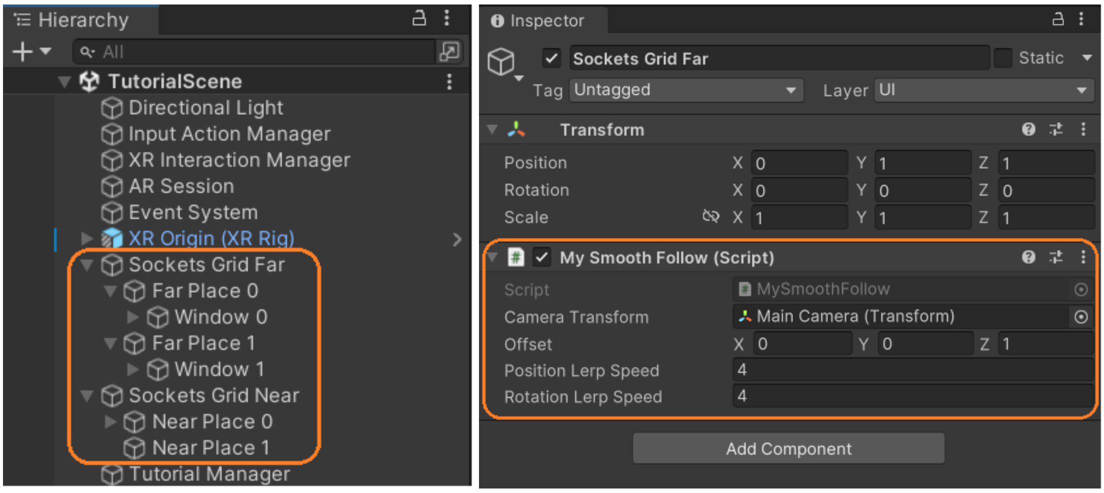

☑ Sockets Grid Far ☐ Static
Tag Untagged Layer UI

Transform
Position X 0 Y 1 Z 1
Rotation X 0 Y 0 Z 0
Scale X 1 Y 1 Z 1

☑ My Smooth Follow (Script)
Script MySmoothFollow
Camera Transform Main Camera (Transform)
Offset X 0 Y 0 Z 1
Position Lerp Speed 4
Rotation Lerp Speed 4

Figure 7 – (a) Socket Grids and Windows hierarchy; (b) Window following script.

Hand Menu Window

During development of this window interface logic, a need was identified for caregivers to have quick access to certain functions, such as viewing patient alerts or adjusting interface settings. Exploring XR Interaction Toolkit, an example of a Hand Menu functionality was found and adapted for this use. This menu is attached to the user's hand and appears only when the arm is raised vertically with the palm facing themselves, providing a portable quick access menu. In AleRa, this hand menu was redesigned as a tabbed interface, to provide:

- A notification list of patient-related alerts.
- Controls to open or close different windows as necessary.

Additionally, the default hand menu position and the threshold values for its activation were adjusted to be more lenient, to reduce strain and fatigue when performing the gesture repeatedly. The final hand menu layout and its configurations are presented in Figure 8.

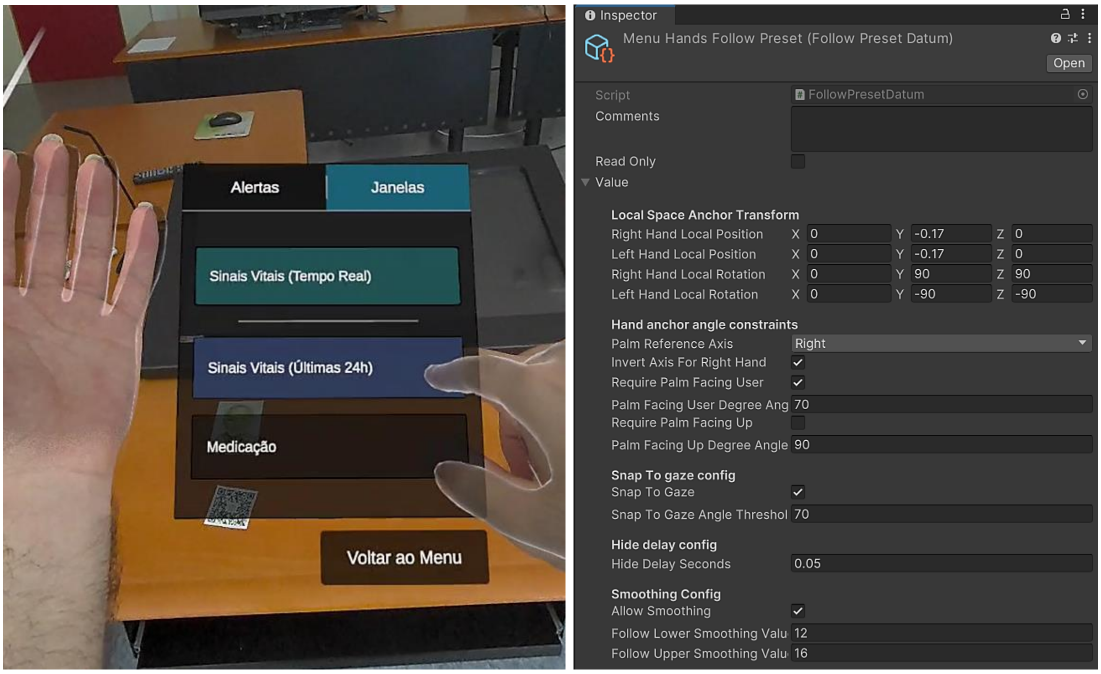

Menu Hands Follow Preset (Follow Preset Datum)

Script: FollowPresetDatum
Comments:
Read Only ☐
Value

Local Space Anchor Transform
Right Hand Local Position X 0 Y -0.17 Z 0
Left Hand Local Position X 0 Y -0.17 Z 0
Right Hand Local Rotation X 0 Y 90 Z 90
Left Hand Local Rotation X 0 Y -90 Z -90

Hand anchor angle constraints
Palm Reference Axis: Right
Invert Axis For Right Hand ☑
Require Palm Facing User ☑
Palm Facing User Degree Ang 70
Require Palm Facing Up ☐
Palm Facing Up Degree Angle 90

Snap To gaze config
Snap To Gaze ☑
Snap To Gaze Angle Threshol 70

Hide delay config
Hide Delay Seconds 0.05

Smoothing Config
Allow Smoothing ☑
Follow Lower Smoothing Valu 12
Follow Upper Smoothing Valu 16

Figure 8 – Hand Menu layout (left) and hand following configurations (right).

4.2.2 Authentication and Device Pairing

The AleRa system was designed to retrieve medical information from a FHIR-compliant server, with authentication and authorization managed by issued JWTs. Given the current limitations of text input in AR environments, identified during the design stage, the authentication process was delegated to the companion mobile application. According to the designed flow, both applications required implementation of UDP and Websocket communication functionalities.

Unity primarily supports scripting in C#, this Unity project specifically using Net Standard 2.1, so many of its standard libraries were directly usable. For UDP implementation, the default System.Net and System.Net.Sockets namespaces were sufficient. For Websocket communication, the external library WebSocketSharp was adopted due to its simpler API. Following the Single Responsibility principle, different functionalities requiring Websocket communication were handled through different services, with different paths, on the same socket: user authentication used the path "ws://<ip>:8888/login" while QR code reading for patient selection (see section 3.7) used the path "ws://<ip>:8888/qr-code".

A key concept that had to be considered is that, in Unity, only the main thread can safely interact with scene objects and update visual interface elements. Since WebSocketSharp handles sockets on a separate thread, the main thread's context needed to be stored and invoked whenever messages were received that required UI updates. Following the Information Expert principle, each WebSocket was managed by two classes with distinct roles:

- Behavior classes (e.g., LoginBehavior), responsible for sending and receiving WebSocket messages, and triggering Manager class functions when needed.

- Manager classes (e.g., LoginManager), responsible for controlling the UI flow, such as transitioning between window states when triggered by a Behavior class.

Figure 9 demonstrates the steps needed to be taken by a caregiver, using both applications, to authenticate in the AleRa system, namely the selection of the HMD device to connect to (found via UDP broadcasting), and the inputting of the confirmation PIN generated by the main HMD app.

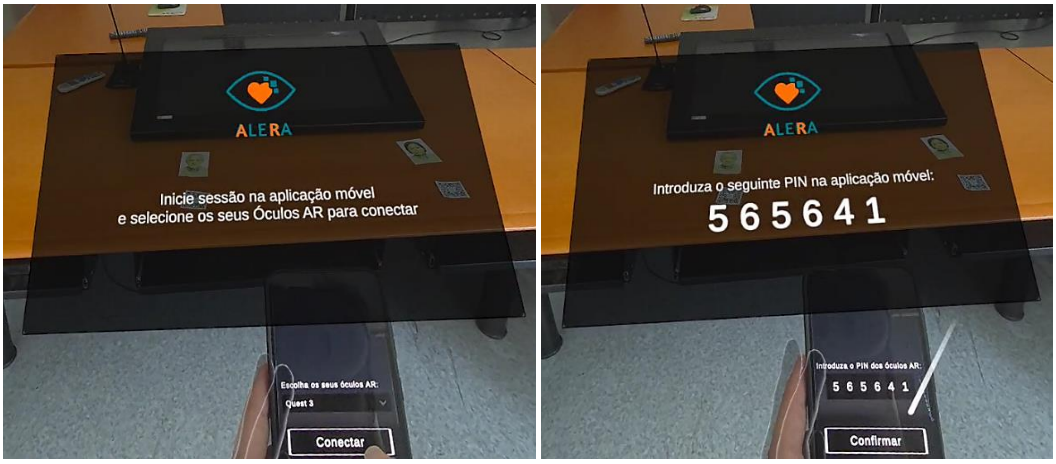

Inicie sessão na aplicação móvel
e selecione os seus Óculos AR para conectar

Escolha os seus óculos AR:

Introduza o seguinte PIN na aplicação móvel:
5 6 5 6 4 1

Introduza o PIN dos óculos AR:
5 6 5 6 4 1

Figure 9 – Steps taken by the caregiver to authenticate in the AleRa system.

4.2.3 Text Input Prototyping

Taking advantage of this modular WebSocket setup, an additional feature was prototyped to try to address the broader challenge of text input in HMDs. An additional socket service, with the path “ws://<ip>:8888/text-input” was used to synchronize both applications: when a Input Field UI object was selected in the HMD app, a message was sent to the mobile app, making it display its own Input Field object. As the user typed in the mobile app, the text was mirrored back to the HMD in real-time.

Although the prototype yielded promising results, it ended up not being further refined or integrated into the case study, as it created an even heavier dependency on the mobile companion app, which was assumed to not be ideal for the caregivers. Figure 10 shows a prototype version of this text input feature, being used as a search bar for a medication usage instructions window, another prototype feature that also ended up not being integrated in the case study (see section 4.2.7).

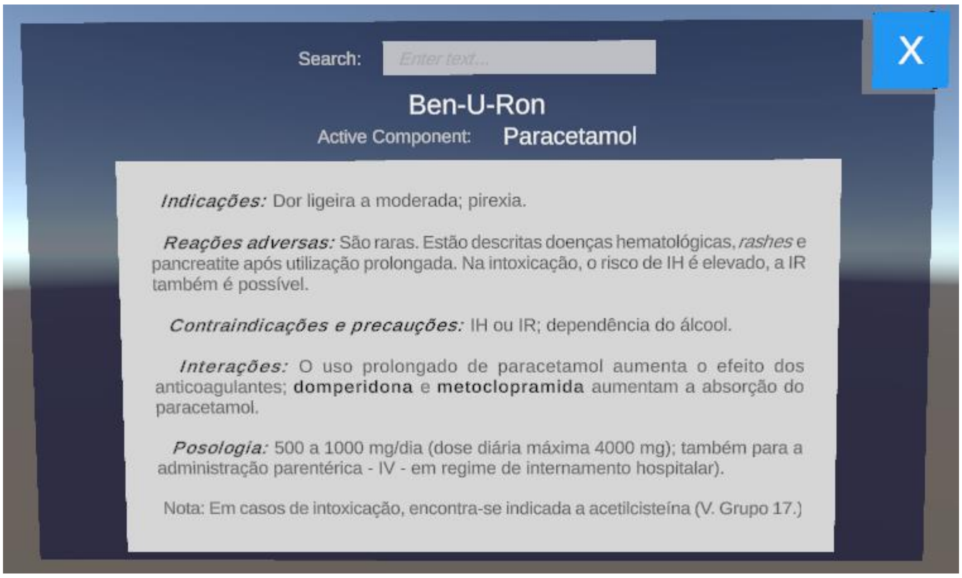

Ben-U-Ron

Active Component: Paracetamol

Indicações: Dor ligeira a moderada; pirexia.

Reações adversas: São raras. Estão descritas doenças hematológicas, *rashes* e pancreatite após utilização prolongada. Na intoxicação, o risco de IH é elevado, a IR também é possível.

Contraindicações e precauções: IH ou IR; dependência do álcool.

Interações: O uso prolongado de paracetamol aumenta o efeito dos anticoagulantes; **domperidona** e **metoclopramida** aumentam a absorção do paracetamol.

Posologia: 500 a 1000 mg/dia (dose diária máxima 4000 mg); também para a administração parentérica - IV - em regime de internamento hospitalar).

Nota: Em casos de intoxicação, encontra-se indicada a acetilcisteína (V. Grupo 17.)

Figure 10 – Prototype window featuring both the synchronized text input and medication administration instructions features.

4.2.4 FHIR Service Integration

The RM4Health system relies on HAPI FHIR, a Java Spring implementation of a FHIR compliant storage server, which exposes its “Create, Read, Update, and Delete” (CRUD) functionalities through an extensive REST API. To enable AleRa to fetch and display its medical information in the main HMD app, a communication layer was required.

Unity provides its own optimized web request package, UnityEngine.Networking, centered around the UnityWebRequest class. However, this is a low-level package, designed for generic requests, and Unity does not include a ready-to-use REST API client tool. As such, a dedicated RESTService wrapper class was implemented for this dissertation project, utilizing C# generics (T) and asyhnchronous Task operations to create a simple and generic communication service.

For serialization and deserialization between JSON and Unity objects, the Newtonsoft.Json external package was used. The RESTService class was implemented to provide high-level methods for the most common CRUD operations: GET, POST, PUT and DELETE. Although only GET and POST were required for this version of the AleRa system, this generalized design aimed to maximize reusability in future Unity projects where integration with REST APIs is a requirement.

4.2.5 Context-Aware Patient Selection

To facilitate rapid access to patient information without requiring navigation through multiple menus, the AleRa system was designed to incorporate QR code scanning as a method for patient selection (see section 3.7), with the scanning functionality itself being delegated to the companion mobile app, and the resulting information regarding the selected patient being then transmitted to the main HMD app using the previously established WebSocket system.

To implement QR code scanning in the mobile app, the external C# library ZXing was utilized, and a dedicated interface in the mobile app was created, as shown in Figure 11.a. Once a QR code is detected, the embedded JSON data is read and sent to the main HMD app via a dedicated WebSocket channel, "ws://<ip>:8888/qr-code". The JSON data consists only of two fields: a property identifying the type of data being read from the QR code, and the patient's FHIR ID, generated automatically by the existing FHIR server, when first registering a patient in the system. Currently, the first property only ever has a value of "patientId", but this field was designed with future expansibility of the QR code reading feature in mind.

The HMD app then follows an event-driven approach, implemented through C# Action delegates. Once it receives the JSON data, it updates the CurrentPatientId property in a PatientContextManager class and invokes the OnPatientChanged() Action. At this point, any window component registered to listen for changes in this Action is automatically notified, retrieves the new CurrentPatientId, and updates itself to display the new selected patient's information. This approach was used, following the Low Coupling practice, to decouple the patient management logic from the individual window's content management, improving modularity and maintainability.

Additionally, to signal to the caregiver that the QR code reading and patient changing was successful, a small window was implemented that temporarily fades in displaying the selected patient's name and profile picture, visible in Figure 11.b.

Figure 11 – (a) Mobile app QR code scanning; (b) "Patient changed" confirmation window.

4.2.6 Vital Signs and Alerts

As per the requirements, the AleRa system should be able to display patient vital signs to the caregivers, as well as alerts when abnormal values are detected. Both types of information were stored in the FHIR server, so during implementation they were retrieved using the previously developed RESTService class (see Section 4.2.4).

Vital sign monitoring in AleRa was separated and implemented according to two commonly seen perspectives:

- Real-time monitoring, where variables are continuously updated to reflect the patient's current state, and alerts are generated when abnormal values are detected. The aim here was to mirror a conventional hospital patient monitor.
- Historical visualization, where 24-hour data is displayed to allow caregivers to identify trends in the vital signs, as well as frequencies and timings of abnormal events. This corresponded to the health dashboards commonly seen in clinical applications.

In both cases, this data is usually displayed using charts, often accompanied by numerical aggregates such as averages, maxima, minima, or alert counts. In AleRa, the vital signs were displayed using the XCharts package for Unity, which provides several types of charts including line, bar, pie, and scatter plots, while accompanying numerical values and alerts were implemented through Unity's built-in UI components. Following the modularity principle designed for the window-based interface, each of these perspectives was created in a dedicated window.

Real-time Window

For the real-time window, line charts were vertically stacked to replicate conventional real-time vital sign monitors, similarly to prior work published by Kimmel et al. [16], analyzed during the literature review stage. Four key variables were shown, based on common ones seen in hospital monitors: heart rate (bpm), body temperature (°C), respiration rate (rpm), and blood pressure (mmHg).

Since this window was intended for passive, continuous visualization, it was classified as a "far" window (see Figure 4) by default, better positioned in the periphery of the user's field of view so as to not obstruct the real world. Nonetheless, as described in section 4.2.1, the socket grids that control window positioning were developed independently of the windows themselves, so this window can become a "near" window simply by reparenting it to the near socket grid. Figure 12 shows the resulting placement of this window in the user's view, and a closeup of its content.

Figure 12 – Positioning and content of the real-time vital signs monitoring window.

Real-time Alerts

In AleRa, alerts were designed to be triggered whenever a vital sign crossed its pre-configured threshold values. Alerts for caregivers are important for early detection of patient deterioration, and several early warning scores (EWS) exist that define such values, including the Modified Early Warning Score (MEWS) [46] and the National Early Warning Score (NEWS) 2 [47]. These EWS tend to define ranges of thresholds, instead of absolute values, since they are dependent on contextual factors such as age, conditions, and medication use.

Still, for this first version of the system the simpler implementation was adopted: absolute threshold values for heart rate, temperature, respiratory rate, and systolic blood pressure were chosen based on the ranges defined in MEWS [48] and NEWS 2 [47] (Table 12). These thresholds were implemented to be easily reconfigured or even replaced by range-based approaches in future iterations of the system.

Table 12 – Alert thresholds by variable and severity.

Variable	**Thresholds by Severity**			
	Warning		**Danger**	
	Above	Below	Above	Below
Heart Rate (bpm)	110	50	130	40
Temperature (ºC)	38.0	36.0	39.0	35.0
Respiratory Rate (rpm)	21	11	30	8
Systolic Blood Pressure (mmHg)	200	90	220	80

Alerts were accessible via the hand-attached menu (see “Hand Menu Window” in section 4.2.1), which listed all emitted notifications. Each notification included its respective patient name, a brief message expressing the abnormality detected, and was colored coded as either yellow or red, depending on if its severity was “warning” or “danger” respectively. Each alert was also accompanied by an audio cue when emitted, different for each severity.

Additionally, a feature was added so that selecting an entry in the alert list redirected the interface to the affected patient, serving as an alternative to QR code scanning for immediate monitoring in cases where the caregiver may be far away from the patient. Figure 13 shows the hand menu listing the alerts, and the user changing patients by selecting one of the notifications.

Mary Ane
Blood Pressure inferior a 90!

John Doe
Heart Rate inferior a 50!

Figure 13 – Alert notifications listed in the Hand Menu.

Historical Window

For the historical vital sign visualization, the dedicated window was implemented as a tabbed interface, showing 24-hour information on each of the four variables listed previously. Each tab contained a line chart of the respective variable, overlaid with two scatter plots used to mark the time of occurrence of “warning” and “danger” alerts. With the three plots superimposed, XCharts configurations were adjusted to allow the caregiver to hover their hand over the chart to visualize detailed information in a tooltip.

Additionally, next to the chart, numerical indicators were implemented to be dynamically calculated, showing the maximum and minimum values of that variable over the last 24 hours, as well as counting the occurrences of each alert severity type over the same period.

Since this window allowed for more active interaction, it was classified as a “near” window by default, being positioned in the bottom center of the user’s view (see Figure 4), but again this could be easily altered if needed, since the window component itself is independent from its positioning. Figure 14 shows the resulting historical vital signs window, and the details tooltip when hovering the chart.

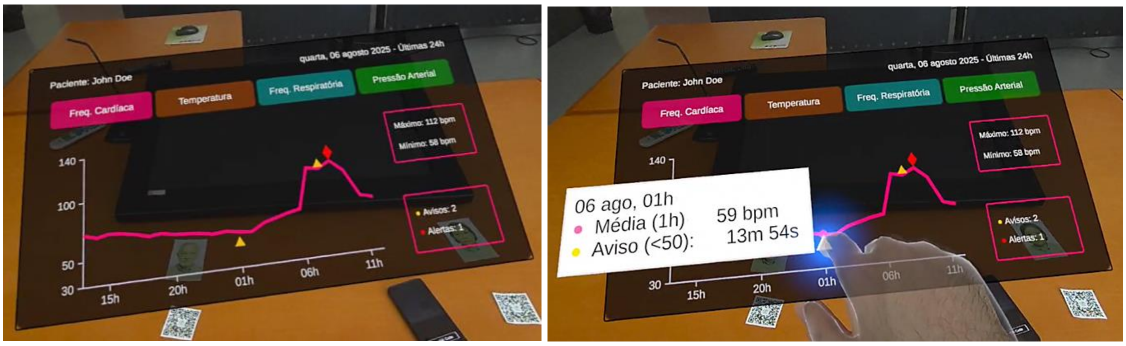

Figure 14 – Historical vital signs window appearance (left) and details tooltip (right).

4.2.7 Medication Intake

The final functional requirement for the AleRa system stated that it should be able to display medication intake information, allowing caregivers to track whether prescriptions are being correctly followed.

The implementation of this feature was inspired by RM4Health, where caregivers from a partner elderly home described that they use weekly blister packs provided by pharmacies. Each blister is filled with medication for seven days, with four daily slots: breakfast, lunch, dinner and before bed. RM4Health integrated a medication device with a top-down camera and an AI model to detect whether each slot was empty or filled, which is then compared against the patient's active prescriptions stored in the FHIR server. This enabled automatic classification of each dose as having been correctly or incorrectly taken (the latter including missed medication).

Based on this mechanism, the interface shown in Figure 15 was implemented in the AleRa system within a dedicated window. The caregiver is presented with a weekly grid where columns represent days, rows represent the four daily times, and each cell is labelled with the medications to be taken in that time slot. Due to each cells reduced size, when there is more than one distinct medication in the same cell, the name of the first one is presented, followed by a "(+X)" notation, where X is the remaining quantity of distinct medications. For example, if a patient was prescribed Zyrtec, Telfast, and Ben-u-ron to the same time slot, the cell presents it as "Zyrtec (+2)".

Additionally, each cell includes a status indicator, color-coded as:

- Green – correct intake.
- Red – incorrect (e.g. medication taken at wrong time) or missed intake.
- White – pending intake, in the future

Paciente: John Doe

04 agosto - 10 agosto

	Segunda	Terça	Quarta	Quinta	Sexta	Sábado	Domingo
Pequeno Almoço	Xanax (+2) 🟥	- 🟩	Zyrtec (+2) 🟥	Valium (+1) 🟩	Zyrtec (+2) 🟩	Xanax ⬜	Pepsamar (+2) ⬜
Almoço	Trifene 200 (+1) 🟥	Xanax (+3) 🟩	- 🟩	Valium (+2) 🟩	Valium 🟥	Telfast 180 ⬜	Xanax (+2) ⬜
Jantar	- 🟩	Xanax (+3) 🟩	Telfast 180 🟩	Ben-u-ron 🟥	Trifene 600 (+1) ⬜	- ⬜	Telfast 120 (+3) ⬜
Antes de Dormir	Telfast 120 (+1) 🟩	- 🟩	Xanax (+2) 🟥	Xanax (+2) 🟥	Ben-u-ron (+1) ⬜	- ⬜	Trifene 600 (+2) ⬜

Figure 15 – Weekly medication intake tracking window.

To provide additional details, a secondary window within this window component was developed, shown in Figure 16. When the caregiver selects a time slot, this new window appears, listing:

- All prescribed medications and quantities for that slot.
- The detected intake status (global to that time slot).
- The image captured by the medication device, allowing caregivers to manually confirm the AI model's analysis, and detect anomalies, for example, objects obstructing the medication blister.

For added usability and immersion, the transition between the two windows was animated using LeanTween. The secondary window appears with a quick popup effect, while the main window fades out to improve readability. In an initial implementation the main grid window only became semi-transparent, but this was discarded as it reduced the legibility window on top, due to the inherent transparency already present in all AleRa windows.

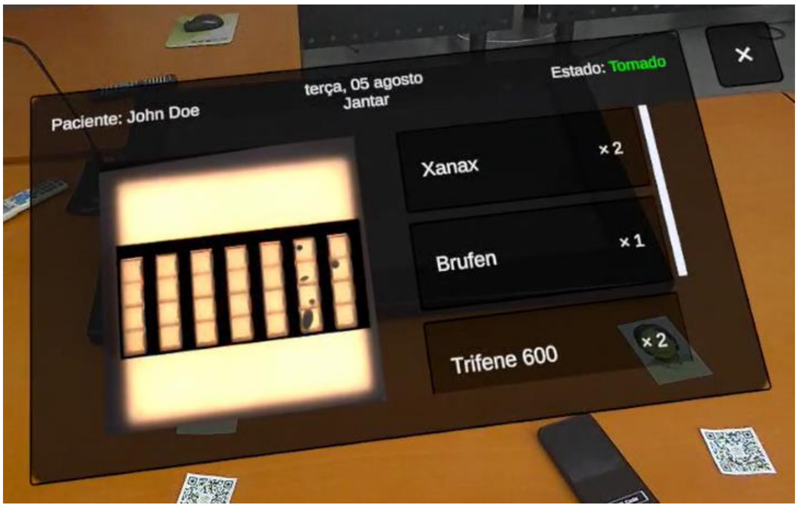

Paciente: John Doe

terça, 05 agosto
Jantar

Estado: Tomado

Xanax × 2

Brufen × 1

Trifene 600 × 2

Figure 16 – Time slot intake details window.

Finally, a prototype for an additional window that would display other details about a selected medication, namely its leaflet and administration instructions, was implemented using the information available on the Infarmed web portal [49]. However, this web portal does not provide a direct API for fetching information. As such, for feasibility testing purposes, the information was extracted by parsing the HTML returned when querying the platform. It showed promising results, as shown previously in Figure 10, but this feature was ultimately suspended due to this not being a very viable or scalable way of obtaining information. Still, this showed the potential for integrating this type of information into the system, in future work, if access to a dedicated API is ever available.

4.2.8 Voice Recognition Prototyping

During this dissertation, voice recognition was studied as a possible alternative to text input and its associated difficulties in HMD applications. However, due to privacy and ethical concerns, voice recognition was only explored as a proof of concept rather than as a realistic or production-ready feature. The EU AI Act, more specifically Article 2(6) of Regulation (EU) 2024/1689 [42], explicitly exempts the use of AI models and its outputs for the sole purpose of scientific research and development activities from strict regulation obligations. This made it possible to prototype voice input in a controlled research setting without aiming for deployment. Still, Article 2(8) clarifies that “testing in real world conditions shall not be covered by that exclusion”, meaning that this feature, even for research purposes, could not be included in the case study.

For this prototype, the Vosk speech recognition toolkit was selected. Vosk is an open-source, offline speech recognition library with multi-language support and bindings for programming languages including C#, making it suitable for Unity integration. The model used was vosk-model-small-pt-0.3 (≈30 MB), chosen since all expected case study participants were native Portuguese speakers.

It showed promising results in preliminary Portuguese voice recognition tests, under controlled conditions. Nonetheless, in addition to the privacy concerns addressed earlier, due to uncertainty about microphone availability and quality across HMDs, and due to expected noise in real-life conditions where caregivers operate, the feature was not further developed. Still, it remains a potential avenue to reduce dependency on the companion mobile app for general text input in future work.

4.2.9 Facial Recognition Prototyping

Face recognition methods were studied as an alternative mechanism for implementing context awareness for patient selection, avoiding the need for QR code scanning via the companion application and for setting up patient specific QR codes in physical locations (e.g. patient’s bedside or room).

However, similarly to voice recognition, privacy and ethical concerns made it unsuitable to be a production-ready feature. Prototyping was nonetheless possible under the scientific research exemptions of Article 2(6) and Article 2(8) of Regulation (EU) 2024/1689 [42], but testing was not included in the case study.

It was developed using the Python programming language (version 3.11.9), using a two-step pipeline of pretrained models. First, a Multi-task Cascaded Convolutional Network (MTCNN), proposed by Zhang et al. [50] was used for face detection, cropping, and alignment. Then, the cropped face was processed by an InceptionResNetV1 network, pretrained according to the FaceNet methodology of Schroff et al. [51], which mapped the face to a 512-dimensional embedding vector. For this implementation, the most important Python libraries were used were: PyTorch and Torchvision for general model development, facenet_pytorch [52] for the provided pretrained MTCNN and InceptionResnetV1 models, and Pillow for image drawing and manipulation.

Figure 17 was created to illustrate this pipeline process and includes two alternative visualizations of the embeddings (a point cloud and a barcode-like representation).

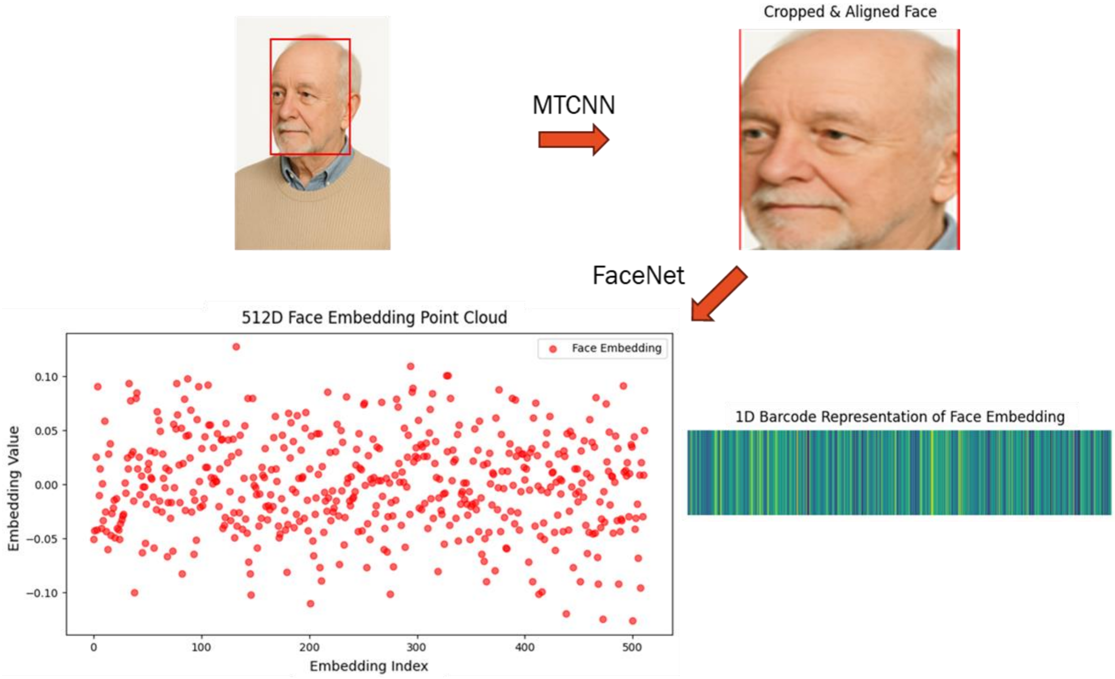

Figure 17 – Prototyped face embedding pipeline.[4]

From here, recognition was achieved by comparing the desired face's embeddings against a database of stored patient embeddings using cosine similarity, which measures the angular distance between two vectors [53], in this case each with 512 dimensions. The formula for calculating cosine similarity is given by:

---

[4] The patient images depicted in the figures relating to this face recognition pipeline are not of real people, being instead generated with AI, to mitigate ethical and privacy concerns.

$$similarity(A, B) = \frac{A \cdot B}{\|A\| * \|B\|} \tag{1}$$

where A and B are the embedding vectors.

For this pipeline, an optimized implementation of cosine similarity was used, provided by the "pytorch" library. From this formula, a higher similarity value attained indicated a closer match between the two vectors, and as such between the input face and stored patient image. In other words, the stored patient image embeddings that, when compared with the target image, resulted in the highest cosine similarity value determined that patient as the one recognized by the developed pipeline. Figure 18 depicts an example of this face embedding comparison process using cosine similarity.

It showed preliminary promising results, and although it was initially planned to operate using the mobile phone camera, after Meta announced developer access to Quest 3 cameras [41] it became technically feasible to run the entire pipeline directly on the HMD, by converting the models used to the Open Neural Network Exchange (ONNX) format and using them in C#.

Despite this, the prototype was not further developed since it would not be realistic as a production-ready feature due to the privacy and ethical constraints, particularly the risk of capturing images of non-consenting patients and caregivers.

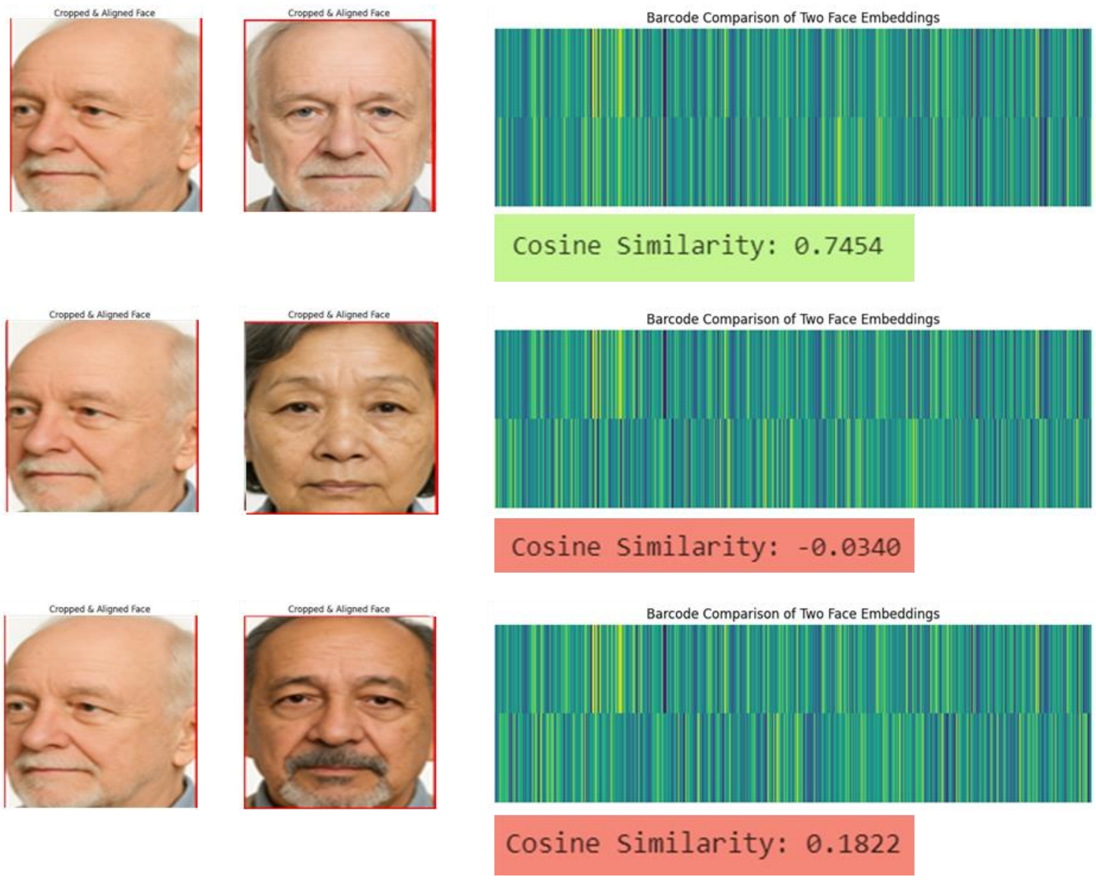

Figure 18 – Face recognition results using cosine similarity between embeddings.

4.2.10 Tutorial

To support first-time users, a dynamic interactive tutorial was constructed, using the previously developed components, to familiarize caregivers with the system interactions. This tutorial combines 3D hand visualizations, interactive windows, and narrated audio instructions, being structured into four steps:

1. Window familiarization: Distant windows appear, and users are instructed to move around and observe how the windows follow their field of view with parallax effects. A progress bar, made using XCharts’ ring chart, tracks the completion of the movement.
2. Ray-based interaction: Buttons appear within the windows, and users are asked to interact using the ray-based pointing and pinching method. An animated virtual hand demonstrates the required motion.
3. Direct touch interaction: One of the windows is repositioned closer to the user, and new buttons are introduced. Users are instructed to activate these buttons by touching them directly with their index fingertip.
4. Hand Menu interaction: Finally, users are guided on how to open the Hand Menu by positioning their hand vertically with the palm facing themselves (see section 4.2.1).An animated virtual hand demonstrates the required movement, and a button within the menu allows users to practice interacting with it.

Throughout the tutorial, short encouraging voice prompts (e.g., "Perfect", "Nice job") are provided by the system to maintain user motivation and engagement. Figure 19 illustrates the appearance of these four steps, from the perspective of a user following the tutorial.

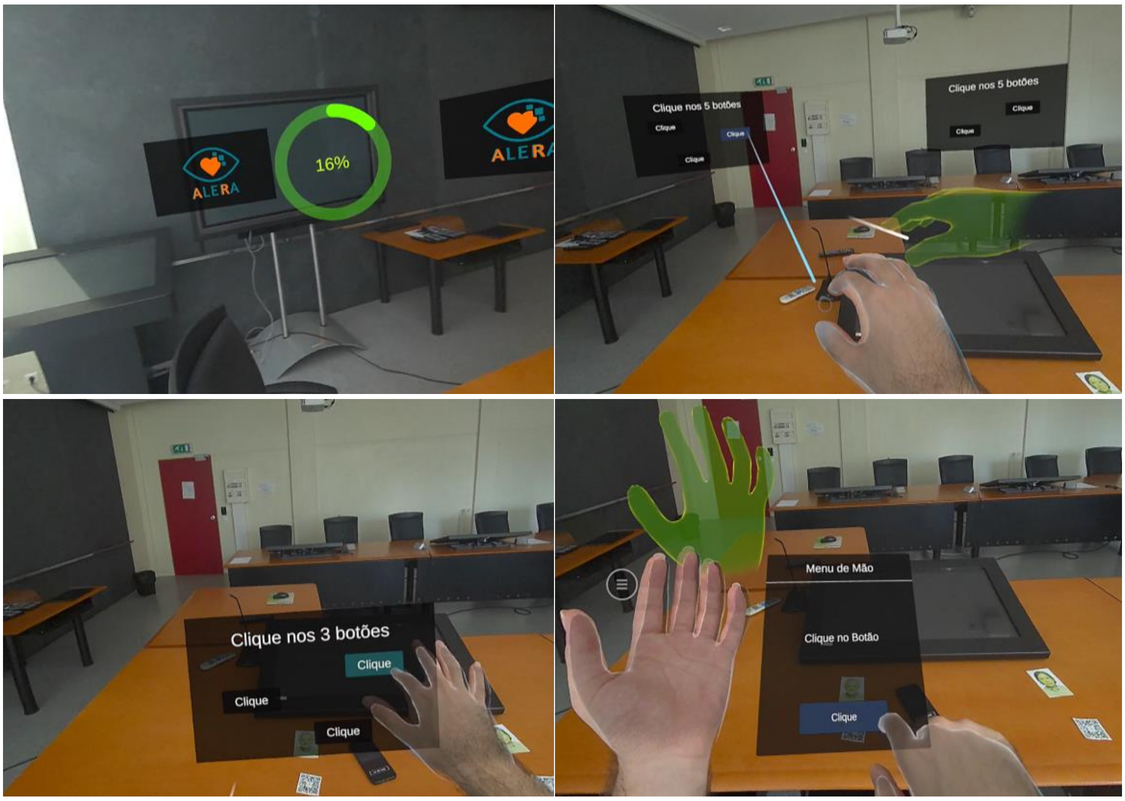

Figure 19 – Tutorial steps. (a) Window following; (b) Ray interactor tool; (c) Fingertip interaction; (d) Hand Menu.

Solution

5 Case Study

An empirical study was conducted to evaluate the users' reception and overall experience of the developed AleRa system as an assisting tool for patient monitoring, with Figure 20 showing an example of a participant's session. Two use cases were simulated by constructing dedicated interfaces using the individual features described in the Solution section (see section 4), to assess usability, intuitiveness, and responsiveness of the system, while also validating the content choices and respective assumptions made for each use case interface.

The first use case is aimed at nurses working in the General Ward of a hospital, with potential applicability also in the Intermediate Care Unit (IMCU). In these settings, patients' vital signs require constant real-time monitoring, and the main assumption was that nurses are often occupied with tasks or carrying equipment, making complex interactions with the system less feasible. As such, this uses case's interface was designed to be more passive, prioritizing data visualization instead of active interaction.

Meanwhile, the second use case targeted domiciliary care nurses attending patients in their homes. In this scenario it was assumed that nurses may need to consult more detailed information on the patient, such as historical vital sign data or medication intake schedules, and that, on the other hand, they have more freedom to actively engage with the system. Consequently, this interface was designed to support more complex interactions with menus and windows in the AR environment.

Figure 20 – Participant interacting with the AleRa system during the case study.[5,6]

5.1 Participants

This study was conducted in a laboratory of the GECAD research center, on ISEP's campus, with a total of 17 participants initially contacted for potential recruitment. Ten participants were recruited from the local university, while the remaining 7 were contacted due to their experience or education in medical or health-related fields. Of the latter, one was unable to participate due to scheduling constraints, leaving a total of 16 participants who completed the experiment and evaluated the current version of the system.

The 10 participants recruited from the university, with the main purpose of testing the system's usability, were evenly divided and randomly assigned to one of two groups: "No Tutorial" and "Tutorial" respectively. Both groups carried out the exact same experiment, evaluating both use cases. The first group received no instructions on how to operate the system, only contextual information about the experiment and each use case's goals, with the intention of performing a discoverability test while the second group completed a dynamic interactive tutorial before proceeding with the testing. The remaining 6 participants followed the same procedure as the "Tutorial" group but were assigned to a third group, "Medical", based on their health-related expertise, to allow their feedback to be more easily isolated and analyzed, particularly on the validation of the use cases' utility for nurses. Since the difference between being trained with a tutorial or not had already been tested with the previous two groups, it

---

[5] In the AleRa system, all user interactions are made through the HMD and the companion mobile app. The monitor visible in the figure was used exclusively during the case study to observe the participant's experience and provide guidance as necessary.

[6] This image was included with the explicit consent of the participant.

was chosen not to split the “Medical” group participants, so that their more informed feedback on use case relevance was not skewed by some having not had tutorial training. This recruitment and grouping process has been summarized through a diagram in Figure 21.

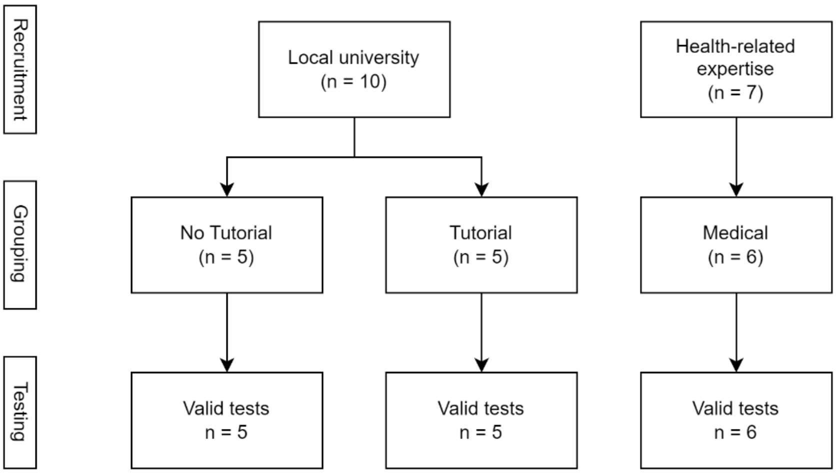

Figure 21 – Case study recruitment and grouping diagram.

Nielsen’s well-known work on usability testing suggests that testing with five users is sufficient to identify approximately 85% of usability issues in a software’s current version of development [54], [55]. However, other studies and researchers find this claim controversial. For instance, Faulkner [56] found through empirical testing that, with 10 users, the lowest percentage of problems found was 80%. Jeff Sauro [57] adds the important nuance that the “5 users rule” only applies to problems that occur more than 31% of the time, and that less common issues need higher quantities of users to be properly detected. Applying Jeff Sauro’s calculations to this study, this set of 16 participants is capable of finding 100% of the problems that affect 31% of users, and approximately 75% of less common problems that only affect 10% of users, as can be seen in Figure 22 (authorship of Jeff Sauro [57]). This user sample was, as such, deemed sufficient to test the usability of this first version of the AleRa system.

In terms of the usefulness of the health-related content, although this evaluation featured a reduced number of users with health-related expertise, the feedback provided by the 6 participants in the “Medical” group was nonetheless considered valuable, serving as preliminary evidence on the system’s potential utility in healthcare contexts.

During recruitment, all participants filled out a pre-questionnaire where they gave informed consent and answered demographic and characterization questions regarding prior technology experience with health applications or with virtual and augmented reality. No names were requested, and each participant was instead randomly assigned a sequential code within their group to ensure anonymity, and to further preserve it, the users’ ages were collected in bin ranges. Table 13 shows a summary of the participant characterization data collected, while the characterization questionnaire, as administered to the participants, can be seen in Appendix A – Characterization Questionnaire.

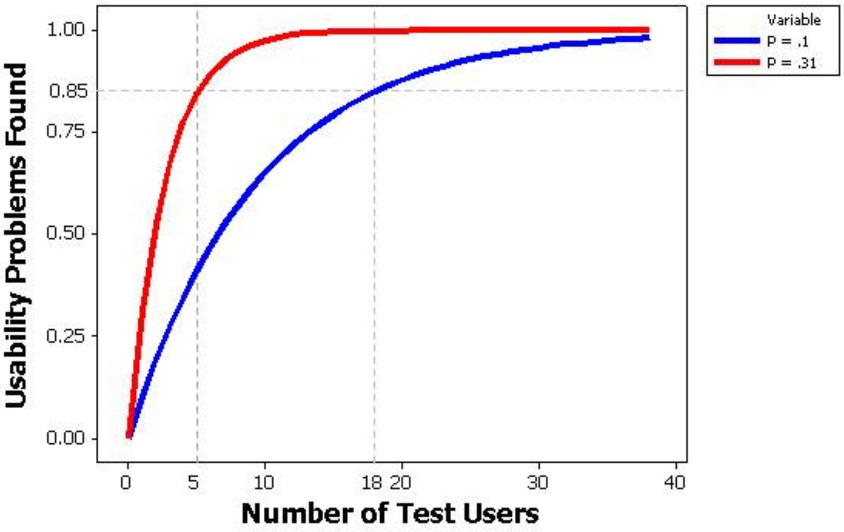

Figure 22 – Jeff Sauro's data on the number of usability problems found for different numbers of users, depending on the percentage of users each problem affects.[7]

The majority of participants were in the age range of 18-34 and all held at least a bachelor's degree. Participants in the “No Tutorial” and “Tutorial” groups specialized in Engineering and Technology, while all in the “Medical” group confirmed their background in Medical and Health Sciences. Among the 16 participants, one reported a temporary physical limitation, being unable to move one hand. Across all participants, only 6 reported having previous experience with virtual reality, and just 3 with augmented reality through an HMD device similar to that used in this dissertation. Only one participant, in the “Medical group”, reported prior experience with a headset-based AR health application, used by patients to simulate scenarios as a therapeutic aid for social phobia (e.g., public speaking simulation). This participant described their previous experience as positive and, like all others, demonstrated interest and curiosity in testing the AleRa system.

---

[7] Source: J. Sauro - “Why you only need to test with five users (explained)” [57]

Table 13 – Case study's participant demographics.

Property	Per Participant Group No Tutorial (n=5)	 Tutorial (n=5)	 Medical (n=6)	Overall (n=16)
Age Range				
18–24	3 (60%)	3 (60%)	0 (0%)	6 (38%)
25–34	1 (20%)	2 (40%)	5 (83%)	8 (50%)
35–44	0 (0%)	0 (0%)	1 (17%)	1 (6%)
45–54	1 (20%)	0 (0%)	0 (0%)	1 (6%)
Gender				
Female	2 (40%)	1 (20%)	4 (67%)	7 (44%)
Male	3 (60%)	4 (80%)	2 (33%)	9 (56%)
Education Level				
Bachelor	3 (60%)	2 (40%)	4 (67%)	9 (56%)
Master	1 (20%)	3 (60%)	2 (33%)	6 (38%)
PhD	1 (20%)	0 (0%)	0 (0%)	1 (6%)
Expertise Fields[8]				
Engineering and Technology	5 (100%)	5 (100%)	2 (33%)	12 (75%)
Medical and Health Sciences	1 (20%)	0 (0%)	6 (100%)	7 (44%)
Exact Sciences	1 (20%)	0 (0%)	0 (0%)	1 (6%)
Management	0 (0%)	0 (0%)	1 (17%)	1 (6%)
Physical Limitations	0 (0%)	1 (20%)	0 (0%)	1 (6%)
Prior Technology Experience				
Health apps for others' care	1 (20%)	0 (0%)	4 (67%)	5 (31%)
Health apps for self care	3 (60%)	3 (60%)	6 (100%)	12 (75%)
Virtual Reality	2 (40%)	2 (40%)	2 (33%)	6 (38%)
Augmented Reality (Mobile)	3 (60%)	1 (20%)	1 (17%)	5 (31%)
Augmented Reality (Headset)	1 (20%)	1 (20%)	1 (17%)	3 (19%)
VR/AR health apps	0 (0%)	0 (0%)	1 (17%)	1 (6%)

5.2 Apparatus

For this study, the main application of the AleRa system was deployed on a Meta Quest 3 headset. As previously described (see section 4.1), the Quest 3 weighs 515 grams, features dual color LCD panels with a resolution of 2064 × 2208 pixels per eye, supports refresh rates up to 120 Hz, and a field of view of 110° horizontal and 96° vertical [44]. Additionally, to improve user comfort, the standard head-fitting strap was replaced with an "Elite Strap", an official Meta accessory which adds more structured support and better distributes the headset weight compared to the default head strap [58].

For the companion mobile application, the same Samsung Galaxy S23 smartphone was used across all participants to avoid possible variability in performance due to differing device characteristics. The Galaxy S23 is equipped with a "Snapdragon 8 Gen 2" processor [59], and 8

---

[8] Participants could select multiple expertise fields, so the total per group may exceed 100%.

GB of RAM [60]. The mobile app handled user authentication and patient selection via QR code scanning, as detailed in the Solution section (see section 4). To complement the latter, two printed QR codes (paired with AI-generated patient photos to preserve privacy) were used to simulate patient selection scenarios. All the apparatus described, used during the case study sessions, is presented in Figure 23.

Although not a part of the actual AleRa system, the Meta Quest 3's "Cast to Web Browser" feature was used during the study to mirror the participant's view over Wi-Fi to a Dell Precision 5570 laptop, allowing to monitor their experience and provide real-time guidance when needed, ensuring study integrity.

Figure 23 – Apparatus used for the case study. From left to right: QR codes and patient photos, Samsung Galaxy S23 running the AleRa companion app, and Meta Quest 3 headset equipped with the Elite Strap accessory.

5.3 Prepared Use Cases

Using the independent features developed and described in the Solution section (see section 4), two interfaces were constructed, to reflect the use cases chosen for this study. Each interface was designed to present information deemed necessary for the corresponding type of nurse of that use case, drawing from the kinds of data that would be available on a FHIR server. It is important to note that, although the AleRa system can connect to and use information from a FHIR server, no real data from RM4Health patients, or any other source, was presented to participants during the case study, in order to comply with ethical and privacy regulations. Instead, mock data within realistic ranges for each variable was used.

The first interface, targeting nurses assigned to a hospital's General Ward unit of a hospital, included the following elements:

- A permanently open window in the top-right of the user's field of view, showing real-time vital signs of the currently selected patient (via QR code).
- A list of notification-style alerts embedded in the Hand Menu, with two levels of severity, “warning” and “danger”. To simulate a patient's deteriorating condition during the case study experiments, alerts were programmed to appear at random intervals, but only after the user had already interacted with the interface for a while and had scanned at least two QR codes to switch patients.

Meanwhile, the second interface, aimed at domiciliary care nurses visiting patients in their homes, in addition to the features from the previous use case, also included:

- An enhanced Hand Menu with a tab system, where the first tab contained the alert notifications list, while the second tab provided a new window management menu. This allowed users to open or close the real-time vital signs window, as well as to access two new windows.
- A vital signs history window, showing the vital signs and alerts of the patient for the last 24 hours, for all four monitored variables.
- A medication intake window, displaying a weekly grid of prescribed medications across four daily periods (breakfast, lunch, dinner, and before bed). Each slot indicated if that medication was correctly taken (green), missed (red), or pending (white). Additionally, by selecting a time slot, the nurse could open a subwindow that lists the medications scheduled for that period, along with a photograph of the patient's medication blister at that time, enabling visual confirmation of the intake.

5.4 Procedure

After recruitment, the participants were asked to complete the previously mentioned consent and characterization questionnaire. Then, since not all trials could be conducted sequentially, the experiment was set up individually for each participant according to their availability.

During setup, first a section of the laboratory's floor (approximately 1.5m × 4m, or 6m^2) was mapped as the usable boundary area for the app, as required by the Meta Quest system. Within this area, two printed QR codes representing patients were placed at different locations, each accompanied by the respective mock patient photograph.

To avoid exposing participants to the Meta Quest 3 operation system’s windows and interaction tools, and thus unintentionally giving them prior experience before using the developed app, the application was launched before handing the HMD to the participant, with a delay being coded to account for the time to transfer the headset to the participant before showing the “welcome” interface animation. This precaution was especially important for the “No Tutorial” group, whose role was to test the intuitiveness and discoverability of the system without prior guidance. At this stage, the casting feature was also activated to mirror the user experience to a laptop, enabling monitoring and guidance during the experiment.

Once ready, the HMD was handed to the participant, who was asked to stand inside the defined boundary. The experiment began with authentication and connection of the main app to the companion mobile app. Then, participants in the “Tutorial” and “Medical” groups were asked to complete the interactive tutorial (see section 4.2.10).

Afterwards, participants were redirected to an interface menu specifically designed for this study, pictured in Figure 24, that allowed them to select which use case to simulate. They were asked to begin with the in-hospital use case, which simulates a nurse in a general ward. Before starting, a brief contextual explanation of the scenario, the participant’s role, and the existing functionalities were given. Apart from scanning QR codes with the mobile app, no instructions on how to access the features were provided, unless participants specifically asked. This approach allowed the “No Tutorial” group to test the natural intuitiveness of the system, particularly the discoverability of the Hand Menu and of the ray interaction tool, which required specific gestures, and allowed the “Tutorial” and “Medical” groups to test the effectiveness of the prepared tutorial as a single source of information to introduce AR-specific interactions to participants with little prior experience.

Internamento /
Cuidados Intermédios

Contexto:
Enfermeiro/a hospitalar a monitorizar o estado de vários pacientes, em constante movimento e com as mãos ocupadas (transporte de medicação, etc.).

Interface:
Apresenta sinais vitais em tempo real e alertas, sem necessidade de interação ativa com o sistema.

Cuidados Domiciliários

Contexto:
Enfermeiro/a em visita domiciliária, com foco na avaliação clínica do paciente e gestão de medicação.

Interface:
Apresenta sinais vitais em tempo real, histórico das últimas 24h e estado de toma de medicação, através de interação ativa com o sistema.

Figure 24 – Use case selection interface menu, created for the case study.

From here, participants were then guided to explore both use cases, interacting with the functionalities described earlier. The complete Unity scene flow that participants had to follow is depicted in Figure 25. After completing the experiment, all participants were asked to fill out a second questionnaire to evaluate the system and providing optional feedback.

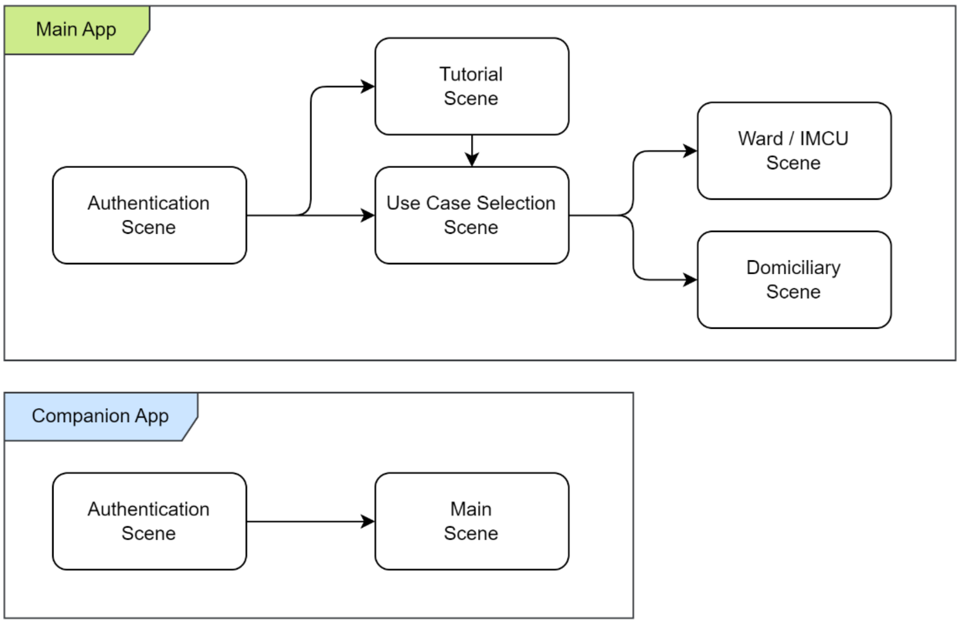

Figure 25 – System scene flow, of both applications, implemented for the case study.

5.5 Evaluation Questionnaire

The evaluation questionnaire was designed to gather the participants' feedback and experiences while using the system. It was equally important to assess both the usability of the software and interface and the perceived relevance of the medical information presented in the use cases. Additionally, attention also needed to be given to physical and mental aspects, such as any discomfort caused by the headset or challenges associated with AR interactions in general, such as fatigue or frustration. To address all these aspects, various recognized evaluation scales were incorporated, resulting in a questionnaire structured in four parts, presented in the following subsections.

The evaluation questionnaire was hosted through the Microsoft Forms platform. Since all participants spoke Portuguese, all scales and questions were translated to their native language for better interpretability. The evaluation questionnaire, as administered to the participants, can be seen in Appendix B – Evaluation Questionnaire.

5.5.1 System Usability Scale (SUS)

The System Usability Scale (SUS) [61] is a widely used instrument, employed in many software and interface evaluation studies to assess their perceived usability. For instance, in the context of AR, Tasfia et al. [62] recently applied it in a study to evaluate various AR-based learning systems for children, and Dutta et al. [63] similarly used it to evaluate an AR app for teaching

Karnaugh maps. Additionally, a systematic review of mHealth applications, by Sousa et al. [64], observed that SUS remains the most commonly employed instrument in the health domain.

In this study, the standard SUS scale was used. The SUS is comprised of 10 questions, presented in Table 14, rated on a 5-point Likert scale from 1 (“strongly disagree”) to 5 (“strongly agree”), with polarity alternating between positive and negative to reduce response bias.

Table 14 – Standard SUS utilized in the evaluation questionnaire.

ID	**Question**	**Strongly Disagree**				**Strongly Agree**
S_1	I think that I would like to use this system frequently.	1	2	3	4	5
S_2	I found the system unnecessarily complex.	1	2	3	4	5
S_3	I thought the system was easy to use.	1	2	3	4	5
S_4	I think that I would need the support of a technical person to be able to use this system.	1	2	3	4	5
S_5	I found the various functions in this system were well integrated.	1	2	3	4	5
S_6	I thought there was too much inconsistency in this system.	1	2	3	4	5
S_7	I would imagine that most people would learn to use this system very quickly.	1	2	3	4	5
S_8	I found the system very cumbersome to use.	1	2	3	4	5
S_9	I felt very confident using the system.	1	2	3	4	5
S_{10}	I needed to learn a lot of things before I could get going with this system.	1	2	3	4	5

The SUS final score is calculated using a set of instructions, due to the questions’ positive-negative polarity. For questions 1, 3, 5, 7, and 9 (positively worded), each item contribution is (score − 1). For items 2, 4, 6, 8, and 10 (negatively worded), each contribution is (5 − score). The sum of contributions is multiplied by 2.5 to yield a score from 0 to 100. This is equivalent to the formula presented below.

$$SUS = 2.5 * \left(20 + \sum(S_1, S_3, S_5, S_7, S_9) - \sum(S_2, S_4, S_6, S_8, S_{10})\right) \tag{2}$$

where S_i is the (1-5) response score to question i.

5.5.2 NASA Task Load Index (NASA-TLX)

Given that AleRa is not a conventional desktop or mobile application, but a HMD AR system that requires physical movement and the learning of novel interaction gestures and tools, a scale developed by the National Aeronautics and Space Administration (NASA), the NASA Task Load Index (NASA-TLX) [65], was used to measure perceived workload. Participants rated their experience across six dimensions: mental demand, physical demand, temporal demand, performance, effort, and frustration, using a 0-100 scale in increments of 5, as presented in Table 15.

This choice is consistent with contemporary usage, as recent studies are applying NASA-TLX to AR/VR interfaces and demonstrating its sensitivity in measuring mental and physical workload. For example, Xi et al. recently used this index to study the effects of AR and VR interfaces on user workload in a shopping-related task, and the challenges that they bring [66]. Similarly, Knudsen et al. [67] validated the use of a VR HMD device to assess the emergency medicine skills of junior doctors, using NASA-TLX as a way to rate their cognitive workload. In addition, a systematic review on AR its cognitive load highlighted that NASA-TLX, by Buchner et al. [68], is the most frequently used measuring instrument.

As such, the standard NASA-TLX dimensions and scales were included as the second step in the evaluation questionnaire, with the dimension names and respective questions translated into Portuguese for better interpretability by the participant base.

However, one temporary modification was made to the “performance” item, to reduce potential confusion when completing the questionnaire. Most NASA-TLX questions are negatively framed, meaning that higher scores indicate a greater workload, and therefore a poorer system. For example, for the question “How mentally demanding was the task?”, a response closer to 100 reflects a more demanding task and thus worse usability. In contrast, the “performance” item (“How successful were you in accomplishing what you were asked to do?”), is worded positively. In the standard NASA-TLX, this item’s scale is reversed so that a higher score corresponds to lower success. While this ensures consistency in scoring (100 always reflects a negative point of the system), it could feel counterintuitive for participants, as the question wording suggests that higher values should reflect greater success. To minimize this issue, in this study the “performance” scale was inverted during the questionnaire administration (0 = failure, 100 = perfect). Later, during analysis, these responses were then re-inverted (100 – response) to align with the standard NASA-TLX scoring method.

Table 15 – Modified NASA-TLX utilized in the evaluation questionnaire.

Dimension	Question	Scale
Mental Demand	How mentally demanding was the task?	0 – Very Low 100 – Very High
Physical Demand	How physically demanding was the task?	0 – Very Low 100 – Very High
Temporal Demand	How hurried or rushed was the pace of the task?	0 – Very Low 100 – Very High
Performance	How successful were you in accomplishing what you were asked to do?	0 – Failure 100 – Perfect
Effort	How hard did you have to work to accomplish your level of performance?	0 – Very Low 100 – Very High
Frustration	How insecure, discouraged, irritated, stressed, and annoyed were you?	0 – Very Low 100 – Very High

5.5.3 Custom Likert-scale questions

Apart from general usability and workload evaluation, it was deemed important to quantitatively measure some specific aspects of the developed system, such as each participant's opinion on the usefulness of the tutorial or on the interface's design and interactivity. This would allow for comparisons between groups as well as an overall analysis of participant feedback.

The previously mentioned systematic review on mHealth applications, by Sousa et al. [64], also reported that most studies, even when not adopting a standardized scale or index, utilized Likert-scale style questions for quantitative measuring. Following this approach, the third step of the evaluation questionnaire consisted of a set of seven additional questions rated on a 5-point Likert scale, as presented in Table 16, designed to capture the participant's views on these specific aspects of the system.

Table 16 – Additional Likert-scale style questions utilized in the evaluation questionnaire.

ID	**Question**	**Very bad**				**Very good**
Design	How do you rate the interface in terms of its design and layout?	1	2	3	4	5
Interactivity	How do you rate the interface in terms of navigation and interactivity?	1	2	3	4	5
Tutorial	Was the tutorial useful and easy to understand?	1	2	3	4	5
Health Info (Ward)	Did the health information on the interfaces for general ward nurses seem adequate to you?	1	2	3	4	5
Health Info (Domiciliary)	Did the health information on the interfaces for domiciliary care nurses seem adequate to you?	1	2	3	4	5
AR Benefit (Nursing)	Do you think the use of augmented reality could be an added value for these areas of nursing?	1	2	3	4	5
AR Benefit (Own Field)	Do you think the use of augmented reality could be an added value for your area of training or activity?	1	2	3	4	5

5.5.4 Open-ended questions

Finally, an optional section invited participants to provide free-text feedback, namely on (i) features they appreciated the most, (ii) features they liked the least, and (iii) suggestions for improvements to enhance usability and practical value of future versions of the system.

These open-ended questions served as a complement to the quantitative scales, allowing participants to highlight aspects that the previous scales might not have captured. This qualitative feedback aimed to identify unexpected issues or positive aspects, and opportunities for future refinement. The three questions presented were:

1. "What did you like most about the system you used? You may mention more than one aspect."
2. "What did you like least about the system you used? You may mention more than one aspect."
3. "How could the system you used be changed to make it easier to use? You may mention more than one aspect."

5.6 Results

All 16 participants were able to follow the procedure steps and successfully completed the experiment, including the participant with the hand mobility limitation.

The initial study design included measuring each group's task completion times for important procedure steps (e.g., first button interaction, first activation of the Hand Menu), in an attempt to observe the differences caused by the existence of a tutorial. However, upon first contact, the participants instead engaged in spontaneous exploration of the system, with most quickly attempting to touch the interface elements directly. The "No Tutorial" group rapidly experimented with hand gesture combinations as well, discovering the Hand Menu before being prompted to search for it. When asked, several participants explained that such interactions felt intuitive, and that they tried to replicate movements inspired by science-fiction media, despite lacking real prior AR experience.

In contrast, the ray-based interaction was observed to be less intuitive. Participants, particularly in the "No Tutorial" group, required explicit instruction on how to aim and perform the pinch gesture select interface elements, and even among tutorial-trained participants some required additional advice. Nonetheless, once accustomed to it, three participants reported preferring the ray-based tool, describing it as more responsive than direct fingertip interaction, and three others considered it a potentially useful interaction alternative given further practice.

Due to the participants' spontaneous exploration behavior, a more flexible testing approach ended up being used, with some freedom being given on the order and duration of feature exploration in each use case, while still maintaining the core steps of the designed procedure. This approach encouraged users, and for the duration they felt necessary to fully explore or get used to that feature. This approach brought the unexpected benefit of encouraging users to comment and give feedback on the system and its features during the active experiment. With the respective authorization by the participants, this feedback was documented and appended to their responses to the open-ended questions in the final section of the system evaluation questionnaire.

Once all participants had tested the system, the responses to the evaluation questionnaire were aggregated and analyzed using the Python programming language (version 3.12.4) and some of its core libraries, such as NumPy and Pandas for data manipulation, and Matplotlib and Seaborn for chart generation.

The following subsections present the aggregated questionnaire results, analyzed across the three scales utilized during the evaluation questionnaire definition (see section 5.5), as well as the most common feedback given in the open-ended questions.

5.6.1 System Usability

Each participant's SUS score was calculated using the standard equation defined in section 5.5.1. These scores were then analyzed for each participant group and overall, the results of which are presented in Table 17 and Figure 26.

Table 17 – SUS scores by participant group, and overall. Scores are presented as mean ± standard deviation.

	Per Participant Group			**Overall**
	No Tutorial (n=5)	**Tutorial** (n=5)	**Medical** (n=6)	(n=16)
SUS score	82.0 ± 12.2	86.5 ± 8.8	79.6 ± 10.6	**82.5 ± 10.3**

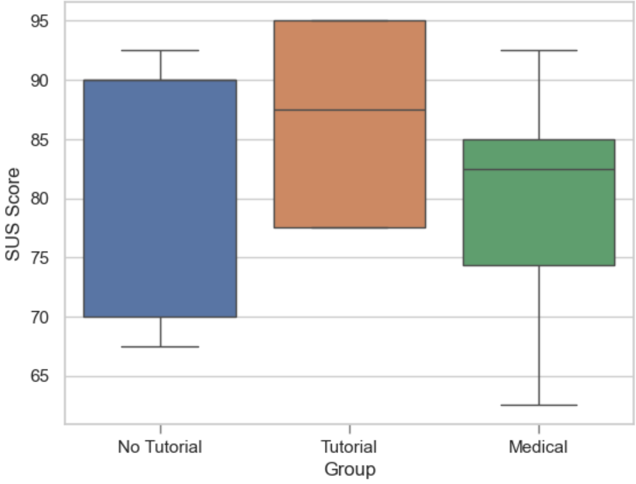

Figure 26 – SUS score plots by participant group.

System usability was rated positively by all groups and received an overall mean score of 82.5/100. The "Tutorial" group reported the highest score, reaching an average value of 86.5, while the "Medical" group rated a slightly lower value of 79.6. To better interpret these results, they were compared against established benchmarks in literature.

An article by Bangor et al. [69] used empirical evidence from various studies to assign adjective ratings along ranges of the SUS score scale. According to this work, the AleRa system obtained an overall rating of "Good", with the "Tutorial" group's score reaching a rating of "Excellent", defined as being above 85/100.

Similarly, in a more recent work, Lewis and Sauro [70] report using data from 241 studies to create a curved letter-grading scale, in which a SUS score of 68 is at the center of the range, corresponding to a "C." Using this benchmark, the AleRa system obtained an overall grading of A, with the "Tutorial" grading reaching A+.

These results suggest that the overall system attained its goal of minding usability, and that exposure to a tutorial generally improved perceived usability. Nonetheless, participants in the "Medical" group appeared slightly less satisfied with the system, possibly reflecting their higher expectations for tool quality and usability when dealing with sensible medical fields.

5.6.2 Perceived Workload

Table 18 and Figure 27 show the workload scores across the six NASA-TLX dimensions and the overall workload index. As mentioned previously, the performance metric was reinverted to become equal to the standard, so a low score near 0 means that the participants felt they had a good performance. Aside from the six dimensions, a general workload score was also calculated utilizing an established simplified alternative to the default NASA-TLX calculation, according to a study by Hart et al. [71], referred to as Raw-TLX. While in the original NASA-TLX the General Workload final score calculation requires asking the participants additional questions about which dimensions they perceive as being more important, and then doing a weighted average, Raw-TLX considers all dimensions equally important to prevent bias, and so the General Workload score is calculated using a simple average. Hart et al. [71] observed that Raw-TLX seems to be as sensitive as the default NASA-TLX calculation, and Said et al. [72] have more recently validated its use to assess perceived workload in patient monitoring tasks.

Table 18 – Raw-TLX workload scores by participant group, and overall. Scores are presented as mean ± standard deviation.

Dimension	**Per Participant Group** **No Tutorial** (n=5)	 **Tutorial** (n=5)	 **Medical** (n=6)	**Overall** (n=16)
Physical Demand	21.0 ± 27.5	7.0 ± 8.4	2.5 ± 6.1	9.7 ± 17.3
Mental Demand	12.0 ± 5.7	22.0 ± 22.0	14.2 ± 22.9	15.9 ± 18.2
Temporal Demand	8.0 ± 17.9	15.0 ± 30.8	4.2 ± 8.0	8.8 ± 19.5
Performance	8.0 ± 11.0	11.0 ± 11.4	27.5 ± 11.3	16.2 ± 13.8
Effort	23.0 ± 29.9	19.0 ± 14.7	25.0 ± 32.7	22.5 ± 25.7
Frustration	21.0 ± 30.1	13.0 ± 26.4	21.7 ± 26.2	18.8 ± 25.9
General Workload score	**15.5 ± 11.5**	**14.5 ± 10.2**	**15.8 ± 10.4**	**15.3 ± 10.0**

In general, workload was perceived as low, with all groups remaining below 30 in all six dimensions. Differences were nevertheless observed between groups: participants without a tutorial reported higher physical demand, effort, and frustration, consistent with the need to discover interaction methods through trial and error. Mental demand was slightly higher in the tutorial group compared to the no tutorial group, though lower in the medical group despite their tutorial exposure. Frustration was also elevated in the medical group, which may reflect unmet expectations about clinical features not available in the prototype. Across all groups, the overall workload index was low (15–20), indicating that use of the system was not perceived as demanding.

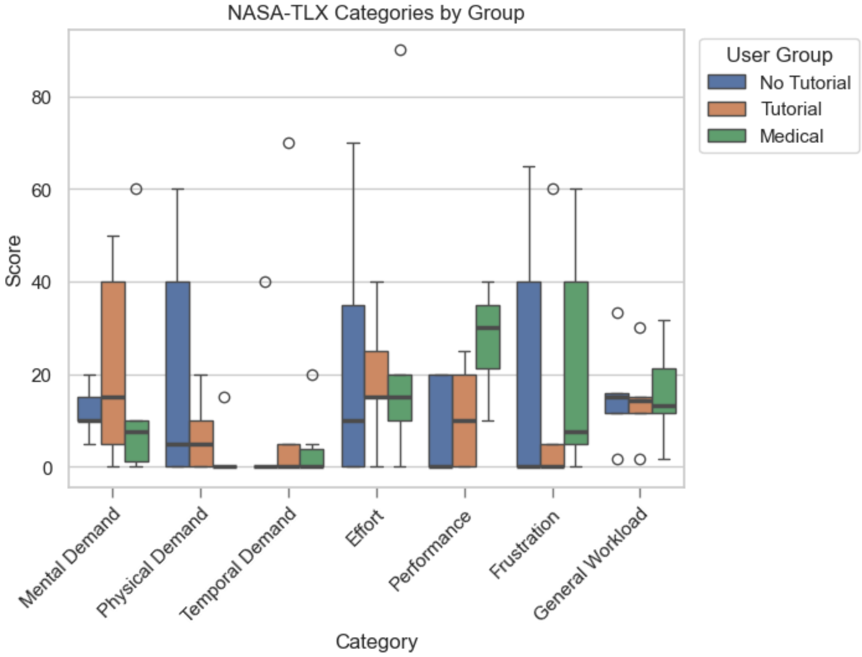

Figure 27 – Raw-TLX score plots, by dimension and participant group.

5.6.3 Quantitative Feedback (Likert-scale questions)

Responses to the seven custom Likert-scale questions, used in addition to the two previously described established scales, were also gathered and analyzed in a similar fashion, with Table 19 and Figure 28 illustrating the score's distribution, averages and standard deviations, per group, and overall. However, in contrast with the established scales, the focus here was not no in deriving a final score, but to examine the quantitative feedback for each question.

Table 19 – Custom Likert-scale questions' scores by participant group, and overall. Scores are presented as mean ± standard deviation.

Question ID	**Per Participant Group**			**Overall**
	No Tutorial	**Tutorial**	**Medical**	
	(n=5)	(n=5)	(n=6)	(n=16)
Design & Layout	4.6 ± 0.5	4.6 ± 0.5	4.3 ± 0.5	**4.5 ± 0.5**
Interactivity	3.4 ± 0.9	3.2 ± 1.3	4.0 ± 0.6	**3.6 ± 1.0**
Tutorial Usefulness	-	4.8 ± 0.4	4.8 ± 0.4	**4.8 ± 0.4**
Health Info (General Ward)	4.0 ± 1.4	4.6 ± 0.5	4.3 ± 0.8	**4.3 ± 0.9**
Health Info (Domiciliary Care)	4.2 ± 1.3	4.6 ± 0.5	4.7 ± 0.5	**4.5 ± 0.8**
AR Benefit (Nursing)	4.8 ± 0.4	4.8 ± 0.4	4.7 ± 0.8	**4.8 ± 0.6**
AR Benefit (Participant Field)	4.4 ± 0.9	4.4 ± 0.9	4.7 ± 0.8	**4.5 ± 0.8**

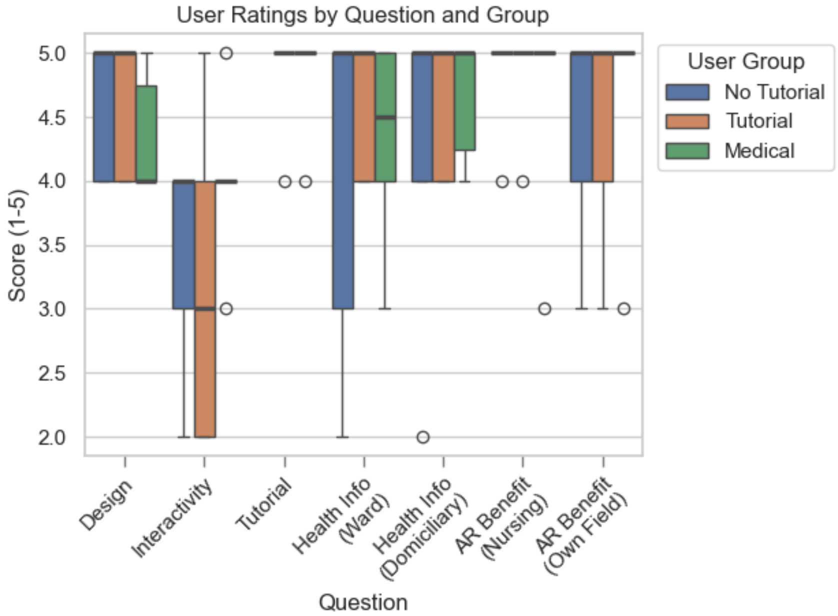

Figure 28 – Custom Likert-scale questions' score plots, by question and participant group.

The overall score for both the design and layout of the system was 4.5 (in a 1-5 scale), and the two groups that were trained with the tutorial rated it a great experience, receiving a 4.8 overall average score.

Focusing on the scores form the "Medical" group, as these participants have a better knowledge of health-related domains, both use cases defined received high scores in terms of the relevance of information displayed to the nurses. This is especially true for the more complex and active-interaction focused use case, the one targeting domiciliary care nurses, which received a score of 4.7, while the use case for general ward / IMCU nurses received a slightly lower score of 4.3.

Additionally, all groups recognized the benefits of AR systems for the nursing use cases defined, and the participants with a more engineering-centered background, those in the "No Tutorial" and "Tutorial" groups, also agreed that it could benefit their own fields of expertise.

Only the Interactivity aspect of the system received notably less positive feedback, although still obtaining an overall average score of 3.6. Surprisingly, here the "Medical" group reported the best feedback out of the three groups, in contrast with the usability scores obtained from the SUS scale, suggesting that the usability of the AleRa system suffers more from other aspects, such as complexity or the amount of knowledge needed to learn before being able to use the system confidently.

5.6.4 Qualitative Feedback (Open-ended questions)

Qualitative feedback was collected both from the open-ended questions in the evaluation questionnaire and from transcribed comments made during the case study session. These responses were then analyzed thematically to identify the AleRa system's strengths, weaknesses, and areas for potential improvement in future work.

Positive feedback

Overall, participants reported a good experience with the system. The most frequently highlighted aspects were its intuitiveness, simplicity of use, visual quality, and strong potential as a helping tool/interface for caregivers. In terms of functionalities, the most valued aspects were the ability to easily access real-time patient data, to do so near the patient rather than being restricted to a computer, and to receive alerts from multiple patients even when not focusing on them directly. The tutorial was regarded as being clear and effective, and the sound feedback when hovering over or selecting UI elements was appreciated as a clear form of receiving interaction confirmation.

Negative feedback

Regarding the HMD device, participants disliked the weight of the HMD, which can become uncomfortable after a while, the distortion / blurring of the real-world environment due to the passthrough cameras, and the added discomfort for users wearing prescription glasses. These issues are device specific and cannot be solved by software alone, but may be addressed by more advanced AR devices, such as the recently announced Meta Orion glasses [73]. In terms of functionalities and the system itself, the main topic was the fixed position of the windows. Some users liked where they were, while others reported preferring that the windows should be other specific places or that they wanted to be able to move and customize the the windows' positions manually, either due to personal preference, or functionally, for example to not be in front of patients in the real world. Another negative aspect reported was difficulty in interacting with a few UI elements, namely scroll bars, and some of the buttons in the "near" windows. Finally, the use of the mobile phone for constant QR code reading was also reported as a negative aspect of the system, while for one time authentication it was seen as tolerable, but better if avoidable, since due to the previously reported distorting of the real world, it was hard to operate with the mobile phone.

Suggestions for improvement.

Regarding potential improvement points, many mirrored from the negative feedback, such as upgrading to a lighter AR HMD, allowing to customize the window layouts (draggable windows) and user-following speed, improving the responsiveness of certain UI elements, particularly scrollbars and a few buttons, and changing patient selection feature so that it does not require the companion mobile app.

Still, participants also suggested various new or enhanced features for the system, including:

- Allow to close windows with a hand-swipe gesture, for faster decluttering of the user's vision.
- Add up and down arrow buttons alongside the existing draggable scrollbars, mimicking familiar UIs, for better list navigation.

- Introduce a Global Positioning System (GPS), or an alternative, to track patient locations inside the hospital.
- If possible, include haptic feedback (vibrations) when interacting with UI elements.

The alert system was highlighted as one of the most promising features, and so it received various suggestions, such as:

- Change the alert sounds to more closely resemble traditional hospital monitor tones, for better familiarity and thus faster recognition and action.
- Include the image of the patient and a severity symbol to each alert, to improve interpretability.
- Display a brief notification in front of the caregiver when an alert is received, even if the hand menu is not open.
- Organize alerts into collapsible groups per patient.

Similarly, the medication intake window also received considerable attention, with participants suggesting the following:

- Add navigation buttons to switch between different weeks, instead of only displaying the most recent one.
- Allow simple management tasks (e.g. adjusting dosage quantities) directly in the AR interface, while still reserving more complex tasks for a traditional computer UI.
- If possible, change the intake status to be per medication, instead of being global per time slot. Otherwise, improve clarity that it is global for that slot.
- Show the number of items in a time slot, instead of a partial list of names, to summarize the window’s content.
- Add access to additional details of each medication, such as its leaflet and administration instructions.

In summary, the detailed feedback obtained from the open-ended questions offered highly valuable insights into the system’s current strengths and limitations, providing a roadmap of potential future improvements.

5.7 Summary of Insights Obtained

This case study showed how the AleRa system achieved its primary goal of delivering a usable and low-workload AR interface for remote patient monitoring. With an average SUS score of 82.5, corresponding to an “A” grade and a “Good” rating in terms of usability benchmarks, and Raw-TLX workload indices consistently below 30, the system was perceived by the participants as intuitive and suitable for the defined nursing use cases.

Intuitiveness, clarity of visualization, and on-site mobility were recognized across groups as major advantages, while negative aspects focused primarily on existing technology constraints (e.g.: HMD weight, passthrough real-world distortion) and on a few responsiveness difficulties. The tutorial improved usability perceptions, in terms of physical demand and effort to use the system, yet also slightly raised mental demand, suggesting that training prior to using the system is beneficial but should be handled carefully to reduce cognitive workload.

Participants showed a strong interest in further development of the system in general, and particularly of both the alerts and the medication intake features, offering suggestions such as patient photos in alerts, draggable windows, per-medication intake status, and support for gesture or haptic interactions.

Overall, this case study validated AleRa’s feasibility as a supporting tool for caregivers, while at the same time gathering ideas on features that could be added, or extended upon, in future iterations.

6 Conclusions

Advances in sensor technology and IoT have made remote patient monitoring an increasingly viable approach to support the evolving needs of healthcare facilities, allowing, for instance, the remote monitoring of vital signs, health anomalies, medication intake. Projects such as RM4Health build on these possibilities, for instance in its Portuguese use case, which focuses on the creation of an intelligent health monitoring ecosystem for seniors who live in residential homes or in domiciliary care.

Still, while such systems can reduce the need for constant in-person supervision, they often require caregivers to remain near a computer to track patient data. This creates a challenge, since caregiving is fundamentally a hands-on role, and caregivers require freedom to move around within their facilities.

AR technologies present a promising solution to this problem. Current literature shows active exploration of AR in diverse fields, such as Industry 4.0 [6], [7] and healthcare [8]. Yet, according to the literature review conducted, its potential for use in patient monitoring systems remains relatively underexplored.

In response, this dissertation project proposed and developed the AleRa system, aiming to verify that using AR to provide context-sensitive overlays of patient information directly in a caregiver’s field of vision could enable continuous unobtrusive monitoring while allowing caregivers to remain actively engaged in direct patient care. The system was designed and implemented as two applications:

- An HMD AR application, that translates familiar healthcare interfaces into the real-world environment, including features such as real-time and historical vital sign visualization, alerts for abnormal conditions, and medication intake tracking.
- A companion mobile application, synchronized with the main app through Websockets, that provides a few essential features, namely user authentication and QR code reading for patient selection, to compensate for limitations in current HMD technologies.

The dissertation focused on the design and development of intuitive AR interfaces for caregivers, that could also serve as a foundation for future applications in healthcare and other fields of study. For this, the system was implemented while adhering to good architectural and coding practices to maximize extensibility. The main features developed that are believed to contribute to the field of AR include:

- A window-based interface with modular windows that follow the user as they move, providing a familiar interaction system to new AR users (see section 4.2.1).
- A window displaying line charts of real-time vital sign data, namely heart rate, temperature, respiratory rate, and blood pressure (see section 4.2.6), fetched from a HL7 FHIR standardized server (see section 4.2.4).
- A dedicated hand attached window, that provides the user with quick access to necessary functions. In this case, visualization of alerts regarding abnormal status of the patient's vital signs (see section 4.2.6), and control over the currently visible windows (see section 4.2.1).
- A window presenting line charts of historical vital signs data, for the four previously mentioned variables, combined with scatterplots for visualization of alerts (see section 4.2.6).
- Two windows showing, respectively, a patient's weekly medication intake schedule, and details about the status of the medication in each time slot (see section 4.2.7).
- A dynamic and interactive tutorial with step-by-step user progress checking, audio narrations and motivating cues, and 3D visualizations and animations, that allows to train caregivers, and users in general, on the use of an head-mounted AR system (see section 4.2.10).

Additional proof-of-concept features were also prototyped, such as a window detailing medication administration information (see section 4.2.7 and Figure 10), voice recognition for text input (see section 4.2.8), and face recognition for patient selection (see section 4.2.9), which, although not tested in the case study due to ethical or privacy constraints, demonstrate promising directions for future research.

In addition to the system development, a case study was conducted to validate its intuitiveness, usability, and usefulness in two specified use cases: The first for nurses assigned to the general ward or IMCU in a hospital, who require continuous patient information visualization, and the second for domiciliary nurses visiting a patient at their home, who need active interaction with the interfaces to access more detailed information. Sixteen participants completed the experiment and praised the benefits of being able to access patient data without being tied to a computer, the intuitiveness of the interface, and the high quality of the layout. The also highlighted this system's potential applicability not only for the prepared health use cases tested, but also in their own fields of expertise.

Overall, this dissertation project achieved its stated objects, demonstrating that AR can be a valuable tool for caregivers joining both the advantages of digital patient monitoring and hands-on care. The AleRa system addressed all set-out requirements, and contributed relevant insights to the specialization area of “Games, Graphics and Interactive Systems” of the Master in Informatics Engineering at ISEP. Finally, the development of this dissertation, along with collaboration in its related projects, also generated contributions to the academic and scientific community, detailed in the next section.

6.1 Scientific Contributions

As this dissertation project was carried out at the GECAD research center, one additional project objective was that the developed work should contribute to the submission of one or more scientific publications to international conferences and journals, indexed in Scopus and Web of Science, related to the research areas of health monitoring and/or augmented reality.

The development of this dissertation project was connected to two international ITEA projects in the healthcare sector, namely Inno4Health, and its successor RM4Health. Collaboration within these two projects were conducted through research grants at GECAD during the three years of the Master in Informatics Engineering, and the resulting artifacts provided the foundation for the medical data logic behind this dissertation.

Inno4Health[9] – Contributions made included protecting the medical data stored in a deployed HAPI FHIR server instance, by developing an Authentication Server for JWT bearer token generation, and management and optimization of the FHIR server itself. Both of these servers were later expanded upon in the successor project RM4Health and eventually utilized by the AleRa system.

RM4Health[10] – Contributions made included the deployment and management of a more recent version of the HAPI FHIR server, as well as the design and implementation, together with other colleagues, of the complete system architecture, seen in Figure 1. Work done on RM4Health also involved defining the content and structure, according to FHIR resource standards, of all patient information needed to be stored. Much of this information model, particularly the vital signs, alerts, and medication information, were later the data considered available to retrieve and display in the AleRa system.

The AleRa system then emerged from the ideas explored in these two projects, by adding a novel AR-based interface that allows caregivers to visualize medical data by integrating with the FHIR-based storage server and the Authentication Server.

Participation and collaboration within the Inno4Health and RM4Health research projects, resulted in the publication of scientific articles that influenced the AleRa system’s logic and

---

[9] The Inno4Health project is presented on the ITEA portal [74] and on its official website [75].

[10] The RM4Health project is presented on the ITEA portal [2] and on its official website [1].

development process, namely 3 articles published to conferences and 1 article submitted to a journal:

- **A Remote Monitoring Platform for the Management of Lower Limb Vascular Diseases**
  Julio Souza, Ana Vieira, Luís Conceição, Rafael Martins, Daniel Rodrigues, Gustavo Corrente, William Xavier, Sérgio Sampaio, Goreti Marreiros, and Alberto Freitas.
  "A Remote Monitoring Platform for the Management of Lower Limb Vascular Diseases," Stud. Health Technol. Inform., vol. 302, pp. 1013–1014, May 2023, doi: 10.3233/SHTI230330.

  This short conference paper describes fundamental components of Inno4Health's system architecture, focusing on its data storage server, a rule-based coaching system for patients, and an alert system that notifies physicians whenever vital sign threshold values are crossed. Although not explicitly stated on this paper, the data storage server was FHIR-compliant, and was the first iteration of the data storage server used by AleRa to fetch clinical data. Meanwhile, the alert system described in this short paper served as direct inspiration to the alert system developed for AleRa (see section 4.2.6). The coaching system was not utilized during AleRa's development in this dissertation project, but since its generated recommendations are also stored in the FHIR server, they could be easily fetched and displayed in AleRa's interfaces in future iterations.

- **Supporting Elderly Care Through an AI-Driven and FHIR-Based Remote Monitoring System**
  Rafael Martins, Hugo Pereira, Gustavo Corrente, William Xavier, Luís Conceição, Alberto Freitas, and Goreti Marreiros.
  "Supporting Elderly Care Through an AI-Driven and FHIR-Based Remote Monitoring System" in Distributed Computing and Artificial Intelligence, Special Sessions II, 21st International Conference, G. Marreiros, L. Grande, J. P. Llerena, L. Conceição, H. Ko, M. Plaza, and M. Ricca, Eds., Cham: Springer Nature Switzerland, 2025, pp. 233–241. doi: 10.1007/978-3-031-80946-0_22.

  This conference article describes RM4Health's initially proposed system architecture, including both the FHIR-compliant data storage server and authentication server that were eventually developed and deployed, and with which AleRa was designed to be able to communicate to obtain all necessary clinical data (see sections 3.4 and 4.2.4). Furthermore, it defines a third server in RM4Health's architecture capable of deriving additional critical information from sensor data, using internally developed AI models, such as medication intake, sleep patterns, and fall detection. Since this newly derived information is also stored in the FHIR server, it can be easily integrated into AleRa's interfaces, as was the case with the medication intake (see section 4.2.7).

- **Enhancing Medication Adherence with Computer Vision: Object Detection Models for Pill Detection**
  Gabriel Pinto, Rafael Martins, Hugo Pereira, Rita Ribeiro, Luís Conceição, and Goreti Marreiros.
  "Enhancing Medication Adherence with Computer Vision: Object Detection Models for Pill Detection" in Emerging Trends in Information Systems and Technologies, Á. Rocha, H. Adeli, A. Poniszewska-Maranda, F. Moreira, I. Bianchi, Eds., Cham, Springer Nature Switzerland, 2025, ch.33, doi: 10.1007/978-3-031-97119-8_33.

  This conference paper details work done on training an object detection computer vision model capable of detecting if and where pills are present in an image. While the version in this paper utilized a public dataset of pills in a person's hand, this model was later refined, during the RM4Health project, to utilize images of a rectangular weekly medication blister, taken by a camera mounted on a medication device, which then allows to monitor medication intake by comparing pill positions to the patient's weekly prescribed medication. The insights derived from this model were directly integrated into AleRa's interfaces, which show a colored grid interface of prescribed medications and their intake status, and even the raw blister images taken by the medication device, to allow manual confirmation of the model insights by the caregivers (see section 4.2.7).

- **A Systematic Review on Wearable-enabled Remote Health Monitoring**
  Rita Ribeiro, Rafael Martins, Hugo Pereira, Vítor Crista, Júlio Souza, Rute Almeida, Diogo Martinho, Luís Conceição, Alberto Freitas, and Goreti Marreiros.
  "A Systematic Review on Wearable-enabled Remote Health Monitoring", submitted to Digital Health.

  This journal article consists of a systematic review done on the topic of remote health monitoring, as part of RM4Health's project development. Based on the main research question "What is the current state-of-the-art regarding the use of wearable technologies for remote health monitoring?", articles were identified, screened, and included following the PRISMA guidelines, and the resulting papers were reviewed and had their main contributions to literature extracted and synthesized. Participation in this larger-scale systematic review positively influenced the literature review carried out during this dissertation (see section 2.1), as it provided experience on good practices to follow and on the operation of the PRISMA guidelines. In turn, the literature review done during this dissertation contributed to the design of various AleRa components.

Furthermore, collaboration and work done within the Inno4Health and RM4Health research projects also resulted in the publication of other scientific articles, which although their topics did not directly influence the development of the current AleRa system, have potential for integration in future iterations. These papers are comprised of 2 articles published and presented in conferences and 1 article submitted to a journal:

- **Multi-class Model to Predict Pain on Lower Limb Intermittent Claudication Patients**
  Rafael Martins, Luís Conceição, Gustavo Corrente, William Xavier, Júlio Souza, Alberto Freitas, and Goreti Marreiros.
  “Multi-class Model to Predict Pain on Lower Limb Intermittent Claudication Patients” in Good Practices and New Perspectives in Information Systems and Technologies, Á. Rocha, H. Adeli, G. Dzemyda, F. Moreira, and A. Poniszewska-Marańda, Eds., Cham: Springer Nature Switzerland, 2024, pp. 44–53. doi: 10.1007/978-3-031-60218-4_5.

  This conference article describes the development, training, and comparation of four multi-class classification machine learning (ML) models capable of detecting the existence and, depending on the model, the location of pain in the legs of patients diagnosed with intermediate claudication. Alerts generated from the output of these models could also be integrated into the AleRa alerts feature, in future iterations.

- **Optimizing AI Models for Fall Detection on Resource-Constrained Embedded Systems**
  Francisco Loureiro, Gabriel Pinto, Rafael Martins, William Xavier, Gustavo Corrente, Luís Conceição, and Goreti Marreiros.
  “Optimizing AI Models for Fall Detection on Resource-Constrained Embedded Systems”, in Distributed Computing and Artificial Intelligence, Special Sessions, 22nd International Conference, Lecture Notes in Networks and Systems, Cham: Springer Nature Switzerland, 2025, in press.

  This conference article details work done on the development of artificial neural network models for detecting falls using triaxial accelerometer and gyroscope signal data, and more importantly, the size optimization of these models to be able to embed the final version in a wearable device, allowing fall detection and alerting with low latency. Although the final model is expected to be embedded in a wearable device, the fall alerts are still planned to be stored in RM4Health’s FHIR-compliant data storage server, and as such could be fetched by and integrated into AleRa’s interfaces in future iterations.

- **Comprehensive Analysis of Machine Learning Models for Five-Class Sleep Stage Classification Using PPG Signals**
  Rafael Martins, Rita Ribeiro, Hugo Pereira, Vasco Silva, Alberto Freitas, Rute Almeida, Goreti Marreiros, and Luís Conceição.
  “Comprehensive Analysis of Machine Learning Models for Five-Class Sleep Stage Classification Using PPG Signals”, submitted to Scientific Reports.

  This journal article discusses the signal processing of photoplethysmography data and model architecture implemented to construct ML models capable of classifying sleep into five stages, well-established in literature: awake, light sleep (stages 1 and 2), deep sleep, and rapid-eye movement. The outputs of these models are expected to be stored in RM4Health's data storage server, and as such could easily be integrated into AleRa's interfaces in future work.

6.2 Limitations and Future Work

The developed AleRa system achieved its set-out objectives and contributed positively to scientific community. Nonetheless, during the system's design, implementation, and evaluation, several limitations were identified, many of which were reiterated by the participants during the case study. These limitations suggest important opportunities for future work both In research and in technical development.

6.2.1 Hardware Limitations

A major limitation found in current HMDs is that most vendors deny or heavily restrict developer access to the device's cameras, which limits the development of context-awareness and image-processing features that could significantly enhance AR experiences. More specific to the Meta Quest 3 HMD, it was observed that the passthrough AR functionality causes distortion and blurring of the real world, which hinders smooth interaction with physical objects. This in turn contributes to other limitations, such as the negative experience participants felt when having to operate the secondary mobile app for certain functionality. Although only observed on the Quest 3, this limitation is assumed to also affect other vendor devices to some extent. Additionally, another limitation detected on the Quest 3, but predicted to be a more general problem, is the considerable weight of the device, which creates discomfort in users after extended use.

Still, the AR field is rapidly evolving. Meta has recently announced both developer access to the Quest 3 cameras [41], and their future product, Meta Orion [73], a new lightweight AR device. Future work on AleRa or similar projects could therefore revisit this system with these new tools and explore how these new functionalities affect usability and feasibility.

6.2.2 Existing Software Limitations

Strong efforts have been made by vendors to provide compatibility with open standards such as OpenXR and the XR Interaction Toolkit. However, these standards mainly provide low-level development tools, with commonly expected interface elements such as windows or menus having to be developed on the spot, which limits the rate at which AR applications for specific use cases can be created.

The AleRa system aimed in part to contribute to this direction by implementing a window-based interface with modularity and reusability in mind. Future work could extend this approach into a completely modular and context independent framework, that could be used as the foundation for multiple future AR applications in healthcare and other fields.

6.2.3 Developed System Limitations

Most negative aspects noted during development and by case study participants were tied to the existing hardware and software limitations, as described in the previous sections, and cannot be fixed by improvements to AleRa alone. Still, a few negative aspects that could be alleviated in future iterations are the responsiveness of some UI elements, and a shift to less reliance on the companion mobile app.

Addressing these issues in future work, in AleRa or similar projects, would further improve AR’s viability in caregiving practice.

6.2.4 Feature Extensions

Finally, with the case study feedback from the participants, various ideas for future functionality were gathered. For instance, a participant with health expertise suggested a patient location tracking functionality within hospitals or other healthcare facilities, while others requested an additional window to be added to the medication intake component, that would allow caregivers to search and visualize detailed medication information such as administration instructions. Additionally, other features from the health projects Inno4Health and RM4Health, such as coaching system recommendations, sleep pattern displaying and fall detection alerts could also be added to the AleRa system in future iterations.

Incorporating such features in future work would expand the system’s framework applicability in different use cases, to further validate AR’s potential in healthcare contexts.

Bibliography

[1] "Homepage - rm4health." Accessed: Sept. 19, 2025. [Online]. Available: https://rm4health.eu/

[2] "ITEA 4 · Project · 21022 RM4HEALTH," itea4.org. Accessed: Sept. 19, 2025. [Online]. Available: https://itea4.org/project/rm4health.html

[3] R. Martins, H. Pereira, G. Corrente, W. Xavier, L. Conceição, A. Freitas, and G. Marreiros, "Supporting Elderly Care Through an AI-Driven and FHIR-Based Remote Monitoring System," in *Distributed Computing and Artificial Intelligence, Special Sessions II, 21st International Conference*, G. Marreiros, L. Grande, J. P. Llerena, L. Conceição, H. Ko, M. Plaza, and M. Ricca, Eds., Cham: Springer Nature Switzerland, 2025, pp. 233–241. doi: 10.1007/978-3-031-80946-0_22.

[4] D. Bender and K. Sartipi, "HL7 FHIR: An Agile and RESTful approach to healthcare information exchange," in *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems*, Porto, Portugal: IEEE, June 2013, pp. 326–331. doi: 10.1109/CBMS.2013.6627810.

[5] 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 Med. Inform.*, vol. 9, no. 7, p. e21929, July 2021, doi: 10.2196/21929.

[6] A. Canito, D. Mota, G. Marreiros, J. M. Corchado, and C. Martins, "Contextual Adaptative Interfaces for Industry 4.0," in *Distributed Computing and Artificial Intelligence, Special Sessions, 17th International Conference*, S. Rodríguez González, A. González-Briones, A. Gola, G. Katranas, M. Ricca, R. Loukanova, and J. Prieto, Eds., Cham: Springer International Publishing, 2021, pp. 149–157. doi: 10.1007/978-3-030-53829-3_14.

[7] G. Marreiros, P. Craamer, I. Garitano, R. González, M. G. Novak, A. Kancilija, Q. Marcelis, D. Martinho, A. Niemelä, and F. Novak, "From KPI Dashboards to Advanced Visualization," in *The MANTIS Book*, River Publishers, 2018. doi: 10.1201/9781003339748-6.

[8] D. Martinho, J. Carneiro, J. M. Corchado, and G. Marreiros, "A systematic review of gamification techniques applied to elderly care," *Artif. Intell. Rev.*, vol. 53, no. 7, pp. 4863–4901, Oct. 2020, doi: 10.1007/s10462-020-09809-6.

[9] B. Kitchenham, "Procedures for Performing Systematic Reviews," *Keele UK Keele Univ*, vol. 33, Aug. 2004.

[10] "About the PICOC," Parsifal. Accessed: Jan. 03, 2025. [Online]. Available: https://parsif.al/help/about-the-picoc/

[11] "What is a PICOC? » CEBMa." Accessed: Jan. 03, 2025. [Online]. Available: https://cebma.org/resources/frequently-asked-questions/what-is-a-picoc/

[12] A. Liberati, D. G. Altman, J. Tetzlaff, C. Mulrow, P. C. Gøtzsche, J. P. A. Ioannidis, M. Clarke, P. J. Devereaux, J. Kleijnen, and D. Moher, "The PRISMA statement for reporting systematic reviews and meta-analyses of studies that evaluate health care interventions: explanation and elaboration," *J. Clin. Epidemiol.*, vol. 62, no. 10, pp. e1-34, Oct. 2009, doi: 10.1016/j.jclinepi.2009.06.006.

[13] "PRISMA 2020 flow diagram," PRISMA statement. Accessed: Jan. 03, 2025. [Online]. Available: https://www.prisma-statement.org/prisma-2020-flow-diagram

[14] M. J. Page, J. E. McKenzie, P. M. Bossuyt, I. Boutron, T. C. Hoffmann, C. D. Mulrow, L. Shamseer, J. M. Tetzlaff, E. A. Akl, S. E. Brennan, R. Chou, J. Glanville, J. M. Grimshaw, A. Hróbjartsson, M. M. Lalu, T. Li, E. W. Loder, E. Mayo-Wilson, S. McDonald, L. A.

McGuinness, L. A. Stewart, J. Thomas, A. C. Tricco, V. A. Welch, P. Whiting, and D. Moher, “The PRISMA 2020 statement: an updated guideline for reporting systematic reviews,” *BMJ*, vol. 372, p. n71, Mar. 2021, doi: 10.1136/bmj.n71.

[15] “Cochrane Handbook - Chapter 5: Collecting data.” Accessed: Jan. 03, 2025. [Online]. Available: https://training.cochrane.org/handbook/current/chapter-05

[16] S. Kimmel, V. Cobus, and W. Heuten, “opticARe - Augmented Reality Mobile Patient Monitoring in Intensive Care Units,” in *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*, in VRST ’21. New York, NY, USA: Association for Computing Machinery, 2021. doi: 10.1145/3489849.3489852.

[17] K. Vock, S. Hubenschmid, J. Zagermann, S. Butscher, and H. Reiterer, “IDIAR: Augmented Reality Dashboards to Supervise Mobile Intervention Studies,” in *Proceedings of Mensch Und Computer 2021*, in MuC ’21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 248–259. doi: 10.1145/3473856.3473876.

[18] T. Lorenz, L. Münch, P. Adair, I. Schiering, and S. V. Müller, “Creating Routes for Landmark-Training with the CompanionApp: A Pilot User Study,” in *Proceedings of Mensch Und Computer 2023*, in MuC ’23. New York, NY, USA: Association for Computing Machinery, 2023, pp. 518–523. doi: 10.1145/3603555.3608557.

[19] L. Albright, W. Ko, M. Buvanesh, H. Haraldsson, F. Polubriaginof, G. J. Kuperman, M. Levy, M. R. Sterling, N. Dell, and D. Estrin, “Opportunities and Challenges for Augmented Reality in FamilyCaregiving:Qualitative Video Elicitation Study,” *JMIR FORMATIVE RESEARCH*, vol. 8. 2024. doi: 10.2196/56916.

[20] S. K. Tanbeer and E. R. Sykes, “*MiVitals-Mixed* Reality Interface for *Vitals* Monitoring: A HoloLens based prototype for healthcare practices,” *COMPUTATIONAL AND STRUCTURAL BIOTECHNOLOGY JOURNAL*, vol. 24. pp. 160–175, Dec. 2024. doi: 10.1016/j.csbj.2024.02.024.

[21] S. Kim, T. Jung, D. K. Sohn, Y. Chae, Y. A. Kim, S. H. Kang, Y. Park, and Y. J. Chang, “The Multidomain Metaverse Cancer Care Digital Platform: Development and Usability Study,” *JMIR SERIOUS GAMES*, vol. 11. 2023. doi: 10.2196/46242.

[22] “Start Building with Meta Quest 3 and 3S.” Accessed: Sept. 17, 2025. [Online]. Available: https://developers.meta.com/horizon/blog/start-building-meta-quest-3s-3-launch-spatial-sdk-2D-mixed-reality/

[23] “Getting Started With Focus Vision, XR Elite and Focus 3 - Developer Resources.” Accessed: Sept. 17, 2025. [Online]. Available: https://developer.vive.com/resources/getting-started-with-xr-elite/

[24] “buildwagon — HoloLens Development Platform.” Accessed: Sept. 17, 2025. [Online]. Available: https://buildwagon.com/hololens.html

[25] MSBryan, “OpenXR - Mixed Reality.” Accessed: Sept. 17, 2025. [Online]. Available: https://learn.microsoft.com/en-us/windows/mixed-reality/develop/native/openxr

[26] MSBryan, “Using WebXR with Windows Mixed Reality - Mixed Reality.” Accessed: Sept. 17, 2025. [Online]. Available: https://learn.microsoft.com/en-us/windows/mixed-reality/develop/javascript/webxr-overview

[27] “Development resources,” PICO Developer. Accessed: Sept. 17, 2025. [Online]. Available: https://developer.picoxr.com/

[28] “Intents | MagicLeap Developer Documentation.” Accessed: Sept. 17, 2025. [Online]. Available: https://developer-docs.magicleap.cloud/docs/guides/features/android-intents/android-intents-overview/

[29] “Overview | MagicLeap Developer Documentation.” Accessed: Sept. 17, 2025. [Online]. Available: https://developer-docs.magicleap.cloud/docs/releases/releases-overview/

[30] “WebXR | MagicLeap Developer Documentation.” Accessed: Sept. 17, 2025. [Online]. Available: https://developer-docs.magicleap.cloud/docs/guides/features/webxr-viewer/

[31] "AR Foundation | AR Foundation | 6.2.0." Accessed: Sept. 17, 2025. [Online]. Available: https://docs.unity3d.com/Packages/com.unity.xr.arfoundation@6.2/manual/index.html

[32] "Developing for Head-Mounted Experiences with OpenXR in Unreal Engine | Unreal Engine 5.6 Documentation | Epic Developer Community," Epic Games Developer. Accessed: Sept. 17, 2025. [Online]. Available: https://dev.epicgames.com/documentation/en-us/unreal-engine/developing-for-head-mounted-experiences-with-openxr-in-unreal-engine

[33] "Setting up XR," Godot Engine documentation. Accessed: Sept. 17, 2025. [Online]. Available: https://docs.godotengine.org/en/stable/tutorials/xr/setting_up_xr.html

[34] "Immersive Web Developer Home." Accessed: Sept. 17, 2025. [Online]. Available: https://immersiveweb.dev/

[35] "Overview of building for Android XR | Android Developers." Accessed: Sept. 17, 2025. [Online]. Available: https://developer.android.com/develop/xr/get-started

[36] "OpenXR - High-performance access to AR and VR —collectively known as XR— platforms and devices," The Khronos Group. Accessed: Sept. 22, 2025. [Online]. Available: https://www.khronos.org/OpenXR/

[37] J. L. Derby, C. T. Rarick, and B. S. Chaparro, "Text Input Performance with a Mixed Reality Head-Mounted Display (HMD)," *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, vol. 63, no. 1, pp. 1476–1480, Nov. 2019, doi: 10.1177/1071181319631279.

[38] S. K. Tanbeer and E. R. Sykes, "*MiVitals-Mixed* Reality Interface for *Vitals* Monitoring: A HoloLens based prototype for healthcare practices," *Comput. Struct. Biotechnol. J.*, vol. 24, pp. 160–175, Dec. 2024, doi: 10.1016/j.csbj.2024.02.024.

[39] W. Xu, H.-N. Liang, A. He, and Z. Wang, "Pointing and Selection Methods for Text Entry in Augmented Reality Head Mounted Displays," in *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Oct. 2019, pp. 279–288. doi: 10.1109/ISMAR.2019.00026.

[40] J. J. Dudley, K. Vertanen, and P. O. Kristensson, "Fast and Precise Touch-Based Text Entry for Head-Mounted Augmented Reality with Variable Occlusion," *ACM Trans Comput-Hum Interact*, vol. 25, no. 6, p. 30:1-30:40, dezembro 2018, doi: 10.1145/3232163.

[41] "Explore a New Era of Mixed Reality with the Passthrough Camera API." Accessed: Sept. 17, 2025. [Online]. Available: https://developers.meta.com/horizon/blog/new-era-mixed-reality-passthrough-camera-api-machine-learning-computer-vision/

[42] "Regulation - EU - 2024/1689 - EN - EUR-Lex." Accessed: Sept. 14, 2025. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

[43] "Unity 2022.3.38f1," Unity. Accessed: Sept. 09, 2025. [Online]. Available: https://unity.com/releases/editor/whats-new/2022.3.38f1

[44] "Meta Quest 3: Next-Gen Mixed Reality Headset." Accessed: Sept. 03, 2025. [Online]. Available: https://www.meta.com/quest/quest-3/

[45] U. Technologies, "Unity - Manual: Build an iOS application." Accessed: Sept. 14, 2025. [Online]. Available: https://docs.unity3d.com/2022.3/Documentation/Manual/iphone-BuildProcess.html

[46] C. P. Subbe, M. Kruger, P. Rutherford, and L. Gemmel, "Validation of a modified Early Warning Score in medical admissions," *QJM Mon. J. Assoc. Physicians*, vol. 94, no. 10, pp. 521–526, Oct. 2001, doi: 10.1093/qjmed/94.10.521.

[47] Royal College of Physicians, *National Early Warning Score (NEWS) 2: Standardising the assessment of acute-illness severity in the NHS. Updated report of a working party*. London: Royal College of Physicians, 2017. [Online]. Available: https://www.rcp.ac.uk/projects/outputs/national-early-warning-score-news-2

[48] K. Jeppestøl, M. Kirkevold, and L. K. Bragstad, “Early warning scores and trigger recommendations must be used with care in older home nursing care patients: Results from an observational study,” *Nurs. Open*, vol. 10, no. 7, pp. 4737–4746, 2023, doi: 10.1002/nop2.1724.

[49] “Prontuário Terapêutico online.” Accessed: Sept. 17, 2025. [Online]. Available: https://app10.infarmed.pt/prontuario/index.php

[50] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, “Joint Face Detection and Alignment using Multi-task Cascaded Convolutional Networks,” *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–1503, Oct. 2016, doi: 10.1109/LSP.2016.2603342.

[51] F. Schroff, D. Kalenichenko, and J. Philbin, “FaceNet: A unified embedding for face recognition and clustering,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 815–823. doi: 10.1109/CVPR.2015.7298682.

[52] “GitHub - timesler/facenet-pytorch: Pretrained Pytorch face detection (MTCNN) and facial recognition (InceptionResnet) models.” Accessed: Sept. 27, 2025. [Online]. Available: https://github.com/timesler/facenet-pytorch?tab=readme-ov-file#references

[53] J. Han, M. Kamber, and J. Pei, “2 - Getting to Know Your Data,” in *Data Mining: Concepts and Techniques (Third Edition)*, J. Han, M. Kamber, and J. Pei, Eds., in The Morgan Kaufmann Series in Data Management Systems. , Boston: Morgan Kaufmann, 2012, pp. 39–82. doi: 10.1016/B978-0-12-381479-1.00002-2.

[54] “Why You Only Need to Test with 5 Users,” Nielsen Norman Group. Accessed: Sept. 02, 2025. [Online]. Available: https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/

[55] J. Nielsen and T. K. Landauer, “A mathematical model of the finding of usability problems,” in *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI ’93*, Amsterdam, The Netherlands: ACM Press, 1993, pp. 206–213. doi: 10.1145/169059.169166.

[56] L. Faulkner, “Beyond the five-user assumption: Benefits of increased sample sizes in usability testing,” *Behav. Res. Methods Instrum. Comput.*, vol. 35, no. 3, pp. 379–383, Aug. 2003, doi: 10.3758/BF03195514.

[57] J. S. PhD, “Why you only need to test with five users (explained) – MeasuringU.” Accessed: Sept. 02, 2025. [Online]. Available: https://measuringu.com/five-users/

[58] “Meta Quest Elite Strap | Meta | Meta Store.” Accessed: Sept. 03, 2025. [Online]. Available: https://www.meta.com/quest/accessories/elite-strap/

[59] “Specs | Samsung Galaxy S23 & S23+ | Samsung US,” Samsung Electronics America. Accessed: Sept. 03, 2025. [Online]. Available: https://www.samsung.com/us/smartphones/galaxy-s23-lite-no-3-party/specs/

[60] “Buy Galaxy S23 | Certified Re-Newed 128GB Smartphone,” Samsung US. Accessed: Sept. 03, 2025. [Online]. Available: https://www.samsung.com/us/smartphones/galaxy-s-series/certified-re-newed-store/buy/galaxy-s23-certified-re-newed-128gb-sm5s911uzkaxaa/

[61] J. Brooke, “SUS – a quick and dirty usability scale,” 1996, pp. 189–194.

[62] S. Tasfia, M. N. Islam, S. A. Nusrat, and N. Jahan, “Evaluating Usability of AR-based Learning Applications for Children Using SUS and Heuristic Evaluation,” in *Proceedings of the Fourth International Conference on Trends in Computational and Cognitive Engineering*, M. S. Kaiser, S. Waheed, A. Bandyopadhyay, M. Mahmud, and K. Ray, Eds., Singapore: Springer Nature, 2023, pp. 87–98. doi: 10.1007/978-981-19-9483-8_8.

[63] R. Dutta, A. Mantri, and G. Singh, “Evaluating system usability of mobile augmented reality application for teaching Karnaugh-Maps,” *Smart Learn. Environ.*, vol. 9, no. 1, p. 6, Jan. 2022, doi: 10.1186/s40561-022-00189-8.

[64] V. E. C. Sousa and K. D. Lopez, “Towards Usable E-Health,” *Appl. Clin. Inform.*, vol. 8, no. 2, pp. 470–490, Apr. 2017, doi: 10.4338/ACI-2016-10-R-0170.

[65] S. G. Hart and L. E. Staveland, “Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research,” in *Advances in Psychology*, vol. 52, P. A. Hancock and N. Meshkati, Eds., in Human Mental Workload, vol. 52. , North-Holland, 1988, pp. 139–183. doi: 10.1016/S0166-4115(08)62386-9.

[66] N. Xi, J. Chen, F. Gama, M. Riar, and J. Hamari, “The challenges of entering the metaverse: An experiment on the effect of extended reality on workload,” *Inf. Syst. Front.*, vol. 25, no. 2, pp. 659–680, Apr. 2023, doi: 10.1007/s10796-022-10244-x.

[67] M. H. Knudsen, N. Breindahl, T.-S. Dalsgaard, D. Isbye, A. G. Mølbak, G. Tiwald, M. B. S. Svendsen, L. Konge, J. Bergström, and T. Todsen, “Using Virtual Reality Head-Mounted Displays to Assess Skills in Emergency Medicine: Validity Study,” *J. Med. Internet Res.*, vol. 25, no. 1, p. e45210, June 2023, doi: 10.2196/45210.

[68] J. Buchner, K. Buntins, and M. Kerres, “A systematic map of research characteristics in studies on augmented reality and cognitive load,” *Comput. Educ. Open*, vol. 2, p. 100036, Dec. 2021, doi: 10.1016/j.caeo.2021.100036.

[69] A. Bangor, P. Kortum, and J. Miller, “Determining what individual SUS scores mean: adding an adjective rating scale,” *J Usability Stud.*, vol. 4, no. 3, pp. 114–123, maio 2009.

[70] J. R. Lewis and J. Sauro, “Item benchmarks for the system usability scale,” *J Usability Stud.*, vol. 13, no. 3, pp. 158–167, maio 2018.

[71] S. G. Hart, “Nasa-Task Load Index (NASA-TLX); 20 Years Later,” *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, vol. 50, no. 9, pp. 904–908, Oct. 2006, doi: 10.1177/154193120605000909.

[72] S. Said, M. Gozdzik, T. R. Roche, J. Braun, J. Rössler, A. Kaserer, D. R. Spahn, C. B. Nöthiger, and D. W. Tscholl, “Validation of the Raw National Aeronautics and Space Administration Task Load Index (NASA-TLX) Questionnaire to Assess Perceived Workload in Patient Monitoring Tasks: Pooled Analysis Study Using Mixed Models,” *J. Med. Internet Res.*, vol. 22, no. 9, p. e19472, Sept. 2020, doi: 10.2196/19472.

[73] “Introducing Orion, Our First True Augmented Reality Glasses.” Accessed: Sept. 27, 2025. [Online]. Available: https://about.fb.com/news/2024/09/introducing-orion-our-first-true-augmented-reality-glasses/

[74] “ITEA 4 · Project · 19008 INNO4HEALTH,” itea4.org. Accessed: Sept. 27, 2025. [Online]. Available: https://itea4.org/project/inno4health.html

[75] “Home. - Inno4Health.” Accessed: Sept. 27, 2025. [Online]. Available: https://inno4health.eu/

Appendix A – Characterization Questionnaire

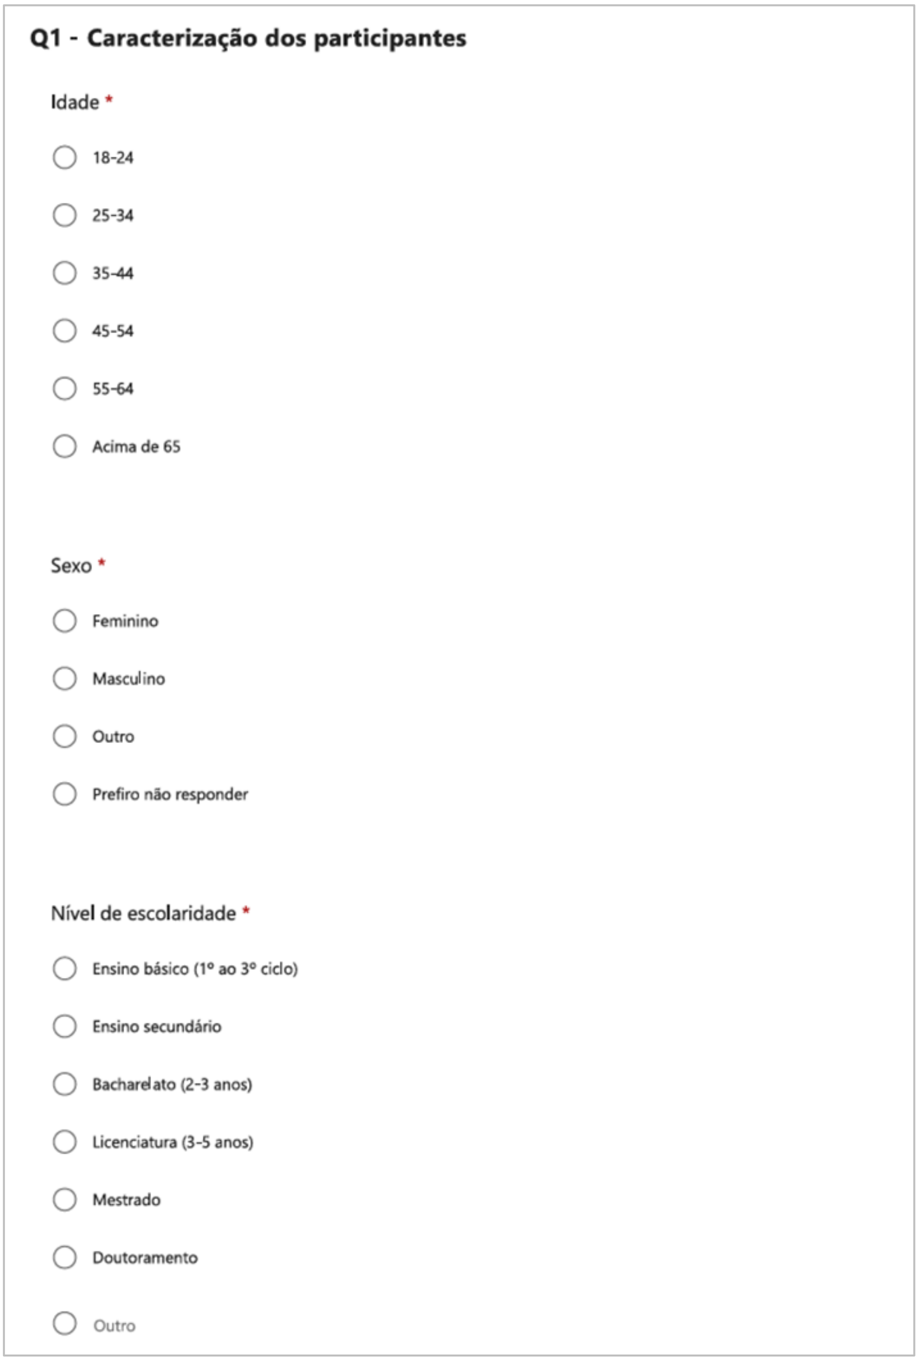

Q1 - Caracterização dos participantes

Idade *

○ 18-24
○ 25-34
○ 35-44
○ 45-54
○ 55-64
○ Acima de 65

Sexo *

○ Feminino
○ Masculino
○ Outro
○ Prefiro não responder

Nível de escolaridade *

○ Ensino básico (1º ao 3º ciclo)
○ Ensino secundário
○ Bacharelato (2-3 anos)
○ Licenciatura (3-5 anos)
○ Mestrado
○ Doutoramento
○ Outro

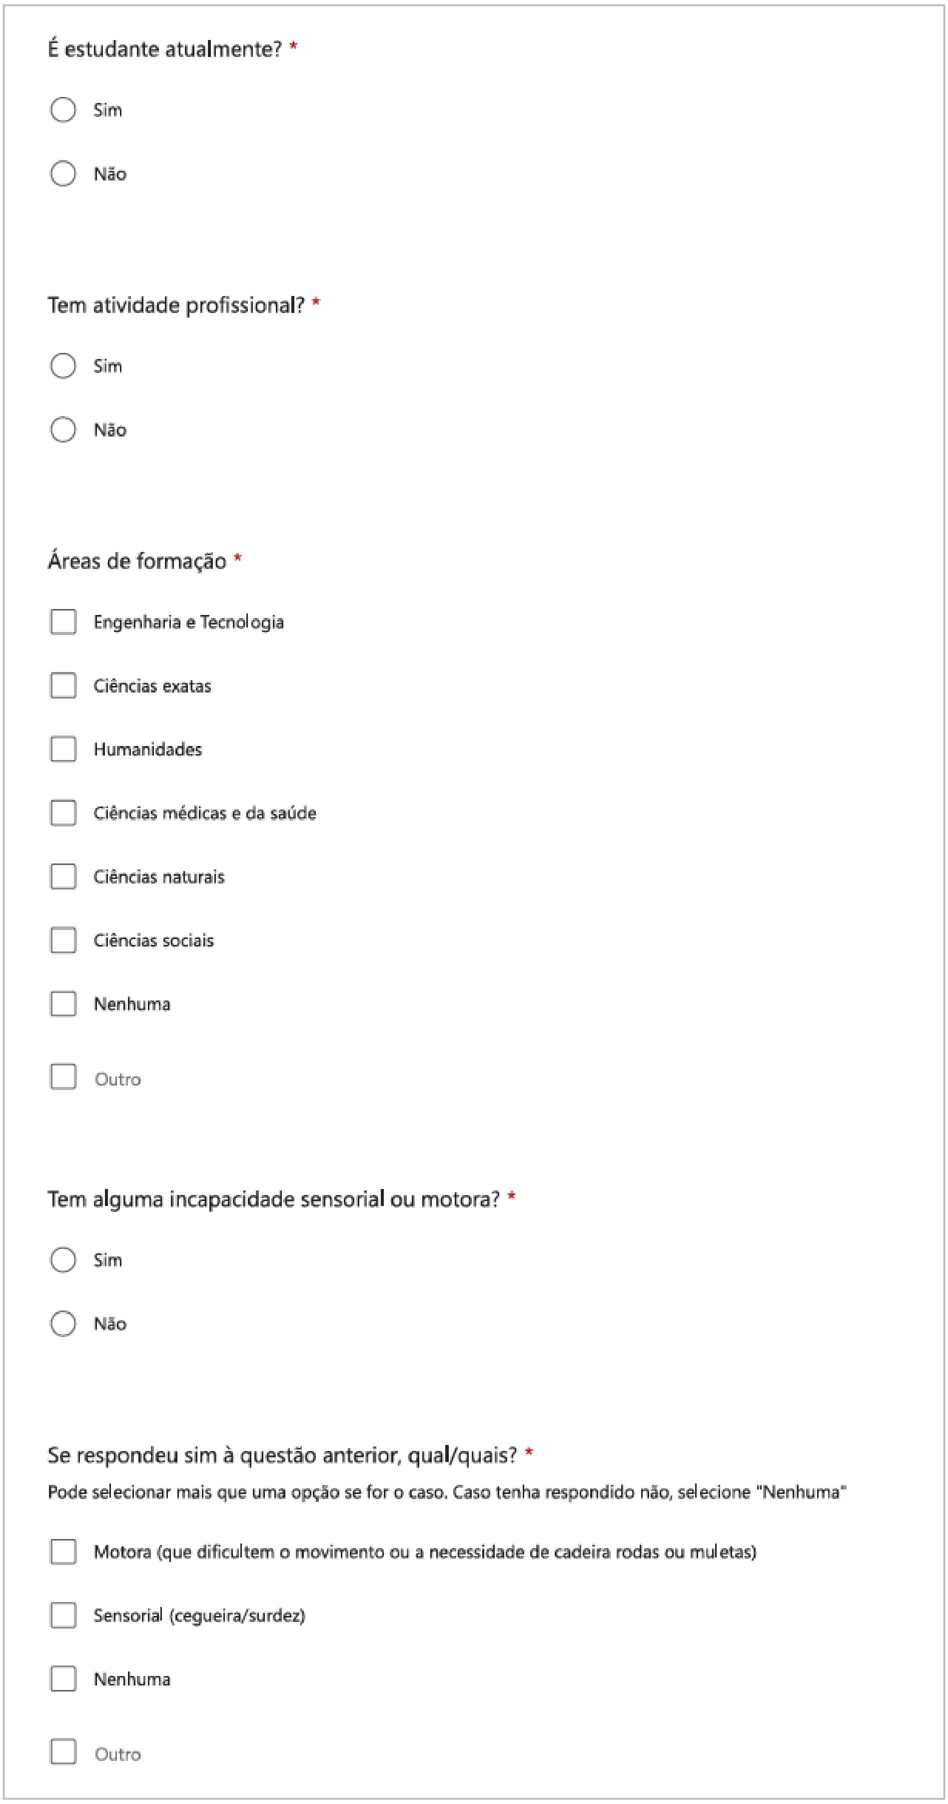

É estudante atualmente? *

○ Sim

○ Não

Tem atividade profissional? *

○ Sim

○ Não

Áreas de formação *

☐ Engenharia e Tecnologia

☐ Ciências exatas

☐ Humanidades

☐ Ciências médicas e da saúde

☐ Ciências naturais

☐ Ciências sociais

☐ Nenhuma

☐ Outro

Tem alguma incapacidade sensorial ou motora? *

○ Sim

○ Não

Se respondeu sim à questão anterior, qual/quais? *

Pode selecionar mais que uma opção se for o caso. Caso tenha respondido não, selecione "Nenhuma"

☐ Motora (que dificultem o movimento ou a necessidade de cadeira rodas ou muletas)

☐ Sensorial (cegueira/surdez)

☐ Nenhuma

☐ Outro

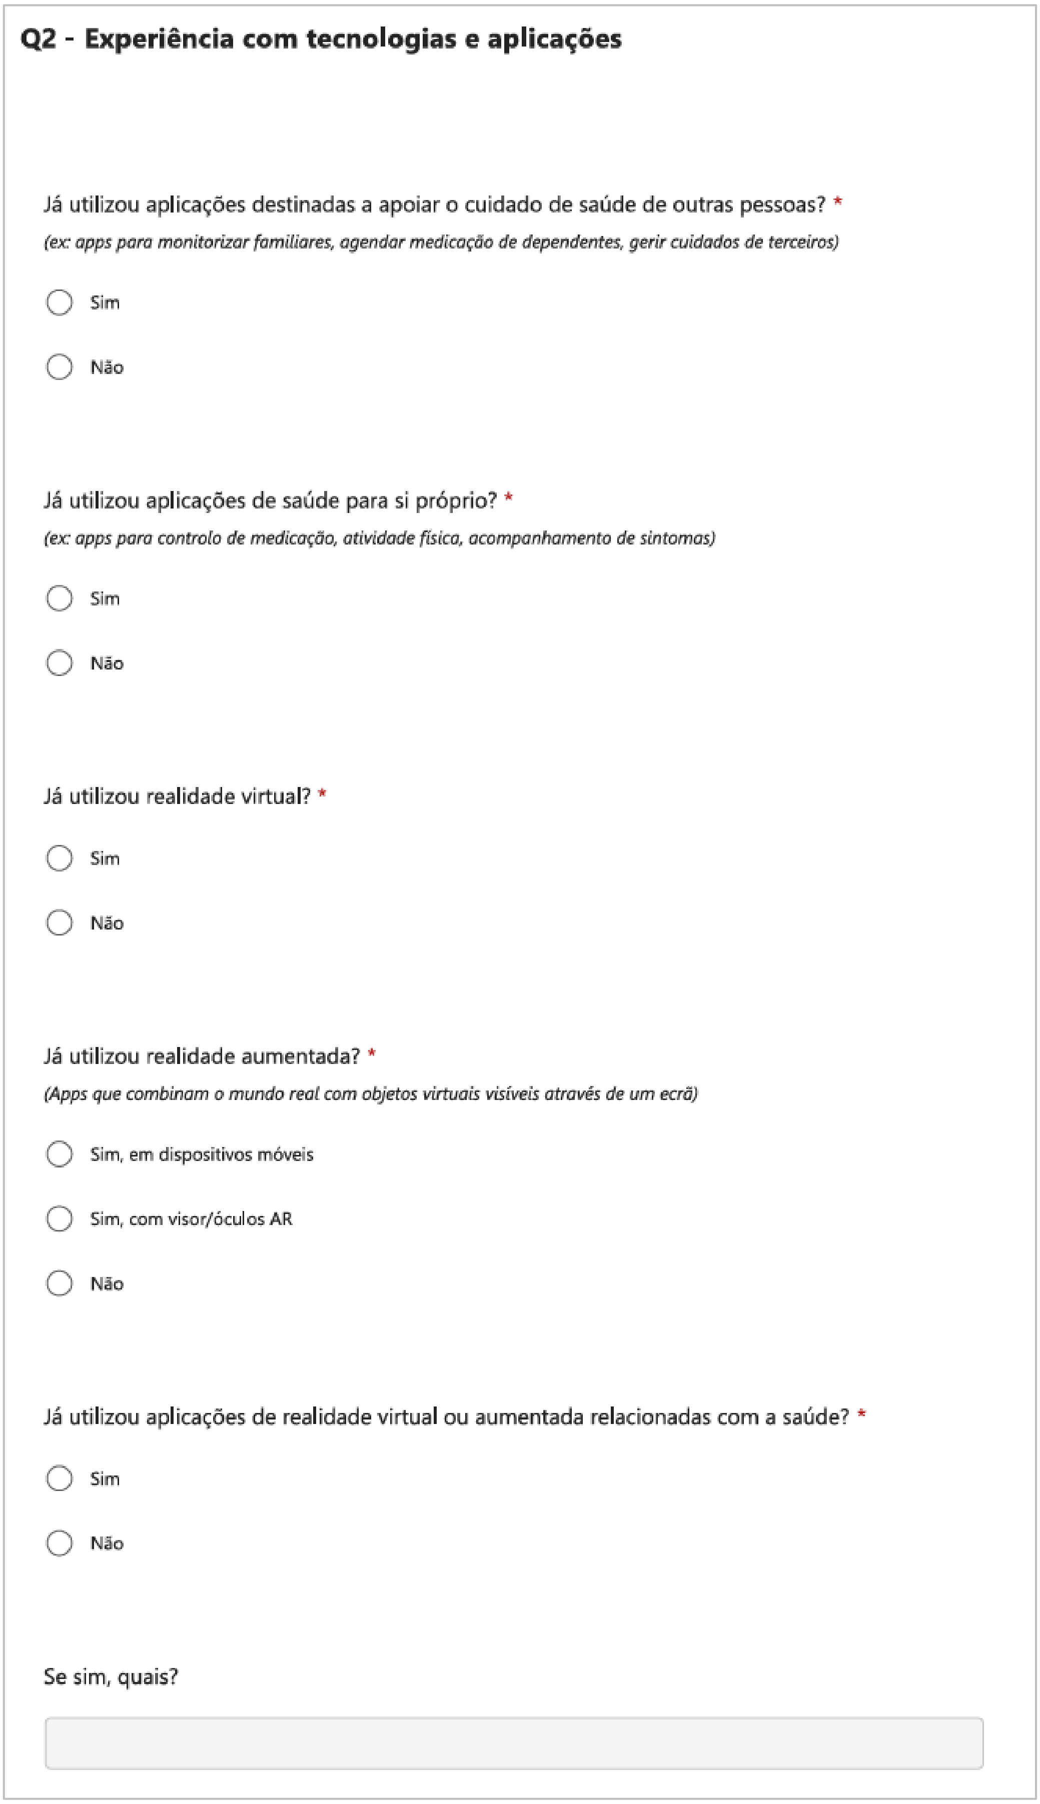

Q2 - Experiência com tecnologias e aplicações

Já utilizou aplicações destinadas a apoiar o cuidado de saúde de outras pessoas? *

(ex: apps para monitorizar familiares, agendar medicação de dependentes, gerir cuidados de terceiros)

○ Sim

○ Não

Já utilizou aplicações de saúde para si próprio? *

(ex: apps para controlo de medicação, atividade física, acompanhamento de sintomas)

○ Sim

○ Não

Já utilizou realidade virtual? *

○ Sim

○ Não

Já utilizou realidade aumentada? *

(Apps que combinam o mundo real com objetos virtuais visíveis através de um ecrã)

○ Sim, em dispositivos móveis

○ Sim, com visor/óculos AR

○ Não

Já utilizou aplicações de realidade virtual ou aumentada relacionadas com a saúde? *

○ Sim

○ Não

Se sim, quais?

Appendix B – Evaluation Questionnaire

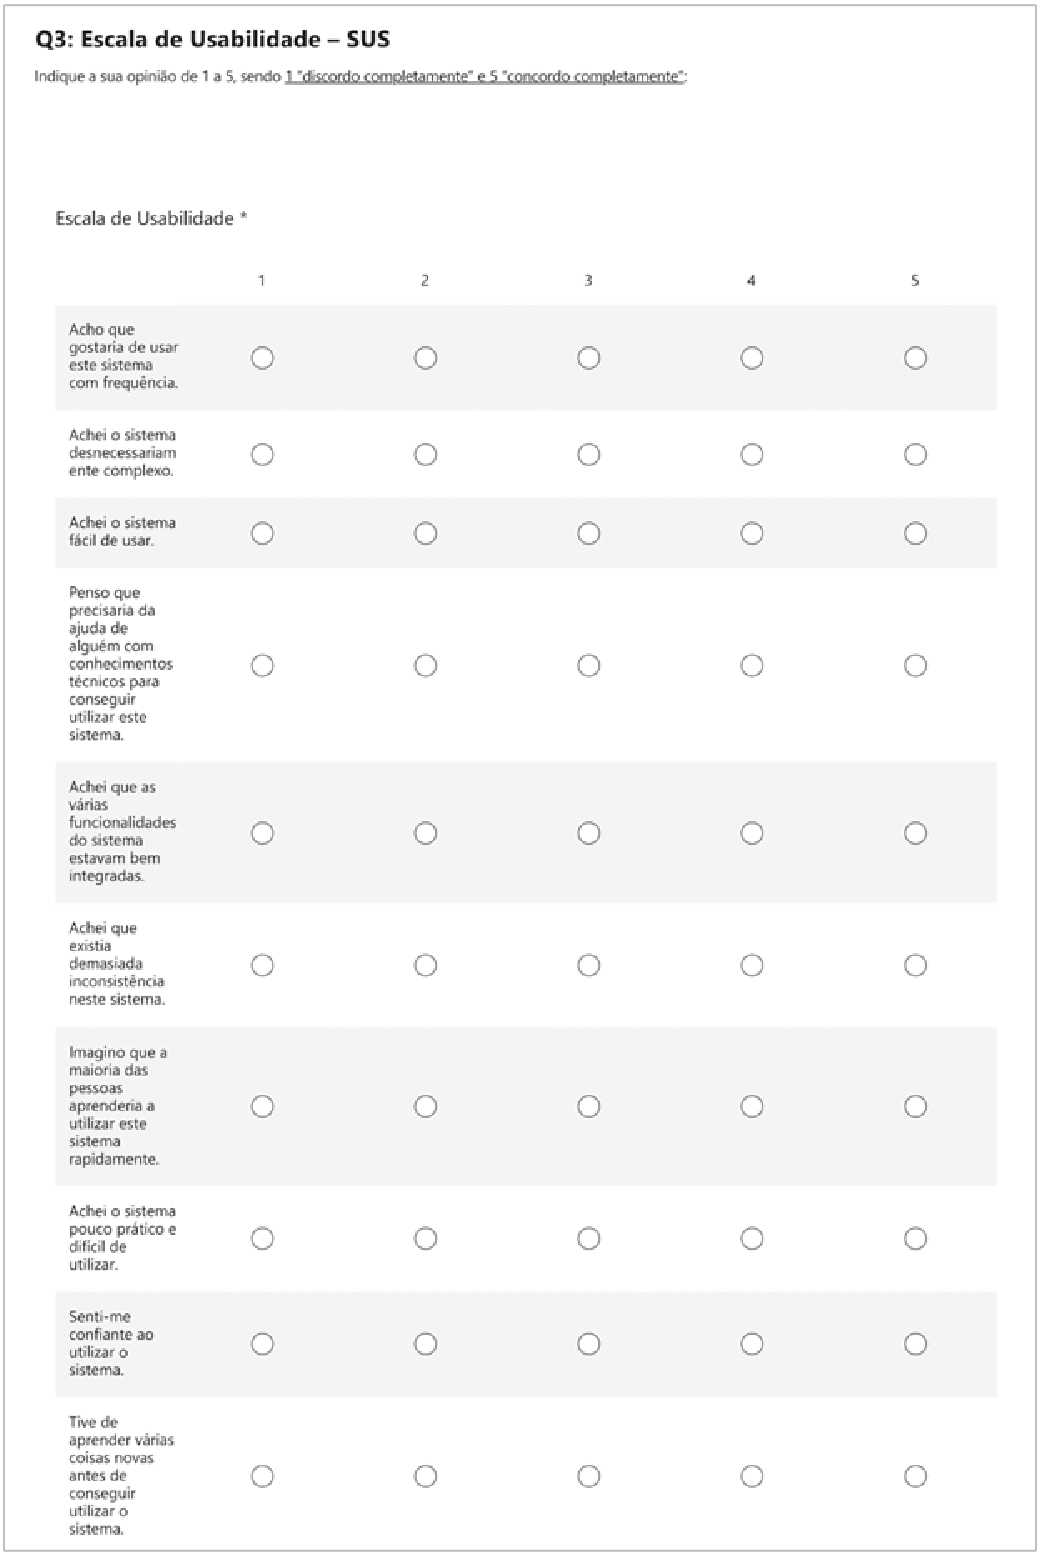

Q3: Escala de Usabilidade – SUS

Indique a sua opinião de 1 a 5, sendo 1 "discordo completamente" e 5 "concordo completamente":

Escala de Usabilidade *

	1	2	3	4	5
Acho que gostaria de usar este sistema com frequência.	○	○	○	○	○
Achei o sistema desnecessariamente complexo.	○	○	○	○	○
Achei o sistema fácil de usar.	○	○	○	○	○
Penso que precisaria da ajuda de alguém com conhecimentos técnicos para conseguir utilizar este sistema.	○	○	○	○	○
Achei que as várias funcionalidades do sistema estavam bem integradas.	○	○	○	○	○
Achei que existia demasiada inconsistência neste sistema.	○	○	○	○	○
Imagino que a maioria das pessoas aprenderia a utilizar este sistema rapidamente.	○	○	○	○	○
Achei o sistema pouco prático e difícil de utilizar.	○	○	○	○	○
Senti-me confiante ao utilizar o sistema.	○	○	○	○	○
Tive de aprender várias coisas novas antes de conseguir utilizar o sistema.	○	○	○	○	○

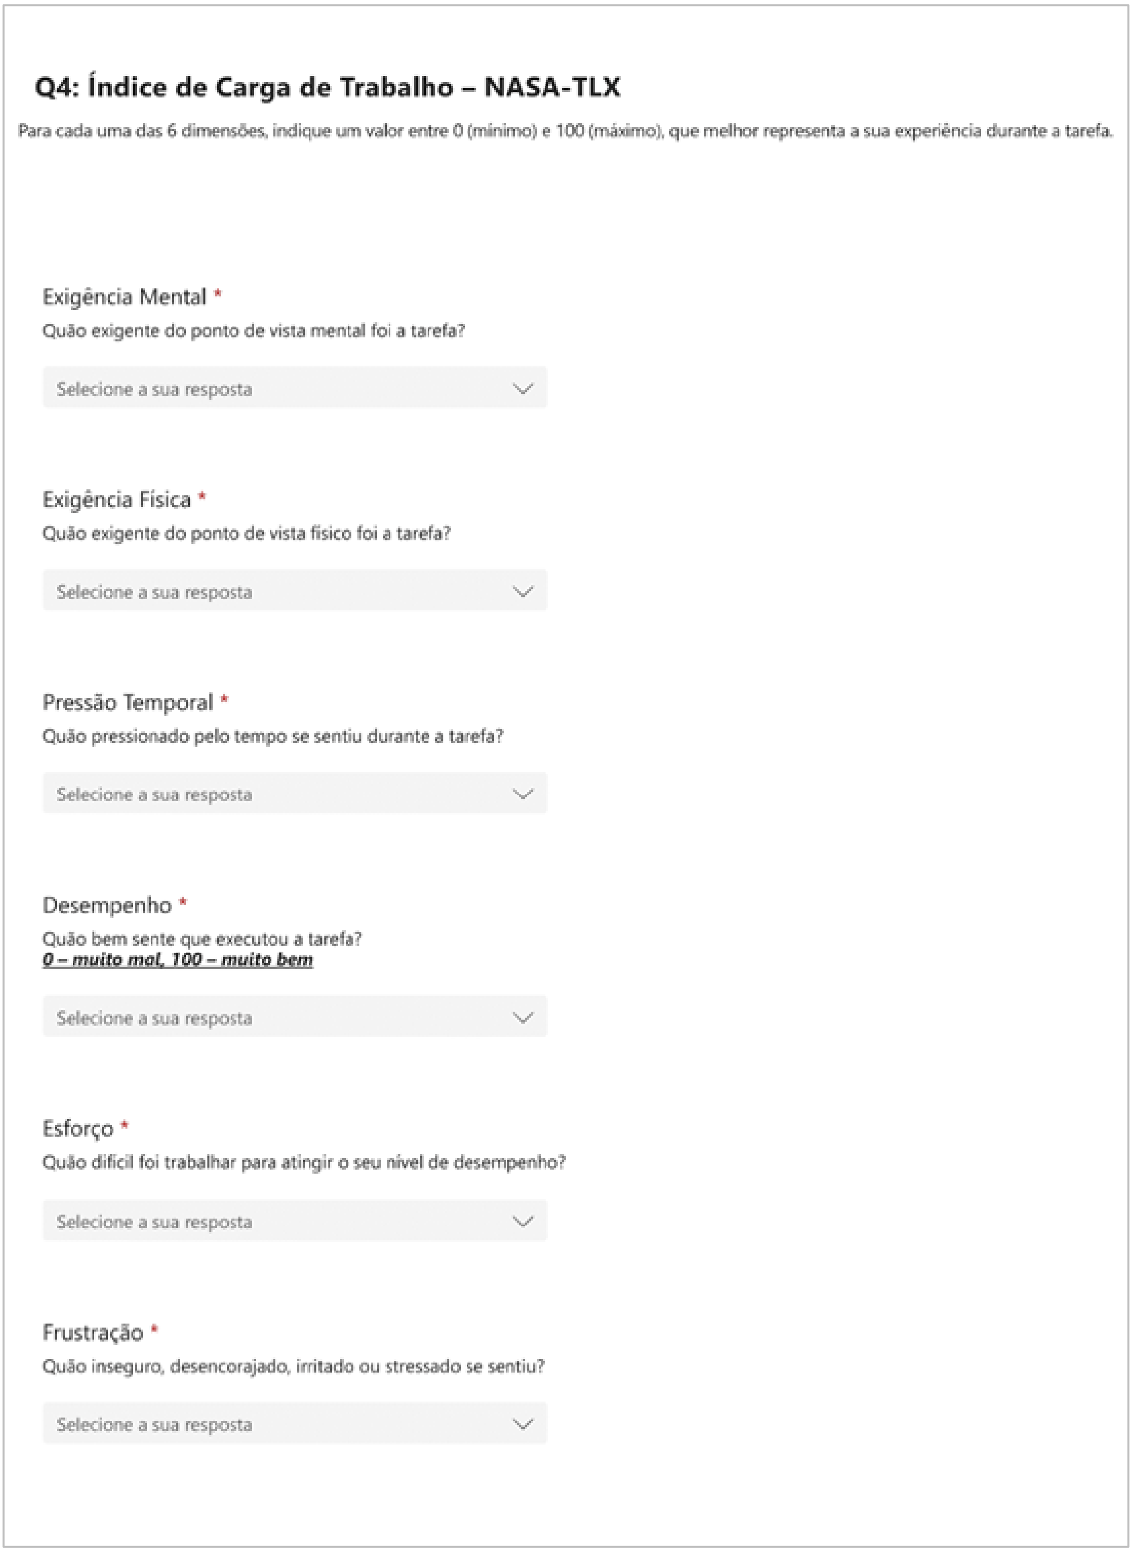

Q4: Índice de Carga de Trabalho – NASA-TLX

Para cada uma das 6 dimensões, indique um valor entre 0 (mínimo) e 100 (máximo), que melhor representa a sua experiência durante a tarefa.

Exigência Mental *

Quão exigente do ponto de vista mental foi a tarefa?

Selecione a sua resposta

Exigência Física *

Quão exigente do ponto de vista físico foi a tarefa?

Selecione a sua resposta

Pressão Temporal *

Quão pressionado pelo tempo se sentiu durante a tarefa?

Selecione a sua resposta

Desempenho *

Quão bem sente que executou a tarefa?
0 – muito mal, 100 – muito bem

Selecione a sua resposta

Esforço *

Quão difícil foi trabalhar para atingir o seu nível de desempenho?

Selecione a sua resposta

Frustração *

Quão inseguro, desencorajado, irritado ou stressado se sentiu?

Selecione a sua resposta

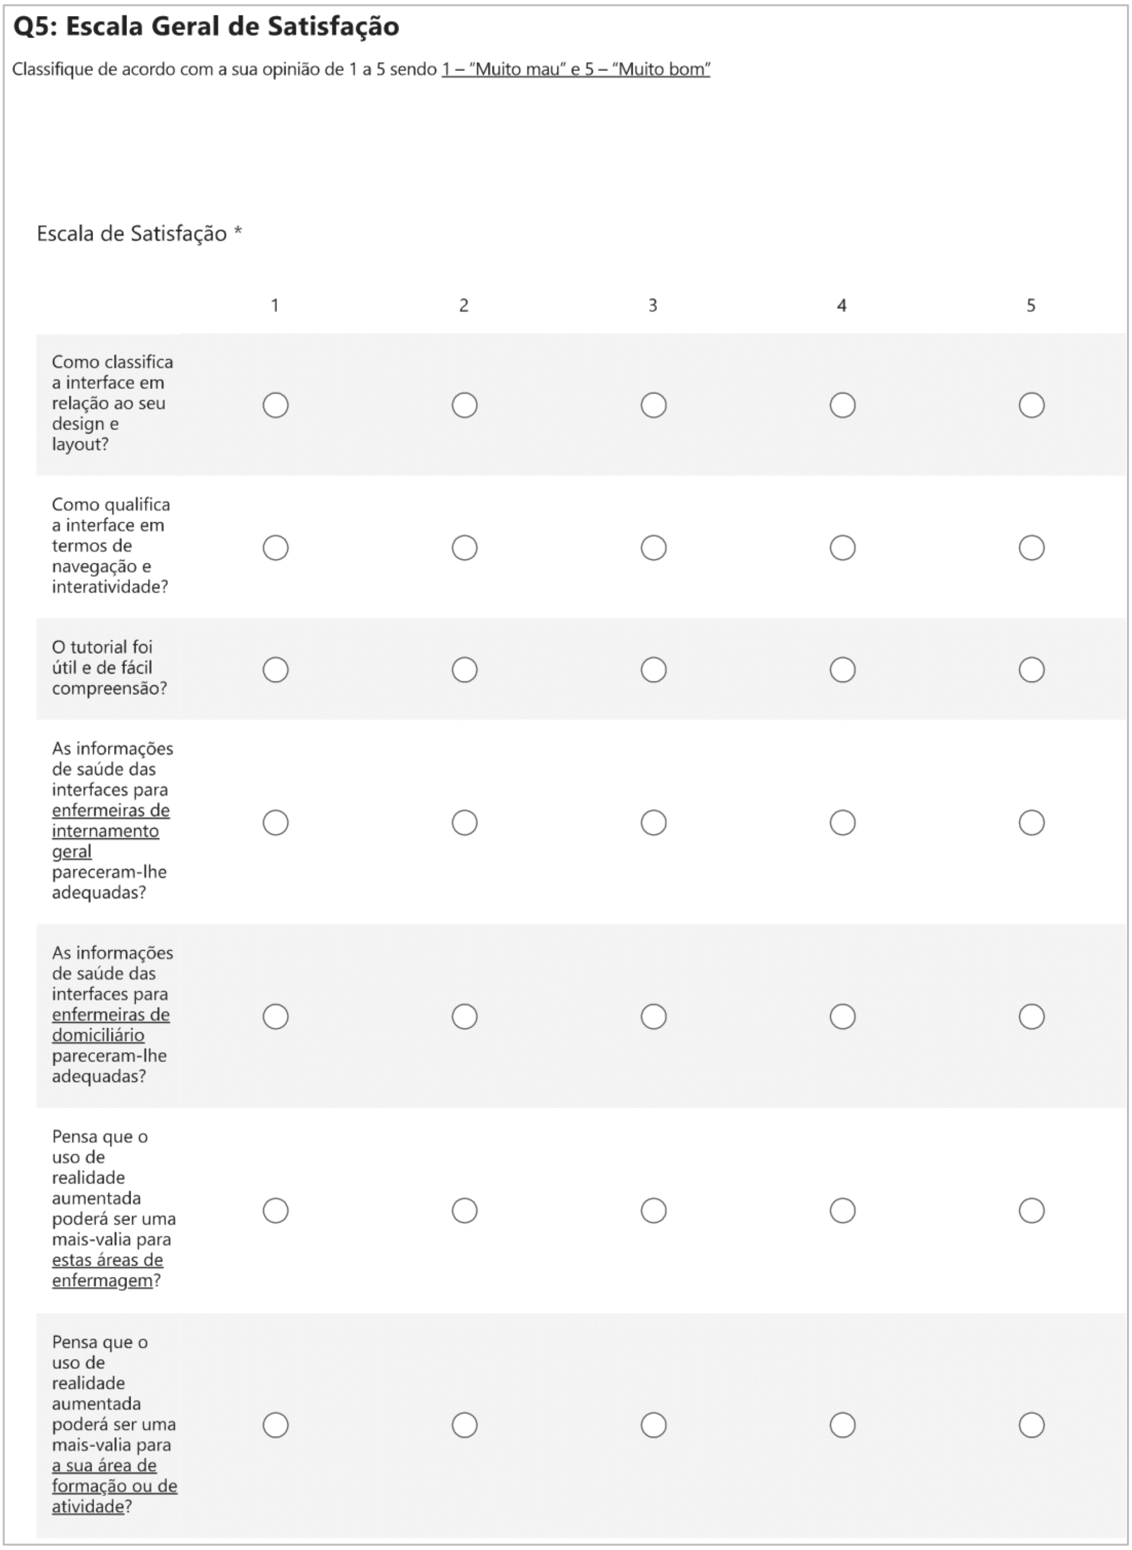

Q5: Escala Geral de Satisfação

Classifique de acordo com a sua opinião de 1 a 5 sendo 1 – "Muito mau" e 5 – "Muito bom"

Escala de Satisfação *

	1	2	3	4	5
Como classifica a interface em relação ao seu design e layout?	○	○	○	○	○
Como qualifica a interface em termos de navegação e interatividade?	○	○	○	○	○
O tutorial foi útil e de fácil compreensão?	○	○	○	○	○
As informações de saúde das interfaces para enfermeiras de internamento geral pareceram-lhe adequadas?	○	○	○	○	○
As informações de saúde das interfaces para enfermeiras de domiciliário pareceram-lhe adequadas?	○	○	○	○	○
Pensa que o uso de realidade aumentada poderá ser uma mais-valia para estas áreas de enfermagem?	○	○	○	○	○
Pensa que o uso de realidade aumentada poderá ser uma mais-valia para a sua área de formação ou de atividade?	○	○	○	○	○

Q6: Questões Opcionais

O que gostou mais no sistema que utilizou? Pode referir mais do que um aspeto.

O que menos gostou no sistema que utilizou? Pode referir mais do que um aspeto.

Como poderia o sistema que utilizou ser alterado de forma a facilitar a sua utilização? Pode referir mais do que um aspeto.