HUMANIZING ROBOT MAPPING

Designing a human-in-the-loop interaction for map understanding in domestic robots

MSc. Thesis Report
Stijn Gruben

> *“We don’t build robots to replace people; we build robots to empower people.”*
>
> - Rodney Brooks,
> Former Chief Technical Officer of iRobot

HUMANIZING ROBOT MAPPING

Designing a human-in-the-loop interaction for map understanding in domestic robots

PROGRAM
Integrated Product Design
Industrial Design Engineering
Delft University of Technology

GRADUATION
April 2026

MASTER THESIS
Stijn Gruben

SUPERVISORY TEAM
Dr. Jordan Boyle
Dr.ir. Olger Siebinga

PREFACE

During my studies in industrial design engineering, I have developed a strong interest in digital and electronic products. Working with technology, combining hardware, programming components, and seeing ideas come to life in functioning systems has often been my favorite part of the projects that I have worked on. The ability to create something that not only works but also behaves as intended continues to be a motivating factor in my development as a designer.

In this graduation project I could once again pursue this interest, as I had the opportunity to work with a physical robotic system. At the same time, I have always been interested in working with people and seeing how we can make interactions with these complex or abstract technical systems more accessible. Exploring how users can interact with and contribute to robotic mapping presented an exciting challenge all throughout this project.

Ultimately, I am pleased with the outcome of this project. I believe it offers an interesting perspective on human-robot interaction and contributes to the exploration of more collaborative approaches in robotics. I hope that this report provides you with new insights and perspectives.

Enjoy,

Stijn

ACKNOWLEDGEMENTS

I would like to thank everyone who contributed to this project and supported me throughout the process. Whether it was through sharing ideas, answering questions, or engaging in valuable conversations, your input has been greatly appreciated and played an important role in shaping this report.

I would like to express special gratitude to my mentors during this project; Jordan Boyle and Olger Siebinga. Thank you for your continuous support, guidance, and critical feedback throughout the whole project. Our discussions and meetings were always insightful and helped challenge and improve my work. Without your coaching, this project would not have been where it is today.

I would also like to thank Arend-Jan van Hilten and Martin Klomp from the TU Delft Cognitive Robotics Lab for providing the necessary robotic equipment and assistance during the initial setup of the robot system. Your willingness to answer any questions and think along was incredibly helpful.

Finally, I would like to thank my friends and family for their ongoing support during this final phase of my studies. Your encouragement as well as the necessary moments of distraction and laughter made this process much more enjoyable.

ABSTRACT

This graduation project explores how a user-in-the-loop approach can improve the usability and reliability of robot-generated maps in domestic environments. While Simultaneous Localization and Mapping (SLAM) allows robots to autonomously create maps of new environments, they are often imperfect due to sensor noise, environmental complexity, and localization errors. Such map errors can negatively impact navigation performance and reduce user trust, particularly when errors are difficult to interpret or correct.

The research phase focused on understanding how humans perceive and interact with robot-generated maps and identifying opportunities for a human-robot collaborative mapping approach. Through literature research and iterative robotics prototyping, insights were gathered and used to create a prototype interface that enables users to interpret, adjust, and refine robot-generated maps. The proposed concept views the user as an active collaborator in the mapping process, combining human spatial reasoning with algorithmic map generation. The design was evaluated through user testing, in which both map understanding and navigation performance were assessed.

The results show that users are able to identify and correct structural inconsistencies in robot-generated maps, leading to more visually consistent representations. Although no statistically significant improvements in navigation performance were found, the robot was able to successfully navigate using refined maps, and the results indicate a positive trend. This suggests that the approach is feasible and does not negatively impact system performance.

This project demonstrates that integrating users into the mapping process can enhance transparency and support more intuitive interactions with robotic systems. While further research is required to validate the effectiveness of the approach, the concept provides a promising direction for collaborative human-robot interaction, contributing to the development of more understandable and user-centered domestic robotics.

LIST OF ABBREVIATIONS

LiDAR	Light Detection and Ranging
IMU	Inertial Measurement Unit
SLAM	Simultaneous Localization And Mapping
HRI	Human-Robot Interaction
ACM	Association for Computing Machinery
IEEE	Institute of Electrical and Electronics Engineers
LOA	Levels Of Automation
SA	Situational Awareness
AI	Artificial Intelligence
GPS	Global Positioning System
EKF	Extended Kalman Filter
KF	Kalman Filter
ROS	Robot Operating System
GUI	Graphical User Interface

TABLE OF CONTENTS

01

INTRODUCTION

This chapter will introduce the context and relevance of the project. It will outline the problem space, the research goals, and the objective and scope of the work. This introduction acts as a foundation for the research and the design development that will aim to solve some of the key problems described in this introduction.

1.1. PROBLEM DEFINITION
1.2. SCOPE & GOAL
1.3. RESEARCH QUESTIONS
1.4. APPROACH

1.1. PROBLEM DEFINITION

1.1.1. Problem

Domestic cleaning robots, like robot vacuum cleaners, rely on autonomous mapping to navigate and efficiently clean your home. Currently, the mapping behavior is done independently, without user interference. As robots become more autonomous, users can experience a lack of control and understanding regarding the robot and the way it perceives its environment. Problems start to arise when the robot's mapping is inaccurate or incomplete, which can occur due to various reasons, like cluttered spaces, layout changes, or technical limitations like noisy sensors or wheel drift. When the robot does not work according to the expectations of the user, this can lead to frustration and reduced trust in the product and brand, especially when there is a lack of transparency and understanding between the robot and user.

1.1.2. Gap

Current mapping systems prioritize robot autonomy over human-robot interaction. This focus on minimizing the user input can also hinder meaningful interactions between the robot and the user (Kaber and Endsley, 2004). While increased autonomy could simplify the cleaning process in some instances, it also overlooks the potential benefits of involving users more actively in the behavior of the robot (Selvaggio et al., 2021). Current interfaces often lack intuitive interaction and require users to navigate abstract maps or complex settings when trying to understand their robot. Adding user participation to the mapping process could reveal an interesting research gap and lead to new insights into how human-robot interaction can enhance the mapping experience for domestic robotics.

1.1.3. Research Goal

This research report aims to explore a user-in-the-loop concept for the mapping process of domestic cleaning robots, focusing on an intuitive, user-centered approach. The goal is to investigate whether giving the user more control over the robot's mapping behavior can improve reliability and overall user experience. By researching, designing, and testing interaction concepts that focus on intuitive user interaction, this project aims to create a seamless human-robot mapping experience. Ultimately, this research seeks to contribute to more intuitive domestic robots that properly balance autonomy with meaningful customer involvement.

1.1.4. Problem Statement

"How might we create a more efficient and intuitive mapping experience for domestic cleaning robots by adding a user-in-the-loop, and how does this approach compare to fully autonomous mapping systems in terms of reliability, ease of use, and user trust?"

1.2. SCOPE & GOAL

The domain of this project is limited to domestic cleaning robots, which will be treated as relatively simple autonomous mobile robots in private households. While cleaning robots can use a multitude of different sensors, for this project the cleaning robot is imagined to have access to a LiDAR sensor, an IMU, and dead-reckoning data based on wheel encoders.

Camera-based data is intentionally excluded, one reason is to avoid privacy issues (see Chapter 2.5) and to explore if, by using limited sensor data, an intuitive but user-friendly interface can be designed. An additional reason is that for this project, the goal is to create a working interface prototype that can be tested on a physical robot rather than using simulations. The robot that will be used for this project is the MIRTE Master robot, see Figure 2.

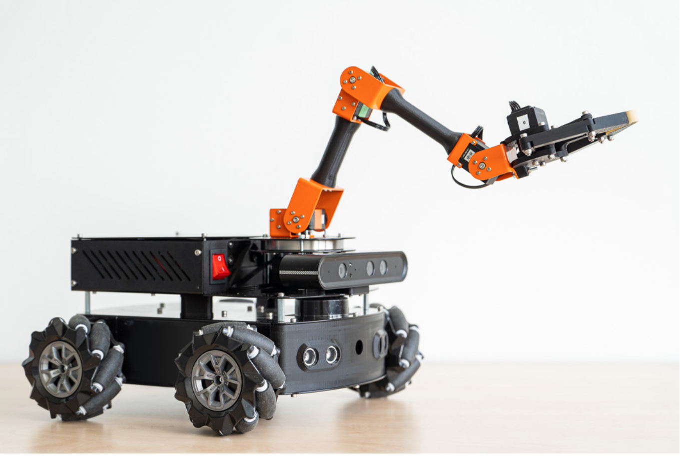

Figure 2: The MIRTE Master robot (MIRTE team, 2025)

1.3. RESEARCH QUESTIONS

Main Research Question

How can a user-in-the-loop mapping interface improve the usability, transparency, and reliability of domestic cleaning robots?

SQ1 - Understanding

How do users currently perceive and understand robot-generated maps?

SQ2 - Design Exploration

How can users effectively contribute to correcting and refining robot-generated maps?

SQ3 - Evaluation

What is the impact of a collaborative human-robot mapping approach on user experience, map accuracy, and reliability?

1.4. APPROACH

The process followed during this project is described in this section and visualized in Figure 3. The design approach is based on the triple diamond framework consisting of an exploration, ideation, and implementation phase. These three phases represent the iterative structure of the design approach. This project also places a strong emphasis on the continuous integration with a physical robot system, utilizing real sensor data, and grounding each phase in real-world exploration.

Exploration
The exploration phase focused on gaining a proper understanding of the context and background information regarding the project, mainly concentrating on human-robot interaction, different mapping approaches, and autonomous systems. The goal of this literature study was to get a broad understanding of the most relevant aspects regarding domestic cleaning robots and the scope of this project.

Along with this theoretical study, the goal for this phase was to gain familiarity with the robot platform that would be used during this project. By exploring the robot and making sure that the possibilities and limitations of this hardware were known, a solid foundation was created for the ideation phase.

Ideation
The ideation phase was a way of structuring the analysis and robot exploration insights into clear design opportunities and directions. By noting the pain points and key opportunities that surfaced during the exploration phase, a clear design vision could be created.

The ideation phase was, along with the necessary idea generation of different interface and interaction concepts, strongly linked with exploration of the physical robot. When general ideas emerged, they were quickly tested on the robot software and by experimenting with the robot new ideas formed. Through these physical iterations, the user-in-the-loop concept came to be.

Implementation
The implementation phase focused on the translation of the design vision into a working interface prototype that could be used in real-life testing. This functional prototype used real mapping data that was created by the robot and let users work with this data in user tests in semi-realistic environments.

Through the testing of the functional prototype insights were gathered, helping the project move along and further iterating the design towards a final design concept for a user-in-the-loop SLAM approach.

Overall, the project followed an iterative and explorative design approach where physical robot integration stood at the center of the project. By combining literature, ideation, hands-on implementation, and evaluation, the project aims to be a valuable asset to the human-robot interaction field.

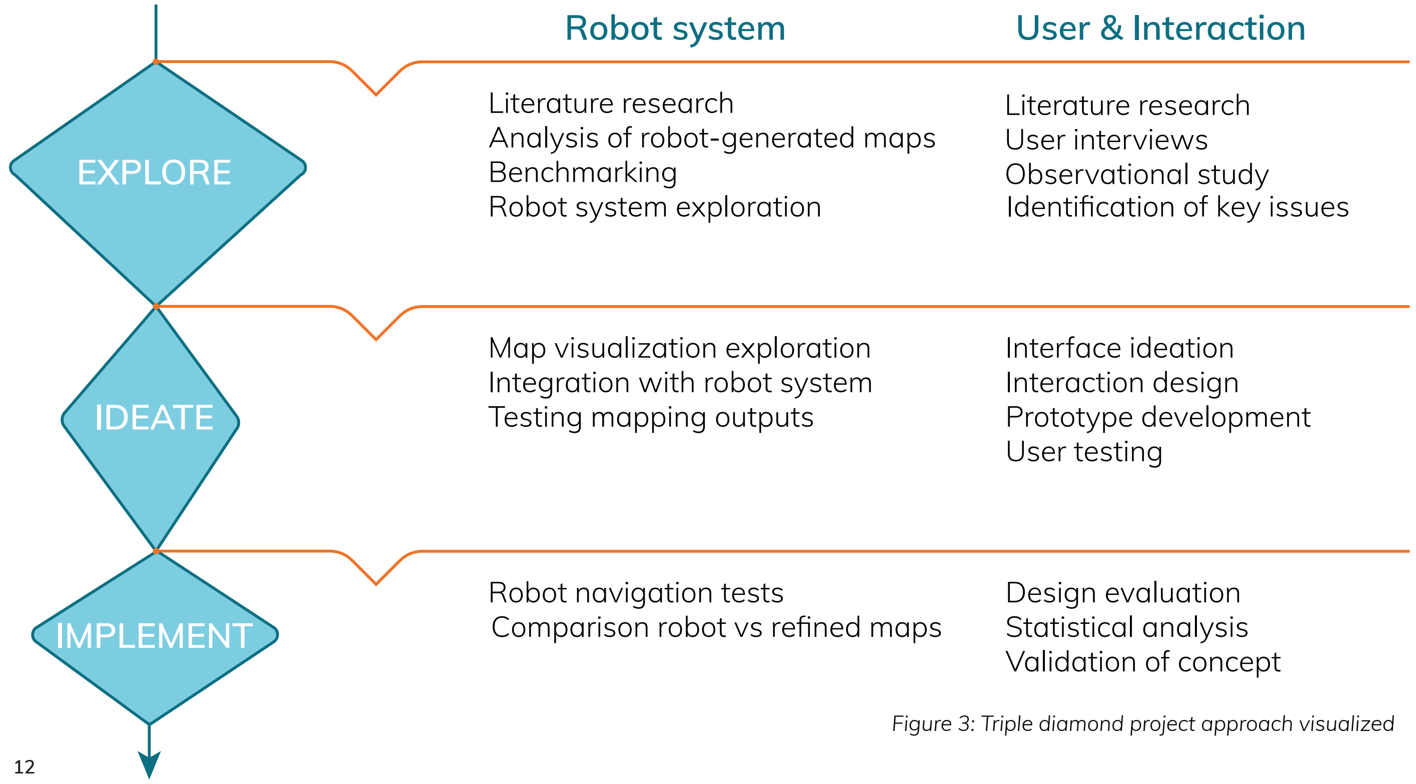

	Robot system	User & Interaction
EXPLORE	Literature research Analysis of robot-generated maps Benchmarking Robot system exploration	Literature research User interviews Observational study Identification of key issues
IDEATE	Map visualization exploration Integration with robot system Testing mapping outputs	Interface ideation Interaction design Prototype development User testing
IMPLEMENT	Robot navigation tests Comparison robot vs refined maps	Design evaluation Statistical analysis Validation of concept

Figure 3: Triple diamond project approach visualized

02

CONTEXT

To gain a proper understanding of the relevant context surrounding this project, an exploratory literature review will be conducted. The chapter will start with research on human interaction with domestic cleaning robots. Secondly, an explanation of SLAM and associated topics will be given. A deeper dive into the fundamentals of interface design will be provided with a focus on robot interfaces. Additionally, benchmarking will give insights into current market trends. Furthermore, the chapter will outline the customer journey that is associated with domestic cleaning robots, and finally address the key ethical considerations relevant to this topic.

2.1. HUMAN-ROBOT INTERACTION

Human-robot interaction (HRI) is a very extensive and diverse research and design field. The literature is ever-expanding, with hundreds of publications released each year, and the topic is relevant across a multitude of technical disciplines, including but not limited to mechanical and electrical engineering, computer and control science, and artificial intelligence. Every year since 2006, the ACM/IEEE has hosted the most influential conference on human-robot interaction and has had an increasing impact within the field (Google Scholar, 2020). The need for research on human interaction aspects and participation in robot research and design is serious. Big opportunities for future research lie in bridging the gap between human factors, engineering, and computer science. Engaging in stronger collaboration between these disciplines might lead to a better overall understanding of human-robot interaction. An approach like this could help the development of robotic systems that are not only technically advanced but also align better with the needs of the user and the context of their use (Sheridan, 2016). Because this project is looking into autonomous robots, the following paragraphs will go into how humans interact with these systems and what this interaction involves.

2.1.1. Human-Robot Interaction with Autonomous Systems

Since the 1940s, humans and robots have been working together. In the beginning, this interaction was mainly unidirectional, using simple switches or controls to operate manipulator joints and remote vehicles. However, as the technical field improved and robots became more autonomous, this relationship changed, becoming more complex, varied, and ultimately more valuable (Sheridan, 2002).

This growing autonomy, however, has not created robots capable of functioning entirely without human involvement across all contexts. Fully autonomous robots are machines that can complete a certain task by sensing, planning, and acting in an environment without any human intervention. While such autonomy has been succesfully achieved in specific, well-defined domains, these systems often operate under controlled or highly predictable conditions.

In more complex and dynamic environments, such as the domestic environment, this full autonomy remains a challenge. Despite recent advancements in automation and artificial intelligence, achieving fully autonomous systems in complex environments remains very difficult (Kok & Soh, 2020).

In most robot applications, autonomy is still limited or constrained; for instance, domestic cleaning robots often stop moving when they detect dynamic objects nearby, such as a human moving towards them. While this is an autonomous behavior, it is based on a predefined response rather than a deeper understanding of a complex situation.

Aside from this, it is common for robots to be operated or supervised by a human operator. This operator is able to aid the robot using superior situational awareness, logic, and problem-solving capabilities. This interaction between human and robot, although sometimes very minor, is needed and often preferred (Selvaggio et al., 2021). Johnson and Vera (2019) even go as far as to say that technology does not work in isolation from people. They claim that technology thrives most when it is successfully woven into human work practices. Machines do not automatically get simpler to use because they get smarter; the opposite is often true. This is where human-centered design comes into play.

2.1.2. Levels of Automation

Through the years, various models and classifications describing different levels of automation (LOA) have been proposed. Because robot systems rarely work totally independently from humans, it is very useful to get an understanding of how the responsibilities and control are distributed between the robot and the user. By distinguishing different LOAs, designers are able to have a clearer overview of the distinct roles of the human and the robot in the system. A framework like this is particularly useful for determining a design context. It enables designers to make informed decisions about which level of autonomy is most appropriate for a given task, user group, and environment, and by testing different levels meaningful human-robot interactions might arise.

The earliest LOA framework was proposed by Sheridan & Verplank (1978). This 10-point scale (see Table 1) describes increasing levels of autonomy and has formed the basis for many later LOA frameworks.

Level of Automation	**Function allocation**
1	The computer offers no assistance; the human must take all decisions and actions.
2	The computer offers little assistance; the human must take almost all decisions and actions.
3	The computer offers a complete set of decision/action alternatives.
4	The computer narrows the selection down to a few alternatives.
5	The computer suggests one alternative.
6	The computer executes the suggestion if the human operator approves it.
7	The computer allows the human a restricted time to veto before automatic execution.
8	The computer executes automatically, then necessarily informs the human.
9	The computer informs the human only if asked to.
10	The computer informs the human only if it, decides to.

Table 1: Levels of decision-making automation (Sheridan & Verplank, 1978)

Based on this taxonomy created by Sheridan & Verplank, Beer et al. (2014) proposed a framework that more specifically defines the levels of robot autonomy within the context of human-robot interaction. As this framework is more in line with the context of this project, it is the framework that will be referenced from here on out (see Table 2).

This project deals with a user-in-the-loop approach and thus will primarily be situated between the levels 6-8 of the LOA framework created by Beer et al. (2014).

Levels of Robot Autonomy	**Function Allocation**			**Description**
	Sense	**Plan**	**Act**	
1. Manual Teleoperation	H	H	H	The human performs all aspects of task including sensing the environment and monitoring the system, generating plans/options/goals, and implementation.
2. Action Support	H	H	H/R	The robot assists the human with action implementation. However, sensing and planning is allocated to the human. For example, a human may teleoperate a robot, but the human may choose to prompt the robot to assist with some aspects of a task (e.g., gripping objects)
3. Assisted Teleoperation	H/R	H	H/R	The human assist with all aspects of the task. However, the robot senses the environment and chooses to intervene with task. For example, if the user navigates the robot too close to an obstacle, the robot will automatically steer to avoid collision.
4. Batch Processing	H/R	H	R	Both the human and robot monitor/sense the environment. The human, however, determines the goals and plans of the task. The robot then implements task.
5. Decision Support	H/R	H/R	R	Both the human and robot sense the environment and generate a task plan. However, the human chooses the task plan and commands robot to implement action
6. Shared Control with Human Initiative	H/R	H/R	R	The robot autonomously senses the environment, develops plans/goals, and implements actions. However, the human monitors the robot's progress, and may intervene and influence the robot with new goals/plans if the robot is having difficulty.
7. Shared Control with Robot Initiative	H/R	H/R	R	Robot performs all aspects of the task (sense, plan, act). If the robot encounters difficulty, it can prompt the human for assistance in setting new goals/plans.
8. Supervisory Control	H/R	R	R	Robot performs all aspects of task, but the human continuously monitors the robot. The human has over-ride capability and may set a new goal/plan. In this case the autonomy would shift to shared control or decision support.
9. Executive Control	R	(H)/R	R	The human may give an abstract high level goal (e.g., navigate in environment to specified location). The robot autonomously senses environment, sets plan, and implements action.
10. Full Autonomy	R	R	R	Robot performs all aspects of a task autonomously without human intervening with sensing, planning, or implementing action.

**Note: H = Human, R = Robot*

Table 2: Taxonomy of levels of robot autonomy for HRI (Beer et al., 2014)

2.1.3. Situational Awareness

One particular concept that is increasingly important in human interaction with autonomous robots is situational awareness (SA). Situational awareness has been formally defined by Endsley (1995) as "the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future." We could say that SA is the ability to sense and interpret your environment. SA is a vital concept in automation, as it grants the ability for decision-making.

Research by Merat et al. (2018) has shown that automation can have a significant influence on SA. When the robot has a high level of autonomy, the human is often given a much more passive monitoring role, which can lead to so-called out-of-the-loop performance problems. These problems typically arise when the human operator has little ability to manually intervene in the system and, due to high levels of automation, has little knowledge about the system, resulting in limited SA. Kaber and Endsley (2004) investigated the effects of different levels of automation on the SA of the human operator. With evidence indicating that leaving the human with some level of involvement helped maintain a better SA and ultimately a better overall system performance.

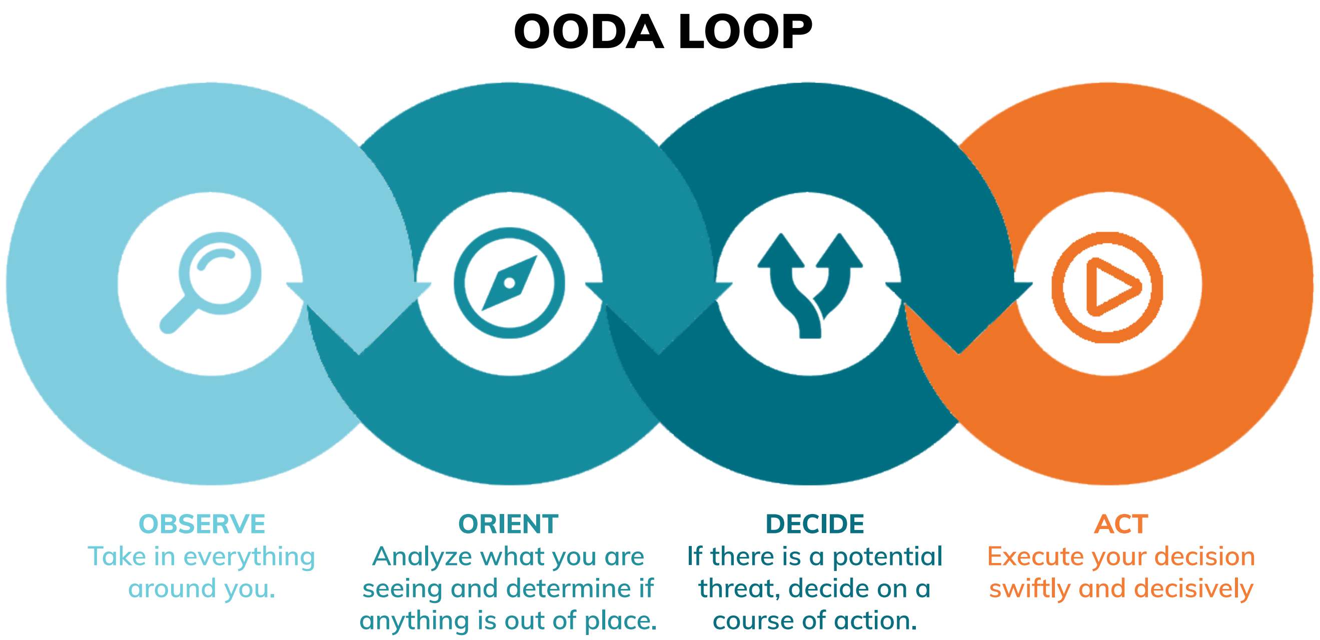

Figure 5: The OODA loop: a framework for situational awareness, developed by military strategist John Boyd (Price, 2023)

2.2. DOMESTIC CLEANING ROBOTS

This chapter will provide the needed background information on domestic robots and the way that users interact with these systems. To create a comprehensive understanding, the chapter starts with a general introduction to domestic robots and their role in the everyday life of people. To gain an understanding of the user group, the user demographics will be determined. Following this, the home environment and cleaning practices, and how these are affected by the introduction of a cleaning robot will be explored. Furthermore, the social and emotional impact of a domestic robot is explained, including the influence on relationships, behavior, and home perception. Finally, the user needs, common pain points, and corresponding solution spaces will be highlighted to address the relevance of this project.

2.2.1. Introduction to Domestic Robots

Domestic robots are among the first autonomous systems to be introduced into the daily lives of humans. Over the last two decades, various brands have designed robots to help relieve homeowners from household chores (Coggins, 2022). Among these, the vacuum cleaning robot is almost certainly the most well-known and also the focus of this report. Even though it is difficult to put a number on the popularity of these robots, iRobot, the manufacturer of the popular Roomba, claimed in 2016 that 20% of the world's vacuum cleaners were robots and their own models accounted for 70% of the market (Etherington, 2016). The first cleaning robots originated from industrial models that have been made smaller, lighter, and less costly to be able to be used at home (Forlizzi & DiSalvo, 2006). The first vacuum robot prototype was presented by a Swedish brand called Electrolux in 1996 and was called "Trilobite" after the extinct arthropod (see figure 6). The Trilobite included very basic navigation features and limited vacuuming mechanisms (Prassler et al., 2000). Nowadays, cleaning robots are a lot more advanced, often consisting of advanced path-planning and automation, internet connectivity, and mobile applications that allow you to remotely run and interact with your robot, and even AI technology for detecting objects and learning about the environment (Yapici et al., 2022).

Due to this continuous development of the technological features of the cleaning robots and these robots being relatively affordable yet useful, these products have quickly gained popularity and are now a great study opportunity for human-robot interaction inside dynamic and personal home environments (Yapici et al., 2022). The integration of these robots into the daily lives of humans not only provides technical challenges related to autonomy and navigation but also adds social, emotional, and behavioral challenges to the mix as human-robot collaboration becomes a center point.

Figure 6: The Electrolux Trilobyte (Cyberneticzoo.com, 2013)

2.2.2. User Demographics and Motivations

To get a better understanding of the user group, the user demographics of domestic cleaning robots will be explored. As Rogers (2019) states, determining the actual audience for a product is crucial to determining the product's overall success on the marketplace.
A result of a survey of 379 iRobot's Roomba users conducted by Sung et al. (2008) offers a general view on the demographics of vacuum cleaning robotics. In Figure 7, the age and gender distribution is shown, and it can be clearly seen that the Roomba users are relatively young, with the majority of the people in this group being in their twenties. The gender distribution of the Roomba user group is almost equally divided, with 194 women and 181 men in total.

When looking at the household composition of the Roomba users, the biggest group consisted of 164 participants (43%) who lived with other adults but not children, 128 households (34%) that included adults and children, and 85 participants (23%) who lived alone. Furthermore, regardless of whether they lived with another person or not, half (49.6%, N=188) of the participants owned one or more pets, including cats, dogs, parrots, and rabbits. Households with children expressed greater satisfaction with the Roomba's performance.

Collectively, the Roomba owners had an above-average level of education, with most having an undergraduate degree (N=229) and many having a graduate degree (N=112). Only 35 people have left education after high school. Moreover, many of those with college degrees had their academic training in technical fields. Among the 341 people with college-level education, 153 had engineering-related degrees. In addition to asking about education, Sung et al. (2008) also asked whether people worked in a technical profession or whether they were a "recreational engineer" with technical hobbies. Based on answers to these three questions, 48% of participants identified as technical (N=182), indicating that a big portion of Roomba users were familiar with technology, and perhaps more so than average. While a large portion of the participants had technical knowledge, about half (43%, N=158) self-identified as being attracted to new technology and gadgets. Lastly the motivation behind the purchase of a Roomba robot was determined by Sung et al. (2008); different motivations can be seen in Table 3.

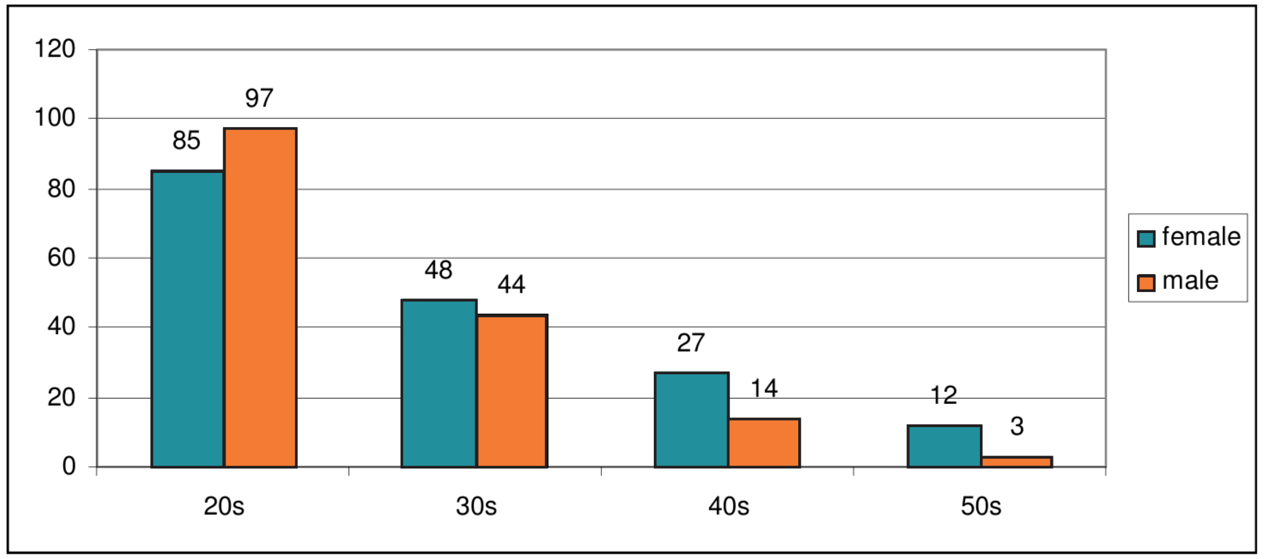

Figure 7: Age & Gender distribution of iRobot Roomba users in study of Sung et al. (2008)

Purchase motivation for First Roomba (N=379)	responses
Through my or other's experience (demonstrated, recomme nded, gifted)	188
Interested in new technology	173
Hate vacuuming	171
Curiosity	152
Always wanted to own robots	79
Overwhelming amount of cleaning. Need assistance	63
To work around physical difficulties	44

Table 3: Purchase motivation of the first iRobot Roomba in the study of Sung et al. (2008)

In recent years, the robotic vacuum cleaner market has grown rapidly, as can be seen in Figure 8. While the user demographics might have changed since the 2008 data report by Sung et al., the study does still provide us with a general profile of the users of robotic vacuum cleaners. Users are often young, active individuals with higher levels of education and an interest in technology. Household compositions vary, although a noteworthy number of users own pets, adding extra cleaning needs that the vacuum cleaner should account for.

2.2.3. Understanding the Environment

Besides the demographic, another important aspect for cleaning robotics is the domain, which, of course, is the commercial home environment. Thus, it is required for researchers and designers to gain an understanding of this unique environment. Unlike labs, where the robots are usually tested, the domestic environment is often way more complex. This environment not only holds material objects but also consists of daily routines and intimate social structures of families, couples, and individuals living in unity with technology (Forlizzi & DiSalvo, 2006).

Forlizzi and DiSalvo (2006) state that homes are not and should not be designed to accommodate autonomous robot technologies. Instead, they propose that autonomous mobile robots should be designed to integrate and adapt to the existing structures, routines, and practices of the home (Suchman, 2002; Taylor & Swan, 2005). For this integration to occur, the user will almost always need a certain control or influence over the robot's behavior.

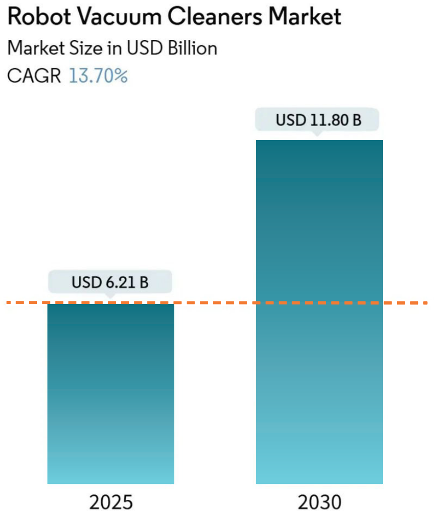

Figure 8: Robot vacuum cleaners market size estimate (Mordor Intelligence, 2025)

A home is a complex environment consisting of people, practices, and products that together create a highly personal social system, within which cleaning has an important role. Cleaning and cleanliness can be described rather practically: to get rid of filth or germs. However, cleaning also reflects the structure and values of what is proper and good. This value will determine the way the house is cleaned and the role of different cleaning products, from cleaning agents to robotics (Shove, 2003).

While this might seem a bit abstract, it is important to notice that every home will have different values on cleaning, meaning that each home will operate its robot vacuum differently, leading to diverse human-robot interactions.

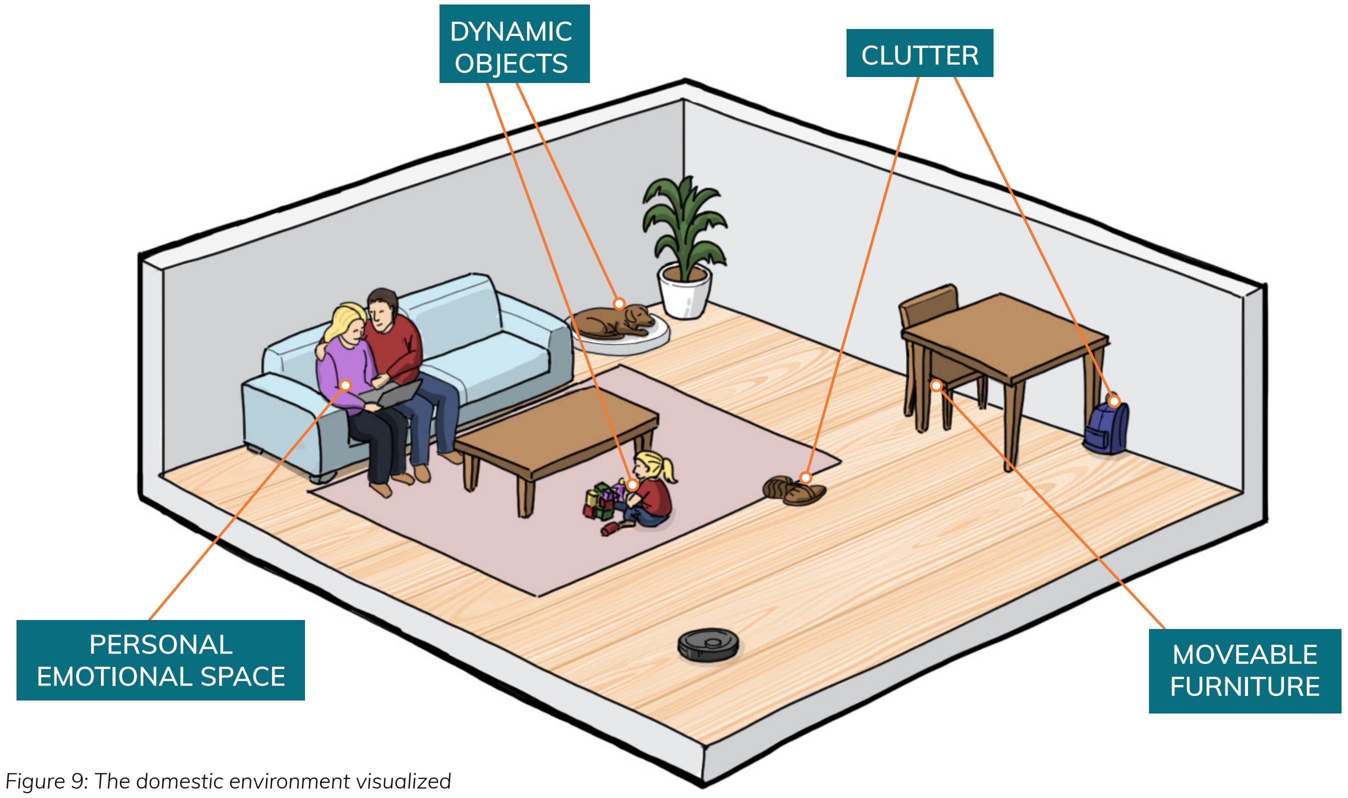

Figure 9: The domestic environment visualized

2.2.4. Robot impact on User Behavior

The introduction of a robot vacuum to a household changes the way that people clean their homes (Coggins, 2022). Foremost, the autonomy of the robot allows for multitasking, meaning that users can do something else while the floor is being vacuumed (Forlizzi & DiSalvo, 2006). The findings of Sung et al. (2008) show that the frequency of cleaning also increased when the robot was introduced to a household (see Table 4). The time of cleaning also shifted with the introduction of a robot vacuum. Before having access to a robot, cleaning routines focused on the days that people could spend time at home: weekends, Friday nights, and work-at-home days. With a robot, however, users can also utilize times they are not home or are occupied by other tasks, like while they are working, while putting their children to bed, or while they are asleep.

	Before Roomba	**With Roomba**
Cleaning Frequency	*# of responses*	*# of responses*
Every day	22	37
Every other day	54	75
Weekly	159	168
Bi-weekly	70	46
Monthly or less	46	22

Table 4: Cleaning frequency before and after owning an iRobot Roomba in the study of Sung et al. (2008)

Besides the frequency of the cleaning and the added multitasking that the robot influences, a different activity, deemed by Sung et al. (2008) as "roombarization," occurs. Roombarization is the activity where householders make physical changes to their homes to help their vacuum robot clean. Common types of roombarization include picking up wires, changing furniture layout, and tucking in the tassels of a rug (Sung et al., 2007).

The amount of roombarization that the user is willing to do, combined with the layout of the space, and the amount of clutter, impacts the effectiveness of the robot. This roombarization can be seen as a human-robot collaboration, which is underlined by a quote from one of the interviewees of Forlizzi and DiSalvo (2006), who characterized it as *"...the way you work with the thing, I have to help it so it can do its job; it is like we are partners. We were working together, you know; that is what I like, because I knew if I exposed an area, it would probably get to it, and if I did not, it probably would not"* This collaboration shows a shift from manual cleaning to a more supervisory role during the cleaning process.

2.2.5. Social and Emotional Aspects of Domestic Robots

In a study by Linjawi & Moore (2018), the different interactions of robots are classified into three categories. The interaction of the robot with the environment and objects in the environment is called physical interaction. The interaction of robots with other intelligent systems and other robots is called cognitive or perceptual interaction, and the interaction of the robot with humans is called social interaction. It has been shown that these social aspects play a crucial role in the adoption of technology (Venkatesh, 2000). Forlizzi and DiSalvo (2006) observed that, besides changing their cleaning patterns or physical arrangement of their homes, people also developed relationships with their cleaning robots by assigning a name or ascribing personality traits to the robot.

Fink et al. (2011) argue that by feeling socially attached to a robot vacuum, people are able to gain more pleasure from cleaning and will invest more effort into fitting the robot into their homes. Scopelliti et al. (2005) conclude that the acceptability of robotic devices in a home setting does not only depend on the functional benefits of the robot but also on the social aspects between the robot and human. A concept that aligns with this topic is anthropomorphism, which Epley et al. (2007) describe as the tendency to "imbue the real or imagined behavior of nonhuman agents with humanlike characteristics, motivations, intentions, or emotions." Forlizzi (2007) found that users were more likely to anthropomorphize autonomous robots compared to other electronic household products.

Users assign human characteristics, gender, personality, and a name to their robots, forming a relationship with the product. Users display some signs of affection, intimacy, and other kinds of sentiments for the robots in various ways, which in some cases is more important than how functional the product is. For example, in a study by Yapici et al. (2022), an interviewee stated that she and her brother call their vacuum robot Alfonso and added, *"It is not the same as my relationship with other domestic appliances or electronic devices. I enjoy it more when dealing with it humorously."* This leads her to accept product issues more easily: *"I can become annoyed; I said it a few times: This robot is not that smart after all. When we had internet problems, he (the robot) could not find the charging stand because of the unavailable mapping feature. He wandered around too much, and he finally ran out of battery under the bed. I said many times: 'Where is this idiot?' and 'This idiot has gotten stuck in a weird place.' We looked all over for him. Then he finally came out from under the bed."*

This social aspect should be kept in mind when working with intelligent, autonomous systems, as it can be useful to help users create a more solid bond with the robot, strengthening product trust.

2.2.6. User Needs and Pain Points

Despite the increase in research, knowledge, and popularity, even the most unproblematic robots still face technical and functional limitations that can have a direct effect on the user experience. An example of this can be taken from the study of Yapici et al. (2022). A participant of their survey had trouble with the robot vacuum not perceiving the environment correctly and therefore knocking down valuable items such as her guitar which was standing on the floor in her living room. After this happened, she admitted to using the robot less often. Other common issues include the robot getting stuck on clothing, cables, or other items on the floor. Some robots can also get stuck while trying to traverse the legs of dining chairs, or they are unable to find their charging stations (Sung et al., 2007; Yapici et al., 2022). These design flaws require users to modify their homes to accommodate for the robot, which can be very frustrating (Forlizzi, 2008).

These issues can translate into concrete pain points that shape the user's view of the robot and the brand of the robot. Another recurring problem is the lack of understanding of how the robot works and its behavior (Yapici et al., 2022). The autonomous robots determine their movements using various path-planning algorithms, which are optimized to find the most efficient paths. However, Kim et al. (2007) argue that there can be a mismatch between the optimal route of the robot and the user's view on what would be the best route. An example of this could be that the robot would clean the living room and the kitchen together along the path that lets it complete the task in the minimum needed time, but the user would want it to clean the entire living room first before moving on to the kitchen (Forlizzi & DiSalvo, 2006). Users can view the movement of the robot as incorrect when it does not align with their expectations, and there might be a mismatch between the user's needs and a technically optimal design.

For traditional products, the user is often familiar with the product's concept and functions and is generally well aware of their needs based on their previous experience. However, with a relatively novel product like an autonomous robot, it can become more difficult to implement user needs into the product's technical development, as most users have little to no knowledge about robotic technology and are inexperienced with using it (Björnfot et al., 2018). By focusing on user-centered design and transparent, intuitive interfaces, it might be possible to incorporate the user into the behavior of the robot and help the robot and user collaborate more easily with each other.

2.2.7. Conclusion

This literature exploration on domestic cleaning robots shows that these robots are not only autonomous cleaning products but also social technologies that influence the home environment and the people within it quite heavily. The presence of a robot vacuum can change users' behavior through multitasking, new cleaning routines, or even letting people modify their homes through roombarization.

The acceptance of the robot is dependent not only on its functional abilities but also on the emotional and social aspects between the user and the machine. Many users gain a bond with their robot by naming it, assigning characteristics, and developing a collaborative cleaning routine together. At the same time, consistent functional issues, like navigation errors, difficulties with scattered objects, or mismatches regarding user needs, can create distrust in the robot.

A design opportunity that presents itself is to bridge the gap between autonomous decision-making of the robot and more easily understandable behavior focusing on user needs. Potential solution spaces include improving transparency in robot behavior, creating clear communication of the robot's intentions, focusing more on user needs instead of optimal path planning, and enhancing collaboration by giving users more control over the robot's behavior. Furthermore, using the social tendencies that users already exhibit to the advantage of the design, such as anthropomorphism, might create more engaging, trustworthy interactions. By improving reliability, transparency, and social aspects, domestic cleaning robots can create more effective human-robot interactions and integrate better into the home environment.

2.3. AUTONOMOUS NAVIGATION AND MAPPING FOUNDATIONS

This chapter will provide the needed background information on autonomous navigation and mapping, with a focus on simultaneous localization and mapping (SLAM), as this is the current state-of-the-art for domestic cleaning robots. The chapter will discuss what SLAM is, how it works, and the different ways to look at the SLAM problem. Additionally, different mapping representations will be discussed, as well as the pitfalls of data association problems. This chapter should ultimately provide readers with the necessary knowledge to understand SLAM on a surface level and know what is possible to explore future solution spaces.

2.3.1. Autonomous Systems

A robot can be considered autonomous when it can navigate through an environment without any human interference. For this navigation to be successful, the robot must have a proper understanding of its environment and a steady, accurate tracking of its location within it (Taheri & Xia, 2020). When referring to the environment, it is helpful to understand that we make a distinction between the static elements, the fixed layout of the map of the environment and the objects present within, and the dynamic components within an environment, such as moving humans, other robots, or movable objects. SLAM mainly deals with the former, as it will aim to build a map of the static environment, while dealing with dynamic obstacles is usually taken care of during the navigation or path-planning stage.

To interpret and operate within an unknown environment, autonomous navigation systems typically follow a continuous loop of sensing, perceiving, planning, and acting (see Figure 10). The physical world represents the robot's external environment, the source of its sensory input, and the space in which it moves. The robot first collects data from this environment through its sensors, interprets that information to build an internal understanding of its surroundings, plans a path toward a goal, and finally follows that path while constantly updating its understanding of the environment.

2.3.2. What is SLAM?

To navigate autonomously, a robot must first understand its surroundings and determine its position within them. Simultaneous Localization and Mapping (SLAM) is known as the problem of constructing a map of an unknown environment while simultaneously estimating the position of the robot within this map, without relying on a prior map or external referencing systems such as GPS (Grisetti et al., 2010).

When we talk about localization, we refer to the task of a robot estimating its pose, which describes the location and orientation of the robot, in a known environment using its sensors' data (Thrun et al., 2005). Mapping, on the other hand, involves the construction of a model of the environment using sensor data. What makes the SLAM problem difficult is that localization and mapping are dependent on each other. The map of an environment is essential for a mobile robot to determine its location, while mapping requires accurate localization. This is often referred to as the "chicken and egg" problem. SLAM provides a framework that allows mobile robots to solve both problems at once, allowing for safe navigation in unknown and dynamic environments (Durrant-Whyte & Bailey, 2006).

To solve the SLAM problem, various frameworks and algorithms have been developed. These methods differ greatly in terms of sensors used, mathematical models applied, and how uncertainty is handled over time. Many SLAM systems also combine multiple methods, each algorithm addressing different aspects of the problem. Despite their differences, all SLAM models share the same fundamental goal: Estimating the robot's position and constructing a consistent and reliable map of the environment. The following section outlines the main components and working principles of state-of-the-art SLAM models.

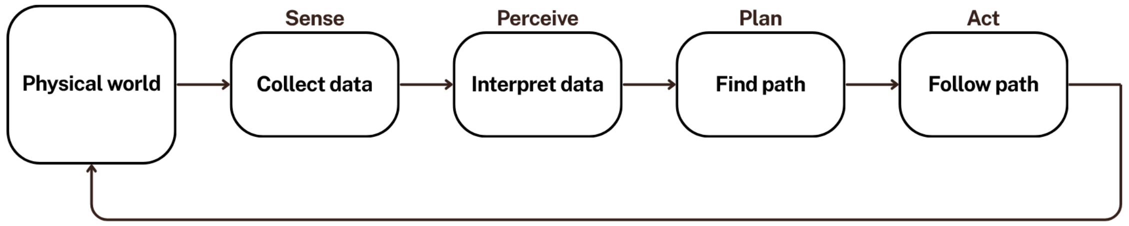

Figure 10: Workflow of an autonomous system

2.3.3. How does SLAM work?

In SLAM, the robot utilizes data gained from different sensors to build a map of its environment and estimate its pose relative to this map. This process can typically be divided into the front-end and the back-end, as can be seen in Figure 11:

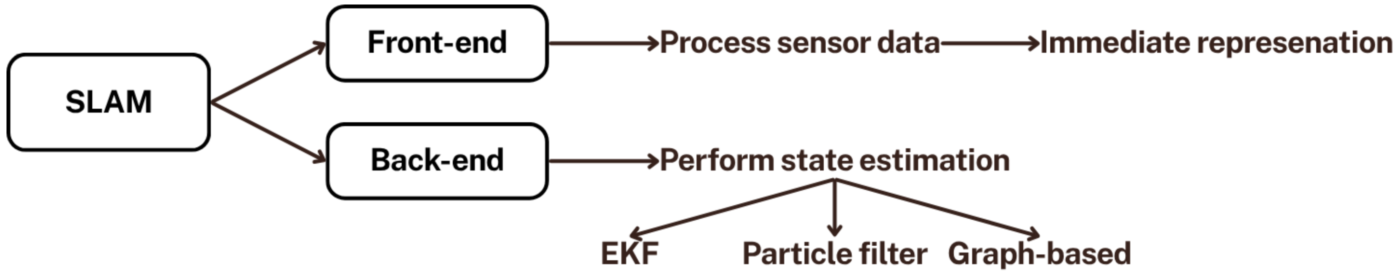

Figure 11: Front-end versus back-end SLAM

The front-end processes raw sensor data to gain an immediate representation of the environment. The back-end performs state estimation to refine the robot's trajectory and the map using mathematical models such as the Extended Kalman Filter (EKF), Particle Filter, or Graph-based optimization methods.

For localization, robots often rely on dead reckoning. Dead reckoning is a method to estimate the robot's pose based on its initial position, combined with measurements of direction, speed, and time. There are several types of dead reckoning models, depending on the sensors used. Common approaches include the use of an Inertial Measurement Unit (IMU) or odometry, where wheel encoders are used to determine motion (Taheri & Xia, 2020; Wang & Zhang, 2020). However, in the absence of a map, dead reckoning quickly accumulates error over time. SLAM helps minimize this error by using a map to "reset" its localization errors by revisiting known areas, a process called loop closure (Cadena et al., 2016).

Different SLAM approaches handle the back-end state estimation in different ways. We can distinguish two main directions, depending on how the state information is defined: Filtering and smoothing (see Figure 12) (Thrun, 2002; Durrant-Whyte and Bailey, 2006). To gain a better understanding of how SLAM operates, both approaches will be outlined below, with an in-depth explanation of the two most popular approaches.

2.3.3.1. Online SLAM

Filter-based methods, such as the EKF or particle filters, model the problem as an online state estimation. This means that the robot is estimating its state, consisting of its position, orientation, and the map, in real time as new sensor data becomes available. The system's state is determined by the most recent robot pose and map of the environment, and as new sensor data arrives, this estimate is continuously updated and refined. Meaning that each new estimate is building on the previous one, so the information from past measurements is already contained in the current state without the need to go back and process all earlier data again. To highlight this incremental nature, filtering approaches are usually referred to as online SLAM methods (Grisetti et al., 2010).

The extended Kalman filter is one of the earliest and most popular model-based algorithms that is being used for online SLAM (Choi et al., 2022). The Kalman filter is a model proposed by Rudolph E. Kalman in 1960. The KF is an algorithm to estimate the evolution of a dynamic system. It is particularly useful for combining noisy data from multiple sensors to gain a better estimate of the system's true state. For example, to gain a reliable estimate of your position, you could combine odometry with LIDAR measurements.

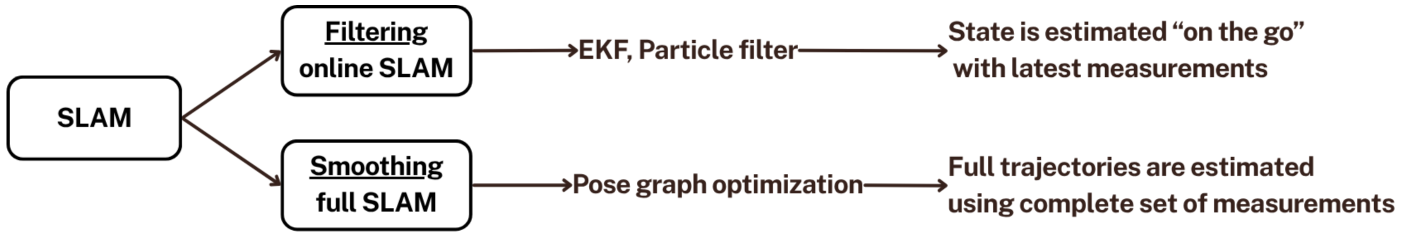

Figure 12: Filter versus smoothing SLAM approaches

The filter works by predicting the future robot state and determining the probability of this prediction. This hypothesis is combined with the measured data to correct and improve the future estimation at each time. This algorithm is suitable for dynamic systems by linking real-time measurement and predicting the state of system parameters through time approaches (Ulin-Avila & Ponce-Hernandez, 2021).

KF-based SLAM approaches are still widely used, as they are relatively simple to implement and can reliably converge to accurate estimates over time. However, this filtering approach also has some clear limitations. KFs do not inherently use loop closures and are sensitive to data association issues, which can lead to failure in estimating the overall system state. Furthermore, KFs are only made for linear systems, but most real and dynamic systems are nonlinear. A solution to still make use of a KF for SLAM is by using the Extended Kalman Filter (EKF), which can linearize nonlinear systems by using first-order Taylor expansion (Julier & Uhlmann, 2001).

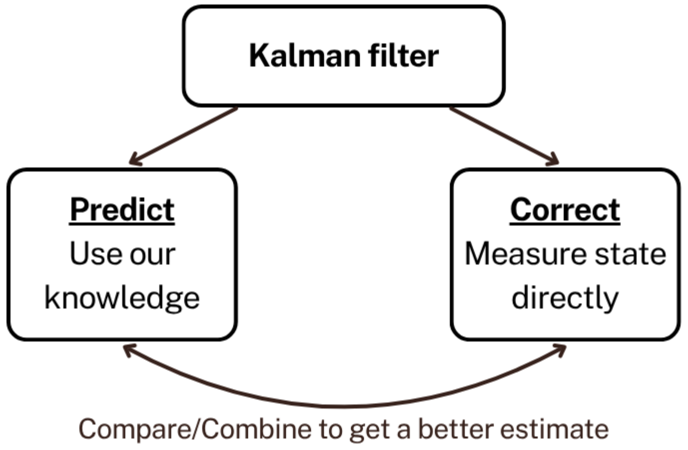

Figure 13: General structure of a Kalman filter

2.3.3.2. Full SLAM

In contrast to filtering approaches like the EKF, smoothing approaches estimate the full trajectory of the robot from the entire set of measurements. These approaches continuously combine predictions from the robot's motion, based on dead reckoning, with new sensor measurements and filter out the noise and uncertainty. This approach results in a gradually more accurate estimate of both the map and the robot's position. Smoothing approaches tackle the so-called full SLAM problem (Dellaert & Kaess, 2006).

An intuitive way to address the full SLAM problem is the so-called graph-based SLAM. The graph-based SLAM formulation was proposed by Lu and Milios in 1997. However, it did not become popular until several years later due to the high complexity of solving the error minimization problem. Later insights into the SLAM problem and advancements in the field of sparse linear algebra created more efficient approaches to the optimization problem that graph-based SLAM presented. Therefore, graph-based SLAM methods gained immense popularity and currently belong to the state-of-the-art techniques in terms of speed and accuracy (Grisetti et al., 2010).

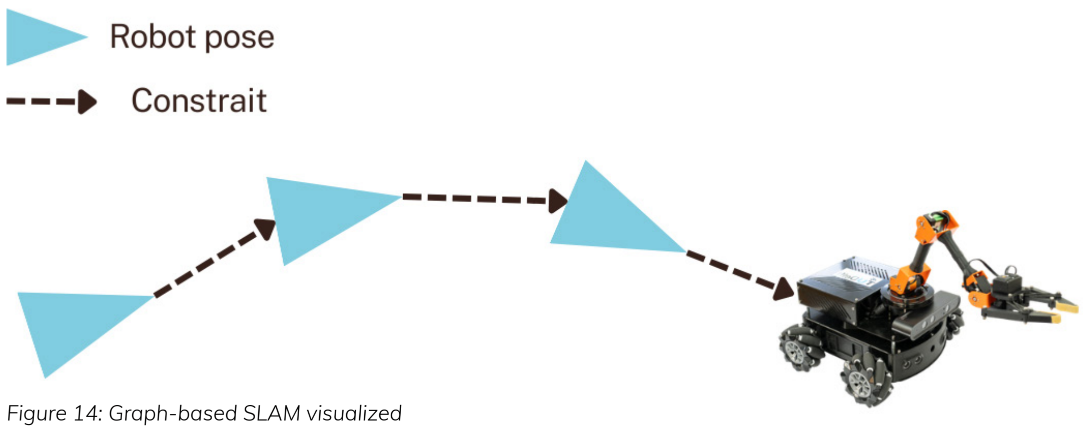

Figure 14: Graph-based SLAM visualized

Graph-based SLAM, also known as pose graph optimization, represents the robot's motion through an environment using a graph. To introduce this approach to SLAM, let's assume a robot is moving through an environment, and as the robot moves, it records its pose, represented by its position and orientation, at different time intervals. Each pose is represented as a node in the graph, and consecutive poses are connected by edges that represent constraints based on the dead reckoning data from the robot. These constraints visualize spatial relationships between different nodes, as can be seen in Figure 14. These constraints should, however, be considered soft due to the uncertainty in the sensor data (Lu & Milios, 1997).

When the robot continues to move through the environment, it may revisit an area where it has already been. This is where a loop closure, a key element of the SLAM problem, can occur. When the robot recognizes a previously visited place, using its sensor data, for example, LIDAR, it can not only create the sequential constraints but also create additional constraints linking the current pose to previous poses associated with the same location (Steder et al., 2007). This loop closure event is demonstrated in Figure 15.

A graph representing the SLAM problem can be created from the collection of interconnected nodes, where each node corresponds to a pose of the robot during mapping, and each edge corresponds to a spatial relationship between poses. The goal for graph-based SLAM is to build this graph and find the optimal node configuration that minimizes the total error introduced by the different constraints. An intuitive way of looking at this error optimization is to imagine each edge as a rubber band connecting two nodes. Every rubber band has an ideal length based on the estimated distance between poses, and it will pull on the nodes to reach its ideal length. The different nodes are being pulled at by multiple rubber bands, and the optimization process tries to adjust the node positions to create the best overall balance, or the configuration of nodes where all the rubber bands are stretched as little as possible.

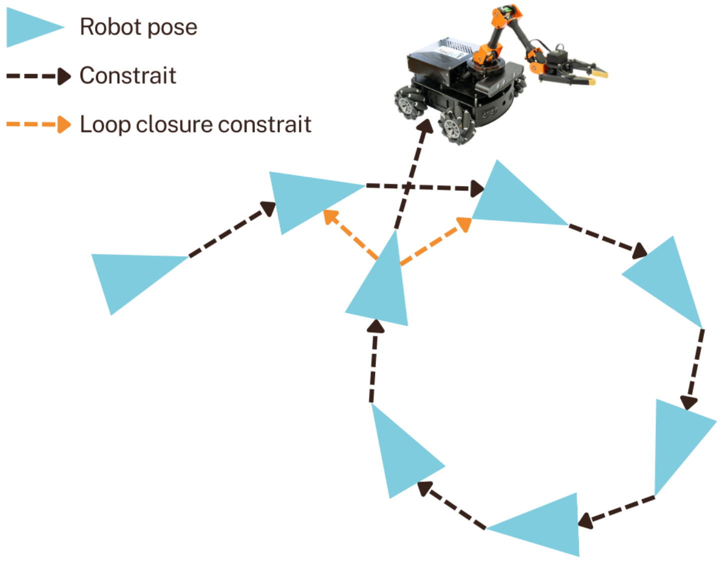

Figure 15: Loop closures using graph-based SLAM visualized

This minimization of the error introduced by the uncertainty in the measurements is also known as the least-squares optimization process (Schubert et al., 2021). In Figure 16, this process is illustrated using real mapping data collected at MIT's Killian Court by Bosse et al. (2004).

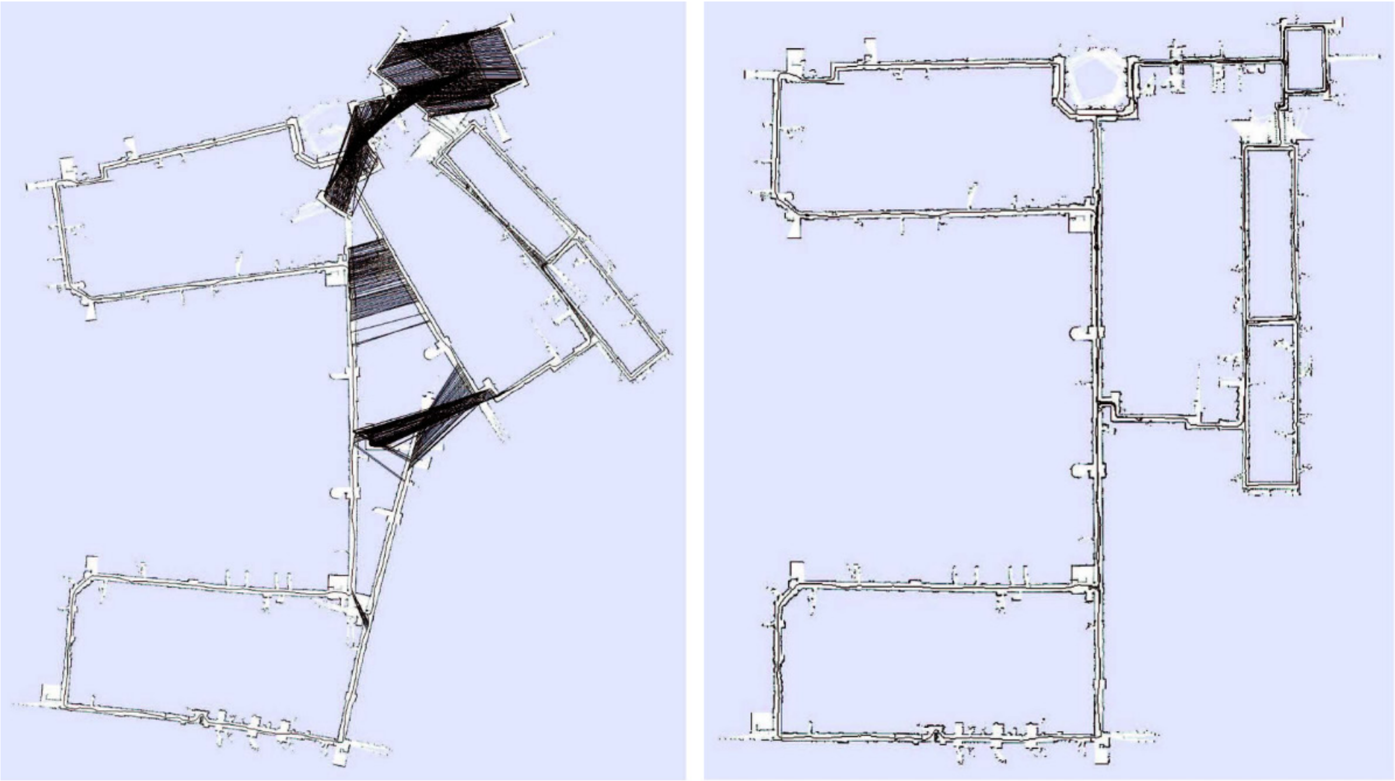

Figure 16: Pose-graph corresponding to a data set recorded at MIT Killian Court before (left) and after optimization (right). The maps are obtained by rendering the laser scans according to the robot positions in the graph (Grisetti et al., 2010).

2.3.3.3. Mapping Representations

For both SLAM approaches, filtering or smoothing, maps can be represented in different ways. It is important to distinguish between the internal representation of the map by the robot and the way that the map will be displayed to a user.

Internally, maps can be simulated as a set of spatially located landmarks or by using dense representations like occupancy grids, where the environment is divided into a grid and each cell is filled with a probability of being occupied. Other potential options are surface maps or simply using raw sensor data. The choice of the map representation depends on the sensors used, the characteristics of the environment, and the estimation algorithm used for SLAM. Landmark maps are often the preferred option in environments filled with easily identifiable and distinct features and are particularly effective when using visual sensors like cameras (Montemerlo et al., 2002). On the contrary, dense representations like occupancy grids are more appropriate when using range sensors like LIDAR (Triebel et al., 2006).

When displaying a map to a user, it can be visualized in various ways, from simple 2D occupancy grids to detailed 3D models. These visualizations can be fully designed for human interpretation and may thus differ completely from the robot's internal map representation.

2.3.4. The Data Association Problem

A fundamental challenge for SLAM algorithms is the data association problem. Data association errors occur when a robot incorrectly identifies two different landmarks as being the same feature observed from a different position (Neira & Tardos, 2001). Data association is crucial, as a wrongly associated landmark can lead to significant localization and mapping errors. The term "data association" also refers to inaccuracy in landmark recognition. This happens when an algorithm is not capable of identifying two slightly different features in an environment. When two different landmarks have a similar shape, it can be difficult for the robot to differentiate between the two. For example, when two different corridors in two different locations within the same environment look the same to a robot, this could have a detrimental effect on the robot's estimation of the environment (Dissanayake et al., 2011; Khairuddin et al., 2015).

Furthermore, due to imperfections in the system, caused by drift, sensor uncertainty, and other noise, the actual movement of the robot will deviate from the estimated motion, which is another source of data association uncertainty to emerge (Ho et al., 2015).

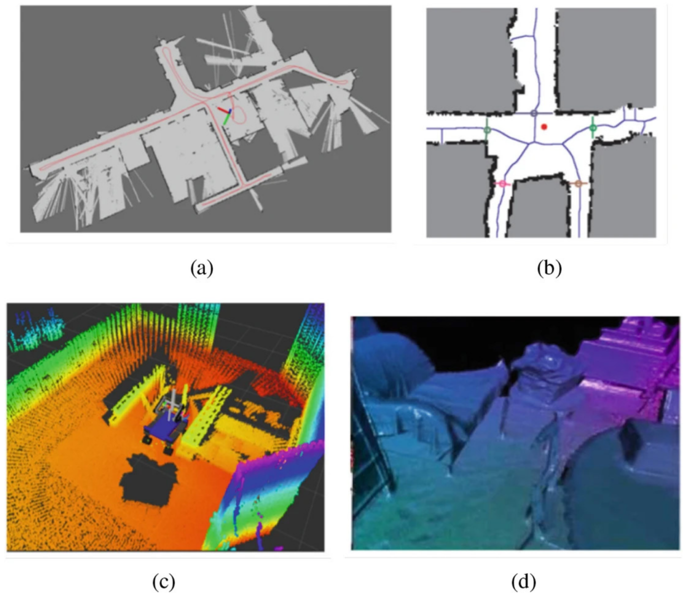

Figure 17: Illustrations of different mapping outputs. (a) 2D grid map. (b) 2D topological map. (c) 3D point cloud map. (d) 3D mesh map (Xu et al., 2024)

2.3.5. Conclusion

SLAM is the foundation for enabling autonomous systems to gain an understanding of their environment and interact with it. Through simultaneous localization and mapping, robots gain the ability to operate safely in a dynamic, unknown context. However, due to the complexity of SLAM, it remains difficult to translate this technology into a form that users can easily understand and confidently interact with. This creates an interesting design opportunity: bridge the gap between complex algorithmic operations and intuitive, user-friendly interaction. Understanding how SLAM works creates a valuable base for the development of a user interface that makes for a more transparent, accessible, and supportive autonomous system.

2.4. BENCHMARKING

To gain a better understanding of the current domestic cleaning robots and their interfaces, the following benchmarking study was conducted. The goal of this analysis is to get an overview of the existing cleaning robots and their mapping approaches, along with the limitations and opportunities currently present.

2.4.1. Mapping Hardware

While SLAM remains the dominant mapping approach in domestic robots, the underlying sensing hardware does have some variations across products. Most cleaning robots rely on LiDAR (light detection and ranging) sensors, camera-based sensing, or a hybrid version combining both (Debeunne & Vivet, 2020).

LiDAR sensors provide the robot with relatively accurate distance measurements of the environment and are widely used in applications that require spatial data. Camera-based mapping approaches provide the ability to gain semantic data of the environment, meaning that the robot can get an understanding of what is a wall and what is, for example, a movable object like a chair. While recent studies are moving more towards hybrid approaches that combine LiDAR and camera data (Zhu et al., 2023), when looking for reliability and consistency, using LiDAR is the best choice.

2.4.2. Domestic Cleaning Robot Apps

To let users interact with their domestic cleaning robot, almost all companies have created a mobile application via which the user can control their robot. To understand how the human-robot interaction is currently shaped, several commercial robot vacuum cleaning apps have been explored.

When looking at the different applications, it quickly becomes clear that almost all work the same way. The app typically presents a simplified floor plan to the user that was created from the SLAM data. The user is often given some interaction features that consist of assigning no-go zones, room labeling, assigning special cleaning zones, and scheduling cleaning. Some newer systems incorporate object recognition, however, this is often still very novel and works sparingly.

While the visual style differs between the different applications, the interaction process is similar. An abstraction of the robot mapping data is used as both the main visual to the user and the primary control interface. Therefore, it is vital that this map is correct, as it is the central element through which the user interacts with and understands the robot.

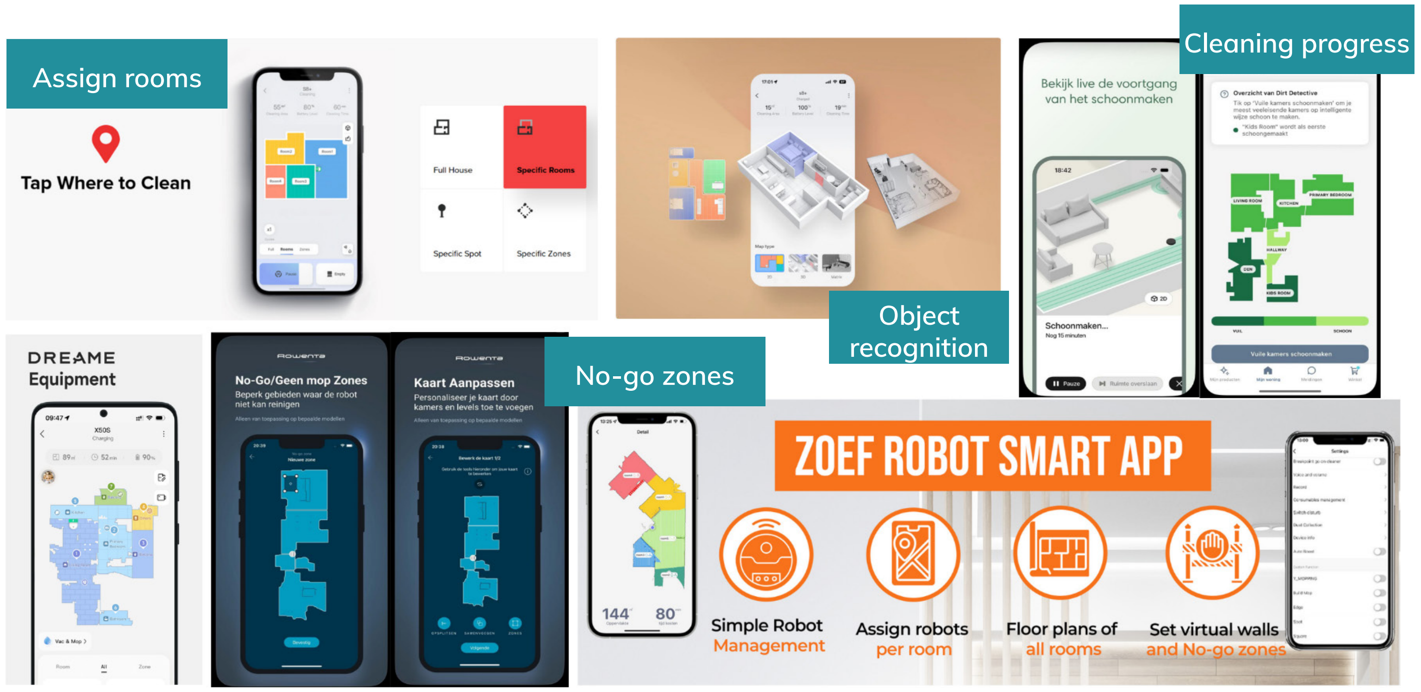

Figure 18: Domestic cleaning robot application benchmarking; applications owned by: Roborock, Rowenta Robots, Zoef robot, DREAME, iRobot

2.4.3. Pain Points in Current Systems

Although domestic cleaning robots have come a long way from early random navigation systems, there are still some recurring usability challenges and issues that can be observed in these products. One major limitation is the lack of direct map editing capabilities offered to the user. In current systems, the mapping process is often fully automated, the user has to start the mapping behavior, but the rest happens automatically, and once completed, the user has little influence on the map. Issues commonly arise when robots have to map spaces with reflective areas, like windows or mirrors, or the robot fails to find all the rooms and gets stuck. Often the only way to fix major errors in the map is to restart the mapping procedure and hope for an improved result.

Another source of frustration is robots getting stuck in cluttered spaces, like between chair or table legs, even when there should be enough space. This issue is often hard for the user to grasp, as the applications offer little transparency on how the robot perceives its environment.

Map readability presents another issue. The generated map abstractions can appear either too messy or too abstract, leaving out vital information or overly simplifying rooms. When users see a certain error in the map that does not align with their idea of their home, it can reduce their trust in the system.

Ultimately the low levels of user control over the map can harm the overall user experience. Especially when errors or unwanted behavior occur, it can lead to significant frustration when the user cannot control or edit the map accordingly.

2.4.4. Conclusion

The benchmarking once again highlights the friction between automation and user control in robotics. While the SLAM systems support relatively high levels of autonomy, this often limits meaningful user interaction, which can damage the system, especially when errors occur. These findings suggest that future robot mapping applications should provide users with more control and perhaps add a certain user-in-the-loop interaction experience. Providing users with intuitive tools to edit and correct the map where necessary could avoid a lot of the problems and limitations currently found during the benchmarking.

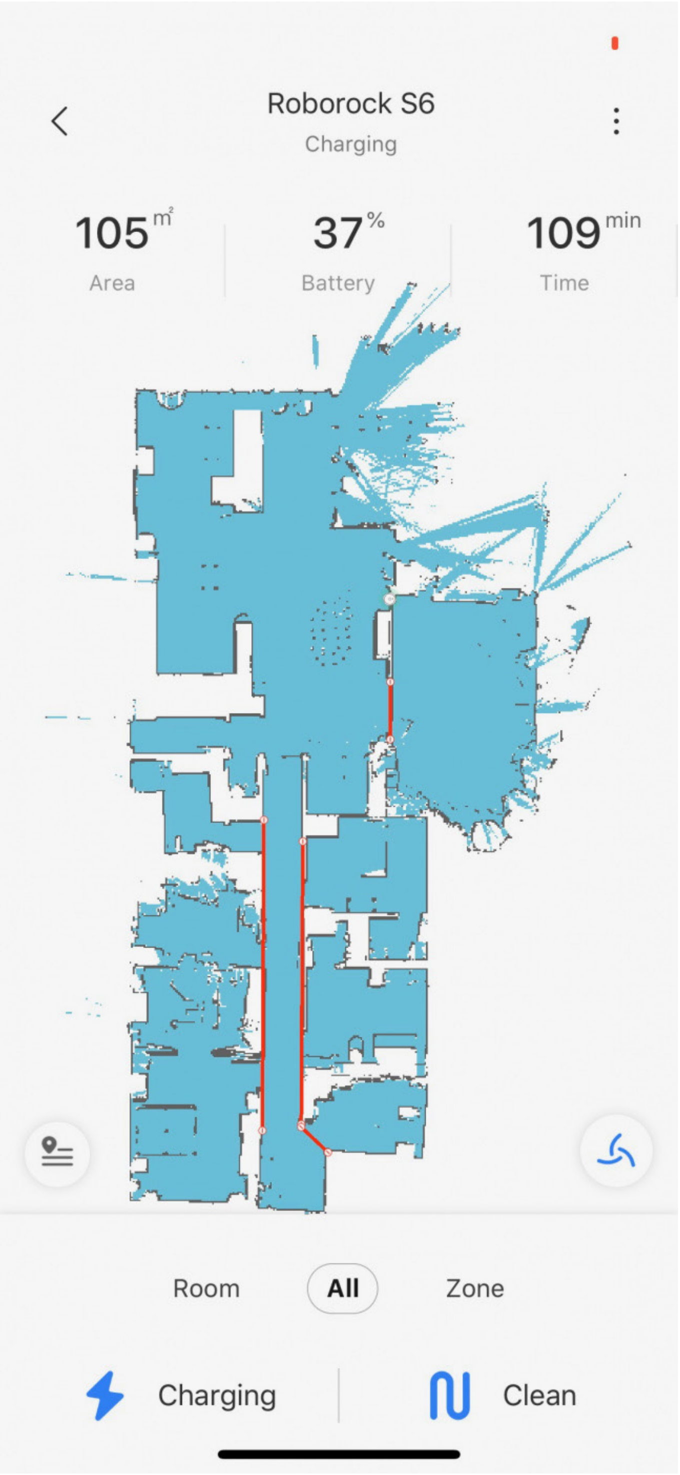

Figure 19: Noisy map visualisation in the Roborock app

2.5. CUSTOMER JOURNEY MAP

A customer journey map was created to gain a proper overview of the key positive and negative experiences of owning a domestic cleaning robot.

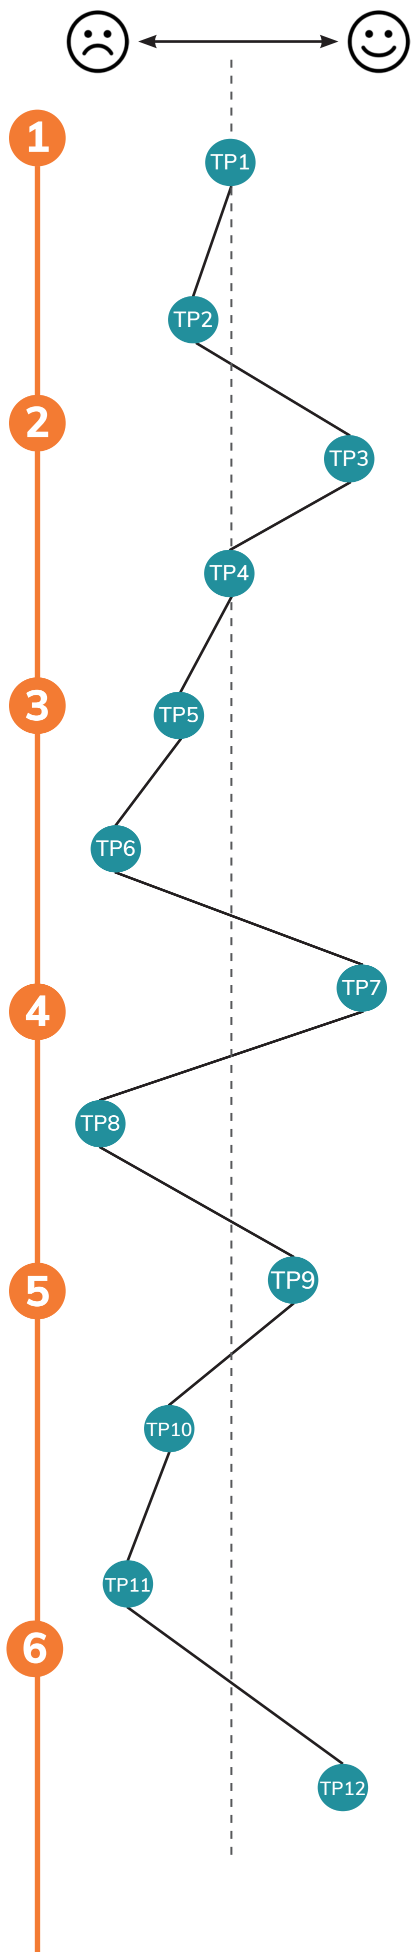

Awareness/Ordering
The user has become aware of Roombas and is now researching whether to purchase one and which model would best meet their needs.

Touch Point 1
First contact with mapping, reading about the smart mapping features online.

Touch Point 2
Confusion about the immense options when choosing a Roomba.

Receiving/Unboxing
The user receives the Roomba delivery, unboxes the product, and gains access to the manual and mobile app.

Touch Point 3
The first time that the user gets the Roomba and can unbox it.

Touch Point 4
Reading the instructions and downloading the mobile app

Initial Setup & Mapping
The user sets up the Roomba, connects it to the app, and begins the initial home mapping.

Touch Point 5
Initial setup of the Roomba and connecting to the app can be quite overwhelming

Touch Point 6
Map shown in app is inaccurate at some points.

First Cleaning Run
The user starts the first cleaning run using the generated home map.

Touch Point 7
Roomba is set out on its first cleaning run which can be roughly followed in the app.

Touch Point 8
The roomba misses certain areas while cleaning due to incorrect mapping.

Manual Customization
The user customizes the map by naming rooms, setting boundaries, and adjusting zones.

Touch Point 9
App helps users to label rooms, split or merge rooms and add no-go zones.

Touch Point 10
Customization can feel clunky and not perform exactly like the user wants to.

Map Adjustments
The user updates the map after moving furniture or changing the room layout.

Touch Point 11
Changing the furniture setup requires updating the map, which may take multiple runs and may not always detect changes correctly.

Touch Point 12
User can edit the map themselves, editing the map to remove inaccuracies or change the layout

2.6. ETHICS & PRIVACY

The use of autonomous robots in private spaces like the domestic environment presents various ethical considerations related to privacy, autonomy, and transparency. As stated before, the domestic environment is an intimate space full of personal and emotional routines and behaviors. Therefore, even a rather technical use case like the spatial mapping of this space can have serious ethical implications.

Domestic cleaning robots create a spatial representation of your private home, this is needed to properly clean. However, these maps may also reveal room layouts, furniture placement, and behavioral patterns, which altogether is sensitive private data that big companies can use for targeted advertisement. With robots becoming more present in our daily lives, concerns about data safety and privacy have become important topics in HRI research. Studies have shown that privacy concerns affect use intention significantly and should always be taken into consideration and treated with the utmost of care (Lutz & Tamò-Larrieux, 2021).

Robot vacuum cleaners are equipped with different sensors and nowadays often utilize cameras for mapping. In 2022, the MIT Technology Review acquired intimate pictures captured from low angles taken by a development version of iRobot's Roomba J7 series. The picture that was taken of a woman on the toilet (see Figure 20) went viral and started huge privacy debates (Guo, 2024). iRobot confirmed that the pictures were taken by a Roomba as part of a development process for training AI in object recognition. Even though the images originated from robots that were still in development rather than consumer products, this incident still shows the potential privacy risks present when introducing robots into your home.

For this project privacy implications were considered in limiting the sensor usage and purposefully staying away from cameras. By avoiding camera-based mapping and instead relying on LiDAR, IMU, and odometry data, the risk of visual privacy being harmed is reduced significantly. The objective is to use human interaction while maintaining the same level of meaning that would be obtained from using cameras for mapping.

At the same time, privacy problems can never really be avoided. This has been displayed by Sami et al. (2020). In a study it was shown that LiDAR sensors can be modified to capture private conversations by recording the vibrations. While scenarios like that are outside of the scope of the project, it does show the complexity of securing privacy in consumer robotics.

This project does not aim to resolve all ethical challenges found in domestic robotics, despite this, it is important to realize that working with robots, interfaces, and private data is a delicate subject. Decisions about data representation, transparency, and user control can have a big influence on the privacy and trust of the user. By focusing on transparency, limited sensing, and user influence, this project tries to make ethical design considerations.

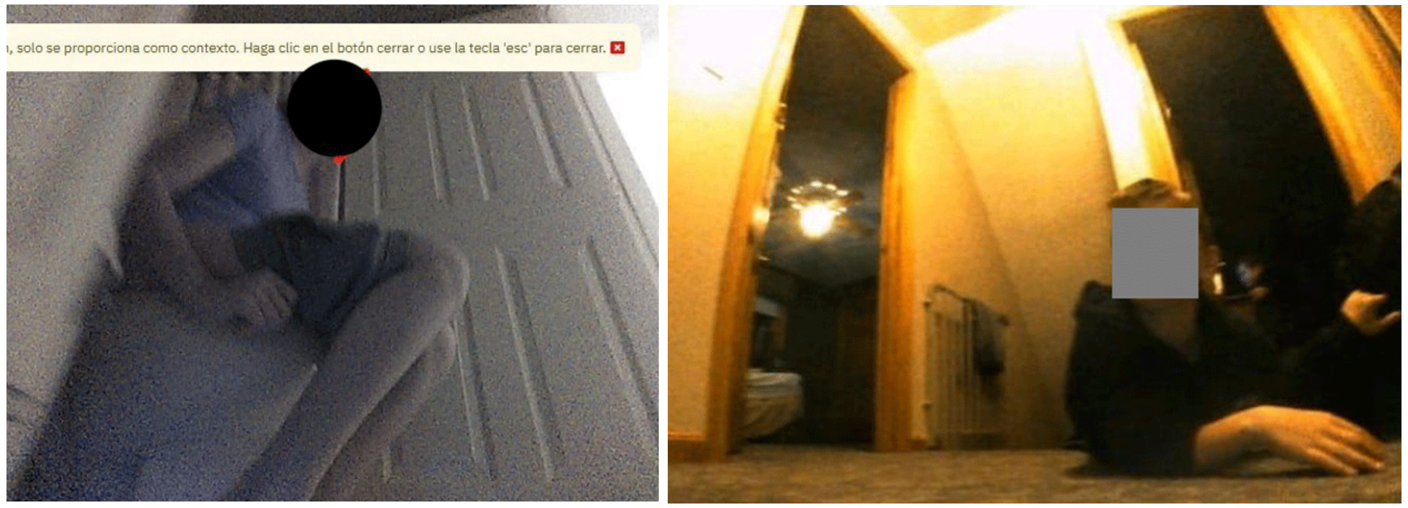

Figure 20: Images captured by iRobot development devices. Faces, where visible, have been obscured by MIT Technology Review.

DESIGN PROCESS

03

This chapter concludes the research phase of this project and introduces the design phase. It outlines the design brief, which will guide the design process, explains the initial interface prototype used in the first testing, and the different design choices that were made to get to this prototype.

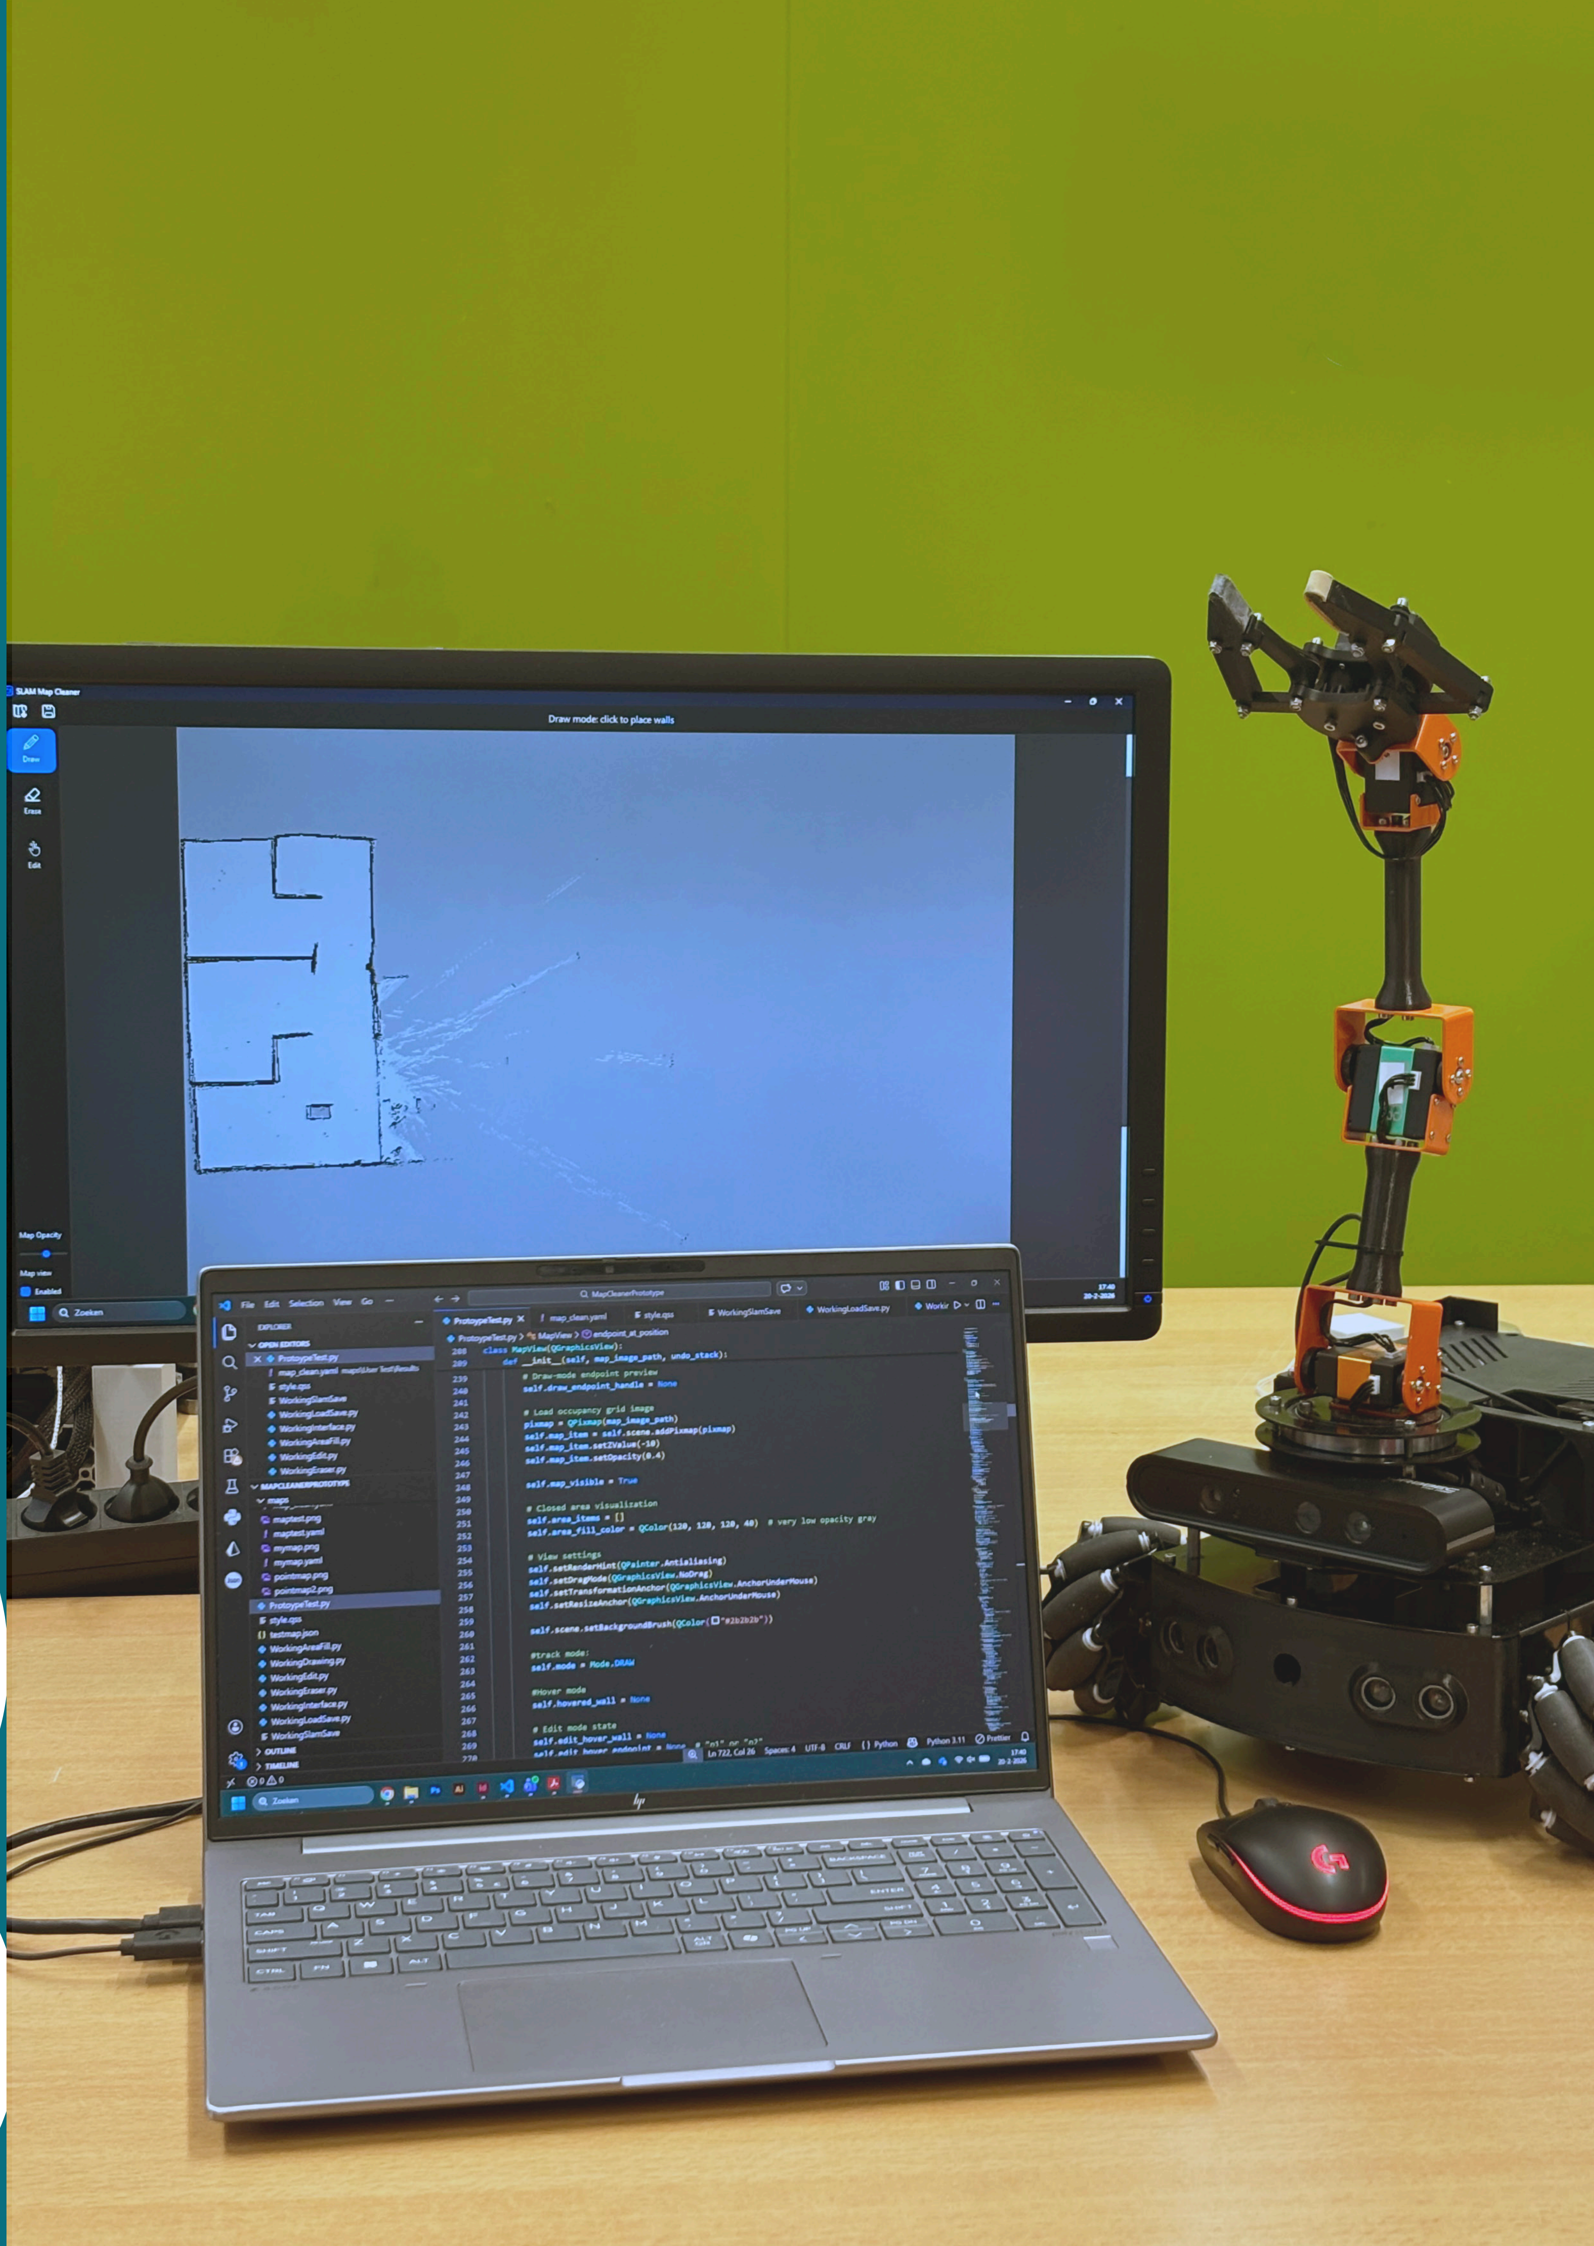

MAIN FINDINGS

Current System

Mapping is robot-driven

- Autonomous SLAM mapping
- User interacts after errors
- Navigation map ≠ Visualisation map
- Robot holds map ownership

The domestic environment

Homes are ever-changing environments

- Dynamic & cluttered
- Layout is constantly changing
- Personal & emotional space
- Users already adapt home to the robot

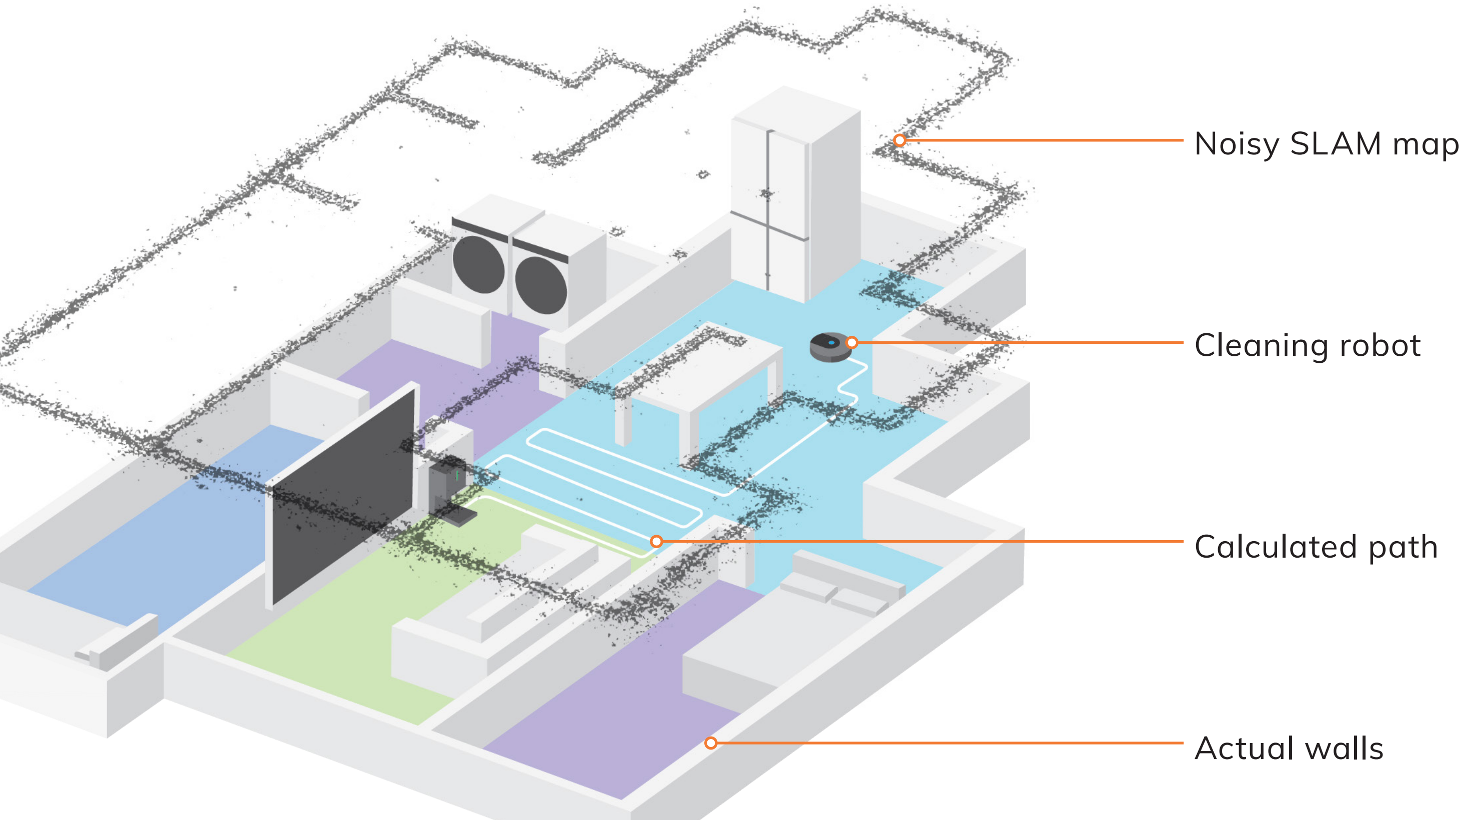

Pain points

Where misalignment emerges

The robot's internal map representation does not align with the user's mental model of the environment.

Robotic decision-making does not always align with users' expectations regarding appropriate system behavior

Consequences

Confusion

Frustration

Reduced Trust

ROBOT-OWNED → SHARED MAPPING

COLLABORATION • TRANSPARENCY • ALIGNMENT • TRUST

3.1. DESIGN BRIEF

The design brief will discuss the main findings of the research phase and combine this with the initial problem definition to clarify the thought process behind the design vision. It will highlight the different insights gathered and translate these insights into a design vision and goal by identifying problems & potential solutions.

3.1.1. Problems & Solutions

The main findings are visualized on the page on the left and give an overview of the big tension points that were revealed when looking into the context of this project. This section will elaborate and summarize what and how these tensions are created and translate them into design directions.

3.1.1.1. The current system

Domestic cleaning robots rely on autonomous SLAM-based mapping techniques. In these systems, the mapping process is primarily robot-driven. The user often needs to initiate the mapping behavior, but after that the process is almost entirely performed by the robot. The robot will create a map of the domestic environment, and after the map is created, the user can sometimes establish some room segmentations; however, the map is already built by this time, meaning that the walls are in place, often containing errors, missed spots, or noisy data, with little the user can do to fix this.

In addition to this, the internal maps that are used for navigation are made for robot understanding and often not created for humans to interpret. Resulting in abstract, noisy visualizations that do not match the mental model of the user.

3.1.1.2. The domestic environment

In contrast to the technical mapping systems, the home environments of users are spaces that are dynamic and intimate, filled with personal routines and emotional meaning. Unlike the controlled labs in which robots are often being tested, the houses in which they will be placed are ever-changing and unpredictable. Research already showed that users adapt their behavior and environment to aid the robot system, creating a human-robot collaboration. However, this collaborative effect is not taken advantage of by the technical systems, and making actual changes to the way the robot works is often difficult or even impossible.

Problem:

The robot's internal map does not align with the user's mental model

The robot uses probabilistic, sensor-based SLAM maps while users interpret their home by routines and functionality:

- Rooms have meaning (kitchen, childrens room, play-zone)
- Objects are temporary and moveable
- Cleaning priorities are personal and differ per household.

Robot internal map ≠ the user's idea of their home

Consequences:

- Maps do not look like home
- Noisy maps lead to confusion (noise due to sensors/windows/mirrors/clutter)
- Perceived errors in the map lead to decrease in trust.

Problem:

Robot decision-making and user's expectations of optimal behavior do not always match

Autonomous navigation makes choices based on efficiency and algorithmic optimality, this does not always align with human logic or preferences.

- e.g. Human wants robot to clean the kitchen first and then move to the living room.
- High levels of autonomy decrease situational awareness, meaning the users has little knowledge of why the robot is acting the way it is.

Due to high robot autonomy the user has no chance of understanding the robot.

Consequences:

- Lack of understanding leads to confusion
- Mismatches in behavior will lead to frustration and decrease in trust.

Key Opportunities		
Collaboration	Alignment	Transparency
Mapping is currently largely robot-driven, leaving little room for the user to interact. Introducing collaborative human-robot interaction to the mapping process could enhance the experience.	User experience issues are largely the fault of misalignment between robot and user. Designing interactions that better align human and robot can reduce frustrations and create a better overall experience.	Lack of transparency into the robot's decision making and perception creates a lack of understanding. Increasing transparency into creation and interpretation of the map can improve user understanding.

3.1.2. Design Vision

"Design a collaborative mapping experience where human and robot work together to reliably and efficiently map domestic environments, while aiming for alignment, transparency, and more intuitive human-robot interaction."

3.1.3. Design Goal

The goal of this project is to design an interactive robot interface allowing users to interact with real robot mapping data and collaborate with the robot to create cleanly refined maps. The interface should be intuitive, and users should have an easy time grasping the concept and learning how to create valuable maps. The final result of the project will be a tested concept that will showcase the added value of introducing a user-in-the-loop approach in robot systems.

3.2. IDEATION

After the design brief was established, the ideation phase could begin to take shape. With the brief defining the design direction towards collaborative mapping, further exploration was needed to determine how this concept coudl be translated to a practical use-case.

In the beginning of the ideation, it was important to keep this view broad, to not focus on one single collaborative mapping approach. However, one of the main requisites of this project was to ultimately develop a prototype that uses real robot sensor data and would let users interact with this collected data. This meant that the ideas should be technically possible considering the robot used, the complexity of the implementation, and the remaining time.

Because of this, it was important to constantly test and experiment with the robot during the ideation phase. To keep a balance between ambition and feasibility. Ideation and prototyping happened parallel to each other, concepts were tested using the robot, and by exploring the robot and sensor data, new ideas came to be. This iterative loop of robot exploration and ideation shaped the future direction of the project.

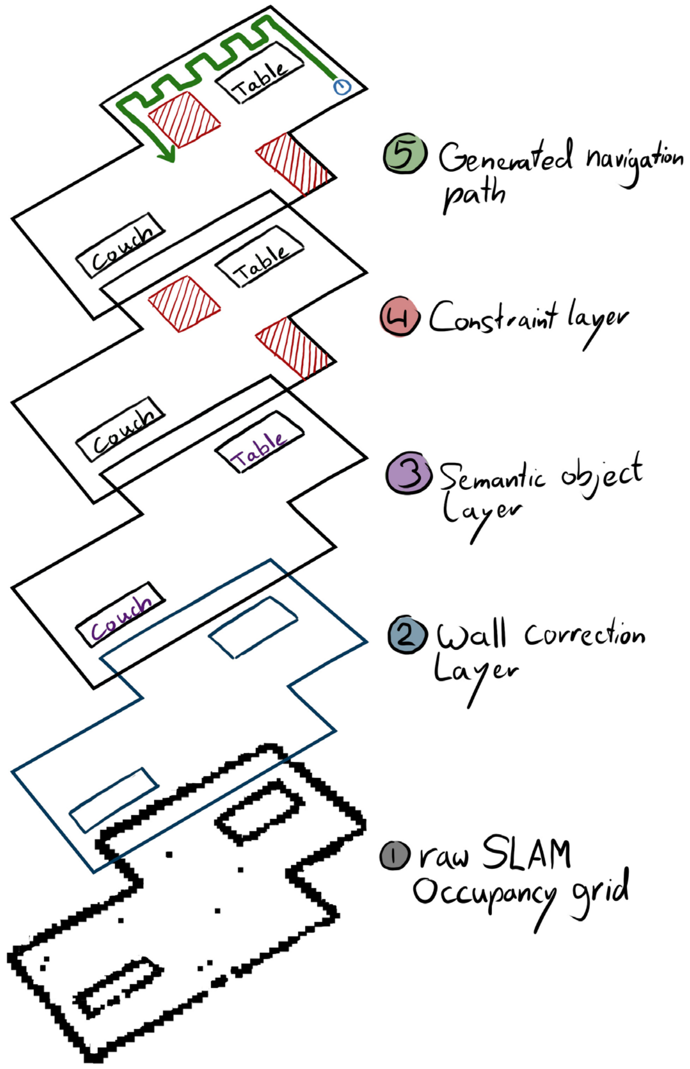

Figure 22: Layered map concept

3.2.1. Layered Mapping

The main outcome of this ideation phase is the layered mapping concept. Instead of treating the robot map as a single static representation, the view of this was changed to multiple informational layers building on top of the original SLAM map.

The core idea behind this framework is to preserve the technical integrity of the original SLAM map but keep adding information and interaction possibilities in every layer. By utilizing this layering approach, the map will gradually become easier to understand and interact with.

The layered map concept consists of five layers:

5) Generated navigation path
The final layer helps visualize the output of the system, it will show the navigational path the robot will take based on all the previous layers. This layer helps the user to understand the robot's reasoning and behavior and gives the user a chance to change this behavior.

4) Constraint layer
The constraint layer adds behavioral rules to the system that will influence the robot's navigation and decision-making. Examples of this include no-go zones, areas requiring additional cleaning attention, and time-specific constraints.

3) Semantic object layer
To give more meaning to the map, this layer will add semantic information. Objects and rooms can be labeled and given attributes, letting the robot use this information for more intelligent cleaning.

2) Wall correction layer
The function of the second layer is to transform the occupancy grid into a refined, readable map. In this layer, errors, excessive noise, and dynamic objects present in the map can be removed.

1) Raw SLAM Occupancy grid
This layer shows the raw output of the SLAM mapping, it consists of an occupancy grid showing free space, obstacles, and uncertainty. This layer forms the foundation of the layer stack.

3.2.2. Focus Point

Although the layered framework offers plenty of design directions, for the remainder of this project it was necessary to define a clear scope that would maximize a meaningful outcome in the remaining time. As not all layers offered the same potential of innovation, feasibility, and relevance to the design brief, a choice was made.

The remainder of the project will focus mainly on layer 2, the wall correction layer. This decision was driven by the following three factors:

1) Meaningful human-robot interaction
In layer two the goal is to translate the noisy, error-filled SLAM map to a cleanly refined, readable map. Researching how well humans can understand the SLAM map, clean it to their liking, and see if this would help the navigation of the robot or maybe even harm it by losing information or adding wrong information could reveal interesting results.

2) High impact on the system
Structural errors in the map will transfer throughout the entire system. Small issues in the walls and sensor data could cause inefficient paths, the robot getting stuck, and user frustration. If an introduction of a user at this stage could eliminate a lot of these errors, the impact on the entire system would be significant. In addition to this, allowing users to intervene at this early stage would also help them gain valuable insights into the robot system.

3) Innovation
Some later stages, for example the constraint layer and the semantics layers, are already partially explored in current commercial products. Manual correction of SLAM mapping by non-technical users is relatively unexplored, making it an interesting and meaningful area for innovation.

3.2.3. Robot Exploration

During the ideation phase, there was a continuous focus on the exploration of the robot, the understanding of the SLAM mapping, and the ROS ecosystem that the robot operates in. Rather than creating an idea or concept and then worrying about the implementation at the end, it was necessary to constantly create prototypes using the robot to see if ideas would be feasible. This way of working also helped gain new ideas by understanding more about the sensor data and visualizations.

An important challenge during this stage was the real-time visualization of SLAM data. This involved exploring how to use this data, modify it to filter out errors and noise, and present it to the user. In Figure 23 the evolution of different SLAM visualizations can be seen, these iterations ultimately led to a clearer understanding of how the technical mapping data could be used and presented to the user and also how we could create valuable output that could later be transferred to the robot for navigation.

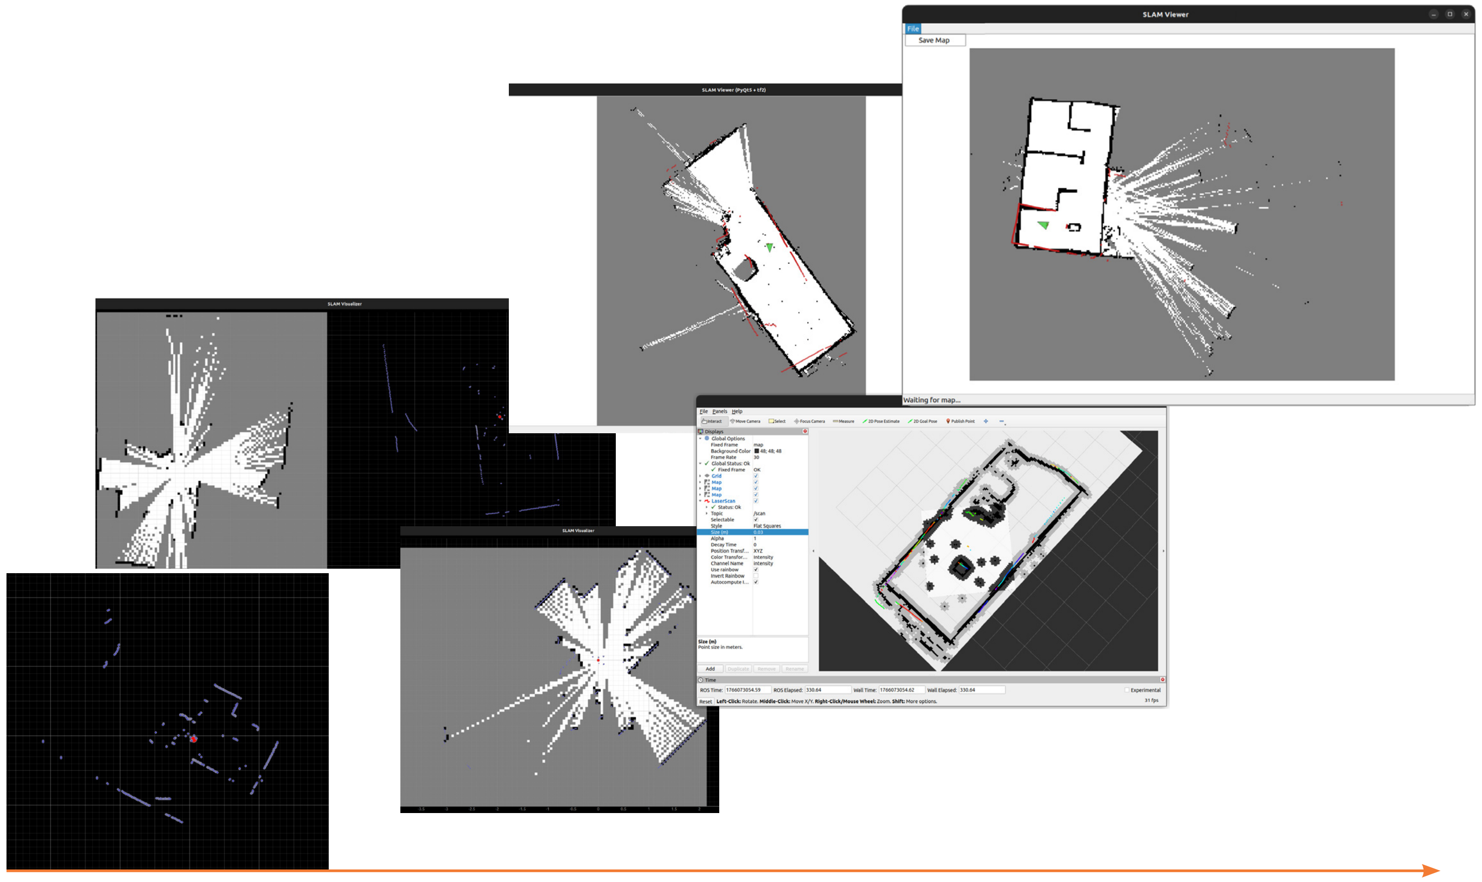

Figure 23: Evolution of the SLAM visualization exploration

3.3. INITIAL PROTOTYPE

The goal for the initial prototype was to create a setup that would be capable of the following functions:

- Connect to the robot and display the required sensor data.
- Control the robot and use the built-in ROS SLAM pipeline to generate a map of a certain space.
- Save this map in the correct format to be able to use this map later.
- Display the robot-generated map to the user via an intuitive interface.
- Let the user draw over the robot map, letting them clean and refine it based on their spatial understanding of the space.
- Save the refined map and use it for localization navigation on the robot.

To achieve all these requisites, two separate interfaces were created. The first interface is used by the researcher and handles the tasks of connecting to the robot, displaying the sensor data, running the mapping process, and ultimately exporting the generated map.

The second interface is designed for the user. It would allow the user to import the robot-generated map, draw a refined version directly on top of this map, and export this refined map. This user-created map could then be uploaded to the robot and used for localization and navigation.

Current interfaces are built for use on a laptop or computer.

3.3.1. The Researcher Interface

The following interface, seen in Figure 24, is used by the researcher to connect to the robot, visualize the ROS sensor data, and export the SLAM map to use in the following stage.

The interface displays:

- The robot's pose (green arrow)
- Real-time LiDAR data (red dots)
- The SLAM-generated map (grayscale occupancy grid).

When saving the map, two files are exported. A .png file containing the occupancy grid and a .yaml file that stores the metadata needed to correctly interpret the map, for example, the resolution and coordinate origin of the map.

Figure 24: The researcher interface

3.3.2. The User Interface

The user interface seen in Figure 25 is meant for the user to interact with and is the main interface for this project. This interface makes it possible to load the robot-generated map and let the user generate their own map on top of it, which can be saved and reused by the robot for navigation. The interface offers different modes to edit the map, which will be explained in the following sections.

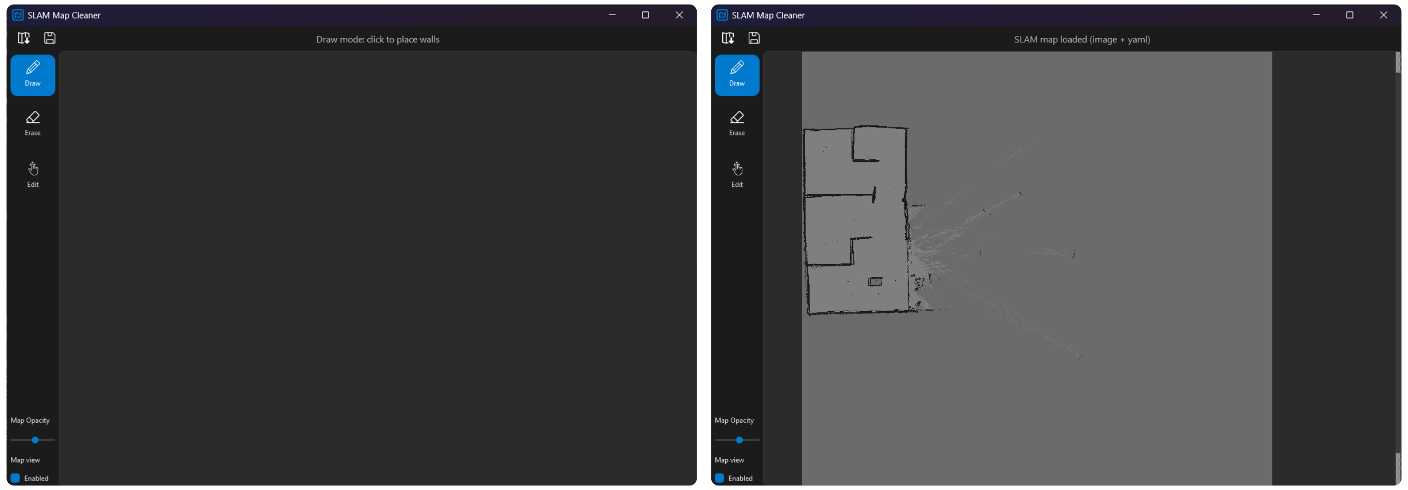

Draw mode: click to place walls

SLAM map loaded (image + yaml)

Figure 25: The user interface, with and without map

Drawing mode
In drawing mode the user can create walls by clicking once to create a start point and clicking again to define the end point. While placing a wall, the line is shown as a dotted preview line, after the second click, the line becomes solid, indicating that the wall is created, see Figure 26. Lines will snap to nearby endpoints or walls, this is indicated by a blue marker.

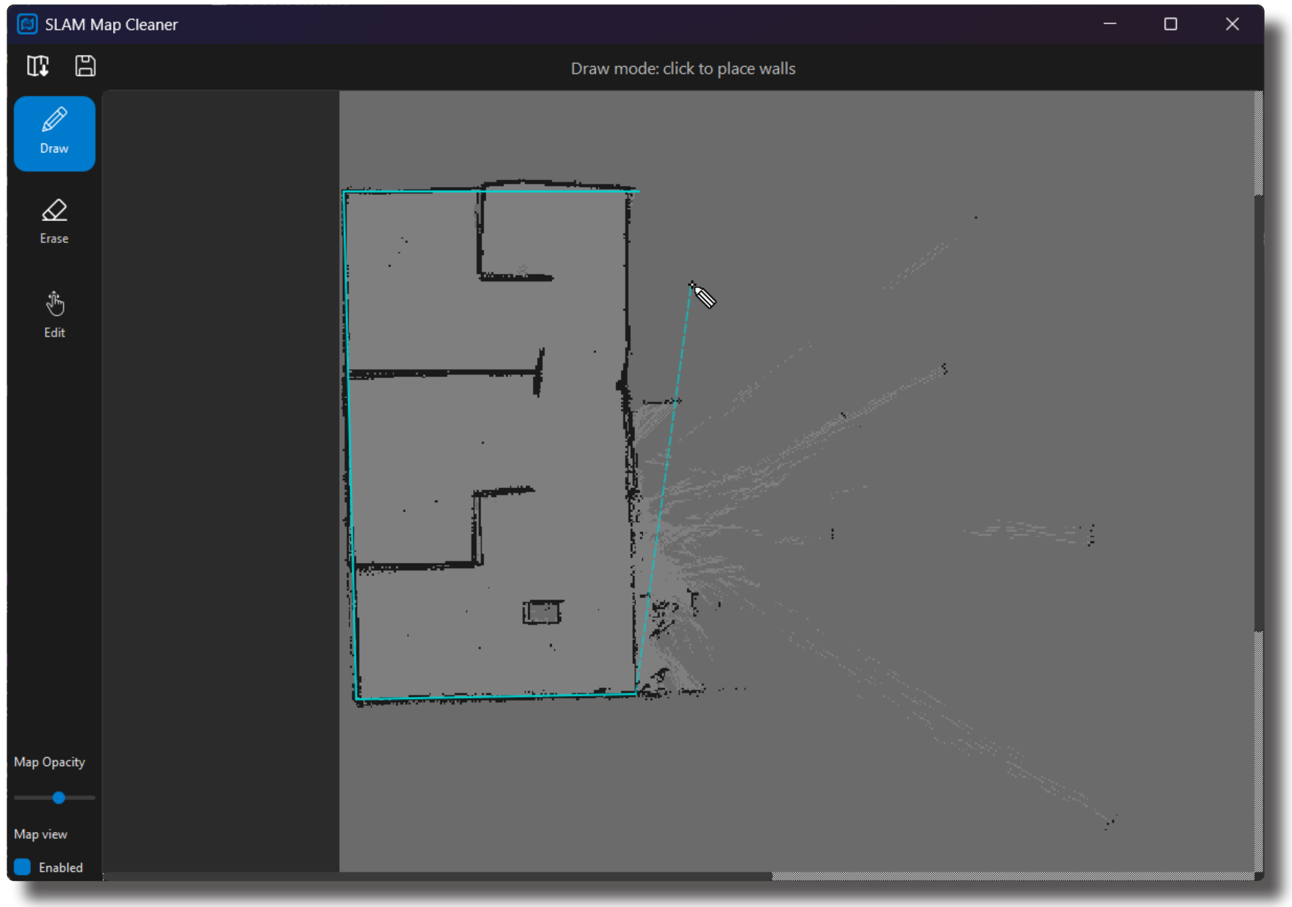

Draw mode: click to place walls

Figure 26: The user interface in drawing mode

Erase mode
Erase mode makes it possible for users to remove previously created walls. When hovering over a wall, it is highlighted in red, see Figure 27, and clicking on it will delete the wall.

Edit mode
Edit mode allows users to adjust walls they have already placed. The endpoints of lines are indicated with a blue circle. When hovering over an endpoint, it is highlighted in yellow. By clicking and dragging an endpoint, the user can alter the line, see Figure 28.

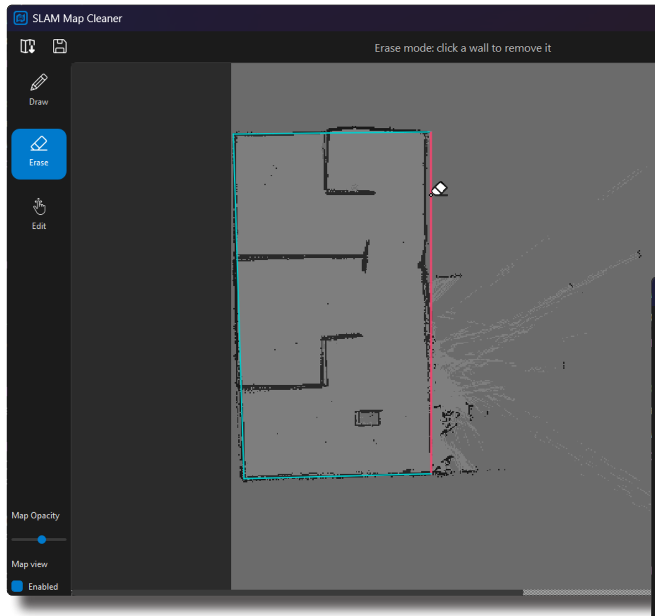

Erase mode: click a wall to remove it

Figure 27: The user interface in erase mode

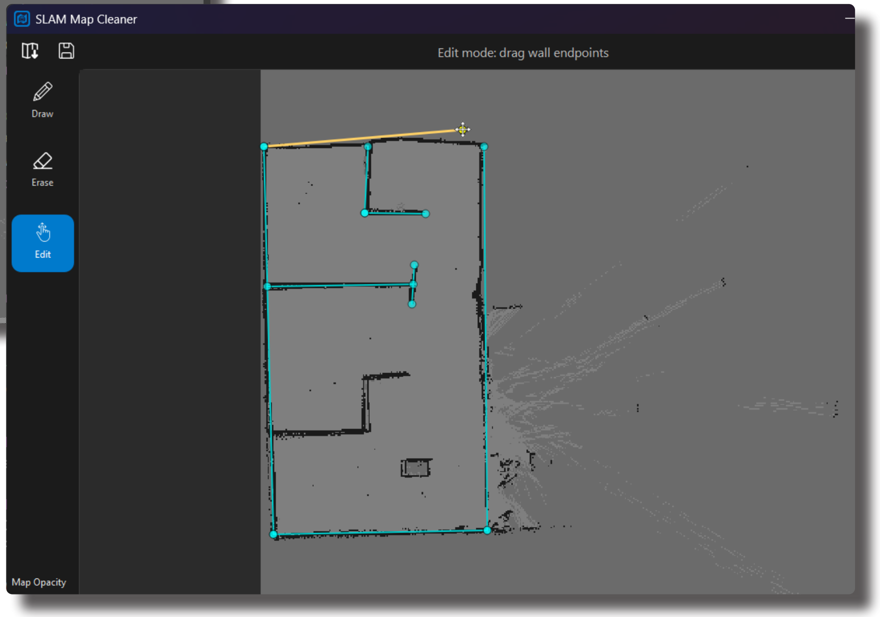

Edit mode: drag wall endpoints

Figure 28: The user interface in edit mode

3.3.3. Design Decisions

The initial prototype was not only developed as a functional tool for testing but also as a way to investigate how users can effectively contribute to refining robot-generated maps. The goal of the interface was therefore to enable clear, structured interaction while keeping the system simple enough to test the core concept.

A key decision was to constrain the interaction to the creation of straight wall segments. This simplification was made intentionally to reduce complexity and focus on evaluating the feasibility of the user-in-the-loop concept. This also aligns with the nature of most indoor environments, where walls are typically straight, allowing users to reconstruct the space in a clean and simplified manner.

To support this, an interaction was chosen where walls are drawn by selecting a start and end point. instead of allowing freehand drawing. This interaction was chosen to guide users toward making larger, more intentional decisions. The idea behind this is that it would result in straighter and more consistent walls. Compared to freehand input, this approach reduces small inaccuracies and leads to a cleaner output, which is important when the refined map is later used for navigation.

Another important design decision made, was to let users redraw the map on top of the robot-generated version, rather than directly editing the existing map. This approach was chosen to avoid carrying over the noise and inconsistencies from the original SLAM output. By reconstructing the map themselves, users can create a simplified and cleaner representation without being constrained by underlying sensor errors. This also reinforces the idea of the user acting as a corrective layer, translating noisy sensor data into a more structured and interpretable map.

Overall, the prototype was intentionally designed to balance simplicity and control, allowing users to make meaningful corrections while minimizing interaction complexity. These choices were made to test whether users can effectively interpret and improve robot-generated maps using their knowledge of the environment.

04

USER TEST

This chapter will present the conducted user test which evaluated the previously introduced interface prototype. The chapter will go into the methodology used in the test, followed by the results gathered from observations, questionnaires, and analysis of the created maps. Following that, the findings are discussed, and the interface is evaluated to determine the next steps in the design process.

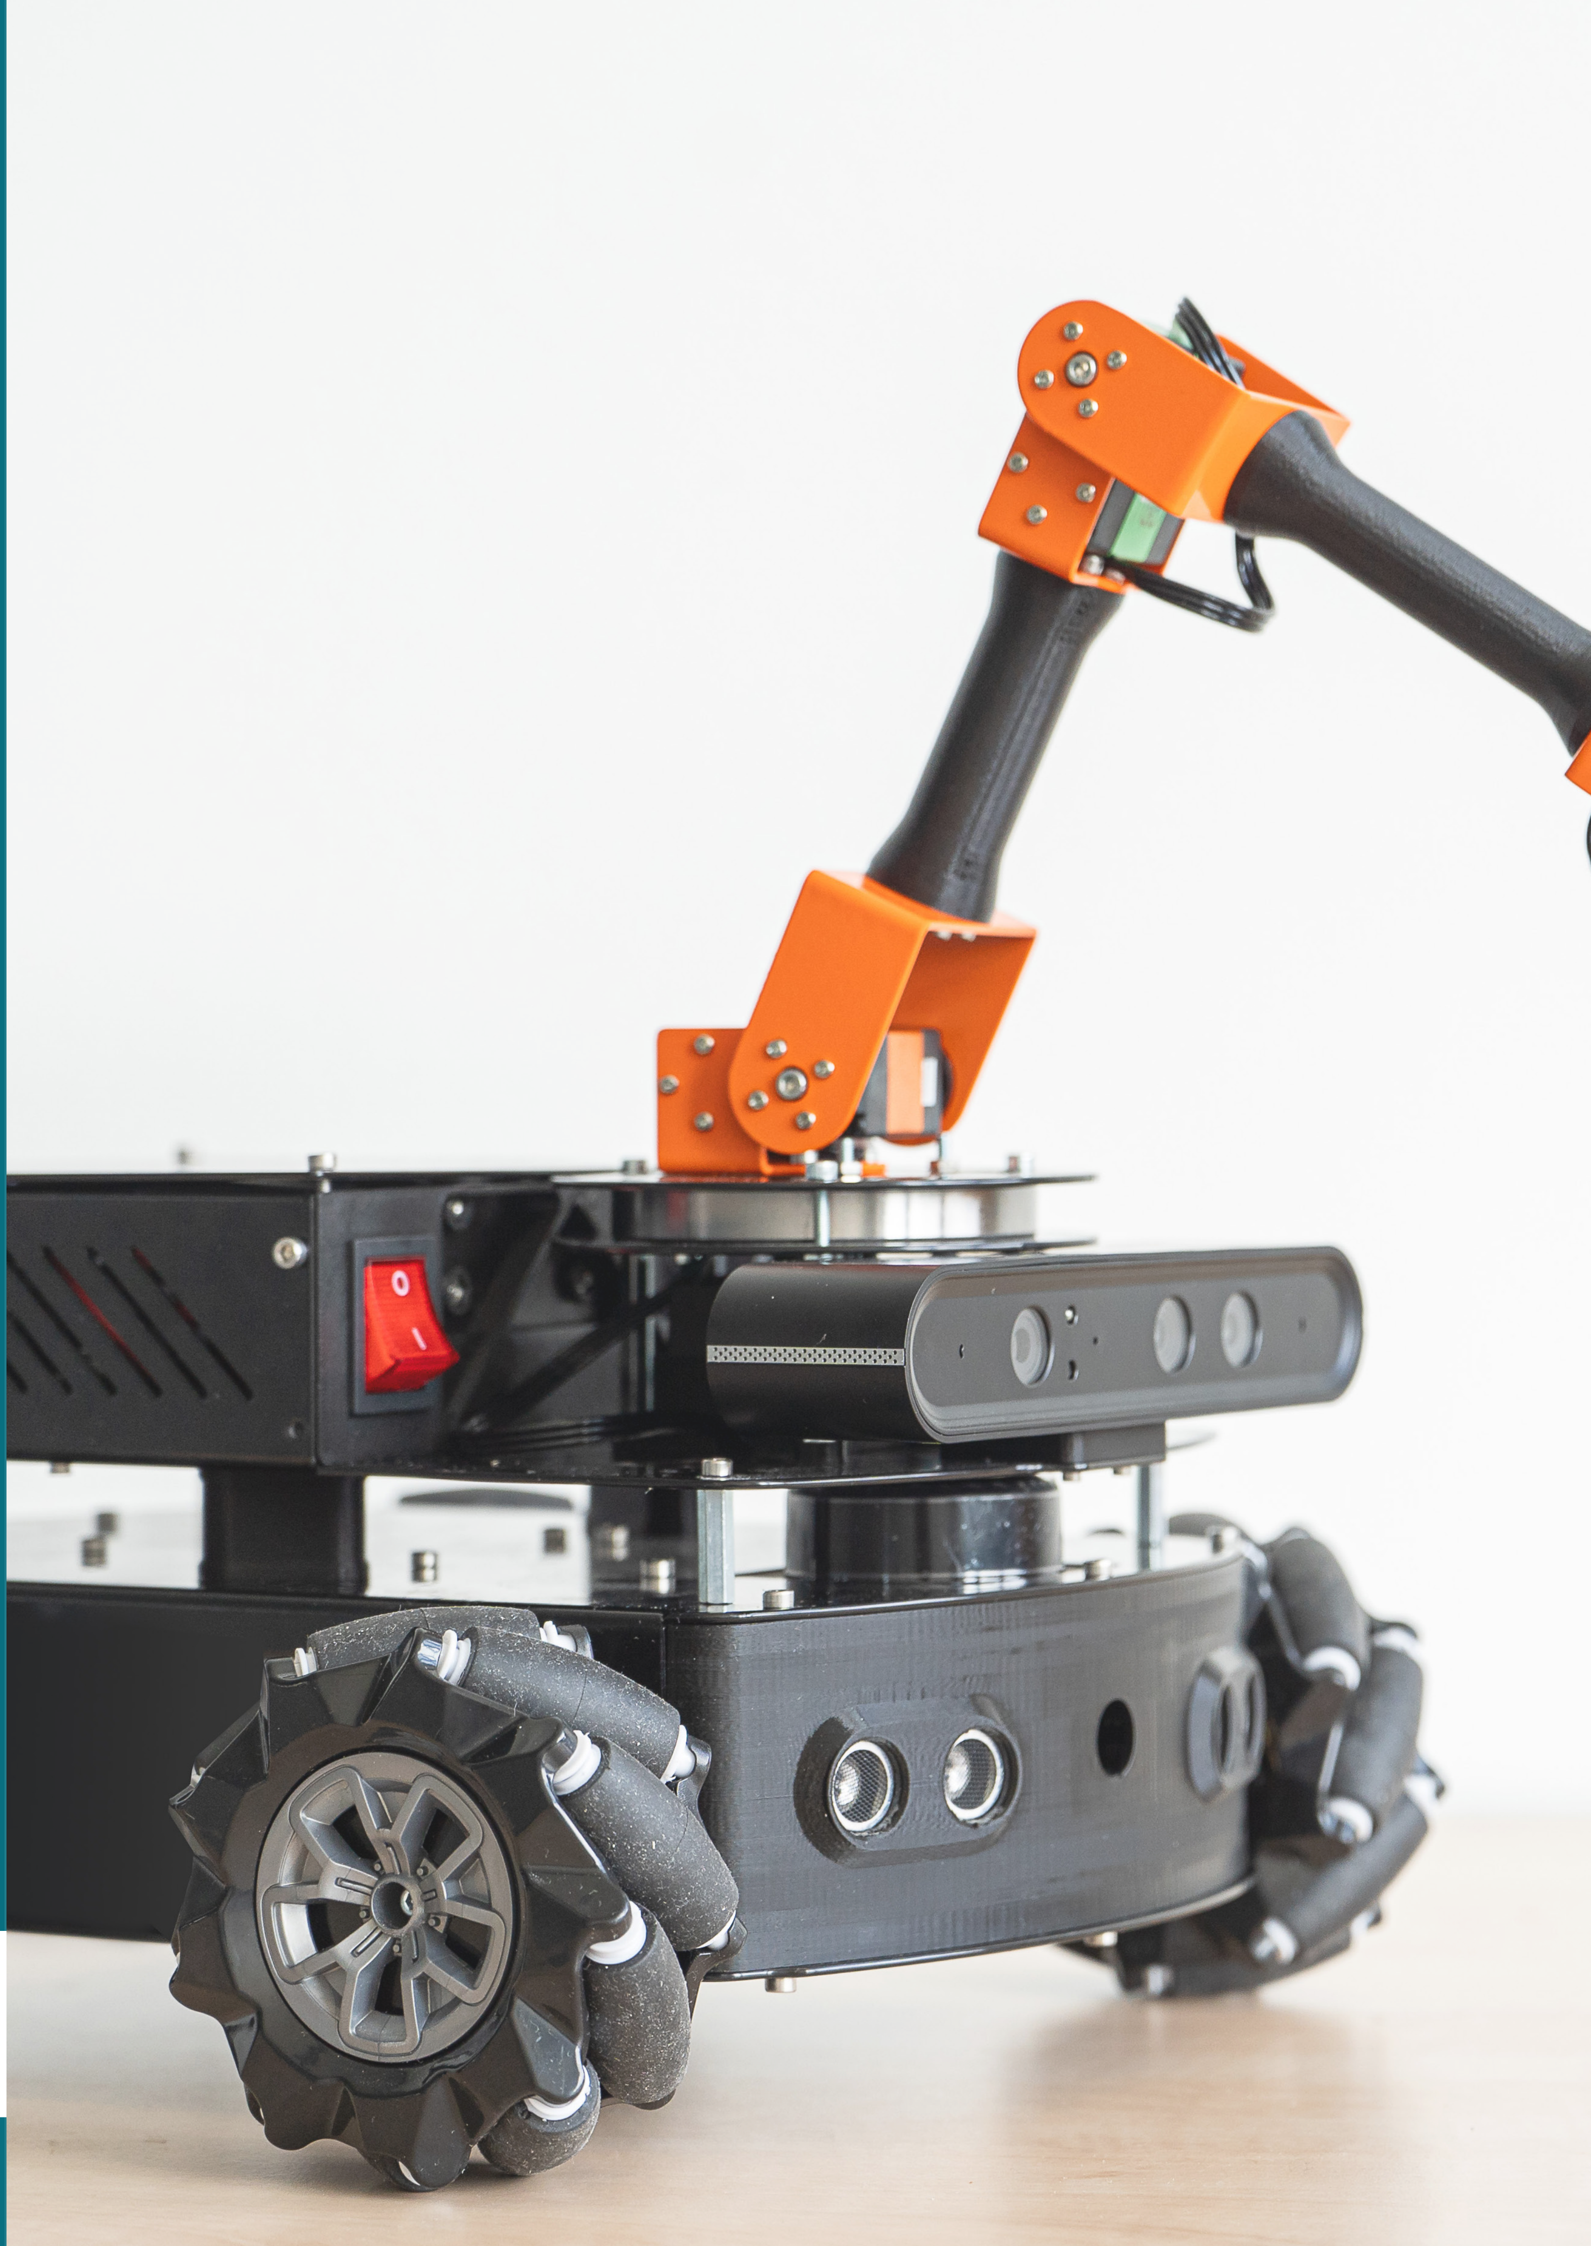

4.1. METHOD

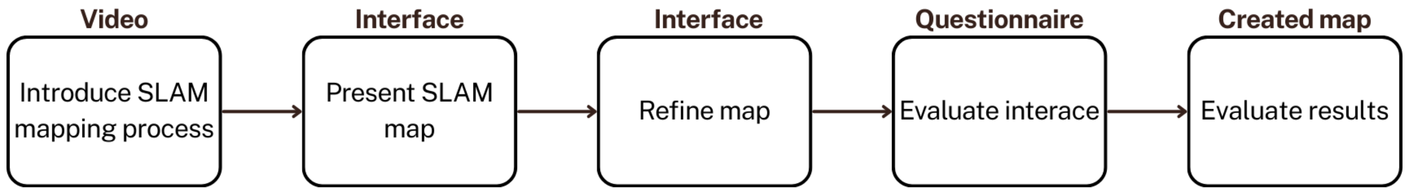

Figure 30: Overview of the user test process

4.1.1. Test Overview

The purpose of this test is to evaluate the effectiveness and usability of the proposed interface design. Furthermore, this test examines how well participants understand the mapping process, and are able to interpret and refine robot-generated maps.

During the test, participants will be introduced to the mapping process via a video demonstration of the mapping of a space by the robot. After this demonstration, the participant will be shown the mapped space in person, and the generated map will be presented within the interface prototype. Participants will be asked to refine this map manually. This process will provide some information on how intuitive participants find it to interpret the SLAM map when having a mental model of the mapped space.

Finally, the participants are asked to provide feedback on the overall interaction experience with the robot and the interface. This feedback will be used to evaluate the interface and identify further places for improvement.

4.1.2. Participants

The user test is conducted with 10 participants recruited using convenience sampling. Participants had no prior experience with the tested interface and were not required to have prior robotics or mapping expertise. Basic demographic information (e.g., age, background, and experience with robotics or mapping tools) is collected for contextual interpretation of results.

Although the sample size is limited, the participant profile aligns relatively well with the demographics reported in the literature. Research by Sung et al. (2008) on early iRobot Roomba users stated that domestic robot users are often quite young and well-educated. While the current user group aligns quite well with the demographics found by Sung et al. (2008), future testing should include a broader range of users to assess the interface across more diverse user groups.

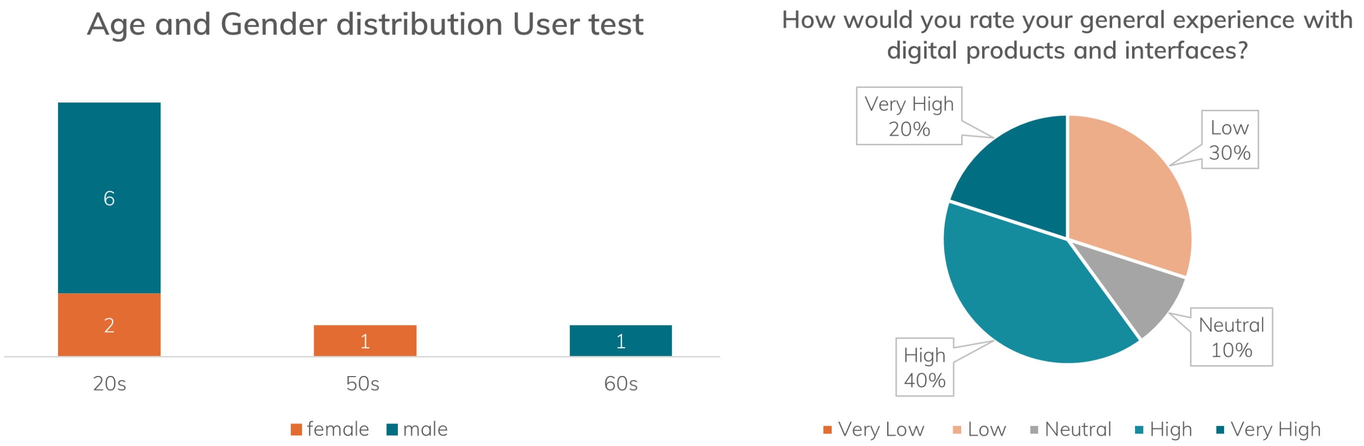

Figure 31: Overview of the demographics of the user test participants

4.1.3. Materials

To conduct the user test, a certain amount of preparation was needed, including creating the correct ethical documentation, including an informed consent form, as well as the setup of the robot, the interface, and the mapping environment. These are described in this section.

4.1.3.1. The Robot

The robot used for this test will be the MIRTE Master robot, an educational mobile robot platform capable of autonomous navigation and simultaneous localization and mapping (SLAM). Although the robot will be shown in video footage to the participants to give them a feel of the mapping behavior, all uses for the robot, like the mapping of the space, will be done beforehand to save time during the testing.

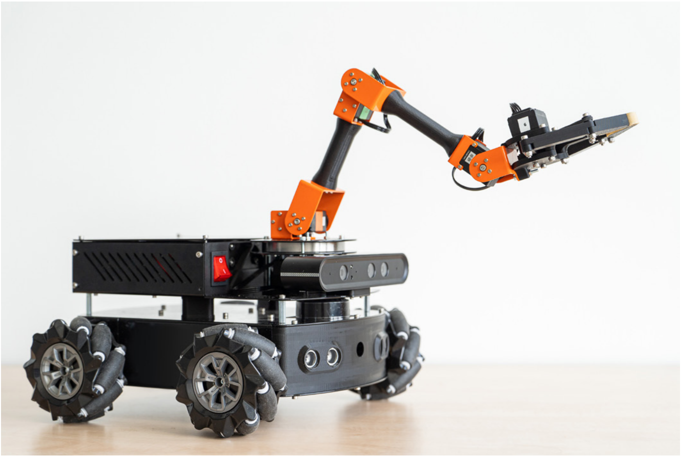

Figure 32: The MIRTE Master robot (MIRTE team, 2025)

4.1.3.2. The User Interface

Participants interact with a laptop-based prototype GUI (see figure 33) that allows them to:

- View the final SLAM-generated map
- Manually refine the map by drawing on top of the SLAM-generated map using drawing and editing tools
- Save the newly refined map in a format the robot can understand for navigation.

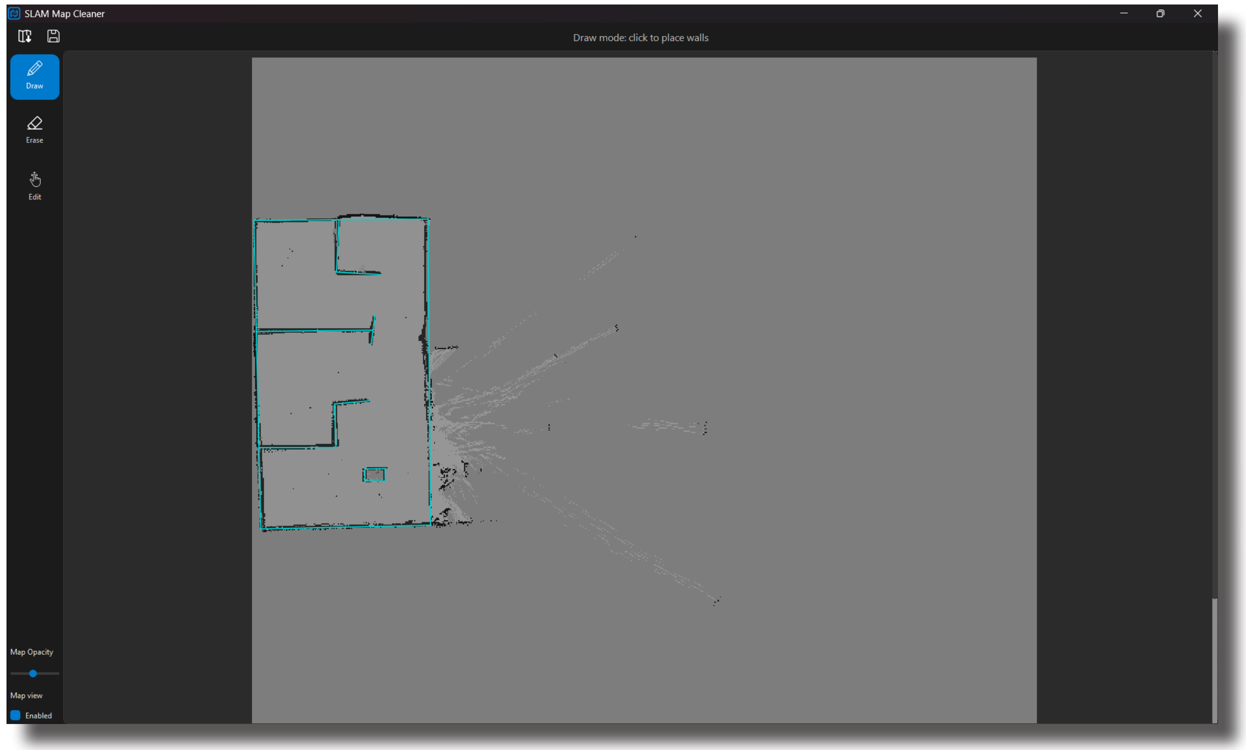

Draw mode: click to place walls

Figure 33: The User Interface with a refined map (in blue) drawn over the SLAM map

4.1.3.3. The Test Environment

The test takes place in a controlled indoor environment in which a simplified floor plan is constructed (see Figure 34). The layout is meant to represent a small-scale indoor space that the robot can navigate and map. The test is done near this test space to allow participants to understand what the map is supposed to look like.

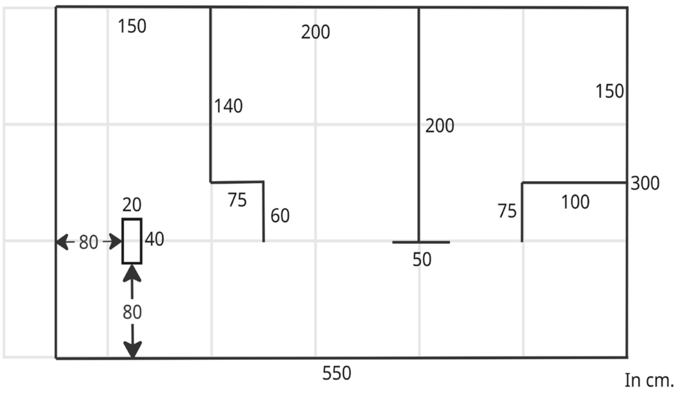

Figure 34: Measurements of the test environment

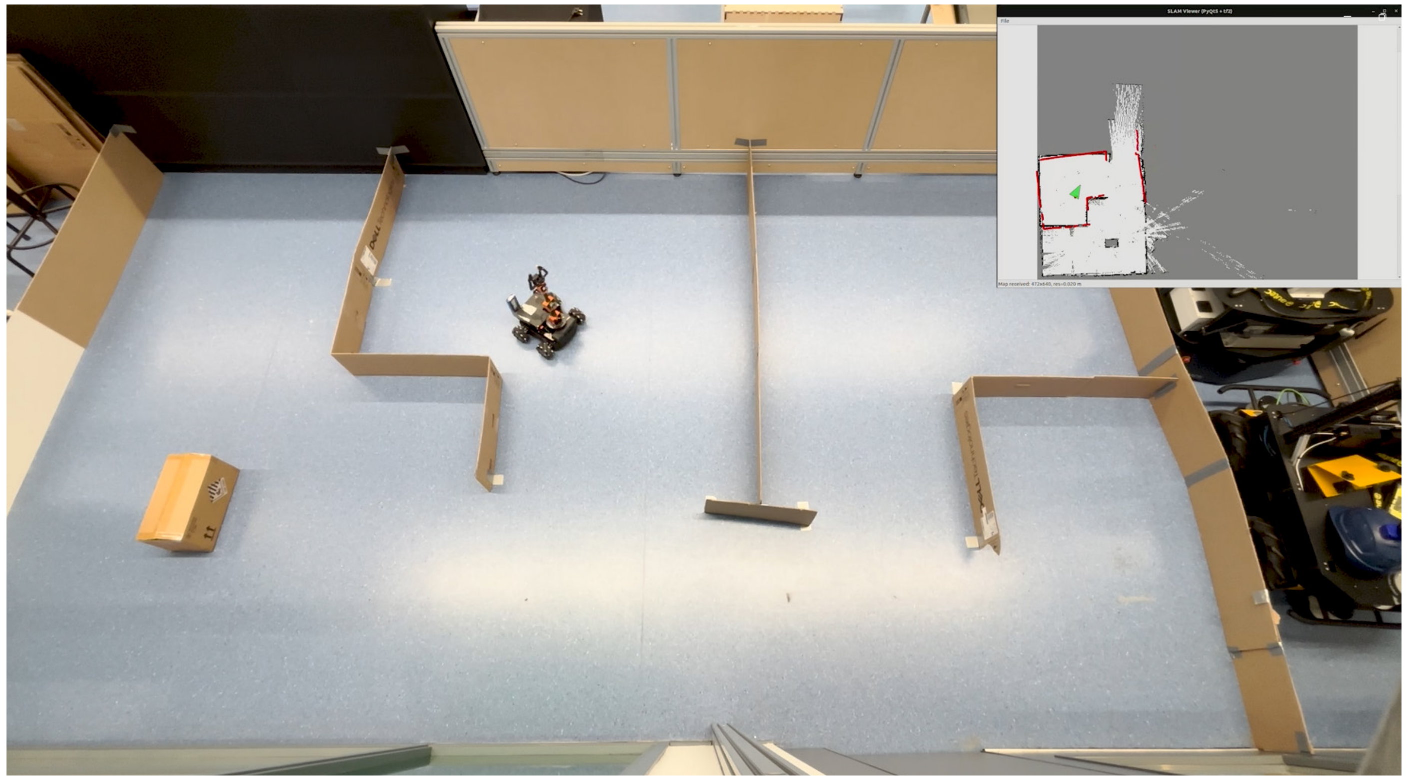

Figure 35: Screenshot of the video shown to the participants during the user test

4.1.4. Procedure

The user test consists of four main phases: introduction, mapping process, map refinement, and post-test evaluation.

1) Introduction

Participants are welcomed and informed about the purpose of the study. The test is quickly introduced, some baseline questions are asked, and informed consent is obtained before the test begins.

2) Mapping process

Participants are shown a video of the robot mapping the space in front of them. The video shows the robot driving through the space and a screen recording of the GUI mapping the space in real-time. This way, the participant can get a feel of the mapping process of the robot. Some questions are asked of the participant about the mapping process

3) Map refinement

After the mapping process is completed, the SLAM-generated map is shown to the user via the interface prototype. Participants are asked to manually clean and refine the map using the interface. The resulting map that is created is saved for further evaluation.

4) Post-test interview

After the participant has finished creating their map, some final questions will be asked focusing on interface clarity, usability, perceived difficulties, and overall experience.

The entire session lasts approximately 15 minutes per participant.

4.1.5. Data Collection

Data is collected through direct observation during the user test, participant responses from the different questions asked throughout the test, screen recordings of the map refinement phase, and the maps that are created during the map refinement process.

4.1.5.1. The Questionnaire

At different points in the test question were asked to the participants. At the beginning of the test to assess a baseline, after showing the video of the robot mapping the space and at the end of the test when the participant has used the interface the refine the SLAM map. The questionnaire used can be seen in Figure 36.

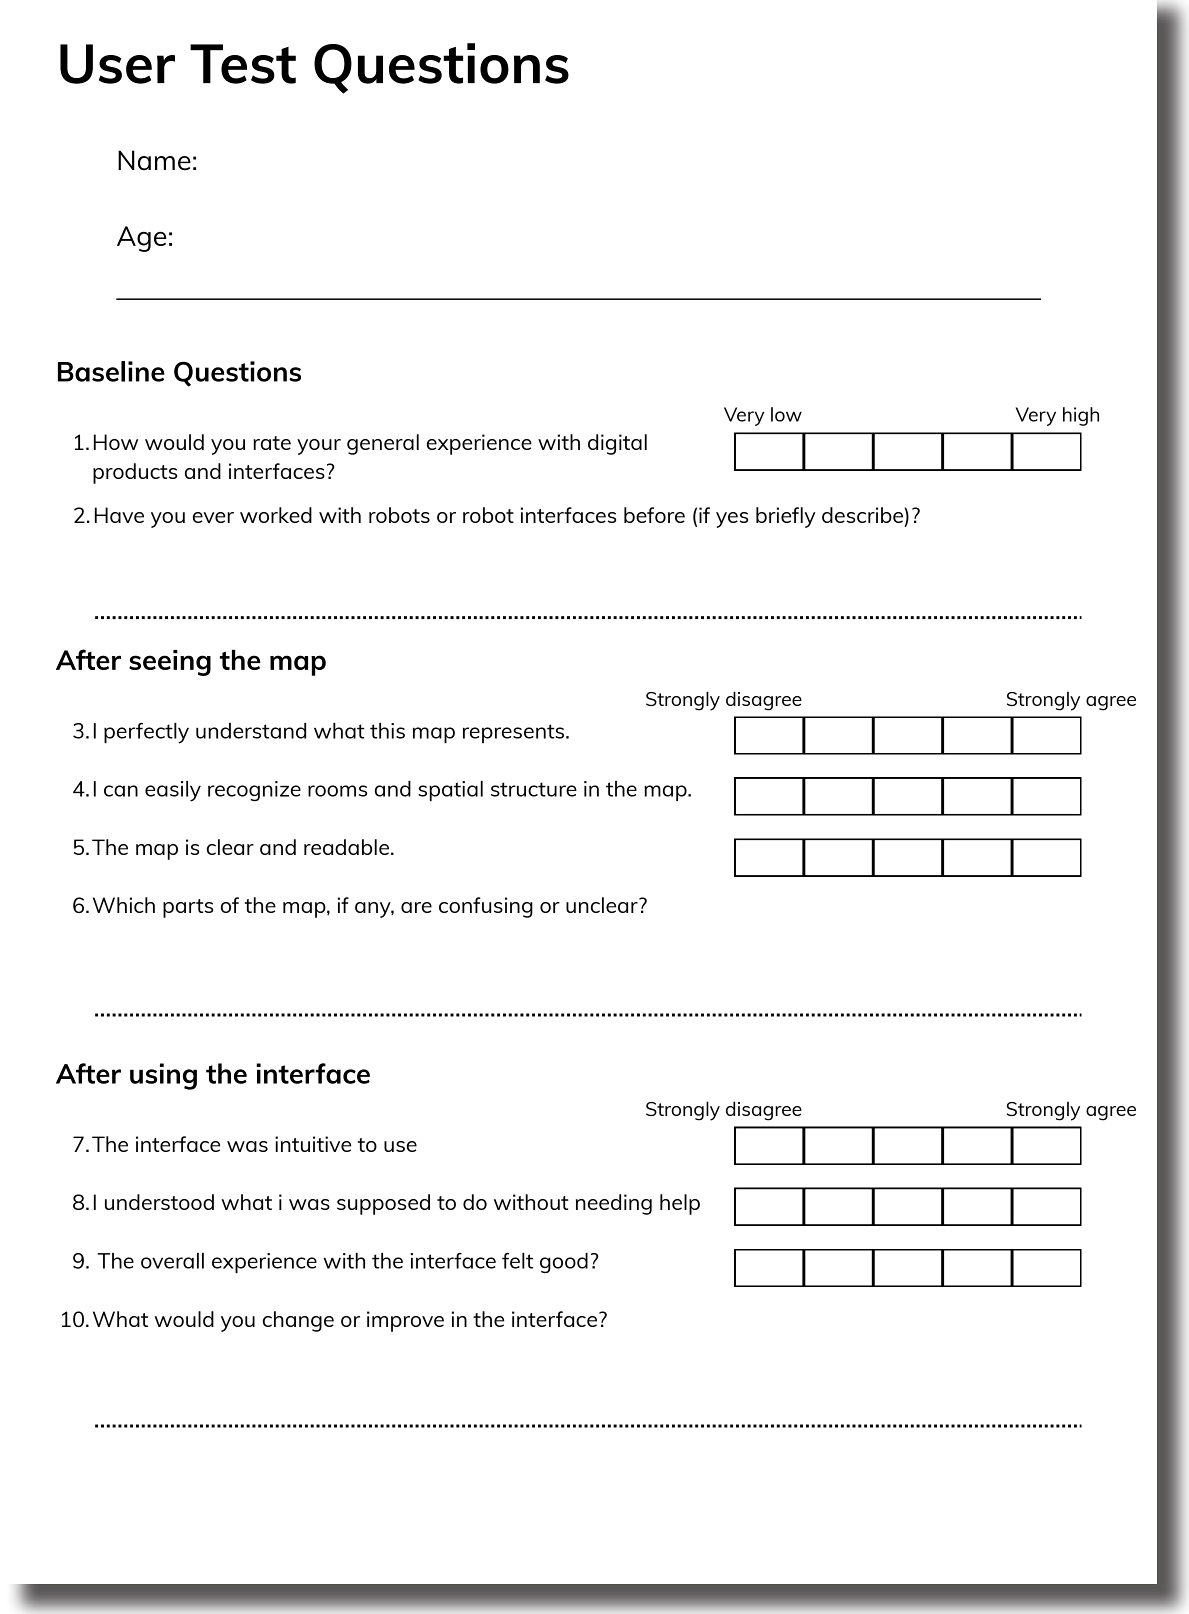

User Test Questions

Name:

Age:

Baseline Questions

Very low ☐ ☐ ☐ ☐ ☐ Very high

1. How would you rate your general experience with digital products and interfaces? ☐ ☐ ☐ ☐ ☐
2. Have you ever worked with robots or robot interfaces before (if yes briefly describe)?

After seeing the map

Strongly disagree ☐ ☐ ☐ ☐ ☐ Strongly agree

3. I perfectly understand what this map represents. ☐ ☐ ☐ ☐ ☐
4. I can easily recognize rooms and spatial structure in the map. ☐ ☐ ☐ ☐ ☐
5. The map is clear and readable. ☐ ☐ ☐ ☐ ☐
6. Which parts of the map, if any, are confusing or unclear?

After using the interface

Strongly disagree ☐ ☐ ☐ ☐ ☐ Strongly agree

7. The interface was intuitive to use ☐ ☐ ☐ ☐ ☐
8. I understood what i was supposed to do without needing help ☐ ☐ ☐ ☐ ☐
9. The overall experience with the interface felt good? ☐ ☐ ☐ ☐ ☐
10. What would you change or improve in the interface?

Figure 36: User test questionnaire

4.2. RESULTS

This section presents the results of the user test evaluating the prototype graphical user interface for robot mapping and map cleaning. The findings are based on questionnaire responses, observations during the test, and analysis of the created maps. The results are presented in three parts: The perceived understanding of the SLAM map, the usability of the interface, and the analysis of the maps.

4.2.1. Perceived Understanding of the Map

After the participants were shown the video of the robot mapping the space and they observed the physical environment, the participants were asked to rate their understanding of the SLAM-generated map, see Figure 37 for the results.

Overall, the participants reported a relatively strong understanding of the robot map. Most participants clearly stated they could see what the map represented and were able to identify distinct features of the map and compare them to the 3D environment that they could see before them. Most participants also found the map quite readable, although some participants did point out that the map was messy, especially near the window in the bottom right corner of the map.

4.2.2. Interface Usability

After the participants used the interface to refine the SLAM map, some questions were asked to evaluate the usability of the interface, see figure 37. The responses indicate that in general most participants found the interface intuitive to use and understandable. Most participants agreed that they understood what they needed to do without any assistance. The overall interaction was generally rated positively.

During and after the test, participants identified several aspects of the interface that could use some improvements. Commonly mentioned points include continuous drawing when creating walls, added shortcuts for drawing straight lines, making the walls and endpoints more visible, and the explanation of the different modes being unclear.

Some uncertainty also occurred for participants when they were wondering how refined the map should be for the robot. Some participants were quickly satisfied, even though they clearly had some errors left in their own map, while others took a long time creating the best walls they could.

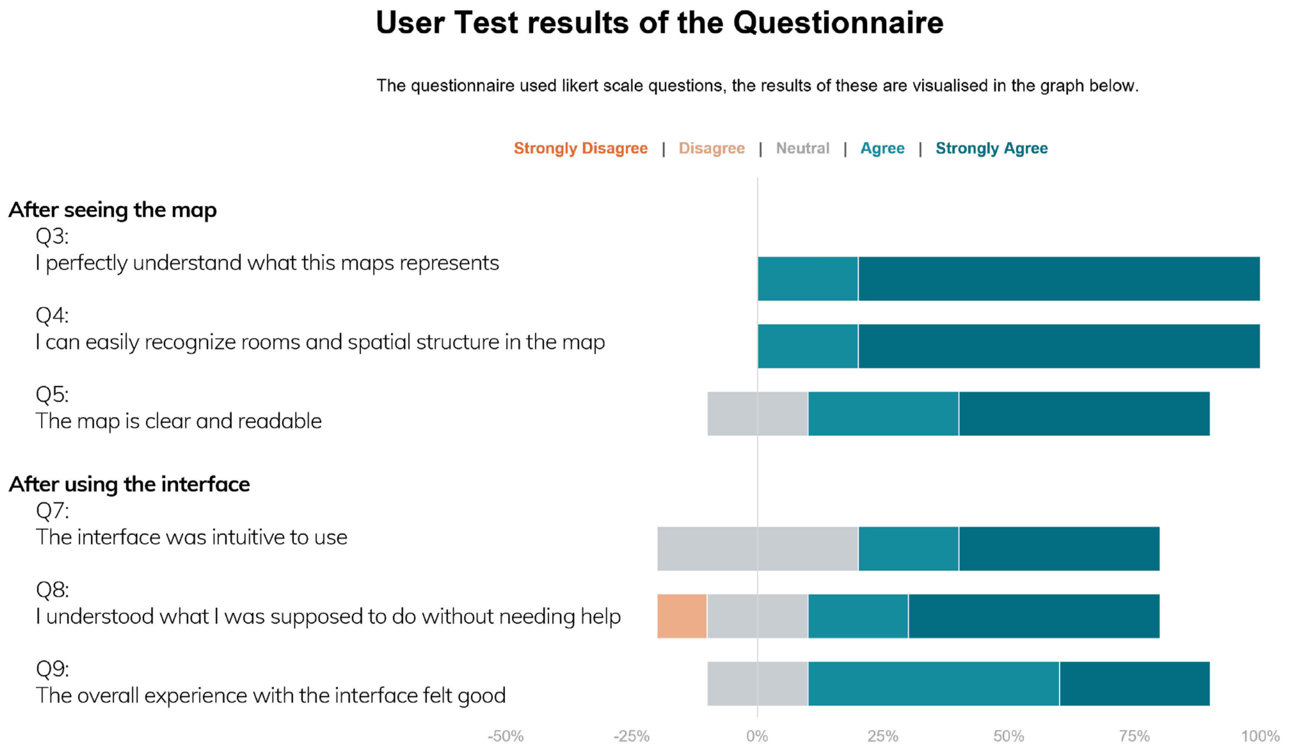

Figure 37: User test results of the questionnaire

4.2.3. Observations during Interaction

During the refinement phase, all participants were able to complete the task of refining the map without direct intervention. Although most participants needed some time to get used to the interface and how exactly walls were supposed to be drawn. Many participants started off by trying to click and drag, like they were actually drawing. Once the participant knew how the interaction worked, they were able to quickly create the map.

Some interesting interactions were observed. Some participants first tried to erase the “wrong” areas on the map, once they noticed they could not directly edit the SLAM map, they started creating their own map.
One participant drew their own map next to the SLAM map, resulting in a map that was proportionally inaccurate, showing that the SLAM map is very important as a reference for the scale.
Another participant also drew some of the lines outside of the map, these lines were picked up by the robot, as it is sometimes able to see through the window.
Another participant only drew the walls that were found to be incorrect and stated, “The other walls are good enough, I will just keep those.” See Figure 38 for some examples of participant-created maps, all maps created by the participants can be seen in Appendix D.

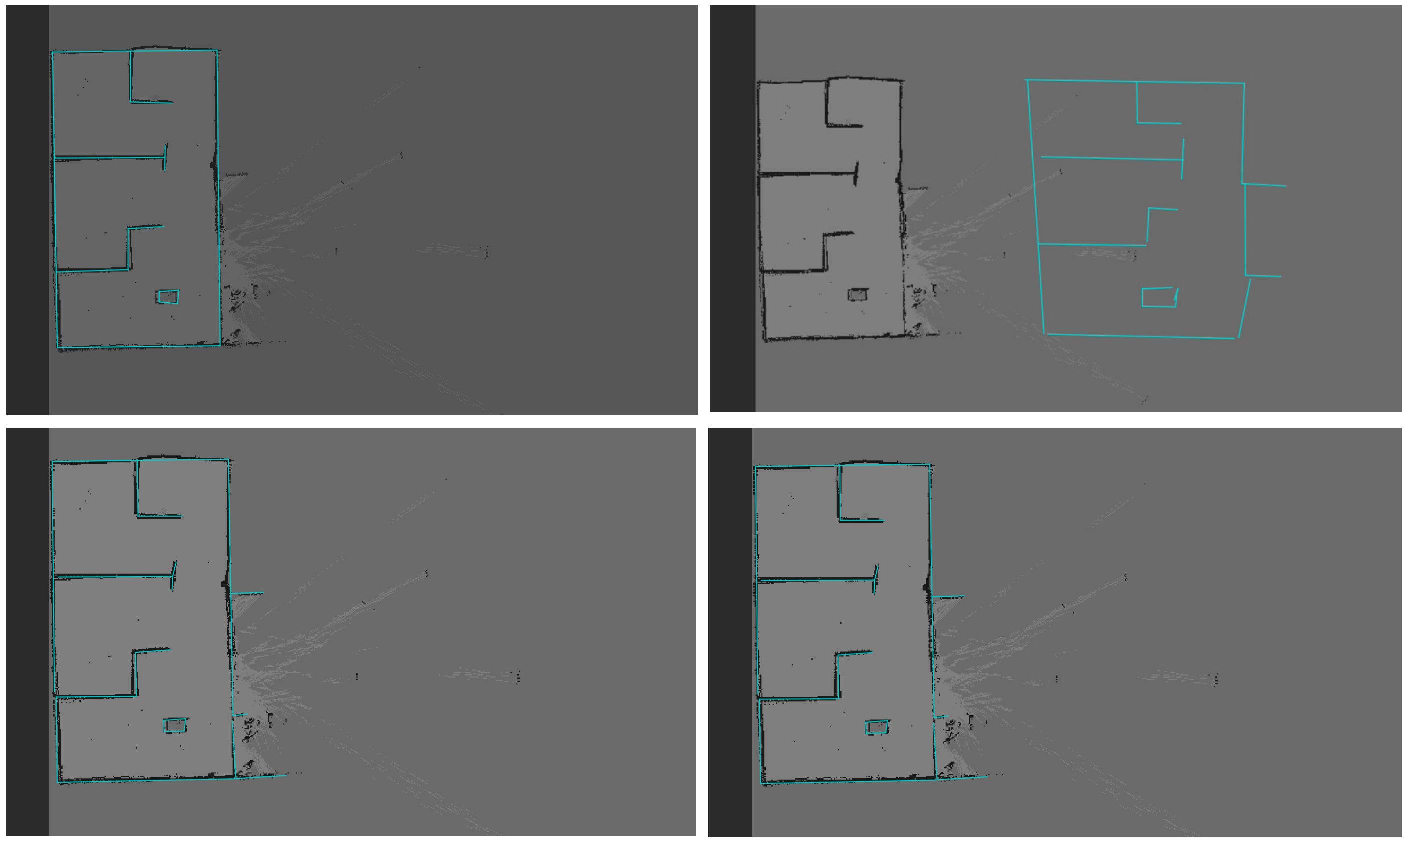

Figure 38: Different participant maps created during the user test

4.2.4. Analysis of the Generated-Maps

The refined maps created by the participants were exported and visually analyzed. The different maps created by the participants were overlaid to see how structurally similar the maps are, this is compared to overlays of three separate robot-generated maps. The results can be seen in figure 39.

The participant-created maps show a high degree of alignment between the different refined maps. When comparing this to the SLAM maps, it can be seen that these display more variation than the refined maps.

To see how these maps would actually function, the different maps were used for robot navigation. A refined map by one of the participants was used for navigation and compared to the robot-generated map that it was based on. The resulting navigation maps are shown in figure 40.

The robot-generated map contains inconsistent walls and visible noise in places where there should not be any obstacles, this harms the navigation of the robot. In comparison, the refined map shows a more simplified representation of the environment where much of the noise is taken out of the system.

Figure 39: Overlay of the robot-generated maps (left) and of the participant-created maps (right)

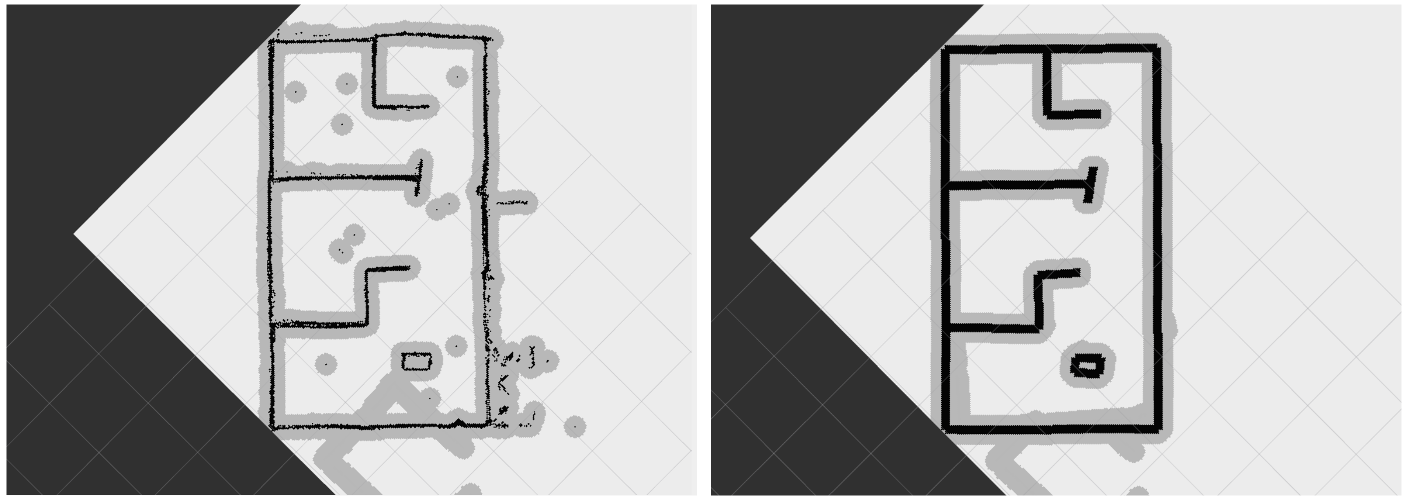

Figure 40: Navigation map of the original robot-generated map (top) and a refined map (bottom)

4.3. DISCUSSION

The results of the user test provide insights into how users perceive, interpret, and interact with robot-generated maps and the user-in-the-loop refinement interface. The different results will now be discussed and used for the next iteration of the project.

4.3.1. Map understanding

Participants showed a high level of understanding of the robot-generated map, and they were able to recognize spatial structures in the environment and relate them to the map. This indicates that participants are capable of interpreting robot-generated map data when they have sufficient contextual information, such as the ability to directly observe the mapping process or a familiarity with the environment.

This aligns with the research on situational awareness introduced in chapter 2. The users' understanding relies on their perception, comprehension, and interpretation of the environment, and if they have enough knowledge about the environment, thus a proper situational awareness, they are able to interpret the robot's representation of the environment more effectively.

However, some participants encountered some issues with interpreting very noisy, cluttered areas in the map, particularly around the window area. This suggests that if the gap between the robot-generated map and the mental model of the user becomes too large, the ability of the user to understand the map can fall short.

4.3.2. Interface understanding

The use of the interface during the user test showed that participants were generally able to navigate the interface without assistance and were able to create a map as intended. This shows that the core interaction concept of the interface was intuitive and easily learned.

However, observations and feedback from the participants highlighted some uncertainty about the level of refinement needed for a proper map. Some participants were unsure about when they were done creating the map and what level of detail the robot would need to properly use the map. While most participants recreated the entire map, some only chose to draw specific walls, because the other walls were perceived as being good enough. This shows that the system currently does not provide clear guidance on what a sufficient map would be. This requires the user to rely on their own assumptions of what a proper map would be like and results in inconsistent outcomes.

These findings suggest that while the interface is usable, it could improve in interaction clarity, feedback, and guidance.

4.3.3. Generated map quality

The comparison between the robot-generated maps and participant-refined maps shows that the maps created by the participants are more consistent, have clearly defined walls, and contain less noise. This indicates that users are capable of identifying errors, particularly those related to sensor noise and structural inconsistencies, and correcting them by creating their own refined map using their knowledge of the mapped environment.

This would suggest that the human input can be a valuable asset and used as a post-processing step in the mapping process. While SLAM algorithms produce maps based on sensor data and calculations, humans apply spatial reasoning and situational awareness to clean up errors and simplify the map.

This finding supports the hypothesis that a user-in-the-loop approach can improve the reliability of robot navigation by creating cleaner and better interpretable maps.

4.3.4. Limitations

There are several limitations that should be considered when evaluating these results.

First of all, while the navigation map and generally navigation with the robot using the refined map seemed to have improved, the way of analyzing this remained qualitative. To actually see the impact of manually refining the map, some quantitative metrics, such as localization accuracy, navigation reliability, and navigation speed, should be linked to the analysis next time.

Another point of discussion is the comparison of the participant-refined maps to the robot-generated maps. While the refined maps seem much more aligned, this could also be because all the maps were based on the same SLAM map, creating an unfair comparison.

Furthermore, the test participant group was relatively young, and most participants reported they had quite a high affinity with technology. Future testing should include some more older participants with lower technological familiarity.

All together this user test created valuable information for the next stage of the project and the creation of the final design concept.

05

FINAL CONCEPT

This chapter presents the final design concept and its evaluation. First, the developed interface design is described in detail, outlining its key features and functionality. This is followed by the evaluation of the design through a user test, including the methodology used. Subsequently, the results of this test are presented and analyzed. Finally, these findings are discussed to interpret the outcomes and assess the performance of the proposed design.

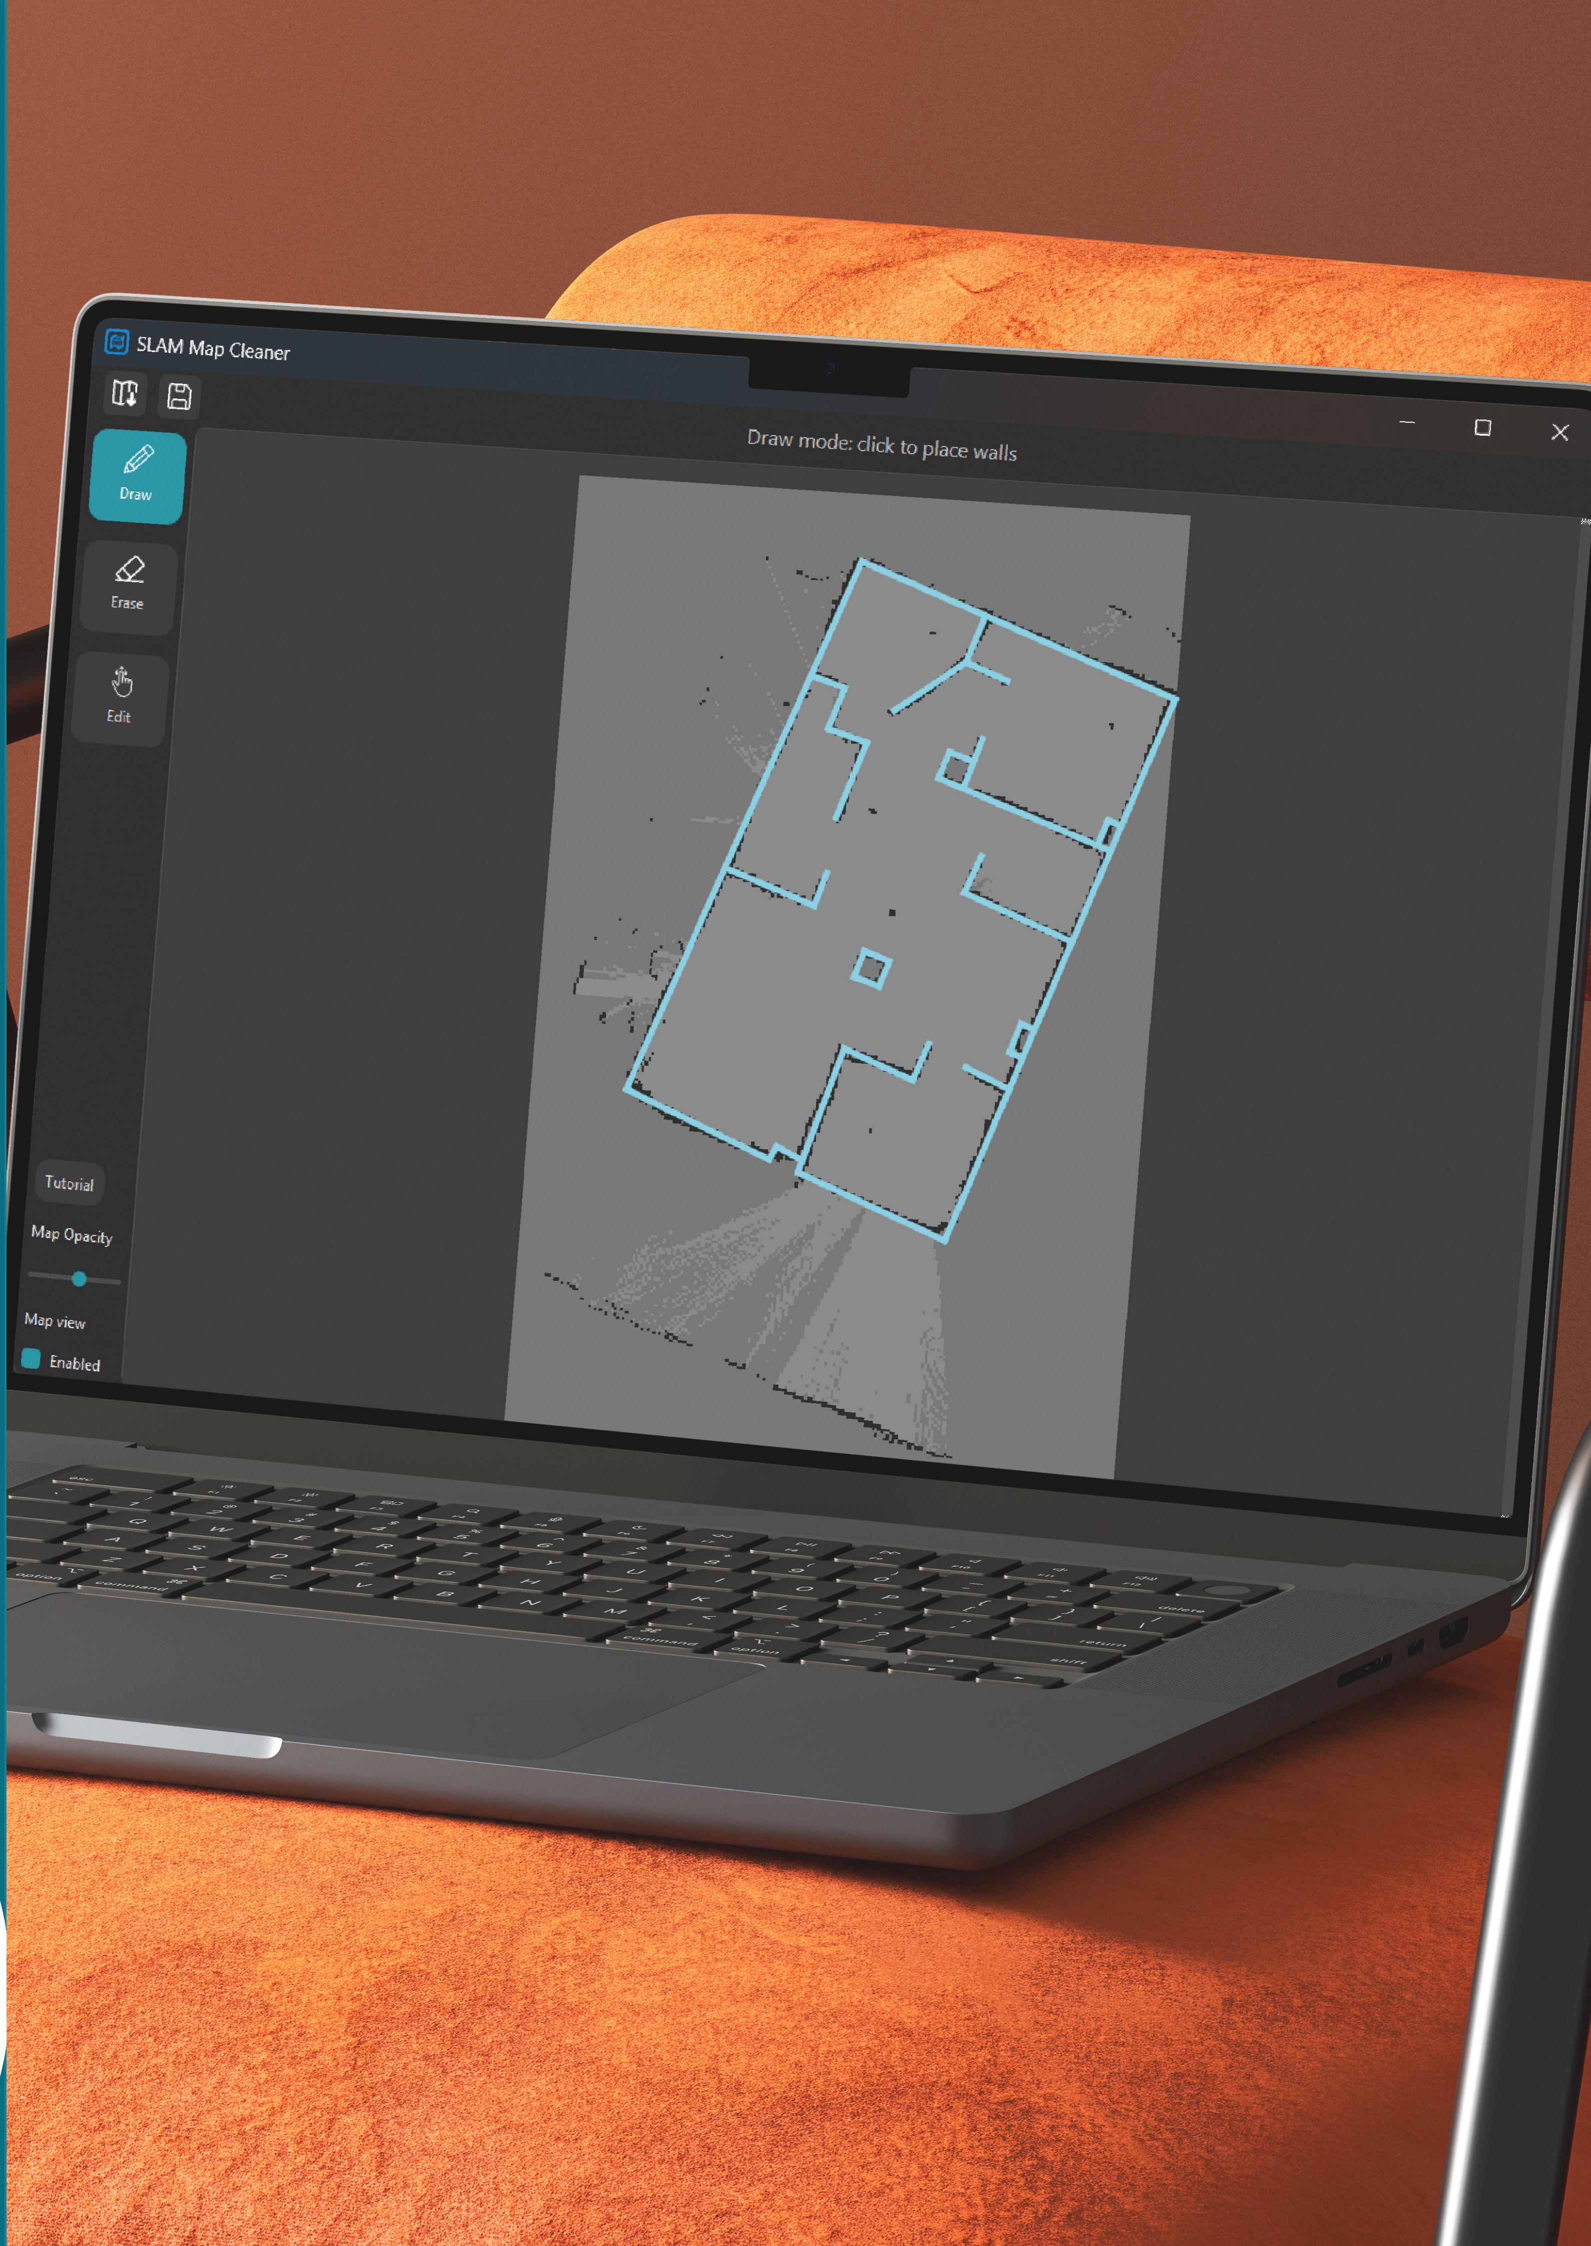

5.1. DESIGN DETAILS

This section presents the final design concept of this project. Building on the insights gathered throughout the research and iterative prototyping phases. The developed interface and the concept behind it are described, along with the main improvements made compared to the initial prototype.

The final concept is a user-in-the-loop mapping approach, in which human input is combined with a robot-generated map to create a more reliable map of the domestic environment. The system aims to combine the strengths of both humans and robots. This is done by utilizing human spatial reasoning and interpretation skills to correct and simplify robot-generated maps while maintaining the robustness and efficiency of the robot's algorithmic SLAM-based mapping. In this way, the mapping process becomes a collaborative effort, where both user and robot contribute to the creation of a more usable and reliable map.

By actively involving the user in a previously autonomous system, the key opportunities identified during the design process, collaboration, alignment, and transparency, are addressed. The collaborative nature of the interface allows users to directly influence the map. Which also helps align the robot's internal representation of the environment with the mental model of the user and reduce mismatches between system behavior and user expectations.

Furthermore, by actively involving the user in the mapping process, the transparency of the system is also increased. The user develops a better understanding of the mapping behavior and limitations of the robot. This improved understanding can translate to increased trust and a more intuitive overall user experience.

The design builds directly on the initial prototype described in Chapter 3, and the core interaction concept remains the same. Users refine a robot-generated map by drawing a cleaner version on top of it, correcting errors and ignoring noisy sensor data by utilizing their understanding of the mapped environment. Since the initial interaction proved to work well during the user test, the final iteration focuses primarily on improving the overall look and feel of the interface. Visual clarity has been enhanced, interaction feedback has been refined, and the system has been made more consistent and intuitive to use.

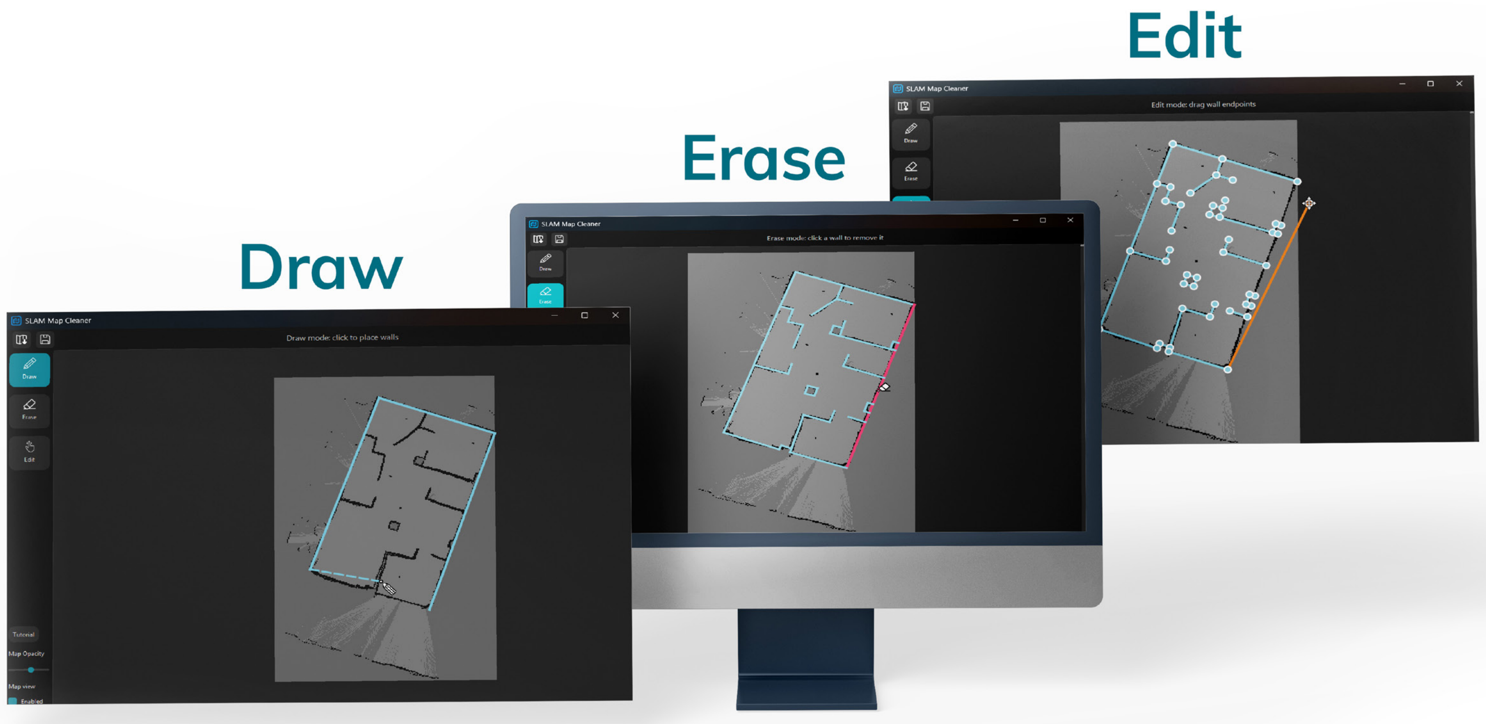

The interface allows users to interact with the robot-generated map through three main modes: drawing, erasing, and editing. In drawing mode, users create new wall segments by defining start and end points, allowing them to reconstruct walls based on their understanding of the environment. Erase mode enables users to remove walls that were incorrectly placed by them, while edit mode allows for fine adjustments of existing wall segments by modifying their endpoints.

Although these modes were already included in the initial prototype, they have been improved through more accurate snapping behavior, added shortcuts for easier use, clearer visual cues, and a more coherent interface layout, resulting in a more consistent and intuitive interface.

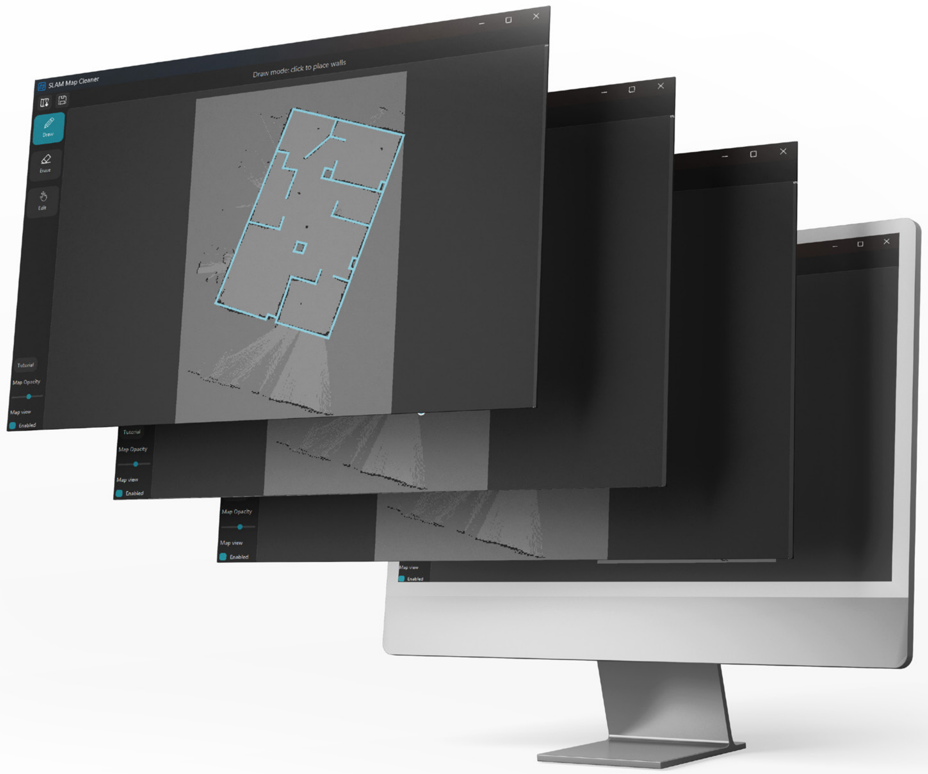

To further support usability and understanding, a simple built-in tutorial was introduced (see figure 45). This tutorial can be launched by first-time users and guides them through the basic functionality of the interface and explains how the different interaction modes can be used. By integrating this tutorial directly into the system, users are able to quickly understand how to work with the interface without requiring additional explanation.

Overall, the final concept refines the initial prototype into a more polished and user-friendly system, creating a collaborative mapping approach where users have clear and accessible control over editing the robot-generated map.

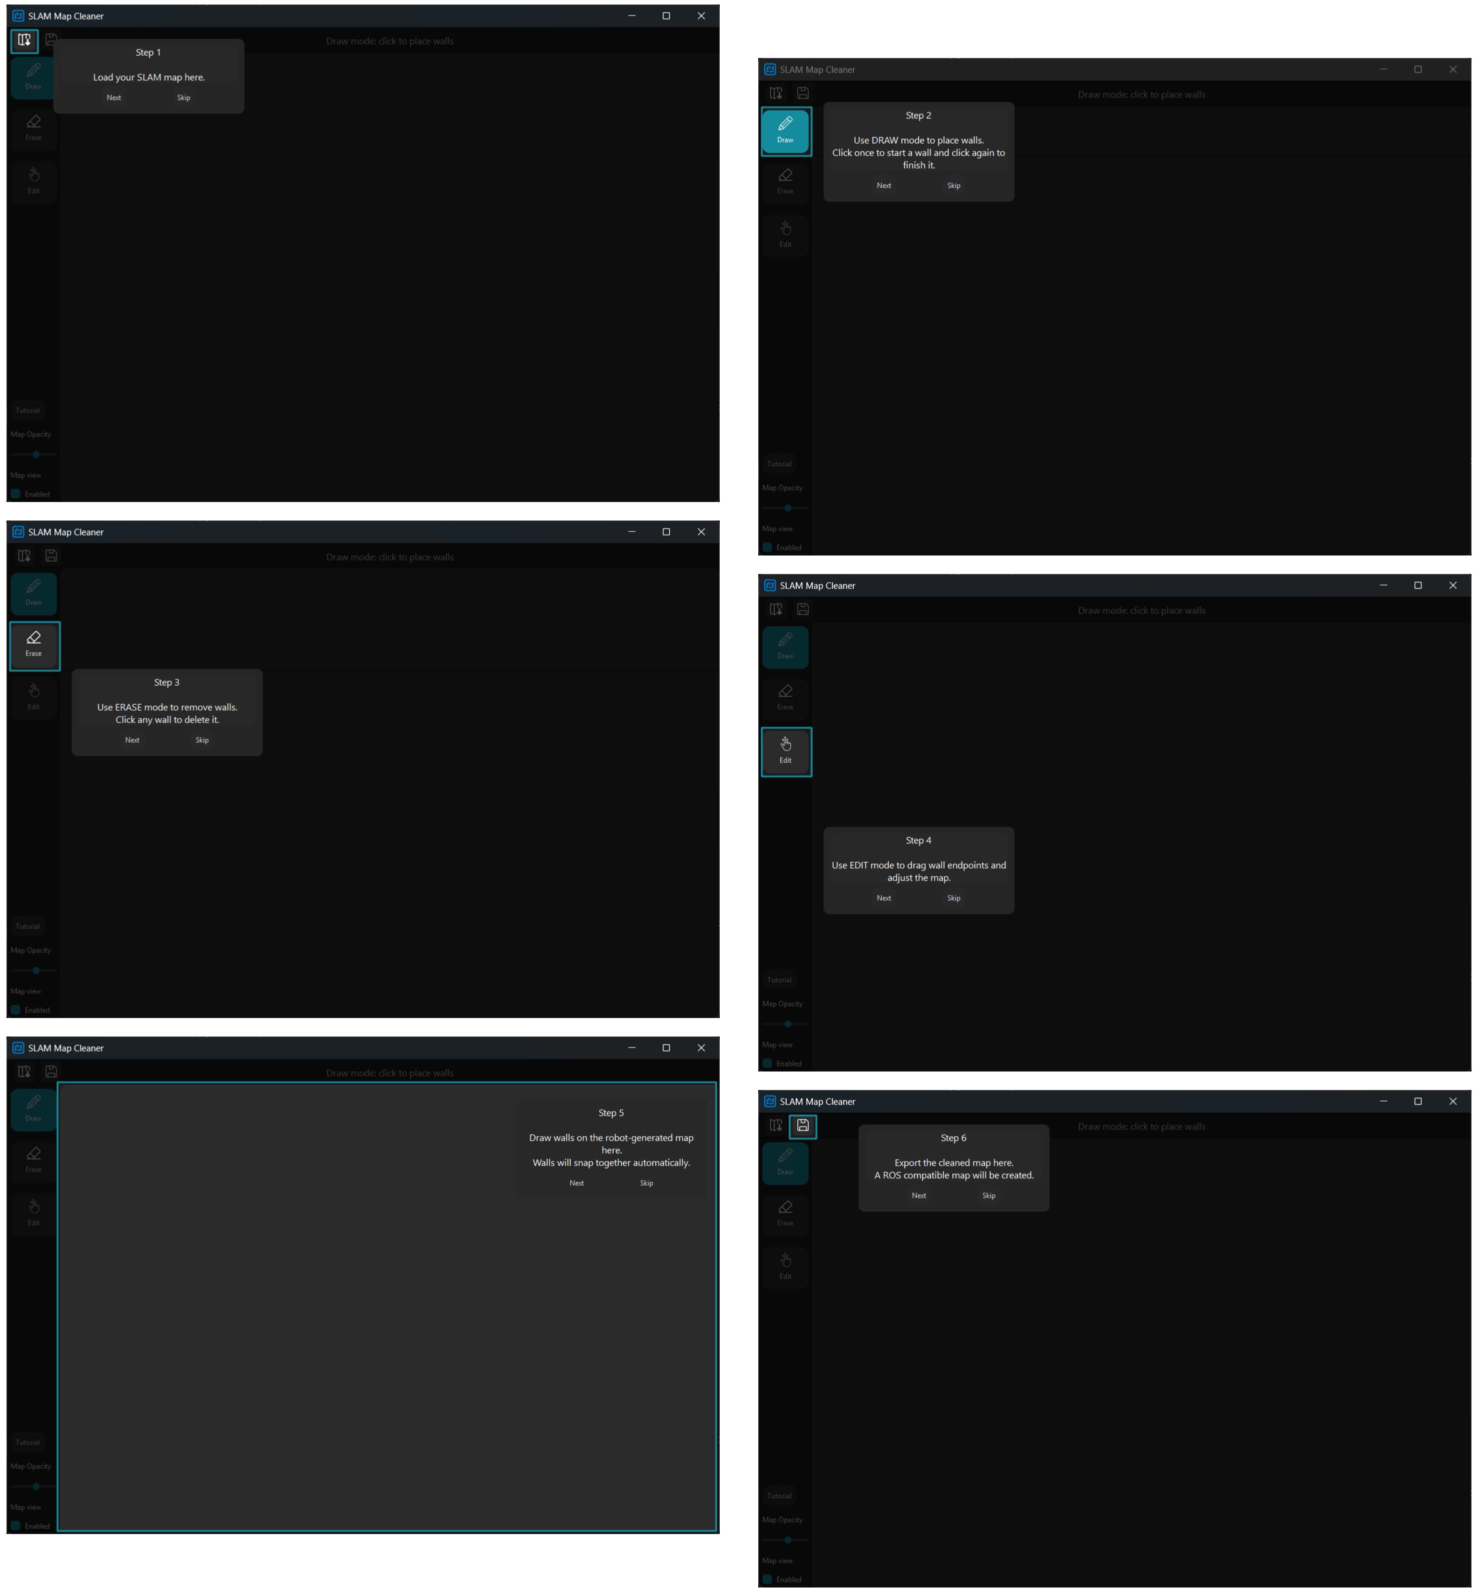

Step 1

Load your SLAM map here.

Step 2

Use DRAW mode to place walls.
Click once to start a wall and click again to finish it.

Step 3

Use ERASE mode to remove walls.
Click any wall to delete it.

Step 4

Use EDIT mode to drag wall endpoints and adjust the map.

Step 5

Draw walls on the robot-generated map here.
Walls will snap together automatically.

Step 6

Export the cleaned map here.
A ROS compatible map will be created.

Figure 45: The tutorial present inside the interface

5.2. DESIGN EVALUATION

For a final design evaluation, a second user test was conducted using the final interface concept. The goal of this test was to quantify the impact of refining the robot-generated maps on the navigation of the robot. While the first test focused primarily on usability and qualitative insights, this test aimed to provide some measurable evidence of the effectiveness of the proposed user-in-the-loop mapping approach.

5.2.1. Method

5.2.1.1. Materials & Participants

The test was conducted using a lot of the same materials as the previous test, mainly the robot, ethical documentation, and ROS (see chapter 4). The things that did change were the interface now being the final concept, a stopwatch for keeping track of time, and the test environment. The user test was conducted with 10 participants recruited using convenience sampling. Participants had no prior experience with the tested interface and were not required to have prior robotics or mapping expertise.

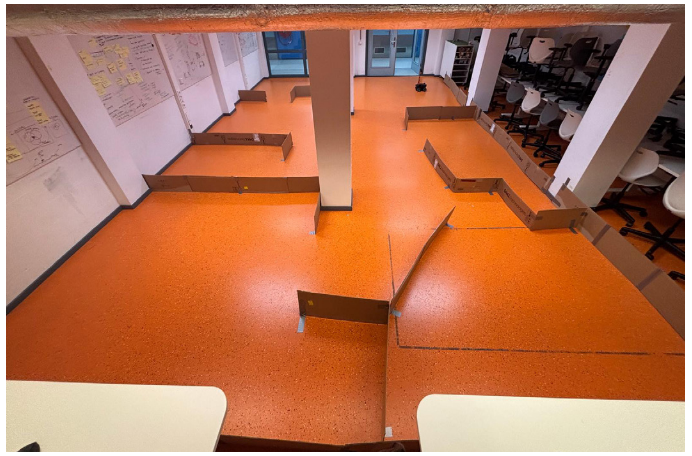

Figure 46: A picture of the test environment

5.2.1.2. Test Environment

The test takes place in a controlled indoor environment in which a simplified floor plan is constructed (see figure 46 & 47). The layout is meant to represent a small-scale indoor space that the robot can navigate and map. The test takes place inside this test space to allow participants to understand what the map is supposed to look like.

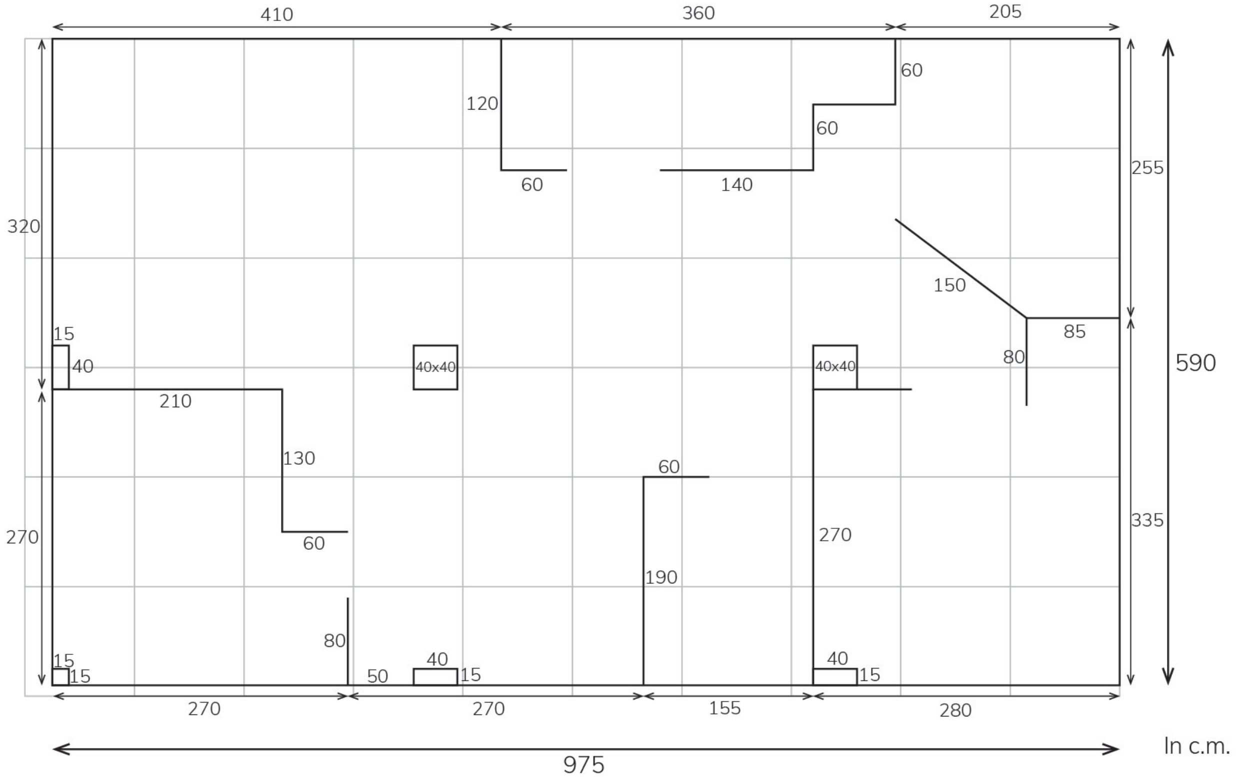

Figure 47: The measurements of the test environment

5.2.1.3. Procedure

Before the test could be conducted, the test environment was mapped a total of five times using the robot. Each mapping run resulted in a different SLAM map of the same space, these multiple runs were used to account for the variation in sensor data.

For the test, participants were asked to refine the robot-generated maps using the interface. Before mapping, they were asked to run through the tutorial present inside the interface to make sure they knew what to do. Each participant was asked to refine all five robot-generated maps using the interface, resulting in a total of 50 refined maps (10 participants x 5 maps). To reduce learning effects, the order in which a participant received the maps was randomized. This ensures that the participants did not systematically improve their performance based on how often they used the interface.

The task for every map was identical: clean and reconstruct the map to best match your understanding of the environment. Each refined map was saved and later used for analysis.

The entire session lasts approximately 15-30 minutes per participant.

5.2.1.4. Data Collection

Similar to the first user test, the refined maps were visually analyzed by overlaying the maps created for each participant. This way we can compare the refined maps with the robot-generated maps

To quantify the effect of refining the map with the user-in-the-loop approach, the performance of the robot-generated map was tested against the refined maps. The goal of the test was to figure out the localization accuracy and navigation performance of the robot using different maps.

To evaluate the localization accuracy, the robot was placed at three different positions in the environment and made to rotate in place until it correctly localized itself on the map. This was tested using both robot-generated maps and refined maps.

However, this method failed due to the software not being fit for localization. The system relies on manual initialization of the robot's position and the start of navigation and was not capable of autonomously determining its location. As a result, localization accuracy could not be reliably measured in this setup. Navigation performance was evaluated by measuring the time required for the robot to move between predefined points in the environment.

Three pairs of start and end locations were predefined: A -> B, C -> D, and E ->F (see figure 48). For each pair the robot was placed at the start location and tasked to move to the corresponding end location. The time taken to reach the goal was recorded. This procedure was conducted using 3 robot-generated maps and 3 refined maps.

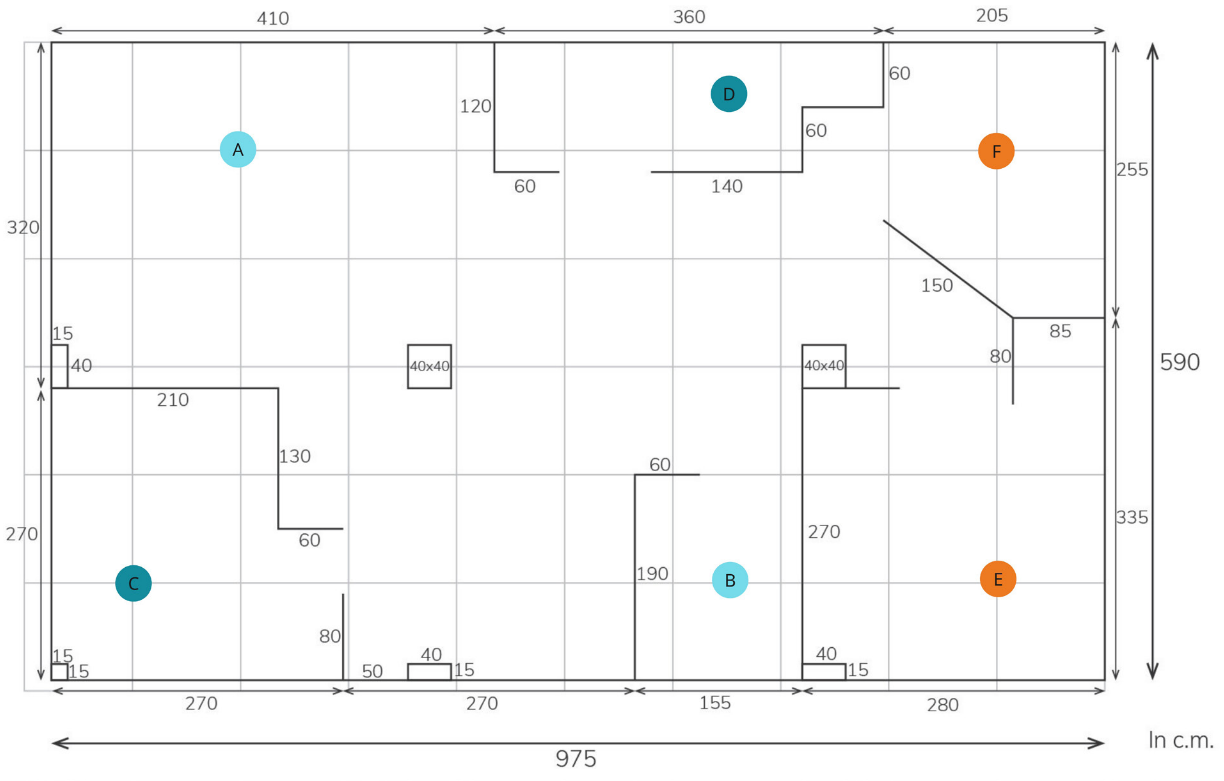

Figure 48: The start and end points for the navigation test

5.2.2. Results

5.2.2.1. Map Consistency

The refined maps created by participants were overlaid per participant to assess structural consistency. The maps per participant got overlaid as well as the five robot-generated maps to visually compare, see Figure 49.

To assess the structural consistency, the refined maps produced by each participant were overlaid and compared against an overlay of the five robot-generated maps (see Figure 50). For all map results see Appendix F

The overlay of the robot-generated maps shows some variation in wall placement and has some noise present in the map. In contrast, most of the participant-refined maps show a higher degree of alignment, with walls appearing sharper and more consistently positioned across overlays. Additionally, most of the apparent noise visible in the robot-generated maps has been removed in the refined versions.

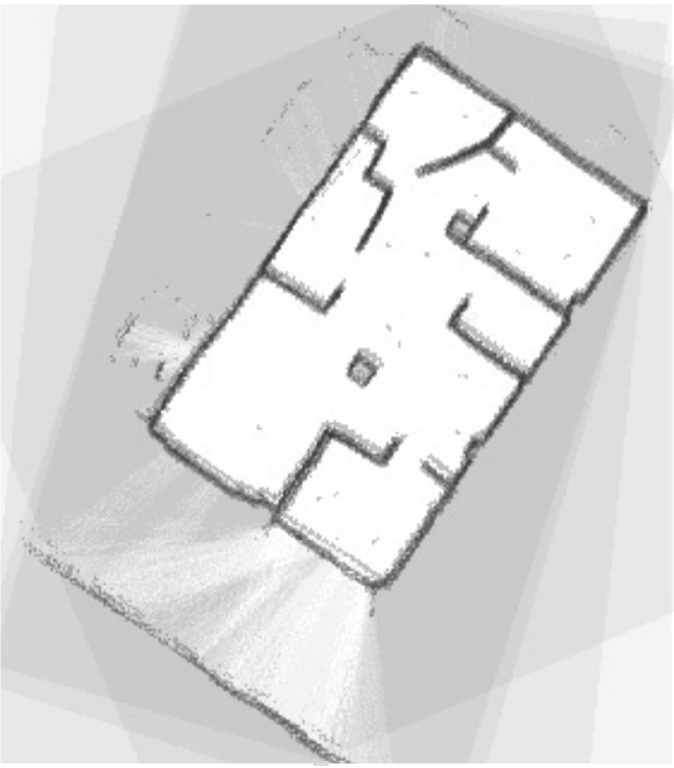

Figure 49: Robot-generated Map overlay

5.2.2.2. Localization Accuracy

An attempt was made to evaluate localization accuracy by placing the robot at different positions and allowing it to localize itself within the map.

However, due to limitations in the software used, the robot required manual initialization of its position and was not able to autonomously determine its location. As a result, no reliable data on localization accuracy could be obtained.

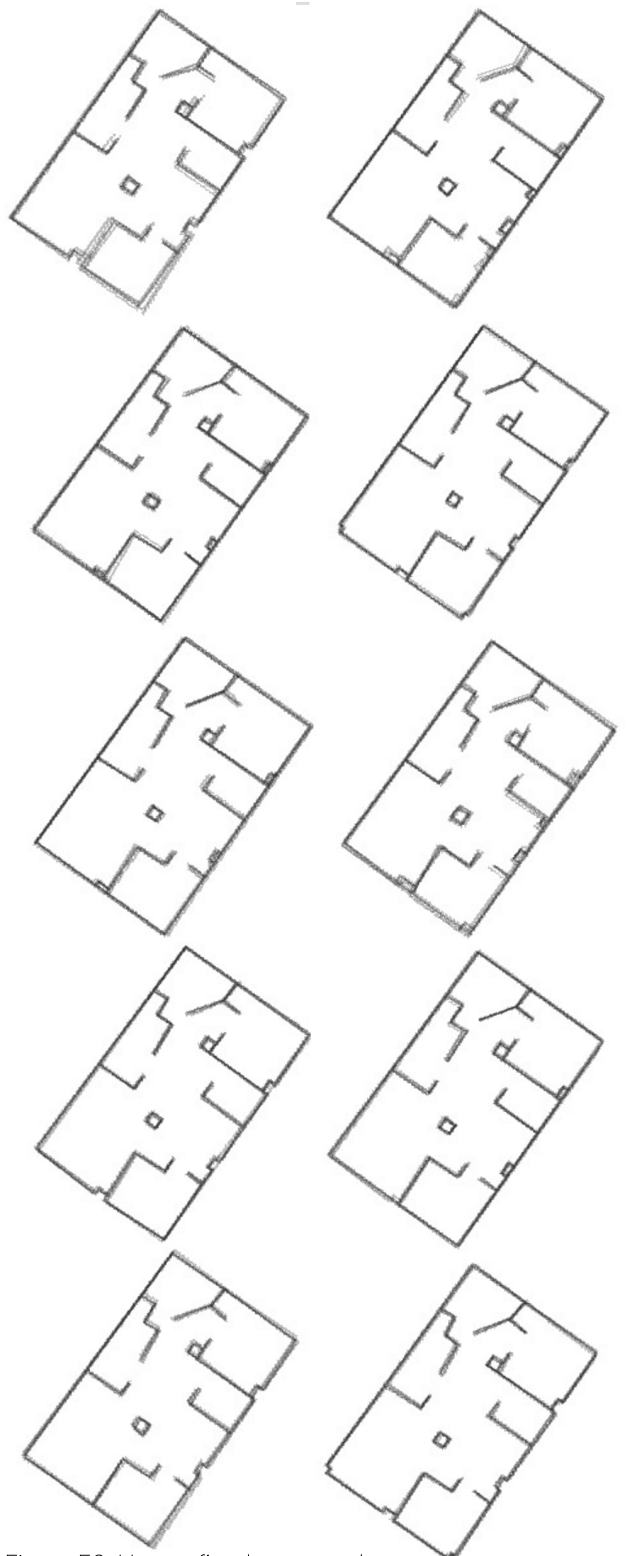

Figure 50: User-refined map overlays

5.2.2.3. Navigation Performance

Navigation performance was evaluated by comparing the time required for the robot to traverse predefined routes using both robot-generated maps and user-refined maps. The recorded completion times, including unsuccessful attempts where the robot became stuck, are presented in Table 5.

Figure 51 presents a scatter plot visualizing the navigation times across different runs. The failed attempts are left out of this graph. The robot-generated maps are represented using different shades of blue while the refined maps are shown in different shades of orange.

Map	A->B (S)	C->D (S)	E->F (S)
Robot map 1	55.4	48.7	**Failed**
Refined map 1	56.9	46.4	76.4
Robot map 2	**Failed**	**Failed**	81.1
Refined map 2	54.9	47.5	**Failed**
Robot map 3	59.3	**Failed**	**Failed**
Refined map 3	**Failed**	48.2	78.4

Table 5: Navigation time (in seconds) for each route across map types

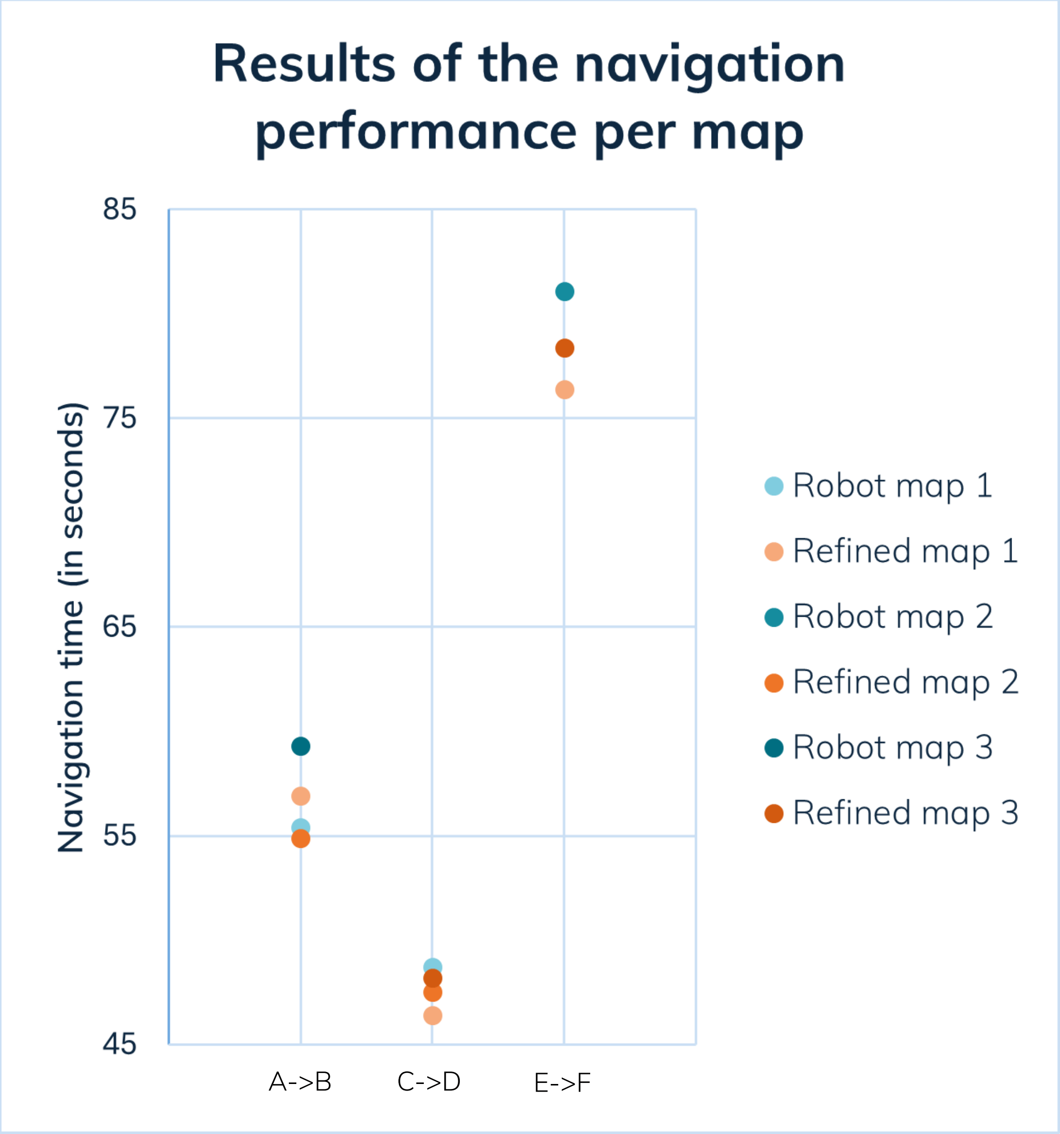

Figure 51: Scatter plot of the navigation performance per map

To evaluate whether map refinement improved navigation performance, a Fisher's Exact test was conducted to see if there is any correlation between the map type and navigation success. The analysis was performed in Python using the SciPy library (see Appendix G for the code). For robot-generated maps, navigation was successful in 4 out of 9 runs (44%), compared to 7 out of 9 runs (78%) using refined maps. This indicates a clear improvement in performance when using refined maps. However, the Fisher's exact test (two-sided) did not show a statistically significant association between map type and navigation success ($p = 0.33$). Despite the lack of statistical significance, the observed effect size was large, with an odds ratio of 4.38, indicating that successful navigation was over four times more likely when using refined maps, although this cannot be said for certain due to the low statistical significance. The results of the Fisher's exact test are summarized in Table 6.

To further evaluate navigation performance, the time required to complete navigation segments was analyzed only for the successful runs. Due to the small sample size and non-normal distribution of the data, a non-parametric Mann–Whitney U test was conducted in Python using the SciPy library (see Appendix G for the code) to compare navigation times between robot-generated and refined maps for each segment. The results showed no statistically significant differences between conditions (see table 7). This is mainly due to the very limited number of successful runs per condition. Across all segments, refined maps consistently resulted in lower completion times compared to robot-generated maps, although these differences are minimal and show no statistical significance.

Segment	Robot mean (s)	Refined mean (s)	P-value
A -> B	57.35	55.9	0.667
C -> D	48.7	47.4	0.5
E -> F	81.1	77.4	0.667

Table 7: Mann-Whitney U test analysis of the successful runs

Map	Success	Failure	Succes rate	Odds ratio	P-value
Robot-Generated	4	5	44%		
Refined	7	2	78%	4.38	0.33

Table 6: Summary of the results of the Fisher's exact test

5.2.2. Discussion

The results of this user test provide insights into the impact of user-refined maps on map quality and robot navigation performance. The different results will now be discussed and used to reflect on the effectiveness of the refinement process and the final design concept.

5.2.3.1 Map consistency

The comparison between the overlaid robot-generated maps and the participant-refined maps reveals an improvement in structural consistency. The refined maps appear cleaner primarily due to the significant reduction of noise in the map. The robot-generated maps also show variation in wall placement, which is a characteristic of SLAM-based mapping where sensor inaccuracies and drift accumulate over time.

The participant-refined maps appear more consistently aligned and contain sharper, more clearly defined walls. This suggests that participants were able to interpret the robot-generated maps and use their understanding of the environment to correct inconsistencies and remove noise. This aligns well with the research done by Selvaggio et al. (2021) and Johnson and Vera (2019), which states that autonomous systems often benefit from human cooperation.

This supports the idea that users can act as an effective post-processing step in the mapping process. While the robot generates maps based purely on sensor input, users apply spatial reasoning and prior knowledge of the environment to simplify and correct the representation.

Ultimately the users eliminated most noise out of the robot-generated maps, however, the participants also avoided making large changes to the original map. Even when they perceived the walls to be straight and the robot drew the walls at a certain angle due to localization drift, the participant often stuck with the angle from the robot-generated map.

5.2.3.2. Localization Accuracy

An important limitation in this test is the inability to measure the localization accuracy. Since the robot required manual initialization and was not able to autonomously determine its position, no conclusions can be drawn about how map refinement influences localization.

Further research should definitely look into quantifying the localization accuracy of the user-in-the-loop mapping approach

5.2.3.3. Navigation Performance

The results of the navigation performance analysis do not show any statistically significant differences between robot-generated maps and refined maps, and therefore no definitive conclusions can be made about the effect of map refinement on navigation performance, however we can observe some trends that the results show.

For navigation success, although refined maps showed a higher success rate (78%) compared to robot-generated maps (44%), no statistical significance was found using the Fisher's exact test between map type and success rate ($p = 0.33$). This means that, based on the current dataset, it cannot be concluded with certainty that map refinement improves navigation reliability. The observed difference may show a positive effect, but it could also be due to random variation within the small sample size.

Similarly, the analysis of navigation times for successful runs did not reveal any statistically significant differences between conditions. While refined maps consistently showed slightly lower average completion times across all segments, these differences were small and not supported by statistical evidence. As such, no claims can be made regarding improvements in navigation efficiency.

The lack of statistical significance across both analyses is mainly due to the limited number of runs and the high number of failed attempts, which reduce the statistical power of the tests. As a result, the current findings should be interpreted with caution, and additional testing with a larger dataset is required to determine whether the observed trends are actual correlations.

Despite this, the observed trends do align with existing literature that states that improving map quality by taking away inaccuracies like noise and misaligned structures can positively influence the navigation ability of the robot (Cadena et al., 2016). Additionally, studies on domestic robots show that navigation failures such as getting stuck or failing to reach a goal are often linked to incorrect perception of the environment by the robot (Yapici et al., 2022). However, based on the current results of this test, it cannot be concluded that map refinement improves navigation performance, additional testing is needed.

Regardless of the absence of statistically significant results, an important observation from this test is that the refinement process did not negatively impact navigation performance. The robot was still able to operate on refined maps, and no decrease in performance was observed compared to the robot-generated maps. Within the scope of this project, this suggests that the refinement approach is at least compatible with the navigation system and does not introduce issues that hinder performance.

5.2.3.4. Relation to Previous User Test

The results of this second test provide a quantitative view on the findings from the initial user test (Chapter 4) but do not offer statistically significant evidence to confirm the observed effects. In the first test, refined maps were qualitatively perceived as cleaner and more usable. While the current test introduces performance measurements, such as navigation success rates and completion times, these did not produce statistically significant differences between conditions.
As a result, the findings of this test cannot conclusively validate the effectiveness of the proposed user-in-the-loop approach. However, the observed trends, such as higher success rates and slightly lower navigation times for refined maps, are consistent with the earlier qualitative observations. These trends may indicate a potential positive effect of map refinement, but further testing with a larger dataset is required to determine whether these effects are reliable.
Within the scope of this project, it can be stated that the user-in-the-loop refinement approach does not negatively impact navigation performance and remains compatible with the system. While its effectiveness cannot be confirmed based on the current results, it shows potential and calls for further investigation.

5.2.3.4. Limitations

There are several limitations that should be considered when interpreting these results.

First, the number of trials is limited, with only three maps per condition. In addition, each navigation path was executed only once per map and route. This results in a very small dataset, which substantially reduces the statistical power of the analyses and limits the ability to detect significant effects. Although some trends can be observed, a larger number of repeated trials would be required to draw more reliable and statistically supported conclusions.

Second, the evaluation in these tests is still primarily qualitative, relying on visual inspection of map overlays and navigation performance assessed through completion time. Although differences can be observed, the quantitative metrics are still fairly limited, especially due to the failure in measuring localization accuracy.

Third, failures in navigation were recorded, with the robot getting stuck and failing to reach its end goal. However, the exact cause of failure was not systematically analyzed. Making it unclear whether failures were caused by map inaccuracies, localization errors, or other factors. This makes it hard to directly relate failed runs to the quality of the map.

Overall, due to these limitations and the absence of statistically significant results, no definitive conclusions can be drawn regarding the effectiveness of user-in-the-loop map refinement. While the observed trends suggest potential improvements in navigation performance, additional testing with a larger dataset and more comprehensive metrics is required to determine whether these effects are reliable. At the same time, the results indicate that the refinement approach does not negatively impact navigation performance, suggesting that it remains a viable direction for further investigation.

06

CONCLUSION

This chapter presents the conclusions of this project by reflecting on the conducted research, design process, and evaluation. The findings are discussed in relation to the defined research questions to achieve a well-structured ending to the project. Lastly, recommendations for future work are provided, outlining possible directions for further development and evaluation of the proposed approach.

6.1. CONCLUSION

This project set out to create a seamless human-robot mapping experience in which a user-in-the-loop can collaborate with the robot to improve the reliability and efficiency of mapping a domestic environment, ultimately creating a smoother overall user experience. While SLAM algorithms provide a strong base for the creation of robot-generated maps, their performance is still constrained by sensor noise and complex environments, especially when containing reflective surfaces. Therefore, faulty maps in consumer products can negatively impact user trust, especially when it is difficult or time-consuming to correct these mapping errors.

The project explored an approach where the spatial reasoning of humans is used to enhance the quality of the robot-generated maps. Through a developed interface, users got the ability to directly interpret, adjust, and improve the robot-generated maps to then utilize them for future navigation.

The results following the user tests show some positive trends. Although no statistically significant results were found, the user tests do show that the approach of the system is feasible. The robot was still able to navigate using the refined maps, and no negative effects on performance were observed. While the results do not allow for definitive conclusions, they at least appear positive and indicate that integrating user input into the mapping process has the potential to improve mapping reliability.

Besides map accuracy, this project also looked into the importance of map interpretability and interaction design. Users were generally able to understand the robot-generated map quite well, especially when they had a clear reference to the physical environment. Furthermore, as found during literature research, involving users in the process helps to improve their understanding of the system, leading to less misinterpretation between robot and user (Kaber and Endsley, 2004).

The final design presents a concept for interactive map refinement that demonstrates a working user-in-the-loop approach and a good base for future research. This concept reframes the role of the user to an active collaborator in the mapping process and suggests a shift toward more collaborative forms of human-robot interaction. However, further research is necessary to validate the positive trends perceived in this project.

Addressing the research questions

To further evaluate the contribution of this project, the findings are discussed in relation to the research questions defined at the start of the project.

SQ1 - Understanding

How do users currently perceive and understand robot-generated maps?

This project investigated how users perceive and understand robot-generated maps. The user test showed that participants were generally able to interpret the content of the maps and identify key elements such as walls, rooms, and obstacles. However, there were some problem areas present in the maps shown, particularly distorted walls, reflective areas like windows or mirrors, or noise.

Users were especially sensitive to structural errors, such as irregularities in walls due to sensor noise, indicating that human spatial reasoning is well-suited to detecting global errors. At the same time, if whole walls were misplaced due to localization drift, it was less consistently recognized, and the robot map was often believed. The understanding of the user appears to be primarily qualitative and pattern-based, with the user putting a lot of trust into the robot-generated map and avoiding making large deviations from the original structure of the base map.

Overall, the findings indicate that users are capable of understanding robot-generated maps to a sufficient degree to work with and edit them. However, the robot map does need a certain level of accuracy and clarity to be able to be used.

SQ2 - Design Exploration

How can users effectively contribute to correcting and refining robot-generated maps?

The second research question explored how users can contribute to correcting and refining robot-generated maps. Through the designed interface, users were able to create their own refined version of the robot-generated map on top and in the process correct structural errors such as wall misalignment, missing sections, and noisy areas.

The results show that users can reliably identify and correct high-level structural errors, particularly when these errors conflict with their knowledge of the mapped environment. At the same time, corrections were mainly focused on easily interpreted errors, staying close to the original map. This highlights that user-in-the-loop systems are most effective when collaborating closely with the robot, rather than attempting to fully replace algorithmic processes.

Even though the results from the user test were not statistically significant, they do demonstrate that the approach is feasible, as the robot was able to navigate using the refined maps. While no definitive conclusion can be drawn on the reliability of the refined maps, the results showed a positive trend and make for a promising direction for future research.

SQ3 - Evaluation

What is the impact of a collaborative human-robot mapping approach on user experience, map accuracy, and reliability?

The third research question examined the impact of a collaborative map ping approach on user experience, map accuracy, and reliability. Based on the results of this study, no statistically significant effects were found, and therefore no definitive conclusions can be made regarding improvements in map accuracy or navigation reliability.

From a usability perspective, participants were generally able to work with the interface and interact with the robot-generated maps. Literature suggests that being actively involved in an automated process increases a human's understanding of how the system works (Kaber and Endsley, 2004). This suggests that incorporating users into the mapping experience can enhance transparency. In addition, user involvement may have a positive influence on their trust in the system. By allowing users to directly interact with and influence the map, the process becomes less of a "black box," making the robot's behavior more understandable and predictable.

In terms of performance, the robot was able to navigate using the refined maps, indicating that the approach is at least feasible and does not negatively impact navigation performance. However, further testing with a larger dataset is required to determine whether the refined map actually has a positive effect on the reliability of the maps.

Overall, while the effectiveness of the proposed user-in-the-loop approach cannot be confirmed based on the current results, integrating users into the mapping process shows potential as a way to improve understanding, engagement, and trust. As such, this collaborative human-robot interaction explored in this project makes for an interesting direction for future research.

Main research question

How can a user-in-the-loop mapping interface improve the usability, transparency, and reliability of domestic cleaning robots?

Based on the findings, it can be concluded that a user-in-the-loop mapping interface has the potential to improve the following:

Usability, by enabling more intuitive interaction with robot-generated maps and reducing the need to interpret abstract system behavior.

Transparency, by allowing users to directly engage with and better understand the mapping process.

Reliability, by enabling users to correct errors that autonomous systems fail to resolve using human spatial reasoning.

The approach does not replace autonomous mapping, but rather complements it by introducing human intelligence where it is most effective. This results in a collaborative system in which the strengths of both humans and robots are combined.

While no statistically significant improvements in performance were found, this project demonstrates a feasible approach to a collaborative mapping experience. Further research is required to determine whether it can lead to measurable improvements in mapping reliability.

Overall, the user-in-the-loop mapping concept shows potential as a collaborative approach that may contribute to more understandable and user-centered robotic systems and therefore represents an interesting direction for future research in domestic robotics.

6.2. RECOMMENDATIONS

The results of this project demonstrate the potential of a user-in-the-loop mapping approach. At the same time, the findings and limitations highlight several recommendations for further development and validation of the concept.

First, the quantitative evaluation of the mapping performance should be improved. In the current study, measuring the localization accuracy failed due to limitations of the robot used. Future work should therefore replicate this test using equipment that supports reliable localization measurement in order to create a more complete assessment of the overall system performance.

Second, more extensive results for navigation performance is required. The current study was based on a limited number of results, which prevented any statistically significant conclusions from being made. Future studies should involve larger and more diverse participant groups and a higher number of repeated trials. This would allow for a more robust quantitative comparison between robot-generated maps and user-refined maps, particularly in terms of navigation accuracy, efficiency, and reliability.

Finally, future works should explore the application of this concept in real-world settings. While this project aimed to closely mimic real-world conditions by working with a physical robot and test environment instead of computer simulations. This still differs from actual users in a domestic environment. Future research should investigate whether users are willing to be actively engaged in editing and refining the map, whether they value the added control, or whether they perceive this process as unnecessary or tedious. Understanding the balance between users' effort, control, and perceived benefit is essential for the success of the collaborative mapping concept.

In conclusion, while this project provides a promising foundation, further evaluation is needed to determine the true potential and feasibility of the proposed concept.

REFERENCES

Beer, J. M., Fisk, A. D., & Rogers, W. A. (2014). Toward a framework for levels of robot autonomy in Human-Robot interaction. Journal of Human-Robot Interaction, 3(2), 74. https://doi.org/10.5898/jhri.3.2.beer

Björnfot, P., Bergqvist, J., & Kaptelinin, V. (2018). Non-technical users' first encounters with a robotic telepresence technology: An empirical study of office workers. Paladyn Journal of Behavioral Robotics, 9(1), 307–322. https://doi.org/10.1515/pjbr-2018-0022

Bosse, M., Newman, P., Leonard, J., Soika, M., Feiten, W., & Teller, S. (2004). An Atlas framework for scalable mapping. IEEE International Conference on Robotics and Automation (Cat. No.03CH37422). https://doi.org/10.1109/robot.2003.1241872

Buonocore, L., Santos, S. R. B. D., Neto, A. A., & Nascimento, C. L. (2016). FastSLAM filter implementation for indoor autonomous robot. 2022 IEEE Intelligent Vehicles Symposium (IV), 484–489. https://doi.org/10.1109/ivs.2016.7535430

Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., & Leonard, J. J. (2016). Past, present, and future of simultaneous localization and mapping: toward the Robust-Perception age. IEEE Transactions on Robotics, 32(6), 1309–1332. https://doi.org/10.1109/tro.2016.2624754

Choi, G., Park, J., Shlezinger, N., Eldar, Y. C., & Lee, N. (2022). Split-KalmanNet: A robust Model-Based Deep Learning Approach for SLAM. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2210.09636

Coggins, T. N. (2022). More work for Roomba? Domestic robots, housework and the production of privacy. Prometheus, 38(1). https://doi.org/10.13169/prometheus.38.1.0098

Cuprik, R. (2023). Gathering dust and data: How robotic vacuums can spy on you. https://www.welivesecurity.com/en/privacy/gathering-dust-and-data-how-robotic-vacuums-can-spy-on-you/

Cyberneticzoo.com. (2013, October 9). 1997/2002 - Electrolux Trilobite Robotic Vacuum Cleaner - Anders Haegermarck, Lars Kilstrom, Bjorn Riise (Swedish) - Cyberneticzoo.com. cyberneticzoo.com. https://cyberneticzoo.com/early-service-robots/19972002-electrolux-trilobite-robotic-vacuum-cleaner-anders-haegermarck-lars-kilstrom-bjorn-riise-swedish/

Debeunne, C., & Vivet, D. (2020). A Review of Visual-LiDAR Fusion based Simultaneous Localization and Mapping. Sensors, 20(7), 2068. https://doi.org/10.3390/s20072068

Dellaert, F., & Kaess, M. (2006). Square Root SAM: Simultaneous localization and mapping via square root information smoothing. The International Journal of Robotics Research, 25(12), 1181–1203. https://doi.org/10.1177/0278364906072768

Dissanayake, G., Huang, S., Wang, Z., & Ranasinghe, R. (2011). A review of recent developments in Simultaneous Localization and Mapping. 2011 6th International Conference on Industrial and Information Systems. https://doi.org/10.1109/iciinfs.2011.6038117

Durrant-Whyte, H., & Bailey, T. (2006). Simultaneous localization and mapping: part I. IEEE Robotics & Automation Magazine, 13(2), 99–110. https://doi.org/10.1109/mra.2006.1638022

Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. Psychological Review, 114(4), 864–886. https://doi.org/10.1037/0033-295x.114.4.864

Fink, J., Bauwens, V., Mubin, O., Kaplan, F., & Dillenbourg, P. (2011). People's perception of domestic service robots: same household, same opinion? In Lecture notes in computer science (pp. 204–213). https://doi.org/10.1007/978-3-642-25504-5_21

Forlizzi, J. (2007). How robotic products become social products (pp. 129–136). https://doi.org/10.1145/1228716.1228734

Forlizzi, J., & DiSalvo, C. (2006). Service robots in the domestic environment: a study of the roomba vacuum in the home (pp. 258–265). https://doi.org/10.1145/1121241.1121286

Google Scholar. (2020). Top publications: Human-computer Interaction. https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng_humancomputerinteraction.

Grisetti, G., Kummerle, R., Stachniss, C., & Burgard, W. (2010). A tutorial on Graph-Based SLAM. IEEE Intelligent Transportation Systems Magazine, 2(4), 31–43. https://doi.org/10.1109/mits.2010.939925

Guo, E. (2024, March 11). A Roomba recorded a woman on the toilet. How did screenshots end up on Facebook? MIT Technology Review. https://www.technologyreview.com/2022/12/19/1065306/roomba-irobot-robot-vacuums-artificial-intelligence-training-data-privacy/

Ho, N. T. S., Fai, N. Y. C., & Ming, E. S. L. (2015). Simultaneous localization and mapping survey based on filtering techniques. 2022 13th Asian Control Conference (ASCC), 1–6. https://doi.org/10.1109/ascc.2015.7244836

Johnson, M., & Vera, A. H. (2019). No AI is an island: the case for teaming intelligence. AI Magazine, 40(1), 16–28. https://doi.org/10.1609/aimag.v40i1.2842

Julier, S., & Uhlmann, J. (2002). A counter example to the theory of simultaneous localization and map building. IEEE International Conference on Robotics and Automation, 4, 4238–4243. https://doi.org/10.1109/robot.2001.933280

Kaber, D. B., & Endsley, M. R. (2004). The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. Theoretical Issues in Ergonomics Science, 5(2), 113–153. https://doi.org/10.1080/1463922021000054335

Kazerouni, I. A., Fitzgerald, L., Dooly, G., & Toal, D. (2022). A survey of state-of-the-art on visual SLAM. Expert Systems With Applications, 205, 117734. https://doi.org/10.1016/j.eswa.2022.117734

Khairuddin, A. R., Talib, M. S., & Haron, H. (2015). Review on simultaneous localization and mapping (SLAM). 2015 IEEE International Conference on Control System, Computing and Engineering (ICCSCE), 85–90. https://doi.org/10.1109/iccsce.2015.7482163

Kim, H., Lee, H., Chung, S., & Kim, C. (2007). User-centered approach to path planning of cleaning robots (pp. 373–380). https://doi.org/10.1145/1228716.1228766

Klein, G., & Murray, D. (2007). Parallel Tracking and Mapping for Small AR Workspaces. 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality. https://doi.org/10.1109/ismar.2007.4538852

Kok, B. C., & Soh, H. (2020). Trust in robots: challenges and opportunities. Current Robotics Reports, 1(4), 297–309. https://doi.org/10.1007/s43154-020-00029-y

Linjawi, M., & Moore, R. K. (2018). Towards a comprehensive taxonomy for characterizing robots. In Lecture notes in computer science (pp. 381–392). https://doi.org/10.1007/978-3-319-96728-8_32

Lu, F., & Milios, E. (1997). Globally Consistent Range Scan Alignment for Environment Mapping. Autonomous Robots, 4(4), 333–349. https://doi.org/10.1023/A:1008854305733

Lutz, C., & Tamò-Larrieux, A. (2021). Do privacy concerns about social robots affect use intentions? Evidence from an experimental vignette study. Frontiers in Robotics and AI, 8, 627958. https://doi.org/10.3389/frobt.2021.627958

Merat, N., Seppelt, B., Louw, T., Engström, J., Lee, J. D., Johansson, E., Green, C. A., Katazaki, S., Monk, C., Itoh, M., McGehee, D., Sunda, T., Unoura, K., Victor, T., Schieben, A., & Keinath, A. (2018). The “Out-of-the-Loop” concept in automated driving: proposed definition, measures and implications. Cognition Technology & Work, 21(1), 87–98. https://doi.org/10.1007/s10111-018-0525-8

MIRTE team. (2025). MIRTE Master — MIRTE Documentation v0.1 documentation. https://docs.mirte.org/develop/doc/robots/mirte_master/index.html

Montemerlo, M., Thrun, S., Koller, D., & Wegbreit, B. (2002). FastSLAM: a factored solution to the simultaneous localization and mapping problem. National Conference on Artificial Intelligence, 593–598. https://doi.org/10.5555/777092.777184

Mordor Intelligence. (2025, August 20). Robotic Vacuum Cleaner Market Size, Growth Report 2030 | Industry Overview. https://www.mordorintelligence.com/industry-reports/robot-vacuum-cleaners-market

Neira, J., & Tardos, J. (2001). Data association in stochastic mapping using the joint compatibility test. IEEE Transactions on Robotics and Automation, 17(6), 890–897. https://doi.org/10.1109/70.976019

Placed, J. A., Strader, J., Carrillo, H., Atanasov, N., Indelman, V., Carlone, L., & Castellanos, J. A. (2023). A survey on Active Simultaneous Localization and Mapping: State of the Art and new frontiers. IEEE Transactions on Robotics, 39(3), 1686–1705. https://doi.org/10.1109/tro.2023.3248510

Prassler, E., Ritter, A., Schaeffer, C., & Fiorini, P. (2000). A Short History of Cleaning Robots. Autonomous Robots, 9(3), 211–226. https://doi.org/10.1023/a:1008974515925

Price, B. R. (2023). Colonel John Boyd’s Thoughts on Disruption: A Useful Effects Spiral from Uncertainty to Chaos. Journal of Advanced Military Studies, 14(1), 98–117. https://doi.org/10.21140/mcuj.20231401004

Rogers, E. M. (2019). Diffusion of Innovations 1 (pp. 415–434). https://doi.org/10.4324/9780203710753-35

Sami, S., Dai, Y., Tan, S. R. X., Roy, N., & Han, J. (2020). Spying with your robot vacuum cleaner (pp. 354–367). https://doi.org/10.1145/3384419.3430781

Schubert, S., Neubert, P., & Protzel, P. (2021). Graph-Based Non-Linear least squares optimization for visual place recognition in changing environments. IEEE Robotics and Automation Letters, 6(2), 811–818. https://doi.org/10.1109/lra.2021.3052446

Scopelliti, M., Giuliani, M. V., & Fornara, F. (2005). Robots in a domestic setting: a psychological approach. Universal Access in the Information Society, 4(2), 146–155. https://doi.org/10.1007/s10209-005-0118-1

Selvaggio, M., Cognetti, M., Nikolaidis, S., Ivaldi, S., & Siciliano, B. (2021). Autonomy in Physical Human-Robot Interaction: A Brief survey. IEEE Robotics and Automation Letters, 6(4), 7989–7996. https://doi.org/10.1109/lra.2021.3100603

Sheridan, T. (2002). Eight ultimate challenges of human-robot communication. Proc. IEEE Int. Workshop Robot-Human Interactive Communication (RO-MAN), 9–14. https://doi.org/10.1109/roman.1997.646944

Sheridan, T. B. (2016). Human–Robot interaction. Human Factors the Journal of the Human Factors and Ergonomics Society, 58(4), 525–532. https://doi.org/10.1177/0018720816644364

Sheridan, T. B., & Verplank, W. L. (1978). Human and computer control of undersea teleoperators. https://doi.org/10.21236/ada057655

Shove, E. (2003). Comfort, cleanliness and convenience : the social organization of normality. In Bloomsbury Academic eBooks. https://doi.org/10.5040/9781474214605

Steder, B., Grisetti, G., Grzonka, S., Stachniss, C., Rottmann, A., & Burgard, W. (2007). Learning maps in 3D using attitude and noisy vision sensors. 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, 644–649. https://doi.org/10.1109/iros.2007.4399414

Suchman, L. (2002). Located accountabilities in technology production. AIS Electronic Library (AISeL) (Association for Information Systems), 14(2), 91–105. http://aisel.aisnet.org/sjis/vol14/iss2/7

Sung, Grinter, Christensen, & Guo. (2008). Housewives or technophiles?: Understanding domestic robot owners. Human-Robot Interaction, 129–136. http://yadda.icm.edu.pl/yadda/element/bwmeta1.element.ieee-000006249425

Sung, J., Guo, L., Grinter, R. E., & Christensen, H. I. (2007). “My Roomba is Rambo”: Intimate home appliances. In Lecture notes in computer science (pp. 145–162). https://doi.org/10.1007/978-3-540-74853-3_9

Taheri, H., & Xia, Z. C. (2020). SLAM; definition and evolution. Engineering Applications of Artificial Intelligence, 97, 104032. https://doi.org/10.1016/j.engappai.2020.104032

Taylor, A. S., & Swan, L. (2005). Artful systems in the home (pp. 641–650). https://doi.org/10.1145/1054972.1055060

The MathWorks, Inc. (n.d.). Https://www.mathworks.com/discovery/slam.html. https://www.mathworks.com/discovery/slam.html

Thrun, S. (2002). Probabilistic robotics. Communications of the ACM, 45(3), 52–57. https://dl.acm.org/doi/fullHtml/10.1145/504729.504754

Thrun, S., & Montemerlo, M. (2006). The Graph SLAM Algorithm with Applications to Large-Scale Mapping of Urban Structures. The International Journal of Robotics Research, 25(5–6), 403–429. https://doi.org/10.1177/0278364906065387

Triebel, R., Pfaff, P., & Burgard, W. (2006). Multi-Level surface maps for outdoor terrain mapping and loop closing. 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2276–2282. https://doi.org/10.1109/iros.2006.282632

Ulin-Avila, E., & Ponce-Hernandez, J. (2021). Kalman Filter estimation and its implementation. In IntechOpen eBooks. https://doi.org/10.5772/intechopen.97406

Venkatesh, V. (2000). Determinants of Perceived Ease of Use: Integrating Control, Intrinsic Motivation, and Emotion into the Technology Acceptance Model. Information Systems Research, 11(4), 342–365. https://doi.org/10.1287/isre.11.4.342.11872

Wang, Y., & Zhang, Z. (2020). Research on dead-reckoning based localization for cleaning robot. Journal of Physics Conference Series, 1449(1), 012108. https://doi.org/10.1088/1742-6596/1449/1/012108

Xu, J., Yang, Z., Liu, Y., & Cao, H. (2024). Understanding Visual SLAM. In Edge Assisted Mobile Visual SLAM (pp. 3–11). https://doi.org/10.1007/978-981-97-3573-0_1

Yapici, N. B., Tuglulular, T., & Basoglu, N. (2022). Assessment of Human-Robot Interaction between Householders and Robotic Vacuum Cleaners (pp. 204–209). https://doi.org/10.1109/temsconeurope54743.2022.9802007

Zhu, J., Li, H., & Zhang, T. (2023). Camera, LiDAR, and IMU based Multi-Sensor Fusion SLAM: a survey. Tsinghua Science & Technology, 29(2), 415–429. https://doi.org/10.26599/tst.2023.9010010

LIST OF FIGURES