



UNIVERSIDADE FEDERAL DE GOIÁS  
INSTITUTO DE INFORMÁTICA  
PROGRAMA DE PÓS GRADUAÇÃO EM CIÊNCIA DA  
COMPUTAÇÃO

LUCAS DA SILVA ASSIS

# **Bridging the Human-in-the-Loop Gap in Robotics Through Language**

**A Multimodal Agentic AI for Scaling Human Supervision  
in Autonomous Systems**

Goiânia  
2025



**TERMO DE CIÊNCIA E DE AUTORIZAÇÃO (TECA) PARA DISPONIBILIZAR VERSÕES ELETRÔNICAS DE TESES  
E DISSERTAÇÕES NA BIBLIOTECA DIGITAL DA UFG**

Na qualidade de titular dos direitos de autor, autorizo a Universidade Federal de Goiás (UFG) a disponibilizar, gratuitamente, por meio da Biblioteca Digital de Teses e Dissertações (BDTD/UFG), regulamentada pela Resolução CEPEC nº 832/2007, sem ressarcimento dos direitos autorais, de acordo com a [Lei 9.610/98](#), o documento conforme permissões assinaladas abaixo, para fins de leitura, impressão e/ou download, a título de divulgação da produção científica brasileira, a partir desta data.

O conteúdo das Teses e Dissertações disponibilizado na BDTD/UFG é de responsabilidade exclusiva do autor. Ao encaminhar o produto final, o autor(a) e o(a) orientador(a) firmam o compromisso de que o trabalho não contém nenhuma violação de quaisquer direitos autorais ou outro direito de terceiros.

**1. Identificação do material bibliográfico**

Dissertação  Tese  Outro\*: \_\_\_\_\_

\*No caso de mestrado/doutorado profissional, indique o formato do Trabalho de Conclusão de Curso, permitido no documento de área, correspondente ao programa de pós-graduação, orientado pela legislação vigente da CAPES.

Exemplos: Estudo de caso ou Revisão sistemática ou outros formatos.

**2. Nome completo do autor**

Lucas da Silva Assis

**3. Título do trabalho**

Bridging the Human-in-the-Loop Gap in Robotics Through Language A Multimodal Agentic AI for Scaling Human Supervision in Autonomous Systems

**4. Informações de acesso ao documento (este campo deve ser preenchido pelo orientador)**

Concorda com a liberação total do documento  SIM  NÃO<sup>1</sup>

[1] Neste caso o documento será embargado por até um ano a partir da data de defesa. Após esse período, a possível disponibilização ocorrerá apenas mediante:

- a) consulta ao(à) autor(a) e ao(à) orientador(a);
- b) novo Termo de Ciência e de Autorização (TECA) assinado e inserido no arquivo da tese ou dissertação.

O documento não será disponibilizado durante o período de embargo.

Casos de embargo:

- Solicitação de registro de patente;
- Submissão de artigo em revista científica;
- Publicação como capítulo de livro;
- Publicação da dissertação/tese em livro.

**Obs. Este termo deverá ser assinado no SEI pelo orientador e pelo autor.**



Documento assinado eletronicamente por **Anderson Da Silva Soares, Professor do Magistério Superior**, em 19/05/2025, às 18:24, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).



Documento assinado eletronicamente por **Lucas Da Silva Assis, Discente**, em 20/05/2025, às 10:17, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).



A autenticidade deste documento pode ser conferida no site [https://sei.ufg.br/sei/controlador\\_externo.php?acao=documento\\_conferir&id\\_orgao\\_acesso\\_externo=0](https://sei.ufg.br/sei/controlador_externo.php?acao=documento_conferir&id_orgao_acesso_externo=0), informando o código verificador **5380873** e o código CRC **F774F961**.

LUCAS DA SILVA ASSIS

# **Preenchendo a lacuna da supervisão humana na robótica por meio da linguagem**

**Uma IA Agêntica Multimodal para escalar a supervisão humana em sistemas autônomos**

Tese apresentada ao Programa de Pós-Graduação em Ciência da Computação (PPGCC) do Instituto de Informática da Universidade Federal de Goiás, como requisito parcial para obtenção do título de Doutor em Ciência da Computação.

**Área de concentração:** Ciência da Computação.

**Linha de pesquisa:** Sistemas Inteligentes e Aplicações.

**Orientador:** Prof. Anderson da Silva Soares

Goiânia  
2025

Ficha de identificação da obra elaborada pelo autor, através do Programa de Geração Automática do Sistema de Bibliotecas da UFG.

da Silva Assis, Lucas

Bridging the Human-in-the-Loop Gap in Robotics Through Language [manuscrito] : A Multimodal Agentic AI for Scaling Human Supervision in Autonomous Systems / Lucas da Silva Assis. - 2025. 87 f.

Orientador: Prof. Dr. Anderson da Silva Soares.

Tese (Doutorado) - Universidade Federal de Goiás, Instituto de Informática (INF), Programa de Pós-Graduação em Ciência da Computação, Goiânia, 2025.

Bibliografia. Apêndice.

Inclui lista de figuras.

1. Robótica. 2. Agentes baseados em LLMs. 3. Human-in-the-loop. 4. Memory Augmented Reasoning. 5. Self-Discovery Reasoning. I. da Silva Soares, Anderson, orient. II. Título.

CDU 004



UNIVERSIDADE FEDERAL DE GOIÁS

INSTITUTO DE INFORMÁTICA

**ATA DE DEFESA DE TESE**

Ata nº **07/2025** da sessão de Defesa de Tese de **Lucas da Silva Assis**, que confere o título de Doutor em Ciência da Computação, na área de concentração em Ciência da Computação.

Aos vinte e dois dias do mês de abril de dois mil e vinte e cinco, a partir das catorze horas, via sistema de webconferência, realizou-se a sessão pública de Defesa de Tese intitulada “**Bridging the Human-in-the-Loop Gap in Robotics Through Language A Multimodal Agentic AI for Scaling Human Supervision in Autonomous Systems**”. Os trabalhos foram instalados pelo Orientador, Professor Doutor Anderson da Silva Soares (INF/UFG) com a participação dos demais membros da Banca Examinadora: Professor Doutor Gustavo Teodoro Laureano (INF/UFG), membro titular interno; Professor Doutor Flavio Henrique Teles Vieira (INF/UFG), membro titular interno; Professor Doutor Aldo André Díaz Salazar (INF/UFG), membro titular externo; Professora Doutora Esther Luna Colombini (IC/Unicamp), membra titular externa. A realização da banca ocorreu por meio de videoconferência. Durante a arguição os membros da banca não fizeram sugestão de alteração do título do trabalho. A Banca Examinadora reuniu-se em sessão pública a fim de concluir o julgamento da Tese, tendo sido o candidato **aprovado** pelos seus membros. Proclamados os resultados pelo Professor Doutor Anderson da Silva Soares, Presidente da Banca Examinadora, foram encerrados os trabalhos e, para constar, lavrou-se a presente ata que é assinada pelos Membros da Banca Examinadora, aos vinte e dois dias do mês de abril de dois mil e vinte e cinco.

TÍTULO SUGERIDO PELA BANCA



Documento assinado eletronicamente por **Gustavo Teodoro Laureano, Professor do Magistério Superior**, em 22/04/2025, às 17:48, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).



Documento assinado eletronicamente por **Anderson Da Silva Soares, Professor do Magistério Superior**, em 22/04/2025, às 17:48, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).



Documento assinado eletronicamente por **Aldo Andre Diaz Salazar, Professor do Magistério Superior**, em 22/04/2025, às 17:52, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).



Documento assinado eletronicamente por **Lucas Da Silva Assis, Discente**, em 22/04/2025, às 17:55, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).



Documento assinado eletronicamente por **ESTHER LUNA COLOMBINI, Usuário Externo**, em 23/04/2025, às 10:38, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).

---



Documento assinado eletronicamente por **Flavio Henrique Teles Vieira, Professor do Magistério Superior**, em 02/05/2025, às 13:56, conforme horário oficial de Brasília, com fundamento no § 3º do art. 4º do [Decreto nº 10.543, de 13 de novembro de 2020](#).

---



A autenticidade deste documento pode ser conferida no site [https://sei.ufg.br/sei/controlador\\_externo.php?acao=documento\\_conferir&id\\_orgao\\_acesso\\_externo=0](https://sei.ufg.br/sei/controlador_externo.php?acao=documento_conferir&id_orgao_acesso_externo=0), informando o código verificador **5314820** e o código CRC **E9FDB334**.

---

**Referência:** Processo nº 23070.011412/2025-11

SEI nº 5314820

Todos os direitos reservados. É proibida a reprodução total ou parcial do trabalho sem autorização da universidade, do autor e do orientador(a).

### **Lucas da Silva Assis**

Mestre em Ciência da Computação e Bacharelado em Engenharia Elétrica pela Universidade Federal de Goiás. Durante a graduação co-fundou o Núcleo de Robótica Pequena Mecânica na UFG e coordenou equipes em competições de destaque, como IEEE Very Small Size Soccer e RoboCup @Home. Além disso, é instrutor certificado do NVIDIA Deep Learning Institute, compartilhando sua experiência em visão computacional e conectando a pesquisa acadêmica com aplicações no setor.

Como cofundador e CEO da Synkar Autonomous, ele alia mais de 15 anos de experiência em liderança de pesquisas em robótica, inteligência artificial e visão computacional para impulsionar inovações em soluções autônomas para entregas.

”To those who believe that humans and machines can work together—not as adversaries, but as partners. May this research serve as a small step toward a future where intelligence, whether human or artificial, elevates rather than replaces human potential.”

---

## Acknowledgements

---

This journey has been one of intellectual challenge and personal growth, and it would not have been possible without the unwavering support of many individuals and institutions.

First and foremost, I would like to express my deepest gratitude to my wife, Ludmille Dias Ribeiro, whose steadfast support, patience, and encouragement sustained me throughout this endeavor. Her belief in me, even during my moments of doubt, has been my greatest source of strength.

I am profoundly grateful to my advisor, Professor Dr. Anderson Soares, for his guidance, trust, and patience. His mentorship has not only shaped this research but has also challenged me to think critically, push boundaries, and refine my vision for the future of AI-driven robotics.

This thesis has also been shaped by the brilliant minds at Synkar, with whom I had the privilege of collaborating. The stimulating discussions, shared passion for innovation, and relentless pursuit of excellence made this journey far more enriching. I am deeply thankful for the opportunity to work alongside such dedicated individuals.

Finally, to those who believed in me—even when I doubted myself—thank you. This thesis stands as a testament to perseverance, curiosity, and the power of collaboration. While this marks the culmination of my PhD journey, I see it as only the beginning of new questions, discoveries, and contributions to the evolving field of AI-driven robotics.

"The real problem is not whether machines think, but whether men do."

**Burrhus Frederic Skinner,**

.

---

## Resumo

---

Assis, Lucas. **Bridging the Human-in-the-Loop Gap in Robotics Through Language**. Goiânia, 2025. 86p. Tese de Doutorado. Programa de Pós Graduação em Ciência da Computação, Instituto de Informática, Universidade Federal de Goiás.

Embora vários sistemas autônomos consigam navegar e evitar obstáculos em condições previsíveis, eles ainda dependem de supervisão humana (Human-in-the-Loop – HITL) para lidar com obstáculos adversos, mudanças repentinas no ambiente ou falhas parciais de hardware. No entanto, as estratégias HITL atuais frequentemente sobrecarregam os operadores com grandes volumes de dados que exigem interpretação em tempo real. Para mitigar esses desafios, propomos uma abordagem baseada em *Agentic AI* que integra memória de longo prazo com técnicas de raciocínio adaptativo, reduzindo a carga do operador e minimizando interrupções nas operações robóticas.

O sistema proposto incorpora subagentes hierárquicos para integrar dados históricos, registros de sensores e técnicas iterativas de resolução de problemas, enfrentando falhas de localização, defeitos de hardware e obstáculos causados por aglomerações — desafios comuns em implantações com múltiplos robôs. Avaliações experimentais entre condições com e sem memória demonstram que seu uso proporciona maior precisão nas soluções e maior satisfação do operador. A recuperação de memória acelera a resolução de falhas recorrentes, enquanto o raciocínio adaptativo melhora a tomada de decisões em cenários novos ou congestionados. Métricas de similaridade textual (Sobreposição de Tokens e Alinhamento Semântico) mostram que a reutilização de linguagem e estratégias validadas no domínio aumentam a clareza e a consistência das ações recomendadas.

Os resultados destacam a viabilidade de um sistema modular baseado em linguagem, que combina diagnósticos orientados por dados, mecanismos robustos de memória e técnicas de *Self-Discovery* para supervisão de autônomos em larga escala. Ao unir as capacidades flexíveis de LLMs com fluxos de HITL, nossa proposta mostra grande potencial para aumentar a eficiência e a transparência na robótica autônoma aplicada ao mundo real.

### Palavras-chave

<Robótica, Agentes baseados em LLMs, *Human-in-the-loop*, *Self-Discovery Reasoning*, *Memory Augmented Reasoning* >

---

## Abstract

---

Assis, Lucas. **Bridging the Human-in-the-Loop Gap in Robotics Through Language**. Goiânia, 2025. 86p. PhD. Thesis. Programa de Pós Graduação em Ciência da Computação, Instituto de Informática, Universidade Federal de Goiás.

While many autonomous systems can navigate and avoid obstacles under predictable conditions, they often rely on a human supervisor (Human-in-the-Loop - HITL) to adapt to adverse obstacles, sudden layout modifications, or partial hardware failures. However, existing HITL strategies frequently leave operators struggling with large volumes of data that demand real-time interpretation. To mitigate these challenges, we propose an agentic AI approach that integrates long-term memory with adaptive reasoning techniques, thereby reducing operator workload and minimizing disruptions in dynamic autonomous robotics operations.

The proposed system incorporates hierarchical subagents to systematically integrate historical data, sensor logs, and iterative problem-solving techniques to address frequent challenges in multi-robot deployments, including localization failures, hardware malfunctions, and crowd-induced obstacles.

Experimental evaluations comparing memory-augmented and baseline (no-memory) conditions reveal that its usage consistently yields higher solution accuracy and operator satisfaction. In particular, memory retrieval accelerates the resolution of recurring failure modes, while adaptive reasoning enhances real-time decision-making in novel or crowded scenarios. Text-based similarity metrics (Token Overlap and Semantic Alignment) further demonstrate that reusing verified domain language and strategies improves the clarity and maintainability of the recommended actions.

The results underscore the viability of a modular, language-based system that combines data-driven diagnostics, robust memory mechanisms, and self-reflective planning for large-scale robot supervision. By uniting flexible LLM capabilities with HITL workflows, our proposal holds considerable promise for improving both efficiency and transparency in real-world autonomous robotics operations.

### Keywords

<Robotics, Agentic LLMs, Human-in-the-loop, Self-Discovery Reasoning, Memory Augmented Reasoning>

---

# Contents

---

List of Figures	<b>14</b>
<b>1 Introduction</b>	<b>15</b>
1.1 Objectives	18
1.2 Literature review	19
1.2.1 Autonomous Robots	19
Foundations of Autonomous Robotics	19
Modern Trends in Autonomous Robots	20
1.2.2 Human-in-the-Loop (HITL) in Robotics	21
1.2.3 Large Language Models	24
Agents using Large Language Models	25
Single Agents vs Multi-Agents	27
Long-Term Memory in Large Language Models	28
Reasoning in Agentic LLMs	29
<b>2 Problem, Methodology and Proposal</b>	<b>31</b>
2.1 Methodology for the Proposed Architecture	32
2.2 The proposed Architecture	33
2.2.1 The Analyst Subagent	34
2.2.2 The Archivist Subagent	35
2.2.3 The Solver Subagent	36
Adapting Self-Discovery.	37
2.2.4 The Supervisor	38
2.2.5 Workflow Integration Example	39
<b>3 Experiments and Results</b>	<b>41</b>
3.1 Dataset Collection	43
3.2 Experimental Goals and Reasoning	44
3.3 Example of the proposed solution in action.	45
3.4 Experiments	46
3.4.1 Scenario Design and Parameter Variations	46
3.4.2 Results Analysis	47
3.4.3 Module Selection Trends	48
3.4.4 Discussion and Future Directions	53
<b>4 Discussion and Considerations</b>	<b>55</b>
4.1 Summary of Findings and Practical Significance	55
4.2 Limitations and Potential Improvements	56
4.3 Scalability and Deployment Prospects	57

4.4	Final Considerations	58
5	Conclusion	<b>60</b>
5.1	Future Works	61
	Bibliography	<b>62</b>
A	Internal messages within the proposed solution.	<b>72</b>
B	Video Showcase	<b>81</b>

---

## List of Figures

---

2.1	The proposed hierarchical multi-agent architecture of this work.	33
3.1	The operation studied as a use case. <b>Top-Left:</b> Collection Stand worker with an order. <b>Top-Right:</b> Robots waiting for deliveries inside the mall. <b>Bottom-Left:</b> Robot collecting orders. <b>Bottom-Right:</b> Driver after collecting orders.	42
3.2	A side by side of the front camera image and top-view rendering of the robot sensors in the example.	45
3.3	Histogram of Module selection on Case 1.	50
3.4	Histogram of Module selection on Case 1 without memory.	50
3.5	Histogram of Module selection on Case 2.	51
3.6	Histogram of Module selection on Case 2 without memory.	51
3.7	Histogram of Module selection on Case 3.	52
3.8	Histogram of Module selection on Case 3 without memory.	52

## Introduction

---

The field of robotics has experienced significant evolution since the early days of mechanistic control, as exemplified by the seminal work of Brooks [Brooks 1986]. Advances in both robotics and artificial intelligence have paved the way for the development of autonomous systems capable of complex tasks such as navigation, manipulation, and multi-robot coordination [Thrun 2002, Katevas et al. 1997]. However, despite these advancements, real-world applications still require continuous human oversight to manage uncertainties, data overload, and unforeseen system failures [Hu et al. 2022].

Many of today’s autonomous robots operate in dynamic environments, such as shopping malls, hospitals, and city sidewalks, where unpredictable human behavior, frequent infrastructure modifications, and hardware vulnerabilities present ongoing challenges [Wong et al. 2018, Sánchez et al. 2022]. Despite advances in localization and sensor fusion, these machines frequently encounter obstructions (e.g., crowds or construction barriers) or face hardware inconsistencies that exceed the capabilities of purely reactive algorithms [Misaros et al. 2023, Thrun 2002]. Due to these environmental complexities, many autonomous systems still rely on naive approaches or rigid, pre-programmed routines that attempt to navigate obstructions and address sensor anomalies without involving human supervisors [Liu 2021, Sánchez et al. 2022, Wong et al. 2018, Misaros et al. 2023].

In practice, these methods struggle to adapt when faced with changing crowd densities or unexpected obstacles—such as impromptu events or playful human interactions—that deviate from the controlled environments for which several techniques were designed. In the absence of human oversight, the robot’s local planner may repeatedly fail to reroute around newly introduced obstructions or misinterpret sensor noise as legitimate obstacles, leading to mission halts and deadlock situations [Sánchez et al. 2022]. Even in more advanced setups, real-time adjustments often rely on simplistic heuristics that quickly degrade under high levels of uncertainty or partial sensor failure. While these autonomous systems can handle straightforward tasks in controlled environments, the lack of a higher-level decision-maker capable of referencing past successes or promptly responding to emergent hazards impedes their reliability in unpredictable indoor con-

texts [Wong et al. 2018].

Meta-learning has emerged as a compelling alternative for adaptive behaviors in real-world robotics, particularly within the context of reinforcement learning (RL). Techniques such as Model-Agnostic Meta-Learning (MAML) and Probabilistic Embeddings for Actor-Critic RL (PEARL) enable agents to learn across distributions of tasks and quickly adapt to new scenarios with minimal additional data [Finn, Abbeel e Levine 2017, Rakelly et al. 2019]. These methods significantly improve the resilience of autonomous robots, allowing them to fine-tune their policies for dynamic or previously unseen environments. This adaptability is particularly valuable in real-world deployments, where the unpredictability of human behavior, environmental conditions, or hardware limitations necessitate rapid and flexible responses. By improving responsiveness and autonomy, meta-learning addresses the inherent scalability issues associated with supporting real-world robotics operations.

While RL-based meta-learning holds great promise for endowing robots with adaptive behaviors, it faces several challenges when deployed in the real world. Foremost among these challenges are issues of sample inefficiency, which often require thousands of interactions to learn a robust policy, and a general lack of interpretability, making it difficult to diagnose failures or verify decisions in safety-critical applications [Kiran et al. 2021, Dulac-Arnold, Mankowitz e Hester 2019]. Furthermore, RL often struggles with safe exploration, especially in environments where trial-and-error may carry operational or financial risks.

Finally, the most common solution is still the application of human supervisors. However, they can become overwhelmed when each system produces large volumes of error logs and sensor data that require real-time interpretation. Manual oversight alone risks delayed interventions or incomplete diagnoses, especially when complex or unusual problems arise. Therefore, a critical challenge in contemporary robotics is scaling human supervision to support large fleets of autonomous robots operating in dynamic environments. Traditional human-in-the-loop (HITL) systems, while effective in controlled scenarios, are limited by their inability to efficiently process heterogeneous sensor data, maintain long-term contextual awareness, and adapt decision-making processes in real time [Leeper et al. 2012, Bringes et al. 2013]. These limitations underscore the need for novel approaches that enhance the efficiency, interpretability and reliability of HITL frameworks.

Deriving from the complexities of multi-robot HITL supervision and the challenges that arise in dynamic indoor delivery settings, this thesis introduces large language models (LLMs) as a means to unify diagnostic data, operator instructions, and real-time context under a single, flexible interface. Rather than relying on rigidly scripted autonomy or narrowly defined algorithms, the integration of LLMs opens pathways for adaptive,

language-driven reasoning that can interpret nuances, recall relevant past experiences, and generate coherent action plans [Brown et al. 2020, Wei et al. 2022]. This approach broadens the scope of human-in-the-loop collaboration, offering a more intuitive bridge between operator directives and the rapidly evolving needs of autonomous robotic fleets.

Recent progress in large-scale language models has given rise to agentic LLM systems—self-directed programs that chain reasoning steps, invoke external tools, and refine their own hypotheses with minimal human guidance [Huang et al. 2024, Kim et al. 2024, Guo et al. 2023, Wang et al. 2024, Abramovich et al. 2024]. In robotics, however, the term agent traditionally denotes an embodied entity whose perception-planning-control loop must satisfy hard real-time and safety constraints while interacting continuously with the physical world. Whereas a mobile-robot agent grounds its decisions in low-latency sensor feedback and direct motor actuation, an LLM-agent reasons over symbolic representations, orchestrates software tools, and maintains context through iterative prompt–response cycles. Leveraging this complementary form of agency, the framework proposed here positions an LLM agent as a transparent, context-aware supervisor that parses heterogeneous robot logs, recalls long-term operational history, and issues high-level corrective actions. By doing so, it alleviates the cognitive load on human operators and improves scalability in dynamic, multi-robot settings, extending earlier work on embodied autonomy [Alami et al. 1998, Thrun 2002, Wong et al. 2018, Liu 2021, Sánchez et al. 2022, Misaros et al. 2023] with an adaptive, language-driven oversight layer.

Given the growing complexity of real-world autonomous robots and the current need for human oversight, this thesis poses a central question: **How can an agentic AI framework, powered by large language models, be leveraged to scale human supervision and enhance real-time problem-solving in dynamic, highly-uncertain environments?** By focusing on memory-augmented reasoning, self-reflective decision-making, and modular subagents—each assigned specialized roles such as diagnostics, memory management, and solution planning—this research seeks to determine whether an LLM-centric approach can reduce operator workload, increase consistency, and maintain robust performance amid the uncertainties of real-world robotic deployments.

The significance of this research lies in its ability to bridge the gap between autonomous operation and human oversight. While existing systems have made notable progress in either autonomy or HITL integration, they often lack the scalability and flexibility required for real-world large-scale deployments. The agentic AI framework introduced herein decomposes complex supervisory tasks into specialized modules that work in concert, ensuring that every aspect of the operational environment is managed effectively. The integration of modular subagents represents a promising step toward more robust, adaptable, and transparent human-robot interaction, aligning with the current tra-

jectory of research in this rapidly evolving field [Vaswani et al. 2017, Devlin et al. 2019, Zhou et al. 2022]. This integrated approach not only enhances the transparency of decision-making but also fosters greater trust and efficiency in human-robot collaboration, as demonstrated in recent studies [Nooralahzadeh et al. 2024, Zhou et al. 2024].

The contributions of this thesis are multifaceted and specifically tailored to address the various challenges that arise from the need for autonomous robots to operate reliably and effectively in harsh and unpredictable environments, where interactions are difficult to model due to dynamics, uncertainty, and complexity [Wong et al. 2018, Sánchez et al. 2022]. It introduces a modular and hierarchical architecture that enhances LLM capabilities through the use of dedicated hierarchical subagents, a design that has been shown to improve both the accuracy and scalability of automated decision-making [Touvron et al. 2023, Yang et al. 2023]. Additionally, it advances HITL strategies by reducing operator workload through real-time subjective reasoning, which is critical in environments where the volume and complexity of sensor data can overwhelm conventional systems [Mirchandani et al. 2024]. Moreover, the thesis incorporates memory-augmented reasoning techniques [Wang et al. 2023] that enable the system to recall and leverage historical data for enhanced consistency in decision-making. Finally, the proposed framework is validated through experiments in realistic scenarios, demonstrating improvements in response time, operational reliability, and overall user satisfaction.

## 1.1 Objectives

Building on the limitations identified in current Human-in-the-Loop (HITL) supervision [Brown et al. 2020, Mirchandani et al. 2024, Royce et al. 2024, Stiefelhagen et al. 2004, Wen et al. 2018], this thesis sets out to tackle five guiding questions: (i) How can we lighten operators' cognitive load during real-time interventions? (ii) Can an automated system interpret heterogeneous log and sensor streams well enough to assist human decision-making? (iii) What mechanisms allow past solutions to be reused when similar failures reappear? (iv) How do we sustain performance and reliability as the fleet size and intervention frequency scale across diverse environments? (v) Most critically, how do we preserve transparency and human control while introducing greater autonomy into the supervisory loop?

The structure of this thesis reflects the progression of these objectives. Following this section, a detailed literature review is presented in Section 1.2 to establish the theoretical foundations of robotics, HITL systems, and agentic LLM architectures. Chapter 2 outlines the proposed methodology and architectural design, while Chapter 3.4 details the experimental evaluation of the framework. A critical discussion of the results is provided

in Chapter 4, and the thesis concludes in Chapter 5 with a summary of the contributions and their implications for future autonomous robotic systems.

In summary, this thesis outlines the overarching goal of developing an agentic AI framework for scalable, robust, and adaptive multi-robot supervision. The continuous narrative presented here serves as a roadmap for addressing the challenges of integrating advanced language modeling with practical robotics, thereby bridging the gap between human oversight and autonomous operation [OpenAI 2023, Touvron et al. 2023].

## 1.2 Literature review

### 1.2.1 Autonomous Robots

The pursuit of fully autonomous robots has motivated decades of research, from simple sensor-actuator loops to sophisticated systems that perceive, decide, and act in real time. Over the years, robotics researchers have addressed key challenges, including robust navigation, distributed coordination among multiple agents, and biologically inspired swarm behavior. This review traces the lineage of these developments, emphasizing the synergy between foundational theories and practical implementations. While the ultimate goal is to minimize human intervention, real-world scenarios often require human oversight to manage technical support, conflict resolution, or safety concerns.

#### Foundations of Autonomous Robotics

Early robots were often controlled by rigid, hierarchical systems that struggled with uncertainty and dynamic environments. A turning point came with [Brooks 1986], who introduced a *subsumption architecture* in which multiple layers of control operated asynchronously, allowing higher-level behaviors to override lower-level actions when necessary. This design ensured robust real-time responsiveness by preserving the functionality of simpler, reactive modules even as more complex modules took charge of tasks like path planning and environment mapping. Around the same time, [Khatib 1987] revolutionized the governance of manipulators by shifting the focus of control from individual joints to the robot's end-effector, an approach known as the *operational space* formulation. By mapping control laws directly to the workspace where tasks are executed, Khatib's framework allowed for finer control over both motion and applied forces, making robots significantly more adept at physically interacting with their surroundings.

As sensors improved and robots began venturing into less controlled environments, the need to systematically handle uncertainty became paramount. [Thrun 2002] popularized a *probabilistic robotics* paradigm that treated all sensor readings and state estimations as inherently uncertain, leading to Bayesian techniques for sensor fusion

and localization. This foundation provided the framework for breakthroughs in *Simultaneous Localization and Mapping* (SLAM). Early theoretical explorations, such as [Dissanayake et al. 2001], established that a robot could indeed localize itself while constructing a map of unknown terrain. Practical scaling was achieved with *FastSLAM*, introduced by [Montemerlo et al. 2002], which factorized the problem so that even large environments with many landmarks could be efficiently mapped in real time. Together, these developments enabled autonomous systems to operate with a high degree of reliability, even when external sensors were noisy or partial.

While foundational work on joint- or workspace-level control aimed to refine the robot's ability to move and act, planning collision-free paths remained a formidable challenge in higher dimensions. [Lozano-Perez 1990] provided a conceptual leap by defining the configuration space (C-space), where each possible robot pose becomes a point in a high-dimensional space. This abstraction transformed collision avoidance into a problem of identifying and circumventing "obstacle regions" within C-space. Building on this framework, [Khatib 1990] employed *artificial potential fields* for real-time obstacle avoidance, where repulsive forces around obstacles and attractive forces toward goals guide the robot through a continuous control process. Although local minima can pose challenges, the responsiveness and simplicity of reactive planning made it appealing for dynamic scenarios, laying the groundwork for more advanced and computationally intensive algorithms.

The push toward higher-dimensional motion planning accelerated as robots evolved from simple mobile platforms to high-DOF manipulators and legged robots. [Kavraki et al. 1996] addressed these challenges by introducing *Probabilistic Roadmaps* (PRMs), which randomly samples the C-space to build a roadmap of feasible paths during a preprocessing phase and then rapidly responds to multiple queries. Later, [LaValle e Kuffner 2001] introduced *Randomized Kinodynamic Planning*, incorporating differential constraints to enhance dynamic feasibility. Over time, asymptotically optimal variants like RRT\* were developed [Karaman e Frazzoli 2011], ensuring that with sufficient sampling, the discovered paths could approach the global optimum. These methods enabled flexible and efficient planning across a wide range of applications, from industrial robotics to self-driving cars.

### **Modern Trends in Autonomous Robots**

As cameras became increasingly affordable and powerful, computer vision emerged as a key enabler of autonomous navigation. [Davison et al. 2007] pioneered a real-time, single-camera SLAM system (MonoSLAM) that simultaneously estimated camera motion and sparse scene geometry. Subsequent developments like [Mur-Artal, Montiel e Tardós 2015, Mur-Artal e Tardós 2017] refined visual SLAM by

incorporating robust feature tracking, loop closure, and lifelong re-localization. To objectively measure progress, datasets such as [Geiger et al. 2013]’s KITTI provided large-scale, real-world benchmarks for perception, odometry, and object detection, thereby accelerating innovation by offering standardized evaluation metrics. Today, vision-based navigation systems integrate monocular or stereo camera inputs with range sensors (such as LIDAR) and inertial measurements to achieve robust performance in diverse conditions.

Once single-robot autonomy matured, researchers turned their attention to teams or swarms of robots capable of sharing tasks and data. [Balch e Arkin 1998] demonstrated how multiple robots could maintain formations by combining behavior-based principles with local sensing and communication. This approach proved especially promising for tasks such as search and rescue, environmental monitoring, and large-scale warehouse logistics, where distributed sensors and parallel task allocation can significantly increase efficiency. Advances in wireless networking, sensor miniaturization, and onboard computing have continued to enhance the scalability of multi-robot coordination, enabling fleets of dozens—or even hundreds—of robots to operate together with minimal central control.

Finally, the field of autonomous robotics has witnessed a remarkable evolution over the past few decades. Layered architectures have revolutionized how robots manage real-time tasks, probabilistic methods have enabled robust navigation in uncertain environments, and advanced planning algorithms have made high-dimensional and dynamic motion feasible. Building upon these foundations, vision-based techniques are further advancing autonomy in the real world, while improvements in multi-robot coordination highlight the effectiveness of distributed approaches.

Despite these achievements, field experience shows that autonomy in the real-world still fails in silent and unpredictable ways that on-board diagnostics cannot always anticipate. When a robot’s localization drifts, a perception stack misclassifies an obstacle, or an actuator degrades, the system may continue operating until it suddenly stalls or behaves erratically—often with only cryptic messages or error codes as clues. In practice, a human operator is still expected to parse logs, recall previous interactions, and intervene to ensure that these sophisticated machines operate safely and responsibly amid the complexities of human environments. This topic is explored in detail in section 1.2.2.

## 1.2.2 Human-in-the-Loop (HITL) in Robotics

Although the ultimate goal of autonomous robotics often remains self-sufficiency, real-world deployments still benefit from careful human oversight. Therefore, HITL robotics has gained increasing prominence as a means of blending human exper-

tise with autonomous capabilities to address the complexities of real-world environments. Building on the need for robustness, scalability, and adaptable human oversight, recent research emphasizes the variety of methods used to integrate human input into multi-robot systems.

Control architecture remains a cornerstone of effective human-in-the-loop (HITL) robotics. Jang et al. [Jang, Jeon e Lee 2012] proposed an adaptive framework for heterogeneous multi-robot teams that modulates the degree of operator involvement according to task complexity, thereby preventing cognitive overload while preserving high mission throughput. However, their design still relies on synchronous operator interventions triggered by low-level status flags, lacks a mechanism for leveraging historical resolutions across missions, and assumes that the operator can parse raw log streams in real time. These omissions limit scalability as fleet size grows and error patterns recur.

While Fernández-Ayala [Fernandez-Ayala 2022] show that control-barrier functions with human feedbacks can keep teams of robots inside certified safe sets, their study leaves two critical questions open. First, the safety constraints are hard-coded into the barrier formulation: when the environment changes (e.g., a new obstacle appears or a temporary zone becomes off-limits) a human must manually retune the barrier parameters before the guarantee holds. Second, the human feedback loop is treated as a binary override; the operator can stop or resume motion, but the framework offers no mechanism for automatically mining past incidents to recommend the most effective corrective action. Consequently, the architecture scales poorly in settings where novel hazards emerge frequently and operators must juggle multiple robots simultaneously.

Deng et al.'s [Deng, Zhao e Pham 2024] dual-loop design successfully tightens the perception-learning cycle by injecting just-in-time human labels, however, its scope remains tightly coupled to single-robot, vision-centric scenarios and assumes that an expert can intervene whenever the active-learning loop requests feedback. The architecture does not address how previously validated fixes are stored for future reuse, nor does it include mechanisms for fusing heterogeneous data streams (e.g., navigation logs, hardware faults) and other data that typically accompany real world robot deployments. Consequently, the operator is still required to supply frequent, low-level annotations rather than receiving high-level, context-aware action plans—leaving a scalability gap when fleet size and task diversity grow beyond a handful of robots.

Mahani and Wang's runtime-verification framework [Mahani e Wang 2018] explicitly incorporates operator trust signals when reallocating tasks among robots, yet it stops short of addressing two intertwined scalability hurdles. First, the approach still presumes that operators can manually parse large volumes of low-level verification data to formulate those trust assessments—an assumption that becomes brittle as team sizes and log traffic grow. Second, the framework does not provide mechanisms for automatically

recalling and re-applying prior successful recovery strategies, forcing humans to reason from scratch each time a familiar fault recurs. These omissions leave a gap between trust-aware task allocation and the deeper need for automated context summarization and memory-driven decision support that would truly relieve operator burden in large, heterogeneous fleets.

New paradigms in distributed learning also hold promise for HITL scalability. Moradi et al. [Moradi e Guastella 2023] demonstrate that federated learning can endow a fleet of robots with a shared, continually updated navigation policy while still letting human operators supply corrective demonstrations on-the-fly. Yet three important gaps remain. First, the method stops at low-level policy refinement; it offers no mechanism for transforming the learned experience into an interpretable knowledge base that an operator—or another learning algorithm—can query when similar faults recur. Second, because each robot exchanges only gradient information, the framework cannot give the supervisor a transparent rationale for why a specific maneuver is safe or unsafe, limiting its usefulness in high-stakes environments that demand auditability and post-hoc explanation. Finally, the approach assumes that periodic connectivity is sufficient for federated updates, overlooking real-world scenarios in which robots must function offline for extended periods and still benefit from long-term contextual memory. Addressing these shortcomings requires a supervisory layer that can (i) store and retrieve past resolutions as explicit memories, (ii) reason over heterogeneous sensor logs to justify its recommendations, and (iii) operate robustly even when synchronous data exchange is impossible.

In manufacturing, Li et al. [Li, Zheng e Lee 2022] show that an AR-assisted digital-twin interface lets an operator monitor multiple production-line robots and inject corrective actions in real time. Yet the study presumes an up-to-date, high-fidelity plant model and well-defined fault taxonomies; unexpected sensor failures or previously unseen process drifts still force the operator to manually diagnose root causes via external tools. Moreover, because the AR overlay presents raw status feeds rather than summarised historical patterns, the cognitive burden rises sharply as task counts or robot variants grow, limiting the framework’s scalability beyond the narrowly scripted scenarios evaluated.

Zhang [Zhang 2023] employs spatiotemporal logic (STL) to let supervisors specify precise timing and spatial constraints for multi-robot missions, ensuring that fleets meet deadlines without collisions. While powerful, the approach depends on handcrafted STL formulas and exhaustive model checking; defining, tuning, and maintaining these specifications becomes increasingly onerous as operating domains diversify and the number of robots scales. The method also treats every violation uniformly, offering no mechanism to rank faults by recurrence or severity, so operators still sift through extensive trace logs before acting.

Hunt et al. [Hunt, Godfrey e Soorati 2024] demonstrate how large-language-model chat interfaces can streamline operator–robot interaction, enabling natural-language queries about task status or simple corrective commands. However, their conversational layer remains reactive: it translates user utterances into API calls without deeper reasoning over system diagnostics or leveraging prior intervention history. The model neither aggregates multimodal evidence (e.g., log clusters plus video cues) nor prioritises repeated failure modes, which means supervisors must switch to specialised dashboards when problems exceed scripted intents—slowing response time as fleet size and error diversity increase.

Aiming to improve the gaps of current HITL systems while maintaining reliability and human control motivates the present work’s focus on leveraging large language models (LLMs) as a vehicle for both human-robot communication and high-level oversight in dynamic, real-world scenarios.

### 1.2.3 Large Language Models

Recent years have witnessed a remarkable surge in Large Language Models (LLMs), dramatically reshaping the landscape of natural language processing. By leveraging vast quantities of training data and scaling model architectures, LLMs now underpin a wide range of tasks, including text generation, summarization, translation, and question answering [Aljohani e Herlofson 2025, Zhu et al. 2025]. Their ability to produce human-like responses has led to successful deployments in areas as diverse as healthcare, legal research, and education [Han, Shi e Tsui 2025]. These developments echo the broader ambition, outlined in earlier chapters, to harness advanced AI for systems that can adapt and collaborate seamlessly with human operators.

The technical evolution of LLMs is closely tied to the introduction of attention-based architectures, most notably the *Transformer*, which enabled models to capture contextual relationships over extended text sequences [Vaswani et al. 2017]. Early language models that relied on recurrent structures faced difficulties in managing long-range dependencies. In contrast, Transformer-based models addressed these challenges with mechanisms such as multi-head self-attention, resulting in significant improvements in performance and scalability. As a result, model families like the GPT series [OpenAI 2022] have steadily increased in size and sophistication, culminating in multi-modal systems that integrate text, images, and other data modalities [Spina et al. 2025, OpenAI 2023].

Beyond general-purpose platforms, domain-specific LLMs illustrate the adaptability of these architectures. Models like Qwen-2.5 demonstrate the power of specialized fine-tuning, achieving state-of-the-art results on medical knowledge benchmarks [Zhu et al. 2025]. A similar trend is evident in other fields where large datasets

and expert annotations can be leveraged to produce domain-tailored tools. These findings underscore the importance of approaches—both in model size and domain adaptation—when addressing the needs of complex, real-world applications. Moreover, training and deploying large-scale models demand significant computational resources, spurring research into methods such as knowledge distillation, model pruning, and hardware optimizations for reduced energy consumption [Liu, Liao e Xu 2025, Daoud e Ben-Hur 2025].

A notable current advancement is the increase in multi-modal capabilities, where LLMs can process text alongside images, audio, or sensor data [Duan, McNeese e Li 2025]. This integrated approach is particularly relevant to robotics, where vision, language, and sensory inputs converge. Studies have demonstrated how LLMs can enhance dexterous robotic manipulation by translating real-time human guidance into improved policy updates [Luo et al. 2024, Wu et al. 2024], a concept aligned with human-in-the-loop paradigms that combine operator feedback with machine intelligence.

In robotics-oriented contexts, these multi-modal LLMs facilitate natural language instructions and interactive dialogues [Royce et al. 2024]. This synergy is evident in frameworks where human operators express goals through text, and the LLM parses, refines, or translates these goals into executable robotic commands [Fan e Zheng 2024, Liu et al. 2022]. By acting as interpreters between human intent and machine execution, LLMs enhance human oversight while expanding the scope of robotic autonomy.

In light of these findings, LLMs stand at the forefront of enabling more scalable and context-aware systems. Their propensity for flexible, natural language interaction holds promise for bridging communication gaps in multi-robot deployments, aligning well with the overarching objective of this research: to promote frameworks where humans guide robots through intuitive language-based control and oversight. The coming chapters delve deeper into how LLMs can be leveraged to enhance the adaptability, safety, and scalability of HITL robotics, thus paving the way for real-world robotic ecosystems capable of handling larger operational demands with minimal strain on human operators.

### **Agents using Large Language Models**

Agentic LLMs represent a pivotal shift in natural language processing, moving from static text generation toward dynamic reasoning and collaborative capabilities. This evolution stems from the demand for more interpretable, context-aware AI systems that can operate autonomously while still aligning with human intentions. Early LLMs primarily served as advanced language predictors, but Agentic strategies increasingly demonstrate decision-making faculties, the capacity to plan multi-step actions, and the ability to articulate rationales behind specific recommendations. These developments

have gained momentum alongside broader challenges in AI, such as trust, explainability, and real-time adaptability, prompting researchers to explore agentic LLMs across diverse application domains.

Recent studies highlight the need for transparency in AI-driven recommendations, particularly in collaborative human-AI scenarios. One investigation describes how providing clear explanations of an LLM's suggestions can strengthen user trust and enhance joint decision-making [Abu-Rasheed, Weber e Fathi 2025]. This emphasis on explainability aligns with the broader requirement to ensure that AI outputs remain interpretable, especially when humans depend on machine-generated guidance for critical or complex tasks.

In parallel, advanced prompting strategies have emerged as a key driver of progress for LLM agents, enabling them to adapt to low-resource scenarios and novel contexts. A framework called INSNER demonstrates how a generative, instruction-based prompting method can address few-shot named entity recognition (NER) more effectively than traditional approaches [Zhao et al. 2025]. This technique underscores the importance of careful prompt design and instruction tuning in maximizing the utility of LLMs, especially in applications where annotated data are scarce.

Researchers have also explored how LLMs can operate in dynamic settings, such as urban sensing, where real-time data integration is crucial. One study integrates LLMs with live traffic and resource data to produce adaptive sensing solutions, promoting efficient decision-making in congested urban environments [Hou et al. 2025]. This same capacity for rapid context assimilation and synthesis could prove vital in robotics, where on-the-fly adjustments are often necessary to cope with unpredictable conditions.

Shifting to human-machine teamwork, another line of research demonstrates how agentic LLMs can manage problem complexity in environments as varied as space missions and aviation [Alexandrov 2025]. These systems dynamically allocate tasks between human operators and AI, assisting in the management of uncertainties in high-risk scenarios. The potential for such collaboration closely aligns with the thesis objective of using LLMs to enable flexible, large-scale robotic ecosystems, where robots autonomously handle routine decisions while human supervisors focus on exceptional or safety-critical situations.

LLMs have also been applied in security and cryptanalysis. Researchers outline how transformer-based models can decode intentions and analyze encrypted text corpora, highlighting novel intersections between AI, intelligence operations, and cybersecurity [Agarwal, Pandey e Mathur 2025]. Such work underscores the dual role of LLMs as both powerful enablers of innovative applications and potential vectors for emerging risks.

For complex social systems, a separate study proposes LLM-AIDSim to model

influence diffusion across social networks [Zhang et al. 2025], highlighting the capacity of LLMs to interpret and forecast group behaviors. Although primarily sociological, these findings reveal how agentic language models can interpret networked states—an ability that translates effectively to multi-robot coordination, where distributed information must be processed collectively.

An additional thread of agentic LLM research focuses on incremental dialogue management for human-robot interaction, exploring how real-time language processing can support fluid, task-oriented exchanges [Kennington, Lison e Schlangen 2025]. Coupled with retrieval-augmented mechanisms that access external databases for specialized knowledge, LLMs can provide context-rich answers specific to an enterprise or industry domain [Jiang, Zhang e Bai 2025]. Together, these advances illuminate a trajectory in which LLMs evolve from mere language models into holistic reasoning entities capable of interpreting, planning, and explaining.

### Single Agents vs Multi-Agents

With the explosion of both research papers and industrial applications of Agentic LLMs, a new debate has emerged regarding whether single or multi-agent systems are better suited for solving complex tasks. In general, single-agent architectures excel when addressing well-defined problems and when user feedback is not strictly necessary, whereas multi-agent architectures are more suitable for solving problems that require collaboration and multiple iterations.

In summary, a single agent can independently achieve its goals and manage multiple tasks within a single domain while performing reasoning, planning, and executing tools autonomously. Given an input prompt, the agent utilizes its available tools and foundational knowledge to execute a task. Common applications of single-agent systems include scientific discovery [Lu et al. 2024, Ghafarollahi e Buehler 2024, Kang e Xiong 2024, Xin et al. 2024], web scenarios [Nakano et al. 2021, Deng et al. 2024, Furuta et al. 2024, Zhou et al. 2024], gaming environments [Yuan et al. 2023, Nottingham et al. 2023], and health-care [Zhang et al. 2023, Abbasian et al. 2023].

Multi-Agent architectures involve two or more agents interacting with one another. Originally inspired by Minsky’s “Society of Mind” [Minsky 1988], which introduces a novel theory of intelligence based on interactions among specialized smaller agents, multi-agent frameworks require careful interoperability between agents, particularly in terms of communication and information sharing. They excel in collaborative tasks that span multiple domains, where each agent is specialized. These agents may share the same LLM backbone or utilize different ones; likewise, they may share tools or employ distinct ones, with each typically assuming a unique system prompt. Multi-agent sys-

tems are widely explored in domains such as scientific discovery [Schmidgall et al. 2025, Baek et al. 2024, Ghafarollahi e Buehler 2024, Swanson et al. 2024, Xiao et al. 2024], software development [Qian et al. 2024, White 2024], and healthcare [Tang et al. 2024, Kim et al. 2024]. While multi-agent systems are powerful for tackling difficult problems in complex environments, communication and interaction among agents present additional challenges compared to single-agent systems.

### **Long-Term Memory in Large Language Models**

Recent studies emphasize that, although large language models (LLMs) demonstrate impressive capabilities in language understanding and generation, they often struggle to retain extensive historical information or process lengthy input contexts. This limitation has prompted researchers to explore new frameworks and mechanisms to equip LLMs with long-term memory.

Wang et al. introduced the Self-Controlled Memory (SCM) [Wang et al. 2023] framework to assist a large language model (LLM) in recalling relevant information over time, thereby reducing the risk of forgetting earlier parts of a conversation or document. In parallel, Wang proposed Language Models Augmented with Long-Term Memory (Long-Mem) [Wang et al. 2023] to address the challenge of fixed-sized inputs by preserving extensive historical contexts for later reference. Focusing on dialogue scenarios, Wang et al. [Wang et al. 2023] argued that recursively summarizing interactions enables an LLM to maintain a coherent understanding of previous exchanges without exceeding context length constraints.

Other researchers have drawn inspiration from human memory models. Liu et al. [Liu et al. 2023] introduced Think-in-Memory (TiM), which stores past "thoughts" and reasoning steps in a dedicated memory space, allowing the LLM to reuse existing insights rather than repeatedly recomputing them. Similarly, Zhong et al. [Zhong et al. 2024] proposed MemoryBank, a mechanism designed to expand an LLM's long-term storage capacity. Deployed in a chatbot setting, MemoryBank demonstrates how specialized memory systems can significantly enhance dialogue consistency and improve the user experience. Meanwhile, Li et al. [Li et al. 2024] explored personalized agents that leverage long-term user information, highlighting the advantages of memory-augmented designs in maintaining consistent persona-based conversations across multiple sessions.

Efforts have also extended beyond text-only tasks. He et al. [He et al. 2024] introduced MA-LMM, a Memory-Augmented Large Multimodal Model capable of referencing historical video frames, thereby enabling continuous understanding of visual scenes over extended durations. This multimodal memory approach expands the applicability of LLMs to domains where video or sensor data are integral.

Evaluating these memory-centric methods has led to the development of new benchmarks and protocols. Maharana et al. [Maharana et al. 2024] proposed an evaluation for very long-term conversational memory in LLM agents, utilizing a pipeline that integrates both machine and human inputs to generate high-quality, context-rich dialogues. Lastly, Wu et al. [Wu et al. 2024] presented LongMemEval, a benchmark aimed at measuring an LLM’s capacity for multi-session reasoning, temporal logic, knowledge updates, and abstention in chat-based interactions.

Collectively, these lines of research reflect a growing emphasis on robust, long-term memory capabilities in LLMs—an evolution that is critical for real-world applications, from conversational agents to robotics. By preserving and dynamically retrieving historical information, future LLM-based systems may exhibit a deeper and more context-aware intelligence, potentially reducing the need for human intervention in tasks that require sustained reasoning over time.

### Reasoning in Agentic LLMs

A growing body of literature [Naidu e El-Gayar 2025] underscores the shift toward agentic LLMs that iteratively evaluate, explain, and refine their own thought processes. These advances are often framed under specific paradigms, such as chain-of-thought prompting [Wei et al. 2022] or self-discovery [Zhou et al. 2024], each highlighting distinct strategies for orchestrating more transparent and robust forms of machine reasoning.

Chain-of-thought prompting refers to an approach where an LLM explicitly outlines its intermediate thinking steps as it processes a query or instruction [Wei et al. 2022]. By revealing a trail of logical inferences, the model becomes more interpretable to human operators, allowing them to validate or correct assumptions midstream. This method not only improves the credibility of LLM-driven decision-making in high-stakes environments—such as medical diagnosis or complex planning—but also facilitates more dynamic forms of collaboration. When an agentic system shares its chain of thought, human collaborators can intervene, offer corrections, or request additional details without losing continuity. Such an interactive loop is especially relevant for multi-robot supervision, where overlapping tasks and unpredictable external conditions often require iterative updates to action plans.

Self-discovery frameworks expand on these ideas by prompting the LLM to critique, refine, and even contradict its own earlier statements as new information emerges [Zhou et al. 2024]. Instead of relying solely on external input to initiate re-evaluations, the model periodically reviews its internal reasoning to identify possible inconsistencies or overlooked nuances. Research in this area suggests that self-discovery can minimize error propagation and reduce cognitive load on human supervisors, as the

LLM autonomously flags potential missteps before seeking operator feedback. In practical terms, self-discovery enhances chain-of-thought by introducing an introspective mechanism: a model that not only reveals its reasoning steps but also continuously audits and corrects them as an integral part of the decision-making process.

Early demonstrations of these paradigms show promise in various contexts. In few-shot and low-resource scenarios, chain-of-thought methods can drastically improve the extraction of relevant insights, particularly when domain-specific prompts are carefully engineered. Self-discovery, meanwhile, has proven advantageous for incremental dialogue management, enabling the system to recalibrate its strategy in real-time if unexpected constraints or user requests arise. These reasoning approaches align closely with the broader narrative that agentic LLMs should offer both transparency and adaptability—two qualities critical to large-scale deployments where humans guide fleets of autonomous or semi-autonomous robots.

As noted in recent studies on agentic AI, including those focused on question answering and incremental dialogue modeling, structured reasoning methods can help mitigate the risks of “hallucinations” or ungrounded outputs. By instructing a model to articulate how it arrives at a conclusion (chain-of-thought) or to iteratively validate that conclusion (self-discovery), developers can reduce the likelihood of spurious responses and build safer, more reliable systems. Within the scope of this thesis, these reasoning paradigms form the conceptual backbone for designing agents that not only interpret operator directives but also confirm whether their own inferences are valid under the environmental and task-specific constraints often encountered in real-world robotics.

---

## Problem, Methodology and Proposal

---

Taken as a whole, the literature illustrates the wide-ranging strategies employed in human-in-the-loop autonomous robotics, from control architectures and safety mechanisms to learning algorithms, trust frameworks, and natural language interfaces. Yet each strand tackles only part of the scalability puzzle; none offers a unified mechanism that can simultaneously interpret multimodal data in real time, reuse hard-won experience across recurrent faults, and converse with operators in transparent everyday language.

Building on the open questions discussed in Section 1.1, this thesis pursues several pathways that together form the backbone of the proposed solution. First, to lower the operators' cognitive load, we adopt a multimodal natural-language interface that ingest logs, video and telemetry and convey only important information back to the operator. Second, we tackle complex reasoning demands through an agentic architecture that delegates subtasks to specialised sub-agents and through a self-discovery workflow that breaks problems into verifiable steps. Third, we embed long-term memory: fine-grained memory tools let each robot recall past incidents, whereas a fleet-level knowledge base allows cross-site sharing of successful fixes. Fourth, we address scale by relying on modular sub-agent composition—so new capabilities can be plugged in without retraining—and on fully automated data processing, which removes latency between a help call and a recommended action. Finally, we preserve transparency via meta-prompted reasoning that surfaces every decision in a step-by-step narrative and by delivering structured outputs whose rationale can be audited by human supervisors. Taken together, these pathways aim to ease operator workload, improve diagnostic accuracy, and ensure that growing robot fleets remain both reliable and explainable as they confront ever more diverse real-world environments.

By anchoring a language-based interface to real-world constraints and robust safety principles, this work aspires to improve the scalability of human oversight while embracing the adaptability and intelligence inherent in agentic AI.

## 2.1 Methodology for the Proposed Architecture

Creating an enhancement—or even a substitute—for a human-in-the-loop (HITL) scenario requires balancing several constraints: the capacity to process large, heterogeneous data sources, the need for rich contextual reasoning, the ability to store and recall long-term information, and, most importantly, the necessity for an easy-to-use, transparent, and intuitive user interface. Modern large language models (LLMs) provide a sophisticated and intuitive natural language interface, positioning them as compelling candidates for supervising multi-agent robotic systems.

Moreover, a natural language interface is key for parsing requests from human operators, generating actionable and transparent plans for interventions, and providing a platform for user questioning and further troubleshooting. Prior works demonstrate that advanced prompting and instruction-based techniques can enhance LLMs in low-data or rapidly changing domains [Liu et al. 2023], suggesting that a linguistic interface can effectively capture operator directives and domain knowledge. In practice, this approach aims to minimize the operator’s workload by simplifying the specification of tasks, rather than resorting to complex user interfaces or extensive retraining whenever the environment changes.

Nonetheless, the use of a simple LLM model is insufficient when prompted with extensive logs, sensor data, and mission updates with the objective of diagnosing and planning a solution in a consistent and efficient way. A single monolithic agent would quickly become overwhelmed by the volume and variety of data, struggling both to prioritize relevant details and to respond efficiently. Recognizing these challenges, the proposed solution adopts an Agentic AI approach that breaks the supervisory system into specialized subagents under a hierarchical structure. Each subagent is equipped with domain-specific functionality—some may manage tool-based data processing (e.g., sensor fusion or path planning), others may maintain long-term memory for extended missions, and still others may focus on adaptive reasoning in evolving contexts.

This modular design offers several advantages. First, it ensures flexibility by allowing new subagents to be introduced or modified as the application scope evolves. Second, the architecture naturally manages complexity: adding more intricate tasks does not necessarily increase the reasoning burden on a single LLM agent, as each subagent specializes in overseeing a specific segment of the workflow. Third, the use of subagents simplifies prompting and tool handling, since each specialized component can focus on its own tools, constraints, and objectives. Finally, modularization reduces the demand for a single large context in the LLM, mitigating known limitations related to context window size and memory retention.

Under this framework, the supervising LLM orchestrates the roles and respon-

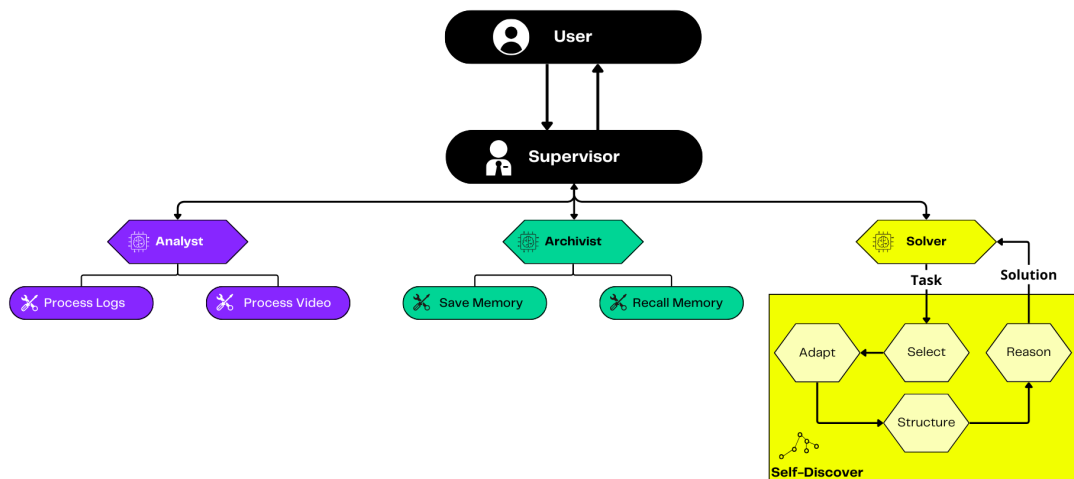
sibilities of the subagents. It interprets operator inputs, breaks down tasks into sub-problems, and issues targeted instructions to the relevant module—whether for data analysis or for a subagent focused on memory retrieval. This approach aligns with recent research highlighting the benefits of hierarchical agent architectures in managing complex decision-making and information processing pipelines. By maintaining an overview of the entire operation, the LLM ensures that the collective efforts of each subagent converge toward consistent goals, enabling reliable and adaptive supervision of autonomous robotic operations.

Moving forward, the following sections will detail the base subagents of this framework, their design principles, and their integration with a supervisory agent. Furthermore, this chapter will conclude by demonstrating how the proposed framework can support human operators more effectively, particularly in the high-stakes or rapidly changing conditions common to multi-robot HITL missions.

## 2.2 The proposed Architecture

The proposed architecture in Figure 2.1 presents a foundational framework developed as a hierarchical multi-agent system, wherein specialized subagents collaborate to tackle complex and dynamic problem-solving tasks.

Figure 2.1: The proposed hierarchical multi-agent architecture of this work.



**Source:** Prepared by the Author.

Each subagent operates autonomously yet synergistically within the system, leveraging specialized tools to process data, manage long-term memory, and perform

adaptive reasoning. This modular design ensures flexibility and scalability, allowing the system to handle diverse scenarios, particularly in supporting autonomous robotic operations.

### 2.2.1 The Analyst Subagent

The Analyst functions as the data collection and processing layer of the system, extracting actionable insights from multimodal, time-sensitive data sources. It processes system logs and visual data, deconstructing diagnostic data and analyzing multiple video feeds to identify anomalies or task-related issues. These insights are crucial for the reasoning and decision-making performed by the downstream subagents.

At the core of the Analyst’s workflow is the ability to interact with domain-specific diagnostic tools to query log files—recording anomalies, error codes, and warning messages—and synthesizing a concise report of the issues uncovered. When necessary, it refines its queries to target specific intervals or event identifiers, potentially resorting to recorded video data to better understand navigational difficulties or environmental challenges. Since the Analyst’s scope is strictly limited to data analysis and reporting, it avoids any suggestions for remedial actions, adhering to the principle that diagnostics and solution planning should remain separate within an agentic framework.

The choice to isolate diagnostic functions within a dedicated subagent arises from the substantial volume and variety of sensor data typical of multi-robot systems. Attempting to parse logs and execute repairs within a single, monolithic LLM supervisor could obscure key diagnostics with decision-related processes. By introducing the Analyst subagent as a standalone module, the architecture ensures that each log, warning, or error code is systematically prioritized and contextualized. This design also supports parallelization, allowing for the use of multiple tools specialized for processing different types of data in the multimodal logs.

Once the Analyst completes the data review, they deliver a concise report to the supervisor, highlighting critical warnings, error frequencies, and any discrepancies between the logged data and captured video. The supervisor can then consult other subagents—such as the Archivist for historical context or the Solver for reasoning and action planning—to determine the appropriate resolution path. This modular approach minimizes redundant computations and clarifies responsibility boundaries: the Analyst identifies problems, while other specialized components either store pertinent memory snapshots or propose corrective measures.

Overall, the Analyst subagent embodies a crucial design philosophy of agentic architectures: delegating specialized tasks to subagents that can be independently refined, tested, or replaced without disrupting the larger system. By focusing solely on diagnostics,

the Analyst ensures that actionable information is readily available to the higher-level reasoning modules, facilitating effective human-in-the-loop supervision and enabling more agile multi-robot deployments.

### **2.2.2 The Archivist Subagent**

The Archivist subagent addresses the challenge of preserving and accessing historical knowledge within a language model-based supervisory framework. Its primary function is to establish and maintain a structured repository of contextual data, which includes previous events, implemented solutions, and the reasoning processes behind those decisions. By systematically organizing interactions, the Archivist facilitates the retrieval of both recent and historical insights to inform current tasks, thereby enhancing the system’s decision-making processes with contextually relevant prior knowledge.

A central design objective for the Archivist is the efficient storage of information in a manner that balances concision with comprehensiveness. In accordance with its operational parameters, the subagent generates condensed summaries of each event, capturing essential details such as the nature of the problem, the selected solution, and the Solver’s rationale while avoiding excessive memory consumption. This structured methodology optimizes retrieval efficiency and facilitates cross-referential analysis to ensure consistency across newly recorded data. When historical context is requested—whether to validate past successes or detect patterns in recurring challenges—the Archivist employs specialized retrieval mechanisms to generate coherent, consolidated reports, thereby preventing fragmented or redundant outputs.

The Archivist’s ability to retain prior reasoning processes and track shifts in the Supervisor’s strategic priorities enables adaptive responses to evolving objectives or constraints. For instance, when previous solutions have proven ineffective under similar conditions, the stored data highlights potential risks. Additionally, when historical actions align with current objectives, the Solver can be enhanced using historical data to prioritize proven approaches. This dynamic integration of accumulated knowledge introduces a layer of personalization, enabling the system to refine its outputs based on empirical evidence—a critical capability for real-world applications characterized by dynamic operational environments.

Within the broader agentic architecture, the Archivist operates as a dedicated module for long-term knowledge retention, alleviating the cognitive burden on the Supervisor and other subagents—such as the Solver or Analyst—by removing the need for them to maintain exhaustive contextual records. This modular design reflects a core principle of the system: decomposing complex responsibilities into specialized subagents enhances scalability and reliability. By specializing in memory management, the Archivist ensures

the persistence of critical insights, environmental nuances, and the Solver’s preferences beyond individual sessions. Consequently, the system supports sustained collaboration across multi-robot teams, extended missions, and diverse operators while preserving operational coherence.

In summary, the Archivist subagent serves as a cornerstone for the system’s long-term learning capabilities, integrating structured knowledge retention, efficient retrieval protocols, and context-aware adaptation. Its mechanisms—designed to comply with invocation constraints and post-processing verification requirements—collectively ensure operational clarity, consistency, and minimal computational overhead. By reinforcing the system’s capacity to leverage past successes and circumvent prior errors, the Archivist not only advances autonomous decision-making but also strengthens human-robot collaboration by maintaining a shared knowledge base across dynamic, mission-critical scenarios.

### 2.2.3 The Solver Subagent

The Solver is central to the proposed agentic architecture, managing advanced reasoning and decision-making. It employs a self-discovery strategy [Zhou et al. 2024] to reason about the problem details, existing diagnostics, and relevant contextual data from other subagents. By dynamically operating based on Analyst data and recalled memories, the Solver identifies and refines potential solutions, ensuring that each step aligns with the broader system objectives.

Its workflow begins with the formalization of incoming issues as discrete tasks, assembling key elements such as error descriptions, sensor outputs, and operator directives. The Solver also reflects on memories provided by the Archivist, including prior solution outcomes and known environmental quirks, to enrich its understanding of potential pathways.

Throughout these iterations, the solver revisits earlier assumptions, adjusts strategies based on updated insights, and refines reasoning in a structured manner. This iterative re-checking is crucial for addressing mission-critical tasks and adapting to high levels of environmental uncertainty or changing operational constraints. Each cycle culminates in a coherent plan that details the sub-steps or actions essential for the system to recover from an error or continue navigation.

Because of its capacity to adapt strategies and converge on a solution through multiple rounds of reflection, the Solver can effectively handle complex, multi-step tasks. This approach also accommodates novel scenarios where a static, one-shot reasoning model might fail.

### **Adapting Self-Discovery.**

As discussed in Section 1.2.3, the Self-Discover approach allows large language models (LLMs) to compose atomic reasoning modules into a task-specific reasoning structure. This approach moves beyond relying on single, predefined reasoning techniques, such as Chain of Thought (CoT), by enabling LLMs to select and combine modules like critical thinking and step-by-step analysis. This framework operates in two stages: first, the LLM uses meta-prompts to select, adapt, and implement a reasoning structure. Second, the LLM employs the self-discovered structure to solve each instance of the task by filling in key-value pairs. This method not only improves performance on challenging reasoning benchmarks but also reduces computational needs compared to other inference-heavy approaches.

In refining the original set of reasoning modules proposed by [Zhou et al. 2024] for the Solver Subagent, the primary objective was to eliminate redundancies and reduce reliance on iterative processes that necessitate multiple exchanges. The updated modules now emphasize immediate problem comprehension, critical analysis, and the generation of a fully formed solution in a single pass. Modules specifically designed for experimental testing, extended collaboration, or lengthy data gathering were removed because they relied on continuous user interaction or ongoing investigation, which contradicts the supervision paradigm. Instead, priority was given to concise problem restatement, verification, and feasibility checks, ensuring that each proposed solution remains both context-aware and actionable. Furthermore, a mandatory "Module 0" was added to compel the solver to always consider past experiences when devising solutions to current problems.

Additionally, creative thinking and risk analysis modules were retained to encourage inventive approaches while still acknowledging potential drawbacks, thus balancing innovation with practical constraints. Reflective thinking continues to be crucial for quickly identifying biases or assumptions that could undermine the solution if left unaddressed. By introducing explicit steps for feasibility and verification, the revised set of modules seeks to incorporate a layer of self-assessment that is otherwise absent in a one-shot context, ensuring that any final recommendation remains logically sound and pragmatically viable within given resource constraints. This methodical approach ensures that each module included contributes distinctly to the overall quality and reliability of the single-pass solution. The following list showcases the proposed reasoning modules:

#### **Self-Discovery Modules for the Solver Subagent:**

**Module 0:** Learning from past experiences: Provide a rationale for how similar problems in memory have been solved and what can be learned from these solutions.

**Module 1:** What are the key assumptions underlying this problem?

**Module 2:** How can I break down this problem into smaller, more manageable parts?

**Module 3:** Critical Thinking: Analyze the problem from different perspectives, question assumptions, and evaluate the evidence or information available.

**Module 4:** Creative Thinking: Generate innovative and out-of-the-box ideas to solve the problem, exploring unconventional solutions.

**Module 5:** Systems Thinking: Consider the problem as part of a larger system, identifying feedback loops and interdependencies.

**Module 6:** Risk Analysis: Evaluate potential risks, uncertainties, and tradeoffs of different solutions.

**Module 7:** Reflective Thinking: Step back for introspection, examining personal biases, assumptions, and mental models.

**Module 8:** What are the underlying causes or factors contributing to the problem?

**Module 9:** Are there any relevant data or information that provide insights into the problem?

**Module 10:** Are there any stakeholders or individuals directly affected by the problem, and what are their perspectives?

**Module 11:** How can progress or success in solving the problem be measured or evaluated?

**Module 12:** Problem Comprehension: Restate and clarify the problem in simple terms, noting any constraints or assumptions.

**Module 13:** Feasibility and Resource Check: Consider whether a proposed solution is realistically implementable given time, funding, and expertise constraints.

**Module 14:** Verification & Confidence Check: Re-evaluate the proposed solution for oversights, logical leaps, or missing evidence.

**Module 15:** Contextual Explanation: Provide a concise rationale explaining how the proposed solution addresses the problem and why it is appropriate.

## 2.2.4 The Supervisor

The Supervisor orchestrates interactions among subagents, ensuring that the entire system remains synchronized and responsive to real-time demands. When a support call or user request enters the system, the Supervisor distributes tasks among specialized subagents, each of which reports back on their progress or any obstacles encountered. By monitoring these reports, the Supervisor maintains overall coherence, preventing duplicate efforts and inconsistent decisions.

This supervisory function also streamlines operator interactions. Users can pose targeted queries, provide context-specific constraints, or explicitly request a particular type of solution. In response, the Supervisor interprets these requests, directs them to

the appropriate subagent, and aggregates relevant insights as the solution evolves. This synergy clarifies responsibilities, with the Supervisor maintaining high-level oversight while subagents manage diagnostics, memory management, or detailed reasoning steps.

### 2.2.5 Workflow Integration Example

To demonstrate how a typical support call progresses through the proposed agentic framework, consider the following eight-step sequence, which includes an optional interaction phase at the end. This workflow illustrates how each sub-agent—*Analyst*, *Archivist*, *Solver*, and *Supervisor*—collaborates to diagnose issues, leverage past experiences, generate solutions, and deliver actionable advice to the human operator:

#### 1. Support Request

A request for assistance or an alert (e.g., robot stall) enters the system. The *Supervisor* initially receives this alert and determines which subagents need to be invoked based on the nature of the issue.

#### 2. Data Analysis

The *Supervisor* forwards the problem to the *Analyst*, which queries the relevant diagnostic logs, sensor streams, and video feeds. The Analyst then compiles a concise report, noting any error codes, anomaly alerts, or suspicious sensor readings.

#### 3. Historical Context

If the situation could benefit from prior experiences, the *Supervisor* instructs the *Archivist* to retrieve pertinent records—for instance, comparable past incidents, solutions, and outcomes—and passes this context back to the Supervisor.

#### 4. Reasoning & Planning

Armed with the Analyst's report and the historical context from the Archivist, the *Supervisor* engages the *Solver*. The Solver applies its self-discovery modules—covering critical thinking, risk analysis, feasibility checks, and so forth—to propose a recommended plan or strategy.

#### 5. Solution Assembly

Once the Solver refines its single-pass plan, it returns a step-by-step outline to the *Supervisor*. This proposal may include suggested interventions (e.g., teleoperation, hardware reset) and references to previously successful solutions, as stored by the Archivist.

#### 6. Feedback & Validation

The *Supervisor* reviews the Solver's proposal for completeness and ensures it aligns with operator preferences and real-time constraints (e.g., urgent delivery timelines or safety considerations). If further diagnostics are needed, the Supervisor re-engages the Analyst or Archivist.

### 7. Operator Delivery

Upon validation, the *Supervisor* compiles a final response -including logs, pertinent historical precedents, and the proposed action plan -to present to the human operator. The operator may confirm, amend, or implement the solution as appropriate.

### 8. Execution & Logging

If the plan is accepted, the *Supervisor* orchestrates the necessary actions for the robots. Meanwhile, the *Analyst* and *Archivist* record the outcome—whether success or failure—enriching the knowledge base for future support calls.

9. **Optional Extended Interaction** If the operator wishes to delve into certain aspects more deeply, they can chat directly with the *Supervisor*. In such cases, the *Supervisor* creates additional tasks, distributing them among the subagents as needed, to fulfill specific user requests or gather further data. This iterative process of operator engagement ensures that no critical detail is overlooked and that the system adapts continuously to evolving operational demands.

Finally, this agentic AI system represents a novel approach to supporting autonomous robotic operations by combining modularity, memory-augmented reasoning, and adaptive decision-making. The integration of a self-discovery framework within the Solver enables the architecture to effectively address unforeseen challenges, while the collaborative interplay among subagents ensures both reliability and scalability. This design has significant implications for real-world applications, such as logistics and industrial automation, where adaptive and intelligent problem-solving capabilities are essential for success.

---

## Experiments and Results

---

Although the proposed agentic framework in Chapter 2 can be generalized to a range of robotic supervision scenarios, the following experiments focus on a specific use case involving a fleet of indoor delivery robots operating in a dynamic and challenging environment—a mall. The robots are equipped with sensors such as RGB-D cameras, ultrasound, and ToF (Time of Flight) sensors, leveraging an appearance-based visual SLAM system for localization and mapping, with ROS as the underlying framework. Their navigation stack employs a global planner that utilizes the SLAM map to reach designated goals and integrates local perception modules into a costmap to enable real-time dynamic obstacle detection. The specific operational context for this system is an indoor shopping mall, where the robots carry out the first mile for a delivery app. They pick up items from stores and restaurants and transport them to designated collection stands, as showcased in Figure 3.1.

Managing an entire fleet of delivery robots within this bustling indoor environment poses significant challenges for human supervisors, who must juggle multiple real-time responsibilities. They oversee robot scheduling, coordinate high-priority deliveries, and respond to system alerts, while also parsing logs whenever technical anomalies arise. In practice, the sheer variety of data to be analyzed from sensors, camera feeds, and navigation modules can overwhelm operators if no structured triage mechanism is in place. Furthermore, the unpredictable nature of shopping mall traffic—such as abrupt layout changes or large crowds—exacerbates the challenge of quickly diagnosing and resolving navigational impasses. Human supervisors often lack the time or resources to analyze intricate data streams on the fly, making it increasingly difficult to implement prompt interventions. These challenges are amplified when multiple robots are operating simultaneously in several malls, underscoring the need for intelligent, agentic support to limit cognitive overload and streamline the supervision process.

Figure 3.1: The operation studied as a use case. **Top-Left:** Collection Stand worker with an order. **Top-Right:** Robots waiting for deliveries inside the mall. **Bottom-Left:** Robot collecting orders. **Bottom-Right:** Driver after collecting orders.



**Source:** Synkar Autonomous, 2024

In summary, this indoor delivery robot scenario highlights the primary challenges that this thesis aims to address: dynamically adapting to human-dense environments, managing potential hardware and software errors, and maintaining a seamless operational flow under varying traffic conditions. While this work focuses on shopping mall delivery scenarios, the underlying challenges are common to a variety of indoor robotics applications, including hospital logistics, warehouse automation, and hospitality services. By successfully deploying our agentic framework in this context, we illustrate a model of adaptability and scalability that can be generalized to other similarly dynamic indoor environments.

The following chapter details how this methodology is experimentally evaluated

in realistic testing environments, emphasizing both the quantitative and qualitative outcomes that arise from the agent’s ability to coordinate multiple robots, interpret sensor-driven logs, and enhance human oversight in high-demand commercial settings.

## 3.1 Dataset Collection

In order to evaluate the proposed agentic framework, data were gathered from two different shopping malls, each served by a fleet of two robots. Whenever a robot flagged a situation requiring human oversight, the system automatically captured a one-minute period of relevant sensor and diagnostic logs leading up to the support call. This section provides an overview of how these recordings were made, the types of information stored, and the accompanying human-collected data that contextualize and enrich our experimental dataset.

The foundation of this dataset lies in the *rosbag* format, which stores both sensor streams (such as RGB-D camera feeds, ultrasound/ToF readings, and odometry signals) and diagnostic-level data (system warnings, error codes, and informational messages). When a support call is initiated—whether due to hardware issues, unforeseen or unsafe situations, or navigation impasses—the system captures a snapshot that encapsulates up to 60 seconds of logs prior to the event. This interval often includes the moment when anomalies or drops in performance first manifest.

In parallel, we recorded the solutions or responses provided by the human operator. While the operator could manually guide the robot (teleoperation), reset hardware, or implement a particular fix, they were also asked to log brief comments on why they selected each method. These remarks often reflect real-world trade-offs, such as time constraints, the operator’s familiarity with specific errors, or the perceived severity of the problem. By linking each error call to a corresponding human-executed solution, our dataset compiles both the empirical robotic data and the reasoning behind the operator’s interventions.

Overall, the data collection setup ensures that each support request is documented in a coherent manner: an entire minute of sensor and diagnostic logs matched with human solution choices and remarks. This structure allows for retrospective comparisons of how well the agentic framework’s subagents replicate or enhance the operator’s judgment, and whether memory retrieval or advanced reasoning could have improved outcomes in near-failure scenarios. The following sections detail how this dataset underpins our experimental methods for assessing diagnostic robustness, memory-based adaptability, and task completion reliability.

The following sections present the experimental validation of the agentic architecture introduced in Chapter 2. This evaluation aims to verify the system’s effectiveness

in managing multi-robot supervision, particularly in the high-traffic indoor delivery scenario described earlier. The experiments focus on several key metrics: solution accuracy, user ratings, time to resolution, and textual similarity measures, such as Token Overlap and Semantic Alignment.

## 3.2 Experimental Goals and Reasoning

Having established the data collection process, we now turn to the broader motivations behind the experiments. Building on the HITL frameworks discussed in Chapter 3, the experiments in this chapter validate how long-term memory and self-reflective reasoning can enhance operator oversight and reduce cognitive burden. By documenting measurable gains in both solution accuracy and operator satisfaction, we provide empirical evidence that the design principles introduced earlier—namely, modular subagents for diagnostics, memory, and planning—offer tangible benefits in real-world multi-robot deployments.

A key dimension of these experiments is measuring the framework’s performance against that of human operators. Although skilled operators possess domain knowledge and practical insight, their efficiency can degrade under information overload or time pressure. By comparing the multi-agent system’s responses with those provided by experienced personnel, we assess the efficacy, consistency, speed, and generalization potential of agentic LLM-driven solutions. Specific questions include whether the agent can handle recurrent failure modes more effectively than a human once sufficient historical information is stored, and whether self-reflection can compensate for incomplete or ambiguous data.

Additionally, the experiments emphasize the evaluation of the system’s memory functionality. Traditional LLMs often struggle with prompts that exceed a certain size or complexity, rendering them susceptible to repeated mistakes in sequential tasks. We aim to verify whether the Archivist, in conjunction with the Solver’s self-discovery approach, can overcome these shortcomings by retrieving and applying relevant knowledge from past interventions.

In short, the experimental program aims to validate the ability of the multi-agent framework to combine human-like insight with robust, data-driven processes. By examining both generalization capacity and memory-based improvements in problem-solving, these experiments provide crucial evidence of how agentic AI can reduce the supervisory burden while maintaining or even exceeding the performance and consistency offered by human operators.

### 3.3 Example of the proposed solution in action.

This section demonstrates how the multi-agent system responds to a single support call in a shopping mall environment. The delivery robot encounters an overly excited child who closely interacts with its sensors, leading to repeated path planning security failures, as shown in Figure 3.2. We showcase each sub-agent’s inputs and outputs in Appendix A, illustrating how diagnostics, memory retrieval, self-discovery reasoning, and supervisor coordination work together to generate a coherent final recommendation.

Figure 3.2: A side by side of the front camera image and top-view rendering of the robot sensors in the example.



Source: Synkar Autonomous, 2024

The following sequence outlines how the agentic system addressed the request for support:

1. **Supervisor Initiates Diagnostics:** The Supervisor instructs the Analyst to investigate logs and video feeds to determine why the robot halted its mission.
2. **Analyst Generates Report:** The Analyst processes the logs and video data, identifying obstacle-detection alerts and local planning errors linked to a child obstructing the robot’s sensors.
3. **Supervisor Requests Context from Archivist:** To ground its response in prior experiences, the Supervisor asks the Archivist for memories involving human or child obstacles causing navigation failures.
4. **Archivist Returns Past Solutions:** The Archivist retrieves summaries of previous instances, highlighting the effectiveness of audio warnings, teleoperation, and mission retries in similar scenarios.
5. **Supervisor Tasks Solver:** The Supervisor hands over the diagnostic findings and contextual memories, requesting a recommended plan to handle the child’s interference.

6. **Solver Self-Reflects:** The Solver ponder on the diagnostics, relevant historical data, and allowed actions, conducting an internal self-discovery loop to propose a stepwise response consistent with past successes.
7. **Solver Delivers Final Proposal:** The Solver suggests using an audio alert to prompt clearance, retrying autonomous navigation, teleoperating if the child remains in the way, optionally playing a second audio announcement during manual maneuvering, and performing a final retry to restore full autonomy.
8. **Supervisor Summarizes for the User:** The Supervisor compiles the Solver’s plan, presenting a concise explanation and justification to the human operator, ensuring clarity and requesting confirmation once the robot resumes operation.

In this example, the Supervisor initiates diagnostic inquiries, the Analyst compiles log-based and video-based insights, the Archivist retrieves similar historical incidents, and the Solver synthesizes a recommended response under the self-discovery framework. Each agent’s output remains specialized, yet collectively contributes to a cohesive solution. By integrating memory-based references and iterative reasoning, the system demonstrates how agentic architecture can intelligently handle dynamic, human-centered obstacles while allowing operators the final discretion to implement or revise the plan.

## 3.4 Experiments

The next sections detail the experimental framework used to evaluate our agentic solution in real-world support call scenarios for an indoor delivery robot. Building on prior research that highlights memory-augmented approaches in LLMs [Liu et al. 2023, Wang et al. 2023], we aim to verify whether the inclusion of the Archivist subagent measurably improves diagnostic accuracy and supervisory efficiency. First, we outline the designed test scenarios and the variations introduced to probe the consistency and accuracy of the proposed architecture. We then present the numerical results of these trials, analyzing the framework’s performance with and without memory augmentation. Finally, we discuss key takeaways, limitations, and potential areas for improving the system’s applicability in large-scale, long-term deployments [He et al. 2024].

### 3.4.1 Scenario Design and Parameter Variations

We devised multiple test scenarios that encompass various common failure modes reported in the context of first-mile delivery within shopping malls. Each scenario utilizes real logs and sensor data from the robot’s operation, prompting the multi-agent architecture to generate recommended actions. To better understand the influence of

memory augmentation, we tested two configurations: one in which the Archivist subagent was enabled (allowing for the retrieval of past solutions) and another in which no memory references were used.

Although short-term interventions have been explored in earlier studies [Zhong et al. 2024], our approach differs by systematically comparing memory-augmented runs to those without recollections of earlier episodes. The parameter variations help ensure that the results generalize to a variety of real-world conditions.

### Sample Scenarios:

- *Case 1: Localization Failures.* The robot’s localization stack is suffering from drift, leading to inconsistent costmaps and aborted path planning.
- *Case 2: Hardware Issues.* An Instance where multiple sensor communication abruptly faults.
- *Case 3: Human Obstruction.* A group of people persistently block the robot’s path, forcing navigation standstills and repeated re-planning attempts, without success.

The experiments included several variations to capture different operational complexities and operator perspectives. First, each scenario was repeated 30 times in both with-memory and no-memory configurations, yielding a total of 60 trials per scenario. Second, certain scenarios featured cluttered logs with multiple warnings or concurrent errors, while others were simpler, allowing us to observe how effectively the framework handled varying levels of data noise. Finally, participants in the final human-labeled evaluation came from diverse operator backgrounds, ensuring that the system’s performance was assessed by a range of skill levels, thus reflecting real-world conditions more accurately.

### 3.4.2 Results Analysis

The evaluation emphasizes the accuracy of solutions, solution time, and text-based similarity metrics, specifically focusing on Token Overlap and Semantic Alignment. Additionally, we draw qualitative insights from experts who reviewed each final proposal and provided ratings. Table 3.1 summarizes the performance across the three scenarios, contrasting runs in which the Archivist was enabled with those using no-memory baselines. Overall, the agentic framework solution is faster than experienced human supervisors and demonstrates similar accuracy. Moreover, the memory augmentation consistently shows improvements in both the correctness of solutions (i.e., how well the recommended actions matched validated outcomes) and user satisfaction, as reflected in the ratings, echoing insights from prior works in retrieval-augmented generation [Woo et al. 2024].

Table 3.1: Comparison of Accuracy, User Rating and Time across Different Cases.

	<b>Accuracy</b>	<b>Rating</b>	<b>Avg Time</b>	<b>Human Time</b>
<b>Case1</b>	93.33%	4.86	66.33	142.82
<b>Case1_noMem</b>	86.67%	4.73	63.71	-
<b>Case2</b>	96.67%	4.79	70.58	254.74
<b>Case2_noMem</b>	80.00%	4.75	64.87	-
<b>Case3</b>	90.00%	4.85	71.24	110.65
<b>Case3_noMem</b>	83.33%	4.68	61.32	-

To gauge how closely the agent’s final instructions align with the reference guidelines and the system’s own historical solutions, the study assessed the overlap in specific domain terminology (Token Overlap) as well as conceptual similarity (Semantic Alignment) in Table 3.2. As observed in prior NLP literature [Jiang, Zhang e Bai 2025, Minato et al. 2004], memory-enabled sessions often exhibit stronger coherence and domain consistency in repeated or intricate tasks. Feedback from human supervisors further underscored the clarity gained when referencing stored context; they reported fewer redundant steps and more cohesive justifications for each action, demonstrating how the stored context can streamline the solution composition process and lower cognitive load.

Table 3.2: Comparison of Similarity Metrics across different cases.

	<b>Token Similarity</b>	<b>Semantic Similarity</b>
<b>Case1</b>	$0.7472 \pm 0.0835$	$0.9471 \pm 0.0144$
<b>Case1_noMem</b>	$0.7394 \pm 0.0737$	$0.9400 \pm 0.0201$
<b>Case2</b>	$0.7243 \pm 0.0729$	$0.9356 \pm 0.0185$
<b>Case2_noMem</b>	$0.7105 \pm 0.0787$	$0.9281 \pm 0.0189$
<b>Case3</b>	$0.7428 \pm 0.0765$	$0.9364 \pm 0.0192$
<b>Case3_noMem</b>	$0.7320 \pm 0.0731$	$0.9271 \pm 0.0258$

### 3.4.3 Module Selection Trends

In addition to solution accuracy and similarity scores, the system’s self-reflection step records which modules it deems most relevant for each scenario. Recent analysis frameworks [Wu et al. 2024] similarly track the inner components of LLM-driven solvers, emphasizing studies on transparency in complex agentic systems. Figures 3.3, 3.5, and 3.7

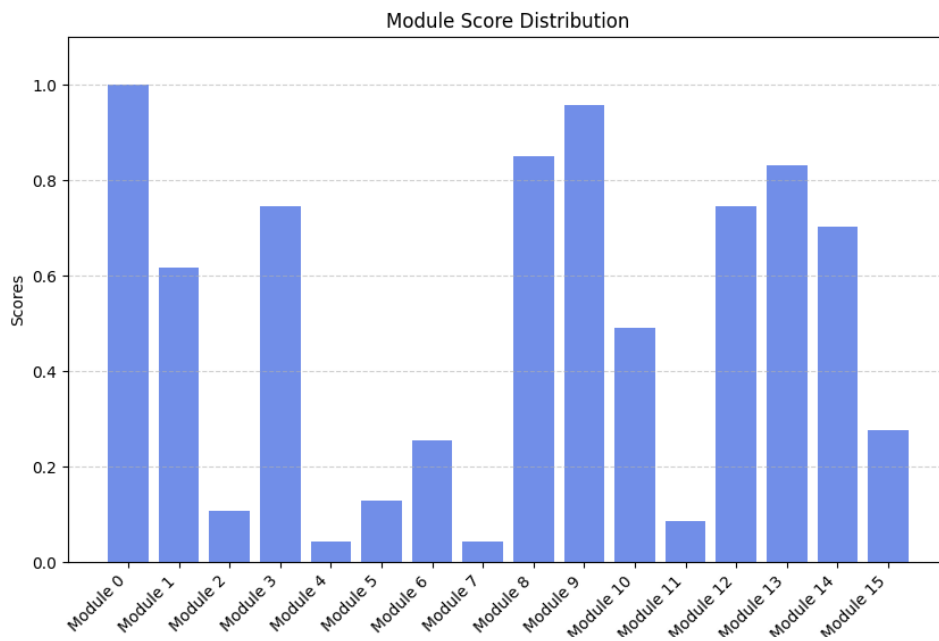
illustrate the frequency of module usage under memory-enabled conditions for each case, while Figures 3.4, 3.6, and 3.8 depict the corresponding no-memory baselines.

A clear observation emerges when comparing the memory-enabled and no-memory histograms. In scenarios with memory access, the Solver performs additional risk or feasibility checks (*Modules 6 and 13*) more frequently. This pattern suggests that the system, when equipped with past information, is more inclined to validate its proposed plan against earlier outcomes. In contrast, sessions lacking Archivist support place a heavier emphasis on general-purpose or exploratory methods, such as *Modules 1 and 2*. This indicates that the Solver compensates for the absence of historical insights by examining problems in a broader and less specialized manner.

Additionally, different failure modes drive unique module priorities. In localization issues (Case 1), the Solver with memory access retrieves and relies on past strategies by frequently invoking *Module 12* (precise problem comprehension) and *Module 13* (feasibility checks), indicating that it anchors its reasoning in documented successes or verified constraints. In the same scenario without memory, there is a spike in *Module 1* (possibly referencing more exploratory steps) and *Module 2* (breaking down the problem in a broader sense), suggesting that the Solver attempts to compensate for the lack of historical context by dissecting the logs more expansively. In hardware issues (Case 2), there is a distinct reliance on modules oriented toward troubleshooting (*Modules 8, 9, 12*). In human obstruction scenarios (Case 3), enabling memory leads the system to consult modules geared toward crowd interactions (*Module 10*) and risk assessments (*Module 6*) at a higher frequency, implying that prior records of teleoperation or audio alert usage are being leveraged. Without memory, the Solver’s reliance on some high-level reasoning steps (*Module 1, Module 2*) increases, but it invokes fewer domain-specific modules, forcing more generalized—and occasionally less refined—solutions. This perceived variability implies that the agent can flexibly tailor its approach: the presence or absence of memory alters which specialized processes are invoked, and the nature of the robot’s difficulties affects which subset of modules is accessed most frequently.

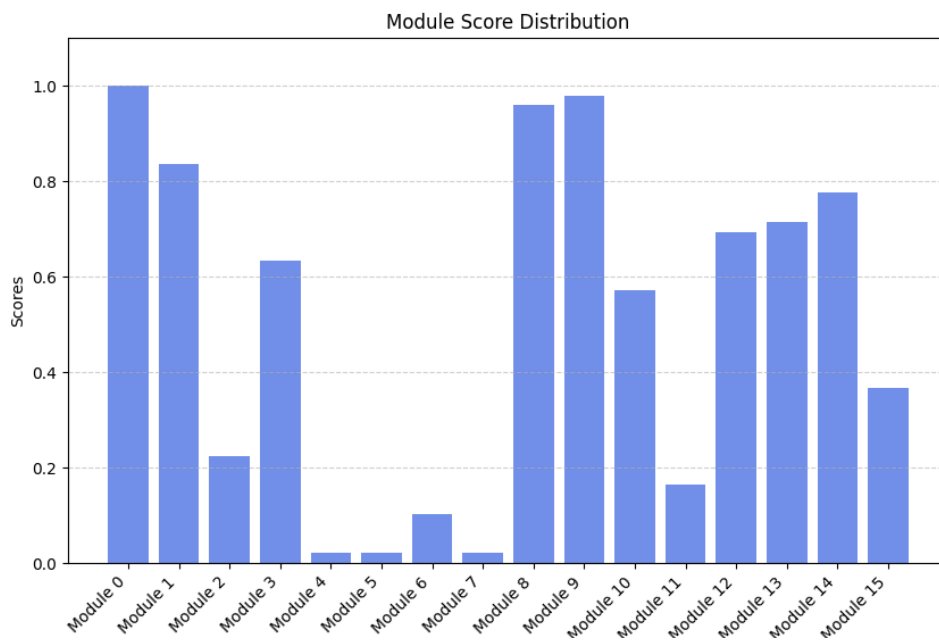
Finally, in memory-enabled runs, the Solver demonstrates a greater use of contextual explanations (*Module 15*), reflecting the system’s aim to maintain transparency by explicitly referencing logs and past successful solutions. Overall, these histograms indicate that activating the Archivist subagent modifies how the Solver balances creative exploration, direct retrieval of prior knowledge, and environment-specific constraints—underscoring the importance of memory-based augmentation for comprehensive and efficient supervision.

Figure 3.3: Histogram of Module selection on Case 1.



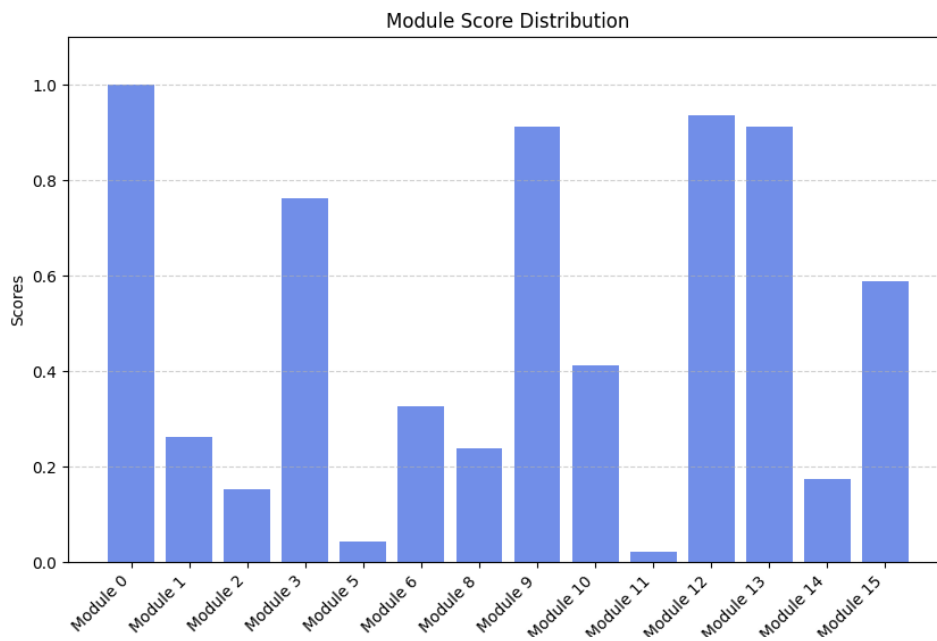
Source: Author, 2025

Figure 3.4: Histogram of Module selection on Case 1 without memory.



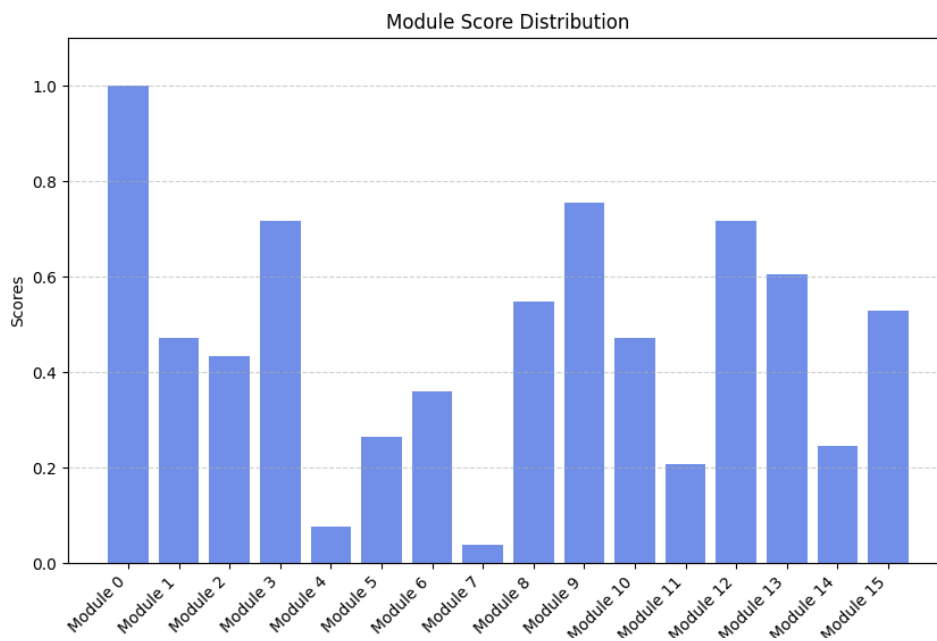
Source: Author, 2025

Figure 3.5: Histogram of Module selection on Case 2.



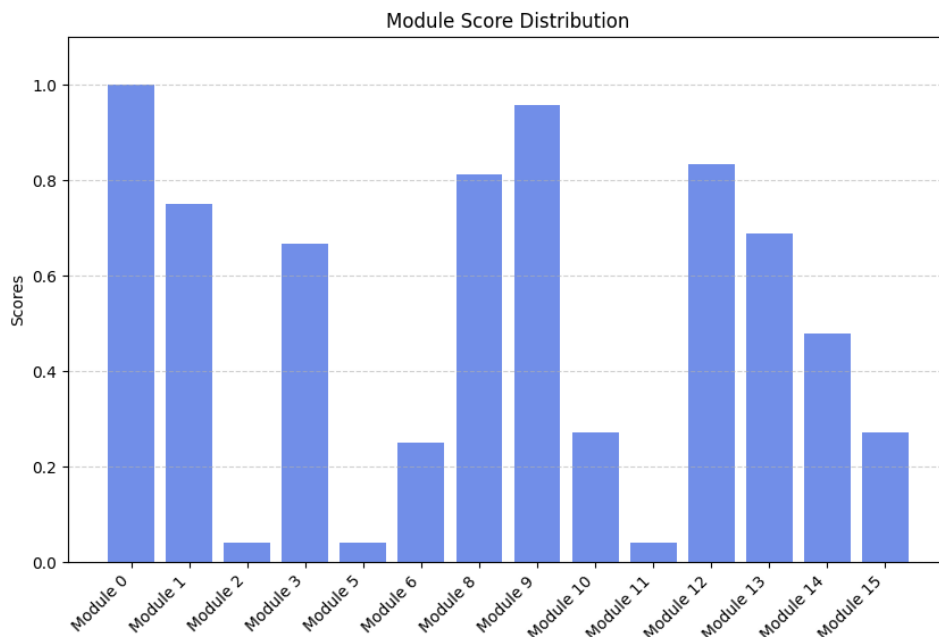
Source: Author, 2025

Figure 3.6: Histogram of Module selection on Case 2 without memory.



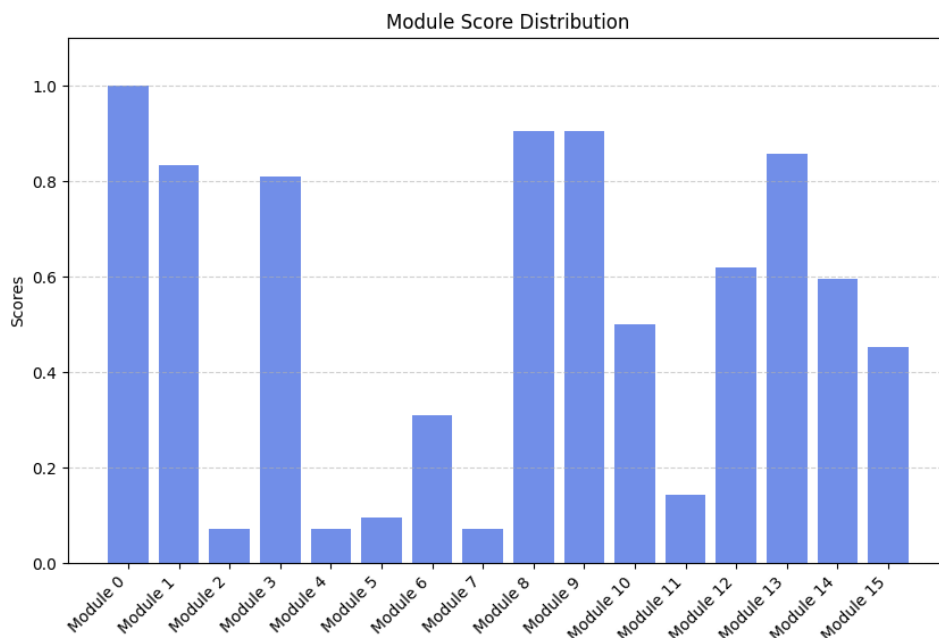
Source: Author, 2025

Figure 3.7: Histogram of Module selection on Case 3.



Source: Author, 2025

Figure 3.8: Histogram of Module selection on Case 3 without memory.



Source: Author, 2025

### 3.4.4 Discussion and Future Directions

These results align with and extend the body of work reviewed in Chapter 3, particularly the emerging emphasis on combining advanced autonomy with targeted human oversight [Jang, Jeon e Lee 2012, Fernandez-Ayala 2022, Li, Zheng e Lee 2022]. While prior studies explored either real-time operator feedback or domain-specific machine learning adaptations, our framework demonstrates that a structured, memory-augmented Agentic approach can unify these strengths. By systematically retrieving past solutions and integrating self-discovery reasoning into a natural language interface, we not only reduce the operator’s cognitive load but also foster higher consistency across repeated failure modes—an outcome that aligns closely with recent insights on trust-based multi-robot coordination [Mahani e Wang 2018] and adaptive LLM-based interfaces [Royce et al. 2024].

Furthermore, we can conclude that the proposed architecture successfully navigates complex scenarios and provides consistent and accurate support for an autonomous robot, comparable to that of a human operator. Additionally, memory augmentation has proven to offer multiple advantages, particularly in scenarios that occur repeatedly over time. Localization drift and hardware failures often necessitate the use of repeated solution strategies that the system can recall and refine, thereby enhancing overall reliability. In dynamically evolving contexts, such as child obstructions, the stored insights regarding prior community interactions assist the system in proposing a combination of teleoperation, gentle audio warnings, and mission retries tailored to recognized patterns.

Nonetheless, the proposed approach relies on extensive system logging and time-stamping, which raises concerns when data is nonexistent, incomplete, or when the volume of recorded interactions is too diverse. While such constraints also appear in other memory-based LLM contexts [Chew et al. 2023], future work might explore advanced memory indexing to prevent changes in the environment from rendering stored data obsolete. Ensuring that the Solver can gracefully discard stale memories or weigh their relevance in real-time remains a vital challenge. Additionally, as environmental complexity grows, concurrency in multi-robot operations will likely demand robust resource and memory management to maintain real-time responsiveness and sharing of experiences within the fleet. Investigating partial caching or distributed archivists could help mitigate potential bottlenecks.

Overall, these findings suggest that a carefully orchestrated agentic solution, in which subagents leverage memory and iterative reasoning, can significantly reduce user workload while enhancing reliability in multi-robot supervision. To fully capitalize on these advantages, attention must be given to data integrity, memory filtering mechanisms, and emergent scenario handling, especially in large-scale rollouts where operators oversee numerous missions simultaneously.

In the next chapter, we explore how these findings inform best practices for scaling agentic AI across multiple deployed robots. We will also refine the interplay between subagents to further reduce the operator's burden while maintaining robust oversight capabilities.

---

## Discussion and Considerations

---

This chapter presents a comprehensive reflection on the implications of the experimental findings of Chapter 3.4, connecting them to human-in-the-loop robotics, memory-augmented large language models, and multi-agent coordination. Following a summary of the proposed system’s operation under various failure scenarios, we explore how design principles, such as modular subagents and iterative reasoning, affect real-world deployments. Furthermore, we investigate potential avenues for scaling this architecture, identify possible pitfalls, and offer targeted recommendations for future work.

### 4.1 Summary of Findings and Practical Significance

The experimental evaluation confirmed that integrating a memory-based Archivist subagent with a modular, LLM-driven architecture achieves human-level performance in robotic supervision. Furthermore, when faced with repetitive or evolving challenges, the memory-enabled system consistently produced more accurate and context-aware solutions compared to a no-memory baseline. Quantitatively, both solution accuracy and user satisfaction improved in memory-enabled trials, while text-based similarity metrics indicated that the framework naturally reused successful vocabulary and strategies from previous instances. Similar advantages have also been observed in retrieval-based frameworks that enhance LLMs with structured knowledge [Liu et al. 2023, Zhong et al. 2024]

These gains extend beyond simple success rates. In practical usage, the Archivist and Solver collaboratively recognized the common ground between new problems and historical resolutions, which sped up diagnostic and planning processes, thereby lowering the operator’s burden. The architecture’s emphasis on iterative reasoning proved critical for adaptive responses, particularly in environments requiring immediate or partial manual intervention (for instance, teleoperation amid crowded walkways). By anchoring proposals in thoroughly documented past cases, the system reduced redundancy and provided

more transparent justifications, reflecting best practices for safe and explainable HITL operations.

The significance of these findings lies in bridging the gap between theoretical LLM-driven solutions and the day-to-day realities of multirobot oversight. Past successes—documented, summarized, and validated—serve as a repository of proven behaviors that the Solver can adapt with minimal input from human supervisors. This approach not only streamlines immediate problem-solving but also cultivates institutional knowledge for repeated missions or robots with shared hardware and software configurations. Consequently, the memory-augmented agentic approach establishes a robust foundation for long-term deployments in high-stakes environments, such as shopping malls or industrial facilities, ensuring that each intervention iteratively refines the system’s understanding of how best to operate in uncertain or changing conditions.

## 4.2 Limitations and Potential Improvements

The proposed memory-augmented, agentic LLM framework demonstrates clear benefits. However, several constraints remain that could be addressed in future work.

First, the system’s reliance on coherent and complete log data leaves it vulnerable to inaccuracies if sensor readings, timestamps, or error codes become corrupted or misaligned—an issue noted in related retrieval-augmented LLM applications [Wu et al. 2024]. Situations in which partial logs are available or where messages are lost could lead the Solver to propose suboptimal actions. Implementing a more robust data validation mechanism, potentially coupled with redundancy in sensor reporting, may help mitigate these risks.

While we presented results that contribute to reduced cognitive load and enhanced operator decision-making, our trial also raises important questions of trust and responsibility. In high-stakes scenarios, operators became overly reliant on the AI-generated guidance, inadvertently neglecting their own situational awareness. Moreover, the persuasive clarity of LLM-based responses can obscure underlying uncertainties. To prevent ‘automation bias,’ future iterations of this framework could incorporate explicit confidence indicators, audit trails for system-recommended actions, and user-interface elements designed to encourage critical human review. These measures not only protect operator agency but also help ensure that ethically sensitive decisions remain transparent, fostering an appropriate balance of trust and human oversight.

Another challenge arises from the dynamic nature of the environments in which the robots operate. Shopping malls can undergo frequent layout modifications or experience sudden influxes of pedestrians, limiting the direct applicability of older solutions stored by the Archivist. While the agent currently distinguishes whether older entries re-

main valid, it does not yet incorporate explicit metadata about environmental changes that could expedite or rule out certain recalled strategies. This tension resonates with discussions on stale data in memory-based agents [Maharana et al. 2024], indicating a need for more advanced indexing and data validation mechanisms. Exploring advanced memory indexing or weighting of historical entries—based on context tags such as crowd density or kiosk placement—could significantly enhance the system’s adaptability.

A further consideration is the computational overhead required to maintain and query long-term memory. Although the memory-enabled trials showed improved outcomes, the system’s performance in real-time depends on efficient retrieval and reasoning processes, particularly when diagnostics are cluttered with multiple errors or warnings. Investigating more fine-grained memory storage, such as compressing rarely accessed events, could help preserve historical knowledge while controlling resource usage.

Lastly, while user feedback confirms that referencing past successes increases transparency and trust, there is a risk of overfitting if the Solver relies too heavily on solutions that worked in the past while overlooking novel or emergent problems. Embedding a more explicit mechanism for updating or discarding outdated patterns could help prevent stagnation in solution design. Overall, refining the interplay between memory recall, real-time diagnostics, and self-discovery reasoning can further strengthen this agentic approach to human-in-the-loop robotic supervision.

### 4.3 Scalability and Deployment Prospects

The results suggest that a memory-augmented, agentic LLM framework can significantly reduce the burden of human supervision, particularly in multi-robot contexts. Although our current implementation focuses on a small fleet of indoor delivery robots, the modular structure of the agentic framework facilitates scaling across larger fleets and multiple sites. Each subagent (Analyst, Archivist, Solver, and Supervisor) can be replicated or distributed, enabling workloads to be balanced as the number of robots and requests increases. For instance, the Archivist could operate as a cloud-based service that aggregates logs and historical insights from diverse locations, while the Solver’s self-discovery processes can be parallelized or specialized for different operational domains. By design, this architecture supports incremental growth—new robots and additional human operators can be integrated with minimal reconfiguration—paving a feasible path to real-world deployments that require managing dozens or even hundreds of autonomous devices under human-in-the-loop supervision.

Prior frameworks advocate for communication-oriented structures that enable multiple LLM agents to share partial solutions or domain-specific knowledge [Zhang et al. 2025, Ahmed 2025]. Our design already demonstrates potential for cross-

location synergy through the Archivist; however, fine-tuning strategies for large-scale data curation remain underexplored. Addressing concurrency bottlenecks and ensuring consistent oversight across multiple Supervisor instances are essential next steps for real-world deployment [Liu et al. 2022].

In environments such as malls or large industrial facilities, this approach can be extended to multiple sites or shifts, with each support station operating a consistent set of subagents capable of retrieving domain- or context-specific knowledge. Operators may also benefit from a shared Archivist, where knowledge gleaned from one location can be transferred to another, provided that hardware specifications and operational parameters align. This cross-location synergy fosters a broader institutional memory that can accelerate troubleshooting and minimize operator training time.

Future deployments must address the system’s resource management. As the Archivist accumulates more data over time, it is essential to design effective storage and retrieval strategies to avoid sluggish retrieval processes or excessive memory usage. Methods such as automated pruning of outdated entries—based on time, domain mismatch, or repeated obsolescence—could help maintain the knowledge base’s relevance and compactness.

The ability to manage wide geographic and temporal dispersion relies heavily on robust connectivity. Since subagents frequently query system logs or camera data, it is critical to maintain stable network links to remote sites. In particularly demanding contexts (e.g., underground facilities or areas with spotty connectivity), a local caching approach could mitigate temporary disconnections. Nonetheless, the architecture is flexible enough that each site can host a partial or mirrored set of subagents, ensuring that day-to-day operations continue even if the main memory store is temporarily unreachable.

Overall, the combination of modular subagents and memory augmentation indicates a replicable model for large-scale, human-in-the-loop robotic oversight. By embedding knowledge reusability and advanced, self-reflective planning into an LLM-based system, organizations can scale their autonomous deployments more seamlessly, without overwhelming the limited bandwidth and cognitive capacity of onsite operators.

## 4.4 Final Considerations

In summary, the memory-augmented, agentic LLM framework proposed in this thesis presents a viable approach to enhancing human-in-the-loop robotic supervision. By coupling modular subagents with a self-reflective reasoning strategy, the system effectively integrates historical insights, real-time diagnostics, and adaptive solution generation in challenging environments, such as indoor delivery robots. The experimental

outcomes highlight distinct advantages—ranging from faster resolution of recurring failure modes to more coherent and user-friendly proposals—when historical context is properly captured and retrieved.

Ultimately, by pairing memory augmentation with a self-discovery strategy, this agentic AI solution has demonstrated notable improvements in solution reliability and operator confidence. These findings align with broader evidence that reasoning-rich LLMs, when augmented with domain-specific data and iterative planning, can effectively reduce manual workloads in high-demand settings [Chew et al. 2023]. Future research might explore incremental learning in dynamic environments, multi-modal expansions integrating more sensor data, and further synergy with human-robot interaction strategies that emphasize real-time collaboration [Fujita 2004, Minato et al. 2004, Pham, Nguyen e Tran 2025]. As agentic frameworks continue to mature, they hold the promise of scaling safe and efficient robotic operations in complex, ever-changing domains.

## Conclusion

---

This thesis introduces an agentic large language model (LLM) framework for human-in-the-loop robotic supervision, focusing on indoor delivery robots operating in dynamic environments such as shopping malls. By combining subagents for diagnostics, memory management, and self-discovery reasoning, the system demonstrates the capability to effectively integrate historical data, sensor logs, and user feedback into coherent, actionable proposals. Experimental trials indicate that memory augmentation significantly enhances solution quality and user satisfaction while reducing the repetitive workload placed on human operators.

Notably, our memory-augmented approach improved solution accuracy by up to **20.8%** and **12.1%** on average compared to the baseline condition, while operator satisfaction ratings increased from **4.72** to **4.83** (on a 5-point Likert scale). Additionally, mean resolution time dropped by **59%** when compared to operator time on the same task, reflecting both the framework's efficiency and the value of historical data retrieval. These figures not only validate the practical gains achieved but also underscore how modular subagents and memory-driven reasoning can reduce cognitive load and enhance real-time oversight in complex, multi-robot environments.

The research findings illustrate how structured memory retrieval and iterative reasoning together produce robust outcomes. Recurring problems, such as hardware malfunctions or crowd-related blockages, were resolved more quickly and with fewer missteps when the Solver could consult the Archivist's record of prior successes. The Analyst's role in summarizing logs and video data minimized the cognitive burden on operators, while the Supervisor provided clear directives for each sub-agent, culminating in coherent final solutions. This modular approach addresses multiple fail points in contemporary multi-robot supervision, ranging from incomplete contextual awareness to operator fatigue, by distributing tasks across specialized modules.

Although valuable, the architecture relies on reliable log data and a stable connection. Inconsistencies or missing entries can undermine the effectiveness of memory-driven suggestions. Environmental changes, such as shifting layouts or unexpected behavior updates, may also diminish the relevance of older solutions. Careful metadata tagging

and selective indexing of archived data could help mitigate these drawbacks, ensuring that historical strategies remain valid or are gracefully discarded as conditions evolve.

## 5.1 Future Works

Looking ahead, continued refinements in memory indexing, environment-specific adaptation, and resource management will be critical to supporting large-scale deployments. Incorporating more advanced logic for filtering obsolete entries or dynamically weighing the value of older solutions may prevent the Solver from relying on stale patterns. Furthermore, the ability to exchange archival data among multiple robots—each with potentially distinct hardware configurations—raises the prospect of collective institutional memory, benefiting entire fleets in distributed or multi-site scenarios. These enhancements, coupled with closer integration of multimodal sensor data and sophisticated user interface design, position the agentic LLM framework for broader adoption in real-world autonomous robotics. By unifying robust memory mechanisms, context-aware diagnostics, and flexible reasoning, this approach holds promise for delivering both reliability and transparency in high-stakes, human-centered automation.

---

## Bibliography

---

- [Abbasian et al. 2023]ABBASIAN, M. et al. Conversational health agents: A personalized llm-powered agent framework. *arXiv preprint arXiv:2310.02374*, 2023.
- [Abramovich et al. 2024]ABRAMOVICH, T. et al. Enigma: Enhanced interactive generative model agent for ctf challenges. *arXiv preprint arXiv:2409.16165*, 2024.
- [Abu-Rasheed, Weber e Fathi 2025]ABU-RASHEED, H.; WEBER, C.; FATHI, M. Tell me why i should “not” follow your recommendation: On the role of explainable ai in collaborative human-ai decision making. *OSF Preprints*, OSF, 2025.
- [Agarwal, Pandey e Mathur 2025]AGARWAL, S.; PANDEY, S.; MATHUR, V. Cryptanalysis and intention decoding in encrypted text corpus using transformer models. In: \_\_\_\_\_. *Books on AI & Cryptanalysis*. [S.I.]: Springer, 2025.
- [Ahmed 2025]AHMED, N. Investigation of cybersecurity risks in multimedia applications of chatgpt. *ResearchGate Preprints*, ResearchGate, 2025.
- [Alami et al. 1998]ALAMI, R. et al. Multi-robot cooperation in the martha project. *IEEE ROBOTICS AUTOM. MAG.*, 1998. Disponível em: <<https://doi.org/10.1109/100.667325>>.
- [Alexandrov 2025]ALEXANDROV, N. Problem complexity and llm agents: Enhancing human-machine team reliability. In: *AIAA SCITECH Forum*. [S.I.]: AIAA, 2025.
- [Aljohani e Herlofson 2025]ALJOHANI, S.; HERLOFSON, B. Evaluation of a context-aware chatbot using retrieval-augmented generation for answering clinical questions on medication-related osteonecrosis of the jaw. *Journal of Cranio ...*, Elsevier, 2025.
- [Baek et al. 2024]BAEK, J. et al. Researchagent: Iterative research idea generation over scientific literature with large language models. *arXiv preprint arXiv:2404.07738*, 2024.
- [Balch e Arkin 1998]BALCH, T.; ARKIN, R. C. Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation*, v. 14, n. 6, p. 926–939, 1998.

- [Bringes et al. 2013]BRINGES, C. et al. Determining the benefit of human input in human-in-the-loop robotic systems. *2013 IEEE RO-MAN*, 2013. Disponível em: <<https://doi.org/10.1109/roman.2013.6628447>>.
- [Brooks 1986]BROOKS, R. A robust layered control system for a mobile robot. *IEEE journal on robotics and automation*, IEEE, v. 2, n. 1, p. 14–23, 1986.
- [Brown et al. 2020]BROWN, T. B. et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, v. 33, p. 1877–1901, 2020. Disponível em: <<https://arxiv.org/abs/2005.14165>>.
- [Chew et al. 2023]CHEW, R. F. et al. Llm-assisted content analysis: Using large language models to support deductive coding. *ARXIV*, 2023. Disponível em: <<https://arxiv.org/abs/2306.14924>>.
- [Daoud e Ben-Hur 2025]DAOUD, A.; BEN-HUR, A. The role of chromatin state in intron retention: A case study in leveraging large scale deep learning models. *PLOS Computational Biology*, PLOS, 2025.
- [Davison et al. 2007]DAVISON, A. J. et al. MonoSLAM: Real-time single camera slam. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, v. 29, n. 6, p. 1052–1067, 2007.
- [Deng, Zhao e Pham 2024]DENG, W.; ZHAO, F.; PHAM, D. Learning by doing: A dual-loop implementation architecture of deep active learning and human-machine collaboration for smart robot vision. *Robotics and Computer-Integrated Manufacturing*, 2024.
- [Deng et al. 2024]DENG, X. et al. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, v. 36, 2024.
- [Devlin et al. 2019]DEVLIN, J. et al. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv preprint*, 2019.
- [Dissanayake et al. 2001]DISSANAYAKE, M. G. et al. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation*, v. 17, n. 3, p. 229–241, 2001.
- [Duan, McNeese e Li 2025]DUAN, W.; MCNEESE, N.; LI, L. Gender stereotypes toward non-gendered generative ai: The role of gendered expertise and gendered linguistic cues. *Proceedings of the ACM on Human Factors in Computing Systems*, ACM, 2025.
- [Dulac-Arnold, Mankowitz e Hester 2019]DULAC-ARNOLD, G.; MANKOWITZ, D. J.; HESTER, T. Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901*, 2019.

- [Fan e Zheng 2024]FAN, J.; ZHENG, P. A vision-language-guided robotic action planning approach for ambiguity mitigation in human–robot collaborative manufacturing. *Journal of Manufacturing Systems*, 2024.
- [Fernandez-Ayala 2022]FERNANDEZ-AYALA, V. Control barrier function-enabled human-in-the-loop control for multi-robot systems: Centralized and distributed approaches. *IEEE Transactions on Automation*, 2022.
- [Finn, Abbeel e Levine 2017]FINN, C.; ABBEEL, P.; LEVINE, S. Model-agnostic meta-learning for fast adaptation of deep networks. In: *Proceedings of the 34th International Conference on Machine Learning*. [S.l.: s.n.], 2017.
- [Fujita 2004]FUJITA, M. On activating human communications with pet-type robot aibo. *PROCEEDINGS OF THE IEEE*, 2004. Disponível em: <<https://doi.org/10.1109/jproc.2004.835364>>.
- [Furuta et al. 2024]FURUTA, H. et al. Multimodal web navigation with instruction-finetuned foundation models. 2024.
- [Geiger et al. 2013]GEIGER, A. et al. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, v. 32, n. 11, p. 1231–1237, 2013.
- [Ghafarollahi e Buehler 2024]GHAFAROLLAHI, A.; BUEHLER, M. J. Protagents: Protein discovery via large language model multi-agent collaborations. *Digital Discovery*, 2024.
- [Ghafarollahi e Buehler 2024]GHAFAROLLAHI, A.; BUEHLER, M. J. Sciagents: Automating scientific discovery through multi-agent intelligent graph reasoning. *arXiv preprint arXiv:2409.05556*, 2024.
- [Guo et al. 2023]GUO, J. et al. Suspicion-agent: Playing imperfect information games with theory of mind aware gpt-4. *arXiv preprint arXiv:2309.17277*, 2023.
- [Han, Shi e Tsui 2025]HAN, S.; SHI, L.; TSUI, F. Enhancing semantical text understanding with fine-tuned large language models: A case study on quora question pair duplicate identification. *PLOS ONE*, PLOS, 2025.
- [He et al. 2024]HE, B. et al. Ma-lmm: Memory-augmented large multimodal model for long-term video understanding. *CVPR*, 2024. Disponível em: <<https://openaccess.thecvf.com/CVPR2024>>.
- [Hou et al. 2025]HOU, C. et al. Urban sensing in the era of large language models. *The Innovation*, Cell Press, 2025.

- [Hu et al. 2022]HU, Z. et al. Human-machine telecollaboration accelerates the safe deployment of large-scale autonomous robots during the covid-19 pandemic. *FRONTIERS IN ROBOTICS AND AI*, 2022. Disponível em: <<https://doi.org/10.3389/frobt.2022.853828>>.
- [Huang et al. 2024]HUANG, Q. et al. Position paper: Agent ai towards a holistic intelligence. *arXiv preprint arXiv:2403.00833*, 2024.
- [Hunt, Godfrey e Soorati 2024]HUNT, W.; GODFREY, T.; SOORATI, M. Conversational language models for human-in-the-loop multi-robot coordination. *arXiv preprint arXiv:2402.19166*, 2024.
- [Jang, Jeon e Lee 2012]JANG, M.; JEON, S.; LEE, C. Control architecture for heterogeneous multiple robots with human-in-the-loop. In: *International Conference on Robotics and Automation*. [S.l.: s.n.], 2012.
- [Jiang, Zhang e Bai 2025]JIANG, Z.; ZHANG, T.; BAI, S. Towards enterprise-specific question answering using retrieval-augmented mechanisms. *SSRN*, SSRN, 2025.
- [Kang e Xiong 2024]KANG, Y.; XIONG, C. Researcharena: Benchmarking llms for information collection and organization. *arXiv preprint arXiv:2406.10291*, 2024.
- [Karaman e Frazzoli 2011]KARAMAN, S.; FRAZZOLI, E. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, v. 30, n. 7, p. 846–894, 2011.
- [Katevas et al. 1997]KATEVAS, N. I. et al. The autonomous mobile robot senario: A sensor aided intelligent navigation system for powered wheelchairs. *IEEE ROBOTICS AUTOM. MAG.*, 1997. Disponível em: <<https://doi.org/10.1109/100.637806>>.
- [Kavraki et al. 1996]KAVRAKI, L. E. et al. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, v. 12, n. 4, p. 566–580, 1996.
- [Kennington, Lison e Schlangen 2025]KENNINGTON, C.; LISON, P.; SCHLANGEN, D. Incremental dialogue management and its applications for human-robot interaction. *arXiv Preprint*, arXiv, 2025.
- [Khatib 1987]KHATIB, O. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, v. 3, n. 1, p. 43–53, 1987.
- [Khatib 1990]KHATIB, O. Real-time obstacle avoidance for manipulators and mobile robots. In: *Autonomous Robot Vehicles*. [S.l.]: Springer, 1990. p. 396–404.

- [Kim et al. 2024]KIM, Y. et al. Mdagents: An adaptive collaboration of llms for medical decision-making. 2024.
- [Kim et al. 2024]KIM, Y. et al. Mdagents: An adaptive collaboration of llms for medical decision-making. *Advances in Neural Information Processing Systems*, v. 37, p. 79410–79452, 2024.
- [Kiran et al. 2021]KIRAN, B. R. et al. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, v. 23, n. 6, p. 4909–4926, 2021.
- [LaValle e Kuffner 2001]LAVALLE, S. M.; KUFFNER, J. J. Randomized kinodynamic planning. *The International Journal of Robotics Research*, v. 20, n. 5, p. 378–400, 2001.
- [Leeper et al. 2012]LEEPER, A. et al. Strategies for human-in-the-loop robotic grasping. *2012 7TH ACM/IEEE INTERNATIONAL CONFERENCE ON HUMAN-ROBOT ...*, 2012. Disponível em: <<https://doi.org/10.1145/2157689.2157691>>.
- [Li, Zheng e Lee 2022]LI, C.; ZHENG, P.; LEE, C. Ar-assisted digital twin-enabled robot collaborative manufacturing system with human-in-the-loop. *Robotics and Computer-Integrated Manufacturing*, 2022.
- [Li et al. 2024]LI, H. et al. Hello again! llm-powered personalized agent for long-term dialogue. *ARXIV-CS.CL*, 2024. Disponível em: <<https://arxiv.org/abs/2406.05925>>.
- [Liu 2021]LIU, B. Recent Advancements in Autonomous Robots and Their Technical Analysis. *Mathematical Problems in Engineering*, Hindawi, v. 2021, p. 12, 2021.
- [Liu, Liao e Xu 2025]LIU, C.; LIAO, W.; XU, Z. Staged multi-strategy framework with open-source large language models for natural language to sql generation. *IEEJ Transactions on Electrical and Electronic Engineering*, Wiley, 2025.
- [Liu et al. 2022]LIU, H. et al. Robot learning on the job: Human-in-the-loop autonomy and learning during deployment. *International Journal of Robotics Research*, 2022.
- [Liu et al. 2023]LIU, L. et al. Think-in-memory: Recalling and post-thinking enable llms with long-term memory. *ARXIV-CS.CL*, 2023. Disponível em: <<https://arxiv.org/abs/2311.08719>>.
- [Lozano-Perez 1990]LOZANO-PEREZ, T. Spatial planning: A configuration space approach. In: *Autonomous Robot Vehicles*. [S.l.]: Springer, 1990. p. 259–271.
- [Lu et al. 2024]LU, C. et al. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.

- [Luo et al. 2024]LUO, J. et al. Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *arXiv preprint arXiv:2410.21845*, 2024.
- [Mahani e Wang 2018]MAHANI, M.; WANG, Y. Trust-based runtime verification for multi-quad-rotor motion planning with a human-in-the-loop. In: *ASME Dynamic Systems and Control Conference*. [S.l.: s.n.], 2018.
- [Maharana et al. 2024]MAHARANA, A. et al. Evaluating very long-term conversational memory of llm agents. *ACL*, 2024. Disponível em: <<https://acl2024.org/>>.
- [Minato et al. 2004]MINATO, T. et al. Development of an android robot for studying human-robot interaction. 2004. Disponível em: <<https://doi.org/10.1007/b97304>>.
- [Minsky 1988]MINSKY, M. *Society of Mind*. [S.l.]: Simon and Schuster, 1988.
- [Mirchandani et al. 2024]MIRCHANDANI, S. et al. So you think you can scale up autonomous robot data collection? *8th Conference on Robot Learning (CoRL 2024)*, 2024.
- [Misaros et al. 2023]MISAROS, M. et al. Autonomous Robots for Services—State of the Art, Challenges, and Research Areas. *Sensors*, MDPI, v. 23, n. 10, p. 4962, 2023.
- [Montemerlo et al. 2002]MONTEMERLO, M. et al. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In: *AAAI/IAAI*. [S.l.: s.n.], 2002. p. 593–598.
- [Moradi e Guastella 2023]MORADI, M.; GUASTELLA, D. Experience sharing and human-in-the-loop optimization for federated robot navigation recommendation. In: *Springer International Conference on Image Processing*. [S.l.: s.n.], 2023.
- [Mur-Artal, Montiel e Tardós 2015]MUR-ARTAL, R.; MONTIEL, J. M. M.; TARDÓS, J. D. ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, v. 31, n. 5, p. 1147–1163, 2015.
- [Mur-Artal e Tardós 2017]MUR-ARTAL, R.; TARDÓS, J. D. ORB-SLAM2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE Transactions on Robotics*, v. 33, n. 5, p. 1255–1262, 2017.
- [Naidu e El-Gayar 2025]NAIDU, N.; EL-GAYAR, O. A review of reasoning in artificial agents using large language models. 2025.
- [Nakano et al. 2021]NAKANO, R. et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.

- [Nooralahzadeh et al. 2024]NOORALAHZADEH, F. et al. Explainable multi-modal data exploration in natural language via llm agent. *ARXIV-CS.AI*, 2024. Disponível em: <<https://arxiv.org/abs/2412.18428>>.
- [Nottingham et al. 2023]NOTTINGHAM, K. et al. Do embodied agents dream of pixelated sheep? embodied decision making using language guided world modelling. In: *International Conference on Machine Learning (ICML)*. [S.l.: s.n.], 2023. p. 26311–26325.
- [OpenAI 2022]OPENAI. *ChatGPT*. 2022. Disponível em: <<https://openai.com/blog/chatgpt>>.
- [OpenAI 2023]OpenAI. Gpt-4 technical report. *ArXiv preprint*, 2023.
- [OpenAI 2023]OPENAI. *GPT-4 Technical Report*. 2023. Disponível em: <<https://arxiv.org/abs/2303.08774>>.
- [Pham, Nguyen e Tran 2025]PHAM, H.; NGUYEN, Q.; TRAN, D. Top 2 at alqac 2024: Large language models (llms) for legal question answering. *International Journal of Asian Language Processing*, World Scientific, 2025.
- [Qian et al. 2024]QIAN, C. et al. Chatdev: Communicative agents for software development. p. 15174–15186, 2024.
- [Rakelly et al. 2019]RAKELLY, K. et al. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In: *Proceedings of the 36th International Conference on Machine Learning*. [S.l.: s.n.], 2019.
- [Royce et al. 2024]ROYCE, R. et al. Enabling novel mission operations and interactions with rosa: The robot operating system agent. *arXiv preprint arXiv:2410.06472*, 2024.
- [Sánchez et al. 2022]SÁNCHEZ, C. M. et al. From perception to navigation in environments with persons: An indoor evaluation of the state of the art. *Sensors*, MDPI, v. 22, n. 3, p. 1191, 2022.
- [Schmidgall et al. 2025]SCHMIDGALL, S. et al. Agent Laboratory: Using llm agents as research assistants. *arXiv preprint arXiv:2501.04227*, 2025.
- [Spina et al. 2025]SPINA, A. et al. Tailoring glaucoma education using large language models: Addressing health disparities in patient comprehension. *Medicine*, LWW, 2025.
- [Stiefelhagen et al. 2004]STIEFELHAGEN, R. et al. Natural human-robot interaction using speech, head pose and gestures. *2004 IEEE/RSJ INTERNATIONAL CONFERENCE ON INTELLIGENT ...*, 2004. Disponível em: <<https://doi.org/10.1109/iros.2004.1389771>>.

- [Swanson et al. 2024]SWANSON, K. et al. The virtual lab: Ai agents design new sars-cov-2 nanobodies with experimental validation. *bioRxiv*, 2024.
- [Tang et al. 2024]TANG, X. et al. Medagents: Large language models as collaborators for zero-shot medical reasoning. p. 599–621, 2024.
- [Thrun 2002]THRUN, S. Probabilistic robotics. *Communications of the ACM*, ACM New York, NY, USA, v. 45, n. 3, p. 52–57, 2002.
- [Touvron et al. 2023]TOUVRON, H. et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [Vaswani et al. 2017]VASWANI, A. et al. Attention is all you need. Long Beach, California, p. 5998–6008, 2017.
- [Vaswani et al. 2017]VASWANI, A. et al. Attention is all you need. *Advances in neural information processing systems*, v. 30, 2017.
- [Wang et al. 2023]WANG, B. et al. Enhancing large language model with self-controlled memory framework. *ARXIV-CS.CL*, 2023. Disponível em: <<https://arxiv.org/abs/2304.13343>>.
- [Wang et al. 2023]WANG, Q. et al. Recursively summarizing enables long-term dialogue memory in large language models. *ARXIV-CS.CL*, 2023. Disponível em: <<https://arxiv.org/abs/2308.15022>>.
- [Wang et al. 2023]WANG, W. et al. Augmenting language models with long-term memory. *ARXIV-CS.CL*, 2023. Disponível em: <<https://arxiv.org/abs/2306.07174>>.
- [Wang et al. 2024]WANG, X. et al. Openhands: An open platform for ai software developers as generalist agents. In: *The Thirteenth International Conference on Learning Representations*. [S.l.: s.n.], 2024.
- [Wei et al. 2022]WEI, J. et al. Chain-of-thought prompting elicits reasoning in large language models. In: *Advances in Neural Information Processing Systems*. [S.l.: s.n.], 2022. v. 35, p. 24824–24837.
- [Wen et al. 2018]WEN, Y. et al. Robotic knee prosthesis real-time control using reinforcement learning with human in the loop. 2018. Disponível em: <[https://doi.org/10.1007/978-981-13-7983-3\\_41](https://doi.org/10.1007/978-981-13-7983-3_41)>.
- [White 2024]WHITE, J. Building living software systems with generative & agentic ai. *arXiv preprint arXiv:2408.01768*, 2024.

- [Wong et al. 2018]WONG, C. et al. Autonomous robots for harsh environments: a holistic overview of current solutions and ongoing challenges. *Systems Science & Control Engineering*, Taylor & Francis, v. 6, n. 1, p. 213–219, 2018.
- [Woo et al. 2024]WOO, J. J. et al. Custom large language models improve accuracy: Comparing retrieval augmented generation and artificial intelligence agents to noncustom models for evidence-based medicine. *ARTHROSCOPY : THE JOURNAL OF ARTHROSCOPIC & RELATED SURGERY ...*, 2024. Disponível em: <<https://www.sciencedirect.com/science/article/abs/pii/S0749806324008831>>.
- [Wu et al. 2024]WU, D. et al. Longmemeval: Benchmarking chat assistants on long-term interactive memory. *ARXIV-CS.CL*, 2024. Disponível em: <<https://arxiv.org/abs/2410.10813>>.
- [Wu et al. 2024]WU, Z. et al. Robot control via natural instructions empowered by large language models. Springer, 2024.
- [Xiao et al. 2024]XIAO, Y. et al. Cellagent: An llm-driven multi-agent framework for automated single-cell data analysis. *bioRxiv*, 2024.
- [Xin et al. 2024]XIN, Q. et al. BioInformatics Agent (BIA): Unleashing the power of large language models to reshape bioinformatics workflow. *bioRxiv*, 2024.
- [Yang et al. 2023]YANG, C. et al. Large language models as optimizers. *arXiv preprint arXiv:2309.03409*, 2023.
- [Yuan et al. 2023]YUAN, H. et al. Skill reinforcement learning and planning for open-world long-horizon tasks. *arXiv preprint arXiv:2303.16563*, 2023.
- [Zhang et al. 2023]ZHANG, H. et al. Huatuogpt, towards taming language models to be a doctor. *Findings of the Association for Computational Linguistics: EMNLP*, 2023.
- [Zhang et al. 2025]ZHANG, L. et al. Llm-aidsim: Influence diffusion simulation using llm-enhanced agent-based modeling. *Systems*, MDPI, 2025.
- [Zhang 2023]ZHANG, Y. *Multi-robot coordination and planning with human-in-the-loop under STL specifications: Centralized and distributed frameworks*. Tese (Doutorado) — Diva-Portal, 2023.
- [Zhao et al. 2025]ZHAO, P. et al. Insner: A generative instruction-based prompting method for boosting performance in few-shot ner. *Information Processing & Management*, Elsevier, 2025.

[Zhong et al. 2024]ZHONG, W. et al. Memorybank: Enhancing large language models with long-term memory. *AAAI*, 2024. Disponível em: <<https://aaai.org/Library/AAAI/aaai24.php>>.

[Zhou et al. 2022]ZHOU, D. et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.11756*, 2022.

[Zhou et al. 2024]ZHOU, P. et al. Self-discover: Large language models self-compose reasoning structures. *arXiv preprint arXiv:2402.03620*, 2024.

[Zhou et al. 2024]ZHOU, S. et al. Webarena: A realistic web environment for building autonomous agents. 2024.

[Zhu et al. 2025]ZHU, S. et al. Qwen-2.5 outperforms other large language models in the chinese national nursing licensing examination: Retrospective cross-sectional comparative ... *JMIR Medical Informatics*, medinform.jmir.org, 2025.

---

## Internal messages within the proposed solution.

---

Here are the internal messages within the hierarchical agentic architecture of the proposed framework, as described in Chapter 3.3.

### 1. Supervisor Instructs the Analyst for Diagnostics

```
[SUPERVISOR] -> [ANALYST]
```

```
Initiate a diagnostics report to determine why the robot required support help.  
Assess the information from the robot's logs and camera to understand the problem.
```

The Analyst collects and interprets relevant logs and video feeds, then compiles a concise report detailing the probable cause of the navigation failure:

```
[ANALYST] -> [SUPERVISOR]
```

```
### Problem Diagnosis Report
```

```
Mission Outcome:
```

- The mission was halted due to the inability to execute the path as indicated by: "Mission '39e79514-4a42-49f7-9062-358d4e211585' halted because Could not execute path (outcome 102)."

```
Key Events and Issues:
```

```
1. Obstacle Detection (19:31:35)
```

- Warning: "Obstacle detected inside the global path."
- Video: A group of adults and a child blocked the trajectory.

```
2. Planning Failure (19:32:02)
```

- Warning: "local\_planner was not able to obtain a local plan."
- Video: The child's proximity blocked sensors, preventing path re-evaluation.

```
3. Mission Halt (19:32:09)
```

- Logged warning indicates a critical failure in path execution.
- Video: Child's clothing obscured the frontal camera.

4. Frequent Calculation Warnings
  - The planner needed excessive time to maintain operational frequency.
5. Temperature Warning (19:31:34)
  - Potential overheating, but not central to the immediate navigation issue.

Summary:

The support call was triggered by close interaction with individuals, particularly a child, preventing the robot from navigating effectively.

## 2. Supervisor Instructs the Archivist to Retrieve Relevant Memories

[SUPERVISOR] -> [ARCHIVIST]

Please retrieve any data related to "Obstacle detected inside the global path" or "local\_planner not able to obtain local plan" when it involved people or children physically interacting with the robot, to provide context on previous troubleshooting and solutions.

The Archivist's memory retrieval mechanism identifies past instances of similar problems, illustrating how audio warnings, teleoperation, mission retries, or localization refreshes contributed to their resolution:

[ARCHIVIST]

Here are relevant memories associated with "obstacle detected" or "local\_planner" issues involving human/child interaction:

1. Memory Date: 2024-12-20
  - Diagnosis: Navigation blocked by human obstacles.
  - Actions: Audio alerts, teleoperation, SLAM map refresh.
  - Feedback: Effective.
2. Memory Date: 2025-01-15
  - Diagnosis: Path planning failures from crowd obstruction.
  - Actions: Audio alerts, retry mission, teleoperation.
  - Feedback: Perfect approach.
3. Memory Date: 2024-08-08
  - Diagnosis: Children in crowded environment delayed planning.
  - Actions: Audio clearance requests, teleoperation, hazard sensors update.
  - Feedback: Highly effective.

4. Memory Date: 2025-01-26

- Diagnosis: Dynamic obstacles & map lag caused local planner failures.
- Actions: Teleoperation, map refresh.
- Feedback: Fully resolved.

5. Memory Date: 2024-12-18

- Diagnosis: Local planning errors with obstructive environment, + overheating.
- Actions: Teleoperate to relocate, mission retries.
- Feedback: Successful resolution.

All prior solutions relied on audio prompts, mission retries, or moving the robot manually, consistently yielding user satisfaction.

### 3. Supervisor Directs Solver to Draft a Plan, Emphasizing Child Interaction

[SUPERVISOR] -> [SOLVER]

Based on the diagnostics and previous experiences, create a solution plan to address the issue of the robot being impeded by human obstructions, resulting in path planning failures, emphasizing child interaction considerations.

#### Solver's Task Proposal

[SOLVER] -> [SELF-REFLECTION]

"The logs show 'Obstacle detected...' and 'local\_planner' issues; video reveals a child blocking sensors. Past memory data highlight successes with audio warnings, teleoperation, and retries. Let's use these while prioritizing child safety to fix the path blockade without software/hardware changes."

## Self-Reflection on the Task

[SELF-REFLECTION]: SELECT MODULES

"Given the constraints and the specific details of the robot navigation problem with human interaction, the following modules are crucial:

\* Module 0: Learning from past experiences:

This is mandatory and directly relevant as past successes with audio warnings, teleoperation, and mission retries provide a strong foundation for the solution.

\* Module 2: How can I break down this problem into smaller, more manageable parts?:

Breaking down the problem (e.g., obstacle detection, human interaction, path replanning, communication with users) will allow for a more structured and effective solution.

\* Module 6: Risk Analysis:

Given the involvement of children, careful risk assessment of any proposed solution is paramount. The potential risks of each past solution (audio warnings, teleoperation, mission retries) need to be considered in the context of children's behavior.

\* Module 10: Are there any stakeholders or individuals directly affected by the problem, and what are their perspectives?:

Understanding the perspectives of children and other individuals interacting with the robot is crucial for designing a safe and user-friendly solution.

\* Module 12: Problem Comprehension:

Clearly restating the problem, including the constraints of no software/hardware modifications, ensures the solution is targeted and feasible.

\* Module 13: Feasibility and Resource Check:

This is vital for ensuring the chosen approach (audio warnings, teleoperation, retries, or a combination thereof) is practically implementable within the given limitations.

\* Module 15: Contextual Explanation:

A clear explanation of why the chosen solution is appropriate, referencing the problem specifics and constraints, will strengthen the proposal.

While other modules might offer some benefit, these are the most crucial for developing a robust, safe, and feasible solution within the single proposal opportunity. For example, modules like Creative Thinking (4) might be helpful, but given we have prior successful strategies, focusing on refining and adapting them is more efficient than brainstorming completely new solutions. Similarly, Systems Thinking (5) is less critical here given the focus on a specific, localized problem.",

[SELF-REFLECTION]: ADAPT MODULES

"Let's rephrase the modules to directly address the robot navigation problem with children:

\* Module 0: Leveraging Past Successes:

Analyze the historical data of successful obstacle avoidance using audio warnings, teleoperation, and mission retries. Specifically, examine how these methods performed in situations involving children and identify patterns or best practices.

\* Module 2: Problem Decomposition:

Break down the navigation challenge into sub-problems:

- Obstacle Detection & Classification (esp. children)
- HRI Strategies for children (audio, teleoperation, retries)
- Path Replanning Constraints (no software changes)
- User Communication

\* Module 6: Child-Centric Risk Assessment:

Evaluate the risks of audio warnings, teleoperation, and mission retries with children, ensuring safety is paramount.

\* Module 10: Understanding Children's Perspectives:

Consider how children perceive robots, potential confusion/fear, and development stages.

\* Module 12: Precise Problem Definition:

The robot must navigate around dynamic human obstacles (children) with no software/hardware modifications, using existing functionalities to address 'Obstacle detected' and 'local\_planner' errors safely.

\* Module 13: Feasibility Check:

Assess whether audio warnings, teleoperation, or retries can be implemented reliably within existing resource constraints.

\* Module 15: Justification of Chosen Approach:

Explain why the chosen combination addresses the logs' error messages effectively, referencing risk assessment, feasibility, and children's perspectives.",

[SELF-REFLECTION]: REASONING STRUCTURE

"1. Precise Problem Definition (Module 12)":

"Problem Statement": "The robot needs to reliably navigate around dynamic human obstacles, especially children, without software/hardware modifications, using existing functionalities (audio warnings, teleoperation, mission retries) to address 'Obstacle detected' and 'local\_planner failure' errors, prioritizing child safety.",

"Constraints": "No software/hardware modifications allowed. Must use existing audio warnings, teleoperation, and mission retries.",

"Success Criteria": "Robot reliably navigates around children without errors, prioritizing safety."

"2. Understanding Children's Perspectives (Module 10)":

"Child Perception of Robots": "",

"Child Reactions to Robot Actions (audio, movement)": "",

"Developmental Stage Considerations": "",

"Insights from Child Development Experts/Studies": ""

"3. Problem Decomposition (Module 2)":

"Obstacle Detection & Classification (Children)": "",

"HRI Strategies for Children (Audio, Teleoperation, Retries)": "",

"Path Replanning Constraints Analysis (No Software Changes)": "",

"User Communication Strategies (Children)": ""

"4. Child-Centric Risk Assessment (Module 6)":

"Risks of Audio Warnings (Startling, Confusion)": "",

"Risks of Teleoperation (Unexpected Movements)": "",

"Risks of Mission Retries (Repeated Approaches)": "",

"Mitigation Strategies for Each Risk": ""

```

"5. Leveraging Past Successes (Module 0)":
  "Success Rate of Audio Warnings with Children": "",
  "Success Rate of Teleoperation with Children": "",
  "Success Rate of Mission Retries with Children": "",
  "Patterns and Best Practices from Past Data": ""

"6. Feasibility Check within Constraints (Module 13)":
  "Feasibility of Audio Warnings (Existing Capabilities)": "",
  "Feasibility of Teleoperation (Responsiveness, Range)": "",
  "Feasibility of Mission Retries (System Resources)": "",
  "Reliability of Child Detection and Reaction Time": ""

"7. Proposed Solution (Combination of Audio, Teleoperation, Retries)": "",
"8. Justification of Chosen Approach (Module 15)":
  "Reasoning for Chosen Combination": "",
  "Evidence from Risk Assessment": "",
  "Evidence from Children's Perspectives": "",
  "Evidence from Feasibility Check": "",
  "Addressing Specific Error Messages": ""

```

[SELF-REFLECTION]: OUTPUT

```

"solution": "1. Problem Diagnosis: The robot's navigation system is failing
due to dynamic obstacles, specifically a child interacting near the robot's
sensors, causing 'Obstacle detected' and 'local_planner failure' errors.

2. Reasoning: Logs + video confirm these obstacles. Past successes with audio
warnings, teleoperation, and mission retries are recommended. Children's behavior
warrants a cautious approach.

3. Proposed Actions:
  a. Play Audio: \"Excuse me, please move slightly away from the robot.\"
  b. Retry mission.
  c. Teleoperate if child remains close.
  d. (Optional) Additional audio for reassurance.
  e. Retry mission once path is clear.

4. Justification:
Audio warnings prompt minimal disruption. Teleoperation ensures safe manual
navigation around a child. Final mission retry returns the robot to autonomy
after obstacle clearance."

```

## 4. Solver Delivers the Final Proposal

[SOLVER] -> [SUPERVISOR]

1. Problem Diagnosis:

Robot's navigation fails due to child-sensor interference, causing 'Obstacle detected' / 'local\_planner' errors.

2. Reasoning:

Video/logs confirm the child blocks sensors. Past memory shows audio warnings + teleoperation + mission retries typically work with close human interactions.

3. Proposed Actions:

- a) Play Audio: "Excuse me, please move slightly away from the robot so it can continue its path. Thank you!"
- b) Retry the mission: Attempt autonomous navigation again.
- c) Teleoperate if needed: If audio fails, steer around the obstacle manually.
- d) Play Audio (during teleoperation if appropriate): "I'm carefully navigating around you. Thank you for your patience."
- e) Retry the mission again: Return to autonomous operation once clear.

4. Justification:

- Audio prompts minimal disruption.
- Retrying checks if the path is now viable.
- Teleoperation ensures safe distance, especially near children.
- A second retry re-enables autonomous mode, limiting manual intervention.

## 5. Supervisor Summarizes the Solution for the User

[SUPERVISOR] -> [USER]

1. Problem Diagnosis:  
The robot cannot navigate due to a child in close proximity, blocking the sensors and plan execution.
2. Previous Experiences:  
Memory logs show successful use of audio alerts and teleoperation in similar situations.
3. Reasoning:  
Logs + video confirm a child obstructed the route.
4. Proposed Actions:
  - a. Audio warning to request clearance.
  - b. Retry mission in autonomous mode.
  - c. Teleoperate if the child remains too close.
  - d. Optional audio reassurance during teleoperation.
  - e. Final mission retry after clearing obstacles.
5. Justification:  
Maintains safety and minimal user disruption, reusing proven approaches from past child-interaction cases.

---

## Video Showcase

---

A video presentation (in Portuguese-BR) of the solution is available in <https://youtu.be/ypHCcjBvM1w>. The interactions shown in the video are also available in this Appendix.

### 1. Supervisor Instructs the Analyst for Diagnostics

```
[SUPERVISOR] -> [ANALYST]
Create a diagnostics report using the logs to understand
why the robot asked for support help.
```

The Analyst collects and interprets relevant logs and video feeds, then compiles a concise report detailing the probable cause of the navigation failure:

```
[ANALYST] -> [SUPERVISOR]
### Diagnostics Report: Robot Support Request

# Mission Outcome
The robot reported the mission was halted due to several critical errors
and warnings related to sensor and navigation difficulties.

#Key Events and Issues
1 - Error: Front and Rear Controller Status (14:38:56)
Message: "Controller is not connected"
Context: This indicates a disconnection of vital components
that could affect navigation and movement capabilities.

2 - Warning: Odom Topic Status (14:38:55)
Message: "Frequency too low."
Context: This suggests issues with the robot's odometry, which is
essential for accurate localization within its mapped environment.
```

3- Error: No Path! (14:38:08)

Logged when the robot attempted to navigate but had no feasible path due to internal or environmental constraints.

4- Warning: Robot Stuck (14:38:05)

Message: "Robot is stuck by some obstacle."

Context: Reflects severe navigation issues possibly due to interacting with the dynamic environment.

5- Warning: Gateway Communication is Slow (14:38:04)

Message: "Gateway communication is slow."

Context: Slow communication could hamper performance, especially under busy operational conditions.

6 - Warning: Mission Halted (14:38:58)

Message: "Mission halted because prerequisites 'Odometry is nominal' failed."

Context: The navigation system halted execution due to failed prerequisites related to sensor performance.

# Video Contextualization

Time Window Analyzed: 14:38:55 to 14:39:00

Summary: In the video, the robot attempted to navigate a wide mall corridor with pedestrian traffic. A child was seen standing directly behind the robot (0:57 to 1:03), causing immediate navigation issues as detected by onboard sensors, corroborating the warning about obstacles. This obstructive behavior aligns with the robot's request for support as it struggled to navigate effectively.

#Error Timeline

14:38:04: Warning - Gateway communication is slow.

14:38:05: Warning - Robot is stuck by some obstacle.

14:38:08: Error - NO PATH!

14:38:55: Warning - Odom topic status: Frequency too low.

14:38:56: Error - Front Controller Status: Controller is not connected.

14:38:56: Error - Rear Controller Status: Controller is not connected.

14:38:58: Warning - Mission halted because prerequisites "Odometry is nominal" failed.

14:38:59: Warning - Front Board is not active.

## 2. Supervisor Instructs the Archivist to Retrieve Relevant Memories

[SUPERVISOR] -> [ARCHIVIST]

Retrieve previous instances where the robot exhibited navigation issues due to dynamic obstacles and odometry failures, including reasons for 'no path' errors and memory of obstacles detected near the robot during operations.

The Archivist's memory retrieval mechanism retrieves previous issues related to dynamic obstacles, odometry failures, and 'no path' errors.

[ARCHIVIST]

Here are the most relevant past cases regarding robot navigation issues related to dynamic obstacles, odometry failures, and 'no path' errors:

DATE: 2024-12-20-15-13-33

Problem Diagnosis: Navigation errors caused by human presence resulted in trajectory feasibility warnings and planning failures.

Reasoning: Logs indicated repeated trajectory and planning issues tied to obstacles. Suggested solutions included using audio to clear paths, teleoperation, and reattempting missions after environment relocalization.

Outcome: Implementations were successful, with confirmation from the supervisor.

DATE: 2025-01-04-13-23-38

Problem Diagnosis: Obstacles near the robot induced trajectory errors and delays, halting the mission.

Reasoning: Audio clearance requests and mission retries resolved the impediment in autonomous mode without requiring additional relocalization.

Outcome: The robot completed its mission after applying the proposed solutions.

DATE: 2024-08-08-20-34-24

Problem Diagnosis: Persistent planning errors from ultrasound/ToF sensor interference and human-induced blockages.

Reasoning: Resetting the sensors eliminated the noise issue, allowing smooth mission resumption.

Outcome: Resolved by sensor reset followed by mission retry.

DATE: 2024-08-08-20-38-08

Problem Diagnosis: Dynamic obstruction from a child and system issues like calculation delays.

Reasoning: Recommendations included audio instructions, teleoperation around obstructions, and returning to autonomous navigation.

Outcome: Successfully addressed and the mission resumed.

DATE: 2025-01-15-19-36-20

Problem Diagnosis: Dynamic obstacles from pedestrians led to prolonged planning difficulties and mission stands.

Reasoning: Solutions involved audio alerts for clearance, mission retries, and manual teleoperation as a backup.

Outcome: Audio clearance primarily resolved the issue, validated through successful completion and user acknowledgment.

DATE: 2024-12-20-15-13-33

Problem Diagnosis: Navigation errors caused by human presence resulted in trajectory feasibility warnings and planning failures.

Reasoning: Logs indicated repeated trajectory and planning issues tied to obstacles. Suggested solutions included using audio to clear paths, teleoperation, and reattempting missions after environment recalibration.

Outcome: Implementations were successful, with confirmation from the supervisor.

DATE: 2025-01-04-13-23-38

Problem Diagnosis: Obstacles near the robot induced trajectory errors and delays, halting the mission.

Reasoning: Audio clearance requests and mission retries resolved the impediment in autonomous mode without requiring additional recalibrations.

Outcome: The robot completed its mission after applying the proposed solutions.

DATE: 2024-08-08-20-34-24

Problem Diagnosis: Persistent planning errors from ultrasound/ToF sensor interference and human-induced blockages.

Reasoning: Resetting the sensors eliminated the noise issue, allowing smooth mission resumption.

Outcome: Resolved by sensor reset followed by mission retry.

DATE: 2024-08-08-20-38-08

Problem Diagnosis: Dynamic obstruction from a child and system issues like calculation delays.

Reasoning: Recommendations included audio instructions, teleoperation around obstructions, and returning to autonomous navigation.

Outcome: Successfully addressed and the mission resumed.

DATE: 2025-01-15-19-36-20

Problem Diagnosis: Dynamic obstacles from pedestrians led to prolonged planning difficulties and mission stands.

Reasoning: Solutions involved audio alerts for clearance, mission retries, and manual teleoperation as a backup.

Outcome: Audio clearance primarily resolved the issue, validated through successful completion and user acknowledgment.

### 3. Supervisor Directs Solver to Draft a Plan

[SUPERVISOR] -> [SOLVER]

Based on diagnostics and memory insights, determine actions to resolve issues related to dynamic obstacles and odometry failures encountered during navigation

### 4. Solver Delivers the Final Proposal

[SOLVER] -> [SUPERVISOR]

Problem Diagnosis: The robot's mission halted because it cannot navigate effectively. This is due to a combination of hardware issues: intermittent front/rear controller disconnections and low-frequency/inaccurate odometry leading to poor localization. These problems are exacerbated when the robot interacts with dynamic obstacles like pedestrians, ultimately preventing it from finding a feasible path and meeting navigation prerequisites.

Reasoning: The core issue stems from unreliable hardware performance, specifically the motor controllers and the odometry system (likely involving wheel encoders and potentially the IMU). The intermittent nature of the controller disconnections and the drops in odometry frequency, especially when interacting with pedestrians, points towards potential physical hardware problems like loose connections, unstable power delivery under load/vibration, or components nearing failure. Pedestrian interactions likely introduce vibrations or cause sudden movements/stops that trigger these latent hardware faults. Resetting the controllers is the most direct allowed action to attempt recovery from common controller fault states (often logged as "Board is not active" or similar communication errors), which are critical for movement and contribute encoder data to odometry. Clearing the immediate vicinity of pedestrians removes a known trigger for these issues.

Proposed Actions:

Action 1: Play Audio: "Excuse me, please allow the robot space to maneuver. Thank you for your cooperation."

Action 2: Reset/Power Cycle Controllers

Action 3: Retry the mission

Justification: \* Action 1: This addresses the operational context where pedestrians exacerbate the problem. Asking people to move clears the immediate path and reduces the chance of physical interactions (like bumps or sudden avoidance maneuvers) that might be triggering or worsening the hardware instability (vibrations, connector stress).

Action 2: This directly targets the observed controller disconnection errors. Power cycling the controllers can resolve temporary hardware glitches, communication timeouts, or error states within the controller boards themselves, which are essential for robot motion and often provide encoder feedback for odometry. This is the most relevant hardware reset action based on the symptoms described.

Action 3: After potentially clearing the environmental trigger and attempting to reset the faulty hardware components, this action commands the robot to try navigating autonomously again to check if the immediate problem is resolved.

## 5. Supervisor Summarizes the Solution for the User

[SUPERVISOR] -> [USER]

Problem Diagnosis: The robot encountered issues including controller disconnections and low-frequency odometry which affected navigation, along with dynamic obstacles like nearby pedestrians interfering with its planned path.

2. Previous Experiences: Records suggest similar issues arose in the past, primarily addressed by requesting clearance from pedestrians, resetting hardware, and retrying missions.

3. Reasoning: The issues with navigation are due to intermittent hardware malfunctions and environmental interactions with individuals around the robot. Steps like clearing nearby obstacles and resetting system components were successfully implemented in past scenarios.

4. Proposed Actions:

1. Play Audio: "Excuse me, please allow the robot space to maneuver. Thank you for your cooperation."
2. Reset/power cycle all robot controllers to clear potential faults.
3. Retry the mission to check if the issue is resolved.

5. Justification: This combined approach targets both environmental and internal factors obstructing the robot, enhancing its functionality and navigational ability. Please confirm whether this solution resolves the issue.