Exploring Multi-Agent Systems for Request for Proposals

Master Thesis

Integrated Product Design
Delft University of Technology

Joost Spaanderman

Joost A. Spaanderman

Student number: 4823702

Master Thesis

Integrated Product Design

Faculty of Industrial Design Engineering

Delft University of Technology

Thesis Committee

Chair: Prof. dr. G.W. (Gerd) Kortuem

Mentor: Dr. H.U. (Uğur) Genç

In collaboration with

Schuberg Philis

Mentor: Laurens Eversmann - Innovation Specialist

Mentor: Frank Buters - Data Scientist

Coach: Ilja Heitlager - Chief Innovation Officer

Date

20th of August 2025

Preface

To conclude my time as a student and to obtain a master's degree in integrated Product Design, this thesis feels as a worthy end of a long period of learning and personal growth. Over the years my personal interest changed from designing physical products towards designing processes and digital systems. This project, focused on AI developments and how we can integrate this to our workplace, provided a great fit with my personal ambition to better understand this emerging technology and focus on designing for people.

Soon enough, I experienced how difficult it is to fully understand how we can control systems that can feel like a black box at times. Where I imagined I would be making a completely functional system at the start, that would be of direct use to Schuberg Philis and its employees, I had to steer back and define a scope that would inform a greater good, that steers other people in right direction and to learn from the discoveries I made. This thesis therefore took a more explorative direction to provide insights in my inevitable failures and successes with the aim to inform organizations how they can create multi-agent systems themselves based on my guidelines.

I am thankful for my supervisors, Gerd and Uğur, for their guidance and support during this project. Our discussions were essential and provided me with the confidence to critically assess my work. Your input helped me to combine the sheer amount of input from various activities towards one goal and outcome.

I also want to express my gratitude to everyone at Schuberg Philis for realizing this project. All stakeholders that participated in the interviews and workshop session made a great impact on my work and provided valuable insights. A special thank you goes out to the Lab271 team, you inspired me with your projects, conversations, and workplace. Laurens and Frank, thank you for guiding me in maintaining an overview, providing me with constant updates on the newest AI developments, refining my narrative and connecting me to the right people. Furthermore, I want to thank Ilja for sharing his vision on the unknown and how we can always learn from uncertainties if we dare to dive in.

To my family, Sabine, and friends: thank you for your unconditional support and company throughout this period. I am looking forward to spending more time with you again and to continue talking your ears off about the possibilities of AI!

J.A. Spaanderman
August 2025

Abstract

Organizations increasingly recognize the potential of integrating AI and agents into complex business processes yet lack practical guidance on how these processes can be translated into AI-augmented systems. The proposal development process for Request for Proposals (RFPs) is an example of such a process as it requires multidisciplinary expertise, adherence to strict requirements for both content and formatting, and consistent quality across writing from individuals and teams. Creating proposals consumes a lot of time and resources, while at the same time quality is essential as they can be decisive in winning new business.

This research was conducted in collaboration with Schuberg Philis, a mission-critical IT company, to explore how multi-agent systems can be applied to their RFP response procedure. The study adopted a multi-method approach: a literature review builds the theoretical foundation of MAS, contextual research in the form of interviews and training observations mapped the current challenges in proposal development, followed by iterative prototyping with a domain expert to design and refine agent roles and evaluate their outcomes, and concluding with a workshop on human-AI interaction guidelines. The prototyping phase utilized the workflow automation platform n8n to develop and test various agent configurations, orchestration patterns, and knowledge retrieval and sharing mechanisms.

The findings show that MAS offer the most value in knowledge retrieval from fragmented sources, requirement alignment with RFP specifications, and drafting compliant content following Schuberg Philis' best practices for proposal writing. The proof of concept demonstrated that MAS design requires a careful balance between agent specialization and over-fragmentation, considerations on parallel executions and sequential executions, and defining explicit instructions for agents' tasks and output hand-offs. Human oversight remains essential for adoption, with traceability and controllable levels of agent autonomy as primary factors for trust.

In addition to the proof of concept, the research delivers a set of design guidelines that translate the findings into actionable steps for developing MAS in business contexts. These guidelines cover the selection of a business process, role and task decomposition, orchestration patterns, iterative agent development, combining agents to a MAS, and human-agent collaboration.

Together, the proof of concept and guidelines provide both a concrete demonstration of a MAS tailored to proposal development at Schuberg Philis, and a practical foundation for organizations that aim to implement MAS in real-world business contexts. While the guidelines offer structured and actionable steps, further research is needed to validate scalability and applicability to other situations.

Terminology

SBP = Schuberg Philis
B2B = Business to Business
SaaS = Software as a Service
IT = Information Technology
KPI = Key Performance Indicator
RFP = Request for Proposal
SME = Subject Matter Expert

AI = Artificial Intelligence
GenAI = Generative AI
MAS = Multi-Agent System(s)
CoT = Chain-of-Thought
ToT = Tree-of-Thought
LLM = Large Language Model
ICL = In-Context Learning
RLHF = Reinforcement Learning from Human Feedback
RAG = Retrieval-Augmented Generation
GUI = Graphical User Interface
HAX = Human-AI Experience
SOP = Standard Operating Prototcol

CAQDAS = Computer-Assisted Qualitative Data Analysis Software

Statement on the use of AI

Throughout the research process AI tools have been used to help in rewriting text, enhancing transcripts of interviews and to function as a sparring partner. As the research is related to AI implementation, naturally AI has played a central role in this research, nevertheless the content of the thesis is my own and AI generated outputs were always critically assessed.

The cover image is generated using ChatGPT's GPT-5 model.

List of figures and tables

List of tables

List of figures

Table of contents

1 Introduction

This chapter introduces the background and relevance of this research. It establishes the current position of AI in the workplace and how this evolves into agentic systems. Furthermore, the context of the research in collaboration with Schuberg Philis and their Request for Proposal process is elaborated. By connecting these elements, the introduction shows why we explore design possibilities for multi-agent systems in real-world business contexts.

1.1 AI in the workspace

As businesses expand their digital landscapes, many are turning to intelligent systems to increase efficiency, support decision-making, and maintain competitive advantages. A central question for businesses of all sizes has become: “How can we rethink our business processes with GenAI?” (Deloitte AI Institute, 2024) and “how can we integrate it effectively to stay ahead of the competition?” based on research by (PwC, 2024).

Generative AI is changing the anatomy of work by automating and augmenting activities such that they require less time from humans. The acceleration of using generative AI for work comes from the understanding of natural language of large language models. This allows generative AI to make an impact on knowledge work, which is typically associated with occupations that require higher levels of education. This marks a shift in automation, which previously targeted mostly manual labor (Chui et al., 2023). The potential of implementing generative AI increases productivity in the workplace on average by 14% (Brynjolfsson et al., 2023).

As AI capabilities advance and are increasingly introduced in the workplace, the field of Industrial Design Engineering plays a role in ensuring these technologies are developed with human needs at the core. Rather than focusing solely on technical functionality, an industrial design perspective highlights the understanding of user workflows, societal contexts, and real needs and desires of end-users. This perspective is essential for creating AI tools that are not only effective, but also meaningful and usable.

1.2 Emergence of AI agents

A particularly promising development within generative AI is the rise of AI agents, these are autonomous entities capable of performing tasks, making decisions, and interacting with their environments. For the context of this research, AI agents can be seen as digital assistants, that can not only answer questions but can also take initiative and perform actions.

While the concept of agents has existed in computer science for many years, the capabilities of modern agents have expanded rapidly with the introduction of large language models (LLMs) that show potential in human-level intelligence (L. Wang et al., 2024). These models allow AI agents to understand and generate human-like language, reason through complex tasks, and interact naturally with users and systems. Because of these advancements, research and advisory firm Gartner predicts that 33% of enterprise software applications include agentic AI in 2028, compared to less than 1% in 2024” (Coshow et al., 2024). Apart from the integration within software, Google expects 82% of large enterprises to use agents in their work (Google Cloud, 2024).

With the advancement of AI agents, a growing interest in multi-agent systems (MAS) emerges. These systems consist of multiple interacting AI agents, working towards a shared goal. MAS architectures are increasingly seen as an effective replication of complex human collaboration and decision-making patterns, making them viable for business contexts.

Although organizations know they want to include AI or agentic systems in their workflows, they struggle with where to start and how to build something valuable for their own processes. Literature largely describes the definitions, capabilities, and architectures of agents and MAS, but offers limited practical guidance on how to translate a specific business process into an automated or augmented process with MAS.

This thesis aims to explore how businesses can approach the translation from a (complex) business process into augmentative multi-agent systems. Using the case study of the Request for Proposal response procedure at Schuberg Philis, this thesis describes how organizations can identify opportunities, configure agent roles, and orchestrate them in an effective way.

1.3 Schuberg Philis

The research is conducted in collaboration with Schuberg Philis (SBP), a mission-critical IT company with a primary KPI of 100% customer satisfaction. Their organizational culture emphasizes people and building relationships which reflects in their way of working with dedicated customer teams and by putting experts in the lead.

The commitment of Schuberg Philis for technical innovation is led by their internal innovation team, called Lab271, where emerging technologies are constantly explored for internal initiatives and customer-facing projects. Each year the innovation team creates strategically aligned 'bold moves' which become the focus point for exploration in that year. One of the bold moves is focused on AI, where embedding AI in a business process stands at the core. During the research period, the Lab271 team rolled out a companywide program to increase AI literacy among colleagues. The environment and everyday insights and discussion on the developments of AI formed an interesting dynamic for this research, with a high understanding of the potential of AI and agents from team members. Furthermore, the collaboration with SBP provided access to domain expertise and a real-world business case, which are essential for this research.

The Lab271 team identified the RFP process as a compelling use case for exploring the possibilities of MAS. Their interest in applying MAS to a business process is what initiated this project together with our interest in this research topic. The choice for the RFP process as direction came from an increase in participation in RFP's, therefore increasing the workload for proposal teams. This makes the case for an RFP focused MAS a good candidate from both a complexity perspective and a desirability perspective.

1.4 Business case: Request for Proposal process

One business process that could benefit from automation is the Request for Proposal (RFP) procedure, as it is complex, repetitive and time-consuming, as will be explained below. An RFP is a specific form of procurement, that is also referred to as a tender, which is the overarching terminology for RFI, RFP, and RFQ as explained in Figure 1.

In an RFP, organizations invite external vendors to submit proposals based on a list of detailed requirements. The public procurement procedure is visualized in Figure 1. With total value of contracts awarded through tenders in the Dutch IT services sector reaching €1,5 billion in 2024 (Opentender Netherlands, n.d.), this highlights the importance of writing high quality proposals, as significant business can be won through this process.

Writing proposals is a time-consuming and labor-intensive task and presents several challenges. Proposals must be tailored to the specific needs and formatting preferences of each client, and high-quality responses typically require input from multiple subject matter experts (SME's). This makes the process not only labor-intensive, but also vulnerable to duplicated efforts and inefficiencies (Lauesen & Vium, 2005). Additionally, unsuccessful proposals will lead to lost opportunities and sunk costs (mostly consisting of labor hours). Automating or augmenting this process using AI technologies, such as a MAS, is interesting as it could reduce the time and costs associated with drafting proposal whilst improving consistency of quality.

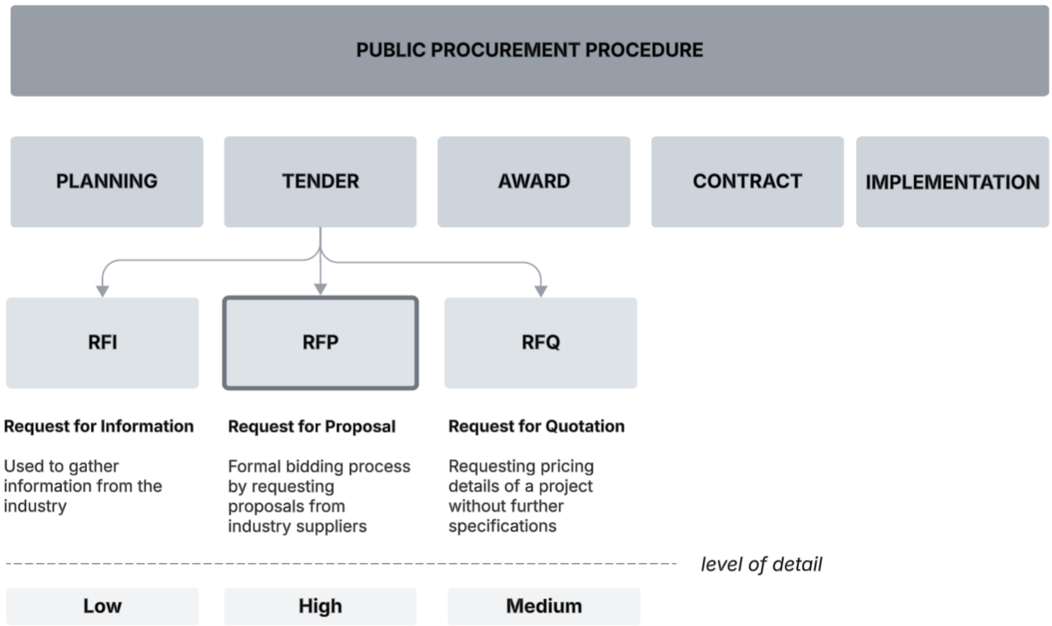

Figure 1: Different stages of public procurement. A tender is the general term for one of three requests (RFI, RFP, or RFQ).

By establishing the relevance and challenges for applying MAS in the business context for proposal development, the introduction formed a foundation for the research objectives.

2 Research approach

This chapter presents the research questions relevant to both exploring business processes and MAS, and a research approach to clarify the methodology, the research activities and report structure. Next to this, the deliverables of the research are described, with the structured overview indicating the foundation for the deliverables and research outcomes.

2.1 Research questions

The complexity of RFP development presents a great test case for exploring how multi-agent systems can address limitations of current AI approaches and offer value to business contexts. This research aims to bridge the gap between the theoretical potential of MAS and practical implementation challenges and insights in a real business process.

Main research question

RQ: *How can Multi-Agent Systems be designed and applied to support the proposal development process in B2B IT request for proposals (RFPs)?*

This main question aims to address the central challenges of translating business processes involving humans into an AI-augmented system whilst remaining similar levels of outcome quality, consistency, and strategic value that are required for successfully supporting proposal writing. Although the research focusses on the RFP process specifically, the aim is to understand how business processes in general can be translated into multi-agent systems.

Sub-questions

To guide and structure the design exploration, we define the following sub-questions.

SQ1: Where in the proposal development process can AI integration offer the most value?

This question aims to form the scope and boundaries of MAS implementation by understanding the RFP process and identifying where automation provides value while the risks remain limited.

SQ2: *What features should a MAS include to make it specific to the company and useful for generating proposal content?*

As a proposal should represent the company through written text, it is important to understand what competitive differentiators are for writing the proposal. This question aims to explore how written text can strategically be used to represent the business.

SQ3: *What design principles and orchestration patterns between agents enable MAS to generate effective output for proposals?*

This question is focused on design choices for AI systems such as prompt engineering, agent orchestration and task dependencies. It aims to effectively address what effects different approaches and strategies have on the system performance and output quality.

SQ4: *How should human-agent collaboration be integrated into a MAS for RFPs to ensure trust in its decisions and output?*

As proposal writing can be viewed as a high-stakes business process, it is important to understand how human oversight and control can be offered and to what extent users desire this.

Together these research questions aim to provide an understanding of how MAS can be effectively designed and implemented in the context of RFP development. While proposal development presents a concrete case study, the findings of the research are intended to inform broader implementation challenges and opportunities in business contexts.

2.2 Research approach and structure

The research adopts a multi-methods approach to capture research background, and practical implementation to address all research questions. Three main phases form the outline of the research and report.

The foundation for understanding large language models, agents and multi-agent system is created by conducting literature research.

The contextual research phase established a foundation of knowledge on both multi-agent systems and the complexity of RFP processes at Schuberg Philis. A combination of semi-structured interviews (n=7) and training observation provides an understanding of the challenges and opportunities at hand.

The prototyping and user research phase involved iterative prototyping through seven co-design sessions with a domain expert to develop and test MAS concepts. A workshop on human-AI interaction (n=8) was held to research what user-centered design principles participants at Schuberg Philis desire for an RFP related MAS.

The research concludes in two main outcomes: a working proof of concept that augments users in proposal drafting, and design guidelines for implementing multi-agent systems. Next to this, a series of suggestions on interface design relating to the multi-agent system is formed as a foundation for potential future work on how interfaces could integrate multi-agent specific features.

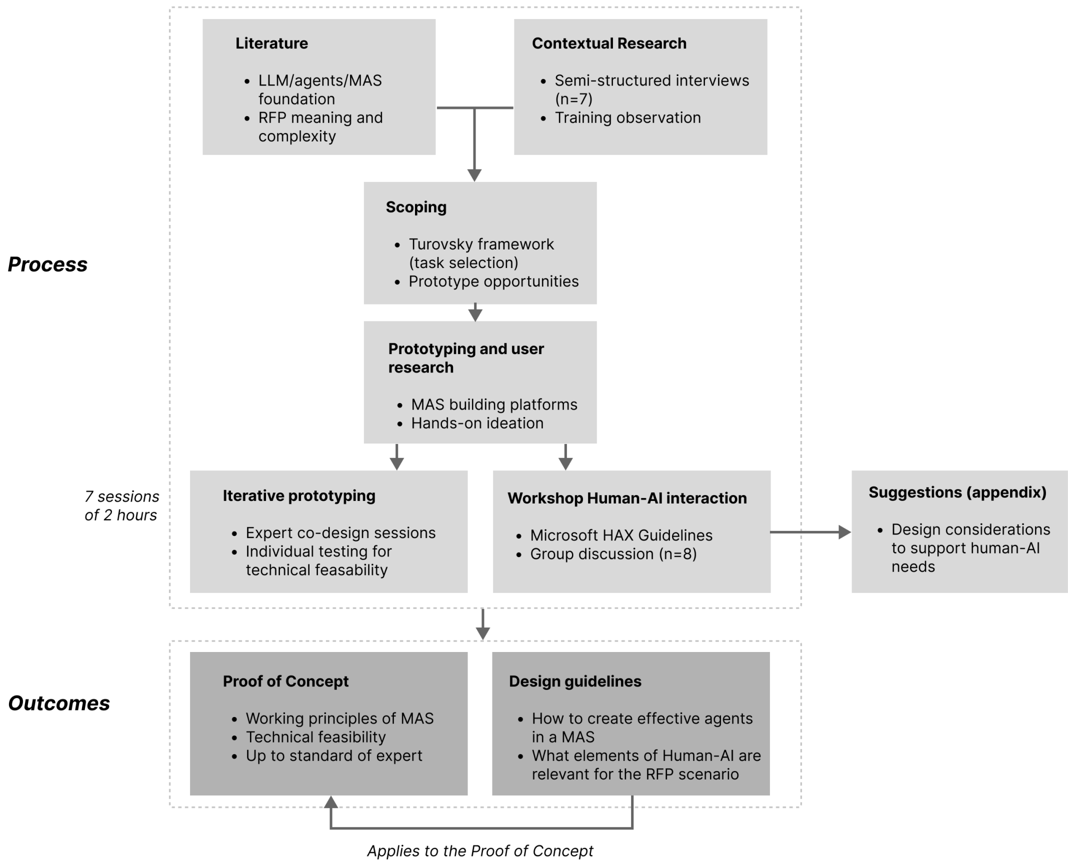

Figure 2: Visual overview of research and report build-up with activities or results.

This page intentionally left blank

3 Background

To explore how multi-agent systems can support the automation of RFP response drafting, it is necessary to establish an understanding of underlying concepts, considerations, and the current state of AI agents. This chapter provides the literature background for the research, focused on artificial intelligence, large language models, agentic systems, and human-AI collaboration. The aim of the section is to form the foundation for research choices described in later chapters.

3.1 Large Language Models

3.1.1 Definition and capabilities

A large language model (LLM) is a large-scale, pre-trained statistical language model based on neural networks (Minaee et al., 2025). The use of natural language in LLMs is part of what makes them suitable for agents. It enables agents to communicate with other agents and with humans in a interpretable conversation that offers clarity (Xi et al., 2023). Because LLMs are trained on large amounts of data and can interpret text, they are increasingly used for a wide range of everyday tasks. A well-known example is ChatGPT (OpenAI, 2025), which offers a chat interface to directly ask questions to the LLM, returning replies based on the trained data and, at the time of writing, by directly scraping the internet.

Fundamentally, LLMs are next-token prediction models trained on large datasets that include information form the internet, books, and other sources for training. Their probabilistic nature causes LLMs to generate text one token at a time based on the probability of the next token (often a word) to come next (Minaee et al., 2025). Understanding this is important for users when using LLMs for any application, as they do not have a sense of 'truth'. This also means that, when a user asks a question to an LLM, the results may vary and lead to potential inconsistencies.

Large Multimodal Models (LMMs) can process data types beyond text, such as images or audio. This can be useful for analyzing graphs, generating images, or generating and understanding audio fragments. Multi-modal models that allow information perception of other input forms than text, allow their perception to come closer to human perception (X. Li et al., 2024). Different LLMs can give other responses as they are trained on different data and have different parameter counts (Brown et al., 2020; Liang et al., 2023). As these LLMs can behave differently, it is important for the user to consider what model should be considered to reach the desired output.

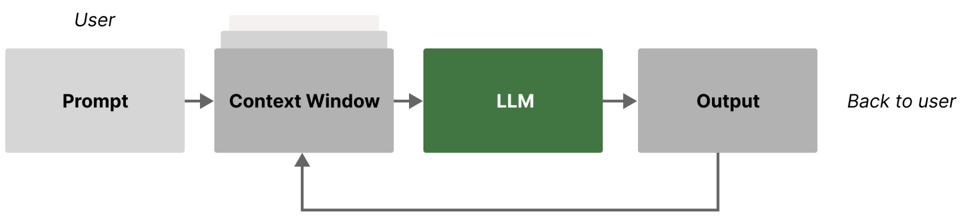

Figure 3: basic interaction with a LLM (often as a chat interface).

3.1.2 Context windows

An important subject in LLMs is the context window. Context windows define the amount of information an LLM can process in a prompt. It can be seen as the total information that can be handled at a time, and it can be seen as short-term or working memory. Context windows are measured in tokens, which are numerical codes that represents sub words or characters of text. Converting text into numbers is called tokenization and each LLM can have their own tokenization process (McKinsey & Company, 2024).

A larger context window allows you to provide more input or context to the LLM, meaning you can provide more examples or generate text based on larger inputs. However, larger context windows increase the required computational power and costs of the prompt. If the context window limit is reached, the LLM will not take new information in consideration, and a new context window must be opened. Most providers of LLM-based chat interfaces cut off the chat and ask to start a new conversation.

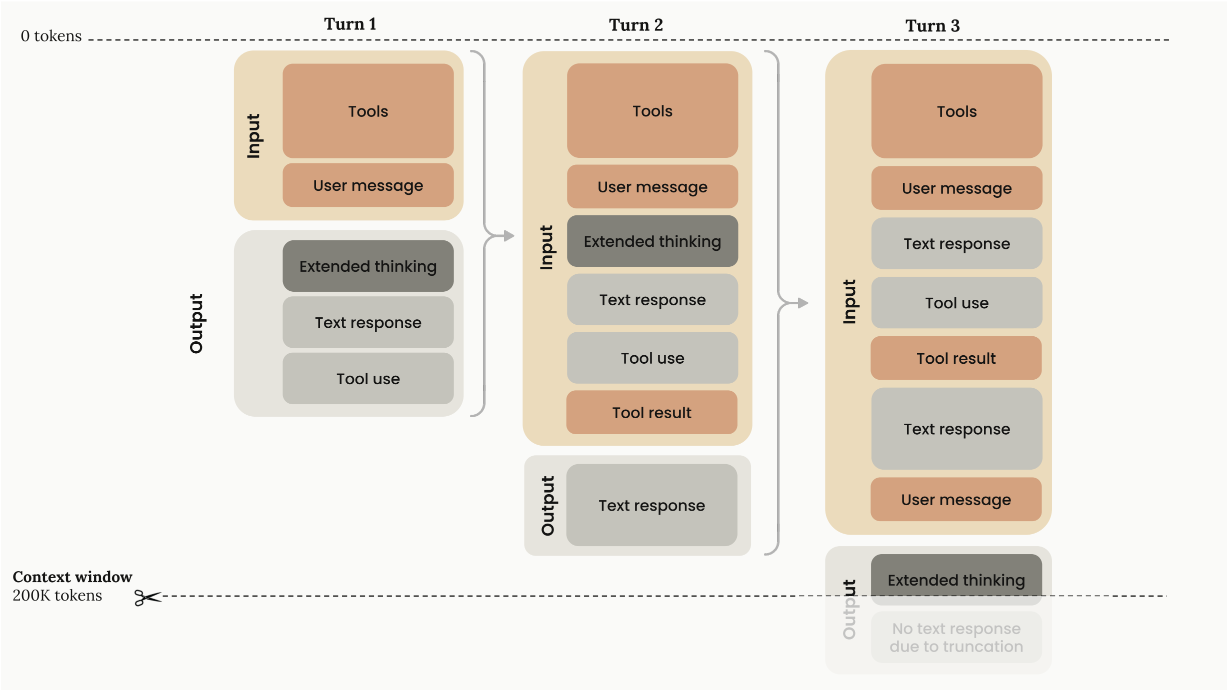

Figure 4: visual representation of how a context window reaches its maximum token count over time, causing a cut-off (Anthropic, n.d.).

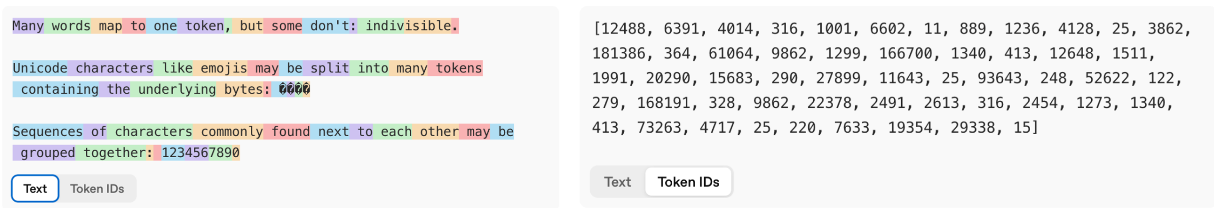

Figure 5: A visualization of tokenization of text. Words or segments of a word are represented by numerical codes that maintain the meaning of these characters/words (OpenAI, n.d.)

3.1.3 In-context learning

LLMs are able to learn tasks when given only a few examples in the form of demonstrations, this is called in-context learning (ICL) (Dong et al., 2023). Unlike traditional machine learning, where parameter updates are required, ICL allows models to adapt to new facts and tasks by learning from examples alone. The ability to learn from examples improves significantly with model size (Brown et al., 2020), which may be one reason why larger models are often considered to perform better at reasoning and adjusting to the user's input.

A popular approach from users to ensure an LLM behaves as intended and reaches the user's goals, is to utilize in context learning (ICL) by implementing feedback, called reinforcement learning from human feedback (RLHF) (Minaee et al., 2025). This involves users suggesting alternatives approaches or corrections to the generated reply, so the model learns to align its behavior to the user's needs.

Although direct instructions to an LLM can influence the behavior and generated content (for example: “use a professional tone”), the most effective way to utilize ICL is to provide examples in the context.

In the context of multi-agent systems, ICL becomes particularly valuable by allowing agents to adapt to new roles or tasks by learning from examples, without the need for retraining the model.

3.1.4 Prompting techniques

While in-context learning focuses on how LLMs adapt form examples, prompting techniques determine how effectively examples and instructions are communicated and handled by the model. The quality of the prompt has a large influence over the output of the LLM, which is why various techniques have been developed to improve prompt effectiveness. The art of optimizing prompts is called ‘prompt engineering’ (Google Cloud, n.d.), and refers to methods for steering LLM behavior to get towards a desired outcome (Weng, 2023). LLMs require careful consideration with prompt engineering as they seem to interpret and understand prompts differently than humans do (Zhou et al., 2022).

The most common prompting techniques are zero-shot, few-shot, and chain-of-thought prompting.

- **Zero-shot prompting** involves directly asking the LLM to perform a task without providing examples (Weng, 2023).
- **Few-shot prompting** involves the user providing the model several examples of input and expected output combinations to clarify what the intention and criteria are of the required output. Few-shot prompting often leads to better performance, but can result in higher cost or reaching the context window limit faster (Weng, 2023).
- **Chain-of-Thought (CoT) prompting**, as introduced by (Wei et al., 2022), is a more advanced technique. By breaking the task down to multiple steps, the LLM is encouraged to break down the question which improves reasoning. CoT shows better results in complicated reasoning tasks, though simple tasks do not benefit from this technique (Weng, 2023). A further advantage of CoT prompting is that it provides an overview of the reasoning process per step, making visible to users where reasoning went wrong, which makes debugging LLM outputs easier (Wei et al., 2022).

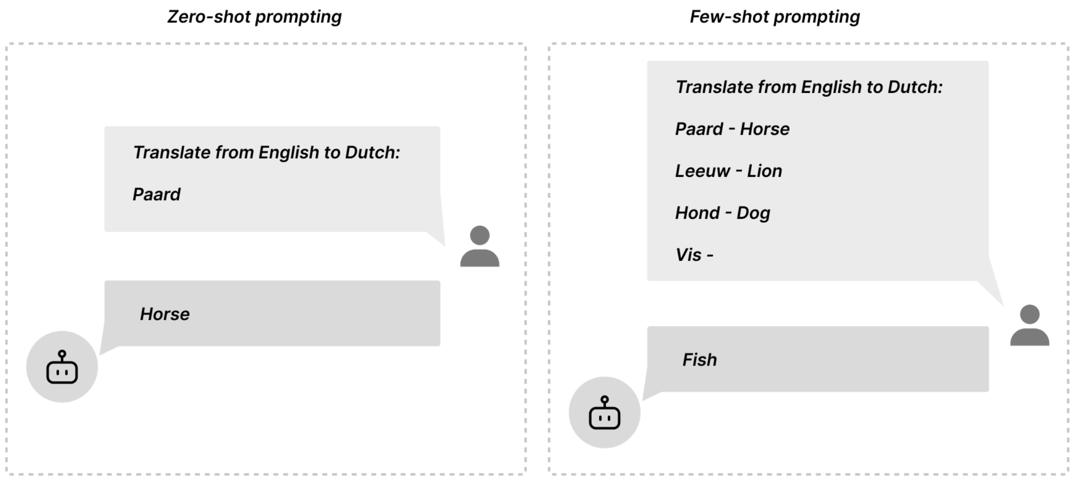

Figure 6: Zero-shot and few-shot prompting technique examples.

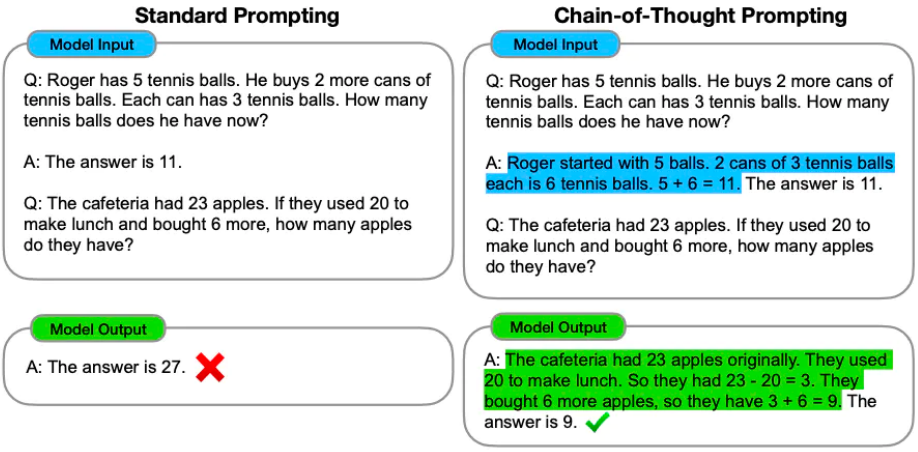

Figure 7: Chain-of-Thought prompting compared to regular prompting (Wei et al., 2022).

Building on CoT, Tree-of-Thoughts (ToT) as introduced by Yao et al. (2023) allows the LLM to explore multiple reasoning paths and choose the most promising solution through self-evaluation. However, ToT prompting is not as easy to implement as above prompting techniques as it requires the setup of a system, rather than providing direct instructions (IBM, 2024).

In the context of creating systems with LLM-based agents (described in section 3.2), robust prompting helps to ensure that agents interpret tasks consistently, which is eventually important for communication between agents as described later in the background for MAS.

3.1.5 Hallucinations

Even with well-designed prompts, LLMs can still produce factually incorrect information and make mistakes. This phenomenon, where LLMs generate contradicting or unverifiable information is referred to as hallucinations.

Hallucinations can appear in different ways: the model can contradict information previously provided by the user (input-conflicting hallucinations), it can generate content that misaligns with the information given in the conversation (context-conflicting hallucinations), or it can generate facts that lack any form of verification (fact-conflicting hallucinations)(Y. Zhang et al., 2023).

More widely adopted terms for hallucination types are intrinsic hallucinations, where generated text contradicts the input, or extrinsic hallucination where generated text cannot be verified with existing information (X. Li et al., 2024).

Because LLM outputs are often written in a confident, well-structured style, it can be hard for users to identify these inaccuracies. This reinforces the importance of incorporating a form of verification, by putting humans in the loop, before implementing generated content in critical workflows.

In the context of multi-agent systems, hallucinations are not just limited to one agent. When a hallucination in the output is communicated to the next agent as input, the incorrect statements can propagate across the entire system. This highlights the importance of implementing options for controlling outputs of AI agents, which are first explored in the following section.

3.2 AI Agents

While LLMs have powerful reasoning capabilities, they remain passive and only respond to the user's prompt. To apply these capabilities in a purposeful way, they must be implemented in AI agents to enable autonomous action taking.

In the field of artificial intelligence, agents are defined as artificial entities that are capable of perceiving the environment through making decisions, and taking action accordingly using actuators (Cheng et al., 2024; Minaee et al., 2025; Russell & Norvig, 2022; Xi et al., 2023). In this thesis, this definition of agents will be maintained.

The goal of implementing agents is often to make repetitive tasks easier to accomplish for users, either by increasing productivity or by reducing resources in time or man hours. Traditionally, agents followed predetermined decision patterns (deterministic), limiting their adaptability to different situation. However, with the introduction of LLMs into agent architectures, agents have the capacity to understand natural language and to learn and improve behavior based on examples and prior experiences by utilizing their ability for ICL (Cheng et al., 2024).

Agents are built up of three core components that enable their abilities (Xi et al., 2023):

1. **Perception:** receiving and interpreting input from the environment.
2. **Brain:** enables reasoning and decision-making, in the case of AI agents this is an LLM.
3. **Action:** based on the chosen response of the brain, it executes a task.

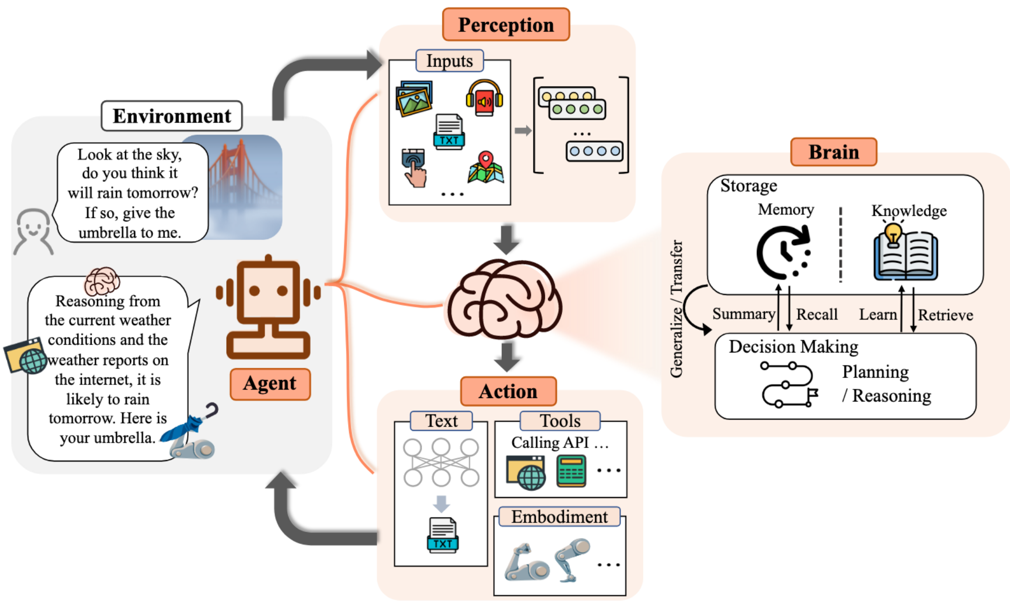

Figure 8: the conceptual framework of LLM-based agents with the components of perception, a brain, and action (Xi et al., 2023).

How agents handle tasks is possible because of multiple capabilities that are explained in the section below.

Planning

The ability of AI agents to plan is what sets them apart and consists of two phases, plan formulation and plan reflection (Xi et al., 2023). Within a single agent, planning is focused on breaking down a task into smaller tasks that are manageable. In a multi-agent system planning also involves deciding on what tasks flow to what agents, based on their role and expertise (Han et al., 2024).

Memory

A critical capability of agents' learning process is their memory. By storing decisions and precepted information in its memory, future decision-making will be based on prior decisions and outcomes of the action (Z. Zhang et al., 2024).

Tool usage

The capability to use external tools or resources enables agents to act effectively in different environments (Guo et al., 2024). The core ability of learning within LLMs enables agents to learn to use new tools. Prompting, as earlier described in the section about large language models, can also be applied to tooling by describing the tool functionalities in a prompt. This allows the agent to utilize new tools (Minaee et al., 2025; Xi et al., 2023). Using tools sets agents further apart from regular LLM chatbots, as they can take actions on their own by utilizing various tools.

Concluding on AI agents' relation to LLMs and MAS

By integrating LLMs into agents, we move from passive language models to active entities that can operate on their own. This progression provides the foundation for scaling to multi-agent systems, where multiple agents collaborate and execute tasks. The following section will provide insight in the background of multi-agent systems.

3.3 From single agents to multi-agent systems

While individual AI agents can effectively automate specific tasks, complex business processes often require coordinated contributions from multiple roles with their own expertise. The business case of drafting high-quality RFP responses, which is the focus of this study, is an example of a situation that involves gathering and integrating inputs from various domains. In such contexts, the abilities of a single agent are not enough. To facilitate support for complex tasks, a system of collaborating agents is needed to manage and execute tasks, share knowledge, and orchestrate communication. This section describes the different aspects that are relevant for developing multi-agent systems.

3.3.1 Foundation of MAS

A multi-agent system (MAS) is a system in which multiple autonomous agents interact with an environment to achieve a common goal. These agents can have individual objectives, but they operate as part of a larger mission, that is often requested from the user. In the context of this study from this point forward, agents always refer to (AI) agents that are based on LLMs.

3.3.2 MAS architecture

MAS architectures and topologies are at the core of how multi-agent systems interact and collaborate. In general, MAS architectures can be categorized in centralized and decentralized systems in terms of communication structure or task allocation (Guo et al., 2024; X. Li et al., 2024; Dorri et al., 2018).

In centralized systems there is often a central agent or agent group that coordinate and orchestrate interactions among agents (X. Li et al., 2024). Often a specific 'facilitator', as a central agent, acts as an intermediary between the user and the system, that routes requests to the right agents (Dorri et al., 2018). Although this can be efficient for coordination, it can create a bottleneck. When the facilitator fails, the entire MAS will fail to function properly (Dorri et al., 2018).

Decentralized systems are built as agent networks without central point, meaning agents can communicate directly to one another and do not rely on one specific agent. This makes decentralized systems more robust, as they are not dependent on one agent (Gutowska, 2024). However, coordination throughout the system can be more complex and there is not one common knowledge base. Potentially, this could lead to inconsistency within the system (Dorri et al., 2018). The decision between architectures will always be a trade-off.

Concluding, centralized systems benefit from the amount of information the central agent receives and can use but are vulnerable to failure as they depend on one agent. Decentralized systems are more robust, but can face challenges in coordination and finding the right agents to retrieve information from (Gutowska, 2024).

Deciding on what architecture to use involves a trade-off and it is important to test what works best for your use specific use case.

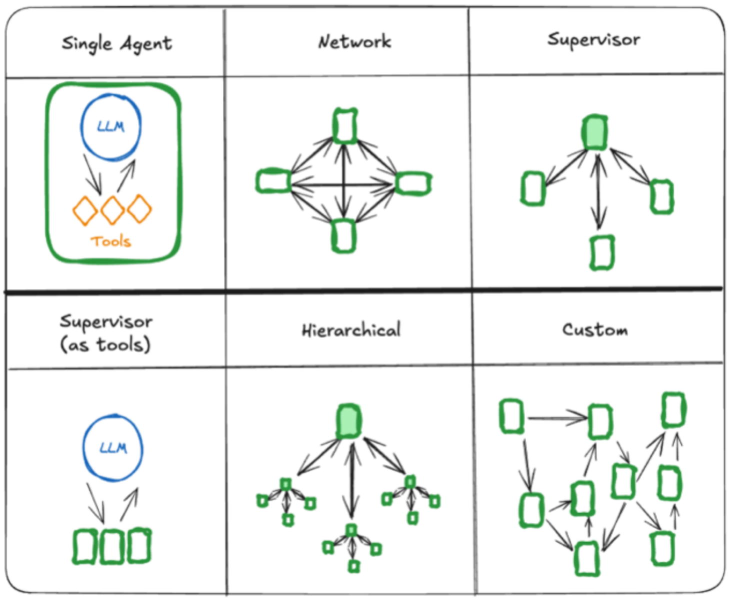

Figure 9: Different options of MAS architectures (LangChain, Inc., z.d.).

3.3.3 Agent communication

The architecture of the multi-agent system defines the overall structure for interaction. The methods of communication between agents determine how effectively they can collaborate and share knowledge and is named ‘message delivery’ (X. Li et al., 2024).

Two main approaches for communications are message passing and the blackboard method, also known as a shared message pool. In message passing, agents have direct communication lines and talk directly to each other through a message. The alternative, blackboard or shared message pool, relies on a central shared data point. Agents can access and retrieve information from the shared data point and write directly to the data point (Dorri et al., 2018).

An important aspect of shared knowledge in communication is memory synchronization among agents. When agents collect information from a knowledge base, this knowledge base should be consistent and synchronized to prevent mistakes in communication and decision making with outdated data information (Z. Zhang et al., 2024).

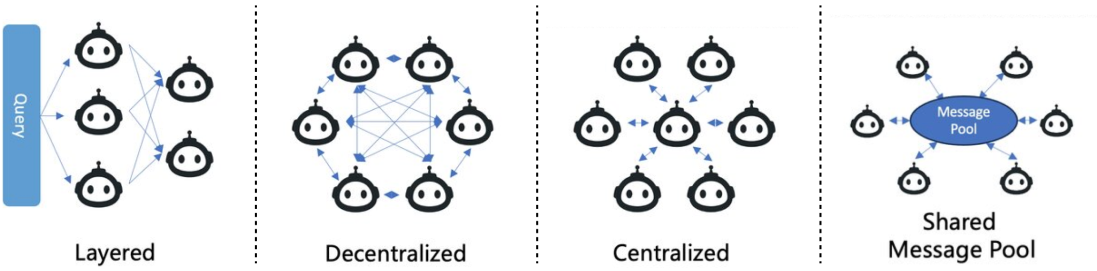

Figure 10: Communication structures with message passing or a shared message pool between agents from (Guo et al., 2024).

3.3.4 Agent profiling

In a MAS the idea is to create diversity in the system to enhance performance. Mirroring the principle of teamwork in human teams, individuals contribute in their own way with their expertise on a specific assigned task. In MAS the tasks are divided over different specialized agents to improve the efficiency and effectiveness of the entire system (Talebirad & Nadiri, 2023). Creating a specific role for an individual agent is called ‘agent profiling’. This profiling is used to give the entity a personalized style, making them able to accomplish specific tasks well (X. Li et al., 2024).

Agent profiling can be categorized in three main methods: pre-defined (handcrafted), LLM-generated, or dataset aligned (Guo et al., 2024; L. Wang et al., 2024). Pre-defined agent profiles are defined by the system designer explicitly. (Guo et al., 2024; L. Wang et al., 2024) split learning-based profiling methods into LLM-generated profiling and data-derived profiling. The LLM-generated profiling method is based on a LLM that defines and generates the right profile for the task. The data-derived method constructs the agent profile based on pre-existing datasets (Guo et al., 2024; L. Wang et al., 2024).

The profiling of an agent is typically done by providing a prompt in natural language to the agent and this is in practice referred to as the system prompt (L. Wang et al., 2024). The system prompts can be written in any way, however, practice teaches that using variations on CoT prompting can be effective in applying desired profile traits (Jiang et al., 2023). The prompts serve as guidelines for LLM based agents and allows the LLM to better understand the context and behavior that would be expected of that role (Xi et al., 2023).

By clearly defining agent roles, MAS developers can ensure that agents work in complementary ways, reducing redundancy, and improving task allocation efficiently. Profiling also enables more predictable behavior, making the system more reliable.

3.3.5 Improved reasoning in MAS

While individual agents can achieve a certain degree of advanced reasoning based on their underlying LLM and prompting strategies, they remain limited in their ability to critically self-reflect on their own outputs. Once an agent has confidence in its own answer, it struggles to reconsider its approach or answer, restricting it from coming up with creative answers and ideas (Liang et al., 2023).

Multi-agent systems can help address this limitation by enabling collaborative reasoning. In collaborative reasoning multiple agents communicate, as Du et al. (2023) states, or introduce a voting system (J. Li et al., 2024). This principle relies on diverse perspectives that are generated by different agents, which allows a multi-agent system to overcome biases or reasoning limitations that can occur in a single agent.

Research by Du et al. (2023) explains how multi-agent debate frameworks can improve the reasoning capabilities of LLMs. In these frameworks, multiple agents engage in a debate where they both present arguments and opinions to each other. By presenting multiple perspectives, the agents show less likelihood to acknowledge their own answer as correct and prevents the agent from becoming trapped in their own reasoning path (Du et al., 2023).

While MAS approaches do show promising aspects for improving reasoning, several challenges remain. The coordination and interaction capabilities between agents become more complex as the system scales. The quality of reasoning is influenced by the information transfer between agents. When developing a multi-agent system, the orchestration strategy and communication patterns should therefore be mapped and structured well to facilitate good knowledge transfer and reasoning of agents.

3.3.6 Evaluation of AI generated text and Multi-Agent Systems

Due to the complexity and collaborative nature of multi-agent systems, evaluating the output of these systems comes with a variety of challenges over the evaluation of single agent assessments. While individual agents can be evaluated directly on quality metrics of the output, MAS require evaluation approaches that capture both individual agent performance and collaborative performance elements such as orchestration success and task hand-off. Effectively evaluating of a MAS can require both quantitative metrics and qualitative assessment to define the performance of the system.

Task completion rate is one of the most common quantitative metrics, in which the percentage of successful executions within defined requirements are measured (Xi et al., 2023). However, effectiveness and usability of task execution also depend on the quality of the outputs produced, which require qualitative evaluation methods.

One approach for evaluating agent-generated content, is the LLM-as-a-judge methodology. As described by (Zheng et al., 2023), LLM-as-a-judge offers two main benefits: scalability and explainability. By reducing the need for continuous human involvement, it enables fast iterations and scalability in benchmarking. Next to this, the LLM does not only score the outcomes quality, but provides explanations as well. Zheng et al. (2023) concludes that strong models can achieve agreement rates of over 80%, which is on the same level as human experts. Building on method of LLM-as-a-judge, ChatEval provides a multi-agent system for evaluating content, using collaborative reasoning (which is one of the advantages of MAS) to produce outputs that are more aligned with human preferences (Chan et al., 2023). Another variant, G-EVAL by Liu et al. (2023), incorporates chain-of-though reasoning in its evaluation process and outperforms many other evaluation frameworks, making it a more scalable alternative for human evaluation.

Despite the advantages of LLM-as-a-judge as evaluation method, human expert evaluation remains essential for considering contextual nuances and domain specific quality criteria. Human annotation, as described by (L. Wang et al., 2024), is a method where evaluators review and score outputs according to domain-specific criteria. Including human experts provides contextual nuances that help to ensure the system meets the expectations of intended users. Combining human annotation with AI-based evaluation techniques becomes important for complex application where standard metrics will not capture the success of the output.

Industry practitioners show examples of evaluating agent and MAS systems through comparative analysis. LangChain explains how their customer, Minimal, tests different prompts by running side-by-side comparisons in different variants (zero-shot vs. few-shot vs. chain-of-thought prompts) (LangChain, 2025). Alternatively, Palantir recommends designing well formulated evaluation plans (Palantir, 2025). Palantir provides a best practice when working with LLM-as-a-judge, where they recommend evaluating the outcome as pass or fail rather than asking for a score. The reason is that it provides the agentic system with actionable feedback to either continue its task or to revise it, rather than thinking about the score of the outcome.

Overall, evaluations serve as quality control of agentic systems and as a feedback tool, ensuring that AI and MAS implementations deliver reliable, accurate, and appropriate results.

Table 1: Comparison of different output evaluation concepts

	Advantages (+)	Disadvantages (-)
LLM-as-a-judge	• Scalable • Fast iterations • Provide a measurable score	• Depend on the quality of the evaluating LLM • May miss contextual nuances • Risks of biases from the LLM
Human annotation	• Capture detailed evaluations • Ensures alignment with human expectations • Interpretability beyond standard metrics	• Time-consuming • Limited scalability • Subjective to annotator variability and preferences

3.3.7 MAS case study examples

There are existing case studies that show how multi-agent systems have successfully created value or achieved their intended goals for real-world business cases. To demonstrate how systems can be designed, this section describes two examples.

ServiceNow (an enterprise SaaS company) has created a multi-agent system in collaboration with Microsoft by using their Semantic Kernel development kit (Lagerkrans-Pandey, 2025). They create a system for tackling P1 (highest priority) incident management with their customers. ServiceNow identified challenges during crisis management, where long and intense communication with their clients lacked documentation.

To address this issue, they created a proof of concept with multiple agents who were assigned to automatically record and transcribe discussions, while other agents interpreted the discussions live. With an orchestrator agent managing the sub-agents, they automatically generate incident reports, reduce the resolutions times and improve their knowledge capture. They concluded with their proof of concept that AI agents can be valuable partners who can work alongside humans.

Avid Solutions created an agentic AI solution with IBM Watsonx Orchestrate (*Boost HR Efficiency with IBM Watsonx*, n.d.). They focused on automating repetitive HR tasks to improve efficiency and accuracy with the aim to boost employee satisfaction. Specifically, they addressed tasks as sending initial recruitment messages, sending mails and creating job descriptions. The created agents could control multiple platforms (LinkedIn, Workday, Outlook). As a result, Avid Solutions saw an increase of 10% in customer satisfaction and reduced project management errors by 10% and new customer onboarding time by 25%.

These examples show the potential of integrating MAS in business processes in different sectors. It remains important to carefully consider what tasks will be handled by a MAS and how humans stay involved, to integrate these solutions in the workplace.

3.3.8 Human-agent interaction

An essential requirement for implementing AI agents in the workplace is to consider the interaction between humans and AI agents. Inconsistency and unpredictable behavior can lead to abandonment of AI systems and should therefore be tackled by design. Humans play an essential role in guiding agents, especially in chat-based applications where their success is heavily dependent on human input of the conversation. Creating good inputs for an agent to can be challenging and time-consuming for the human user (G. Li et al., 2023).

While AI technology is improving rapidly, challenges remain in human-agent interaction. There are two main sources of challenges within designing with AI (Yang et al., 2020). First, there is often uncertainty for humans surrounding the actual capabilities of AI. Second, the output of AI can vary in complexity, making it harder to interpret the results,

Trust can form a critical barrier for adopting AI and is critical to tackle by design (Hwang, 2022; McKinsey & Company, 2025; *People + AI Guidebook*, n.d.; Schwartz et al., 2023; D. Wang et al., 2020). Developing trust can be a challenge and relies on multiple factors. (Schwartz et al., 2023) states two different categories of trust that are relevant for AI agents, cognitive trust and emotional trust. Cognitive trust is based on how one sees an individual in terms of reliability, competence and integrity. Emotional trust is based on emotional bonds between individuals and is more focused on a shared connection and shared values, beliefs, and goals. The main considerations and factors that are considered are reliability, openness, tangibility, immediacy behavior, and task characteristics (Schwartz et al., 2023).

Specific to agentic AI, human-agent interaction can take on different forms. (Xi et al., 2023) names two types: unequal interaction (instructor-executor paradigm), is a form of collaboration where a human hands instructions to the agent, that will perform the given task. Equal interaction (equal partnership paradigm), is performed when the agents is placed on the same level as humans, creating an equal collaboration in terms of decision making and communication. Following the statement of (Xi et al., 2023) that the ultimate goal of agents is not to become more powerful, but to better equip humans with agents, most use cases will benefit from the instructor-executor paradigm.

Effectively implementing multi-agent systems for a business process requires consideration of human-AI collaboration principles to improve user experience and adoption. While Cila (2022) defines 11 design considerations for human-agent collaboration, the most widely adopted framework comes from Microsoft researchers, who developed and validated 18 AI design guidelines for different phases of user interaction (Amershi et al., 2019). These guidelines invoke designers and developers to create AI systems that facilitate user understanding, trust, and effective engagement. These guidelines are utilized later in this research through a workshop to identify critical guidelines specific to the context of assisting in RFP drafting.

3.4 Structured procurement procedures (Request for Proposal)

In structured procurement the term 'tender' is used for procedures where a company asks for a service or product solution for a current need. The other party involved is the bidder (also referred to as supplier), with the goal to win the contract for supplying the product or service. Contracts that require to be handed out through a tender are of a large size financially. The Dutch government procures around €73 billion worth of work every year (PIANOo, n.d.), and in 2024 the IT services sector alone published tenders worth €1,5 billion (Opentender Netherlands, n.d.).

An RFP is a formal document with a description of the context and a set of questions that give an insight in how the bidder would handle to a certain situation or solution offer. The deliverables must often adhere to a very strict formatting style to enable the contracting party to compare offers from different bidders. Next to RFPs, tenders can request other types of information. The most used other requests are the Request for Information (RFI) or a Request for Quotation (RFQ). An RFI is often a first request from the contracting party to get familiar with potential solutions and offerings from the industry. The RFQ is solely focused on receiving pricing details from different bidders for the desired product or service. While an RFP can be similar to an RFI, they typically require more detail and requirements (Rajbhoj et al., 2019). The process of preparing responses to either an RFI or RFP remain similar.

As an RFP is the most detailed request requires the most work and expertise, the scope of this project is focused on RFP response drafting.

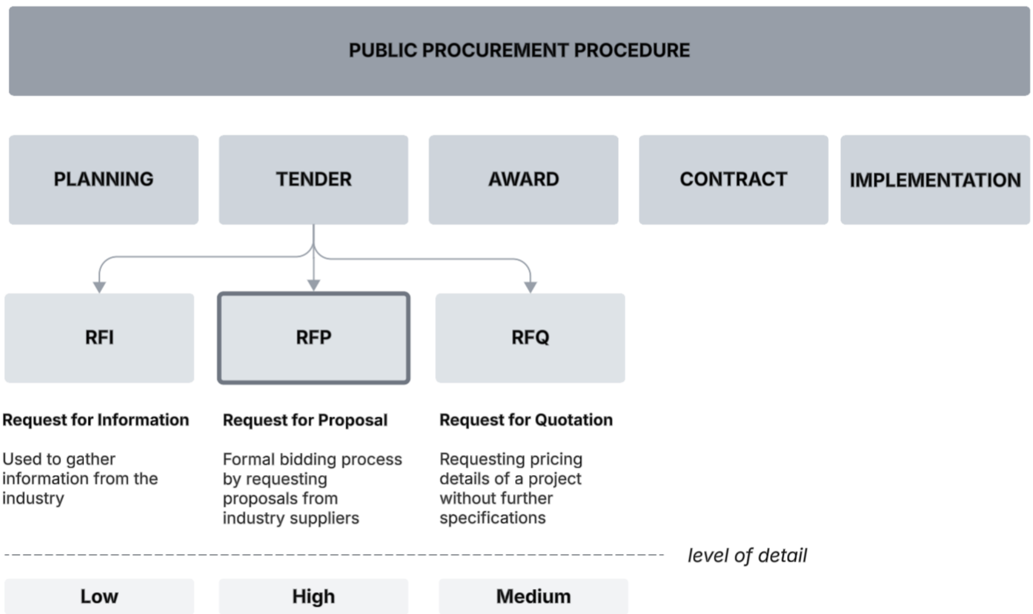

Figure 11: Overview of the procurement procedure and tender phase.

Although few studies address the proposal development process for RFPs from the supplier's perspective, some challenges are pointed out by (Paech et al., 2012). Suppliers face internal challenges in inefficient communication between different roles and often do not identify experts they need to consult. Next to this, there is often no reliable and structured basis for knowledge that could be reused from other responses or projects and RFPs can lack details in their specifications, making it hard to fully understand what work is needed (Lauesen & Vium, 2005).

Organizations have identified the potential in supporting suppliers during the process of proposal development, although organizations tend to provide process management solutions over content development solutions. Recent developments do show AI embedded features in RFP management solutions, which claim to draft secure and accurate proposal content based on the suppliers own data and knowledge (Loopio, n.d.). However, there are no reviewed reports or studies available on their effectiveness and quality.

3.5 Takeaways from literature

Opportunities of LLMs, agents, and MAS

The literature demonstrates that LLMs, AI agents, and MAS have strong potential for automating complex and collaborative business processes. LLMs are the powerhouses that enable reasoning, content generation, and interaction capabilities, but they show limitations with hallucinations, biases (more common in single agents), and are constraint by context windows. Embedding LLMs into agents allows them to act, plan, and learn. Combining multiple agents in MAS orchestrations make it possible to distribute tasks, define different roles and perspectives, and improve reasoning through collaboration.

Challenges and research gaps for MAS

At the same, challenges remain for MAS in evaluating their outcome and creating reliable systems. To evaluate the system, approaches can be found in automated evaluations that suit for scalable situation in the form of LLM-as-a-judge or manual efforts in the form of human annotation to ensure contextual accuracy. As MAS are harder to evaluate then single agents, with the addition of task distribution and knowledge chains, evaluating MAS remains challenging. Next to this, defining human-agent interaction patterns can be challenging and is a critical factor for creating trust and transparency, which are necessary for user adoption. While case studies show promising applications of MAS in business contexts, literature remains theoretical and rarely explores end-to-end integration of MAS in a business context and how this can be created.

Implications for this research

The review highlights opportunities and gaps that inform this thesis. MAS have the potential to provide meaningful support in knowledge-intensive processes, but their success depends on careful orchestration, reliable agent behavior, and their usability by human users. This sets the foundation for the thesis and build on the research gap to explore how we can effectively design and apply MAS to the proposal development process at Schuberg Philis.

Table 2: Choices and challenges in implementing MAS.

Area	Choices	Challenges
Orchestration	Centralized vs. decentralized vs. hybrid structures	Centralized systems create bottlenecks if the facilitator fails (Dorri et al., 2018). Decentralized systems are more robust but face coordination difficulties (Dorri et al., 2018; Gutowska, 2024).
Agent profiling	Pre-defined (handcrafted), LLM-generated, or dataset-derived profiling	Risk of unclear responsibilities and redundancy in the system. Requires careful definitions to create complementary roles (Guo et al., 2024; L. Wang et al., 2024)
Communication method	Message passing vs. shared message pool (blackboard)	Communication between agents is essential for sharing knowledge (Xi et al., 2023). Inconsistent communication can lead to issues with an asynchronous system (Z. Zhang et al., 2024).
Reasoning approach	Single-agent reasoning vs. collaborative reasoning (debate or voting)	Single agents can get stuck in narrow reasoning paths (Liang et al., 2023). MAS debate frameworks improve reasoning but increase orchestration complexity (Du et al., 2023; J. Li et al., 2024).
Evaluation methods	Automated (with LLM-as-a-judge variants) vs. human annotation vs. hybrid	LLM-based evaluations are scalable but can miss contextual nuances (Zheng et al., 2023). Human annotation Is accurate but time-consuming (L. Wang et al., 2024)
Human-agent interaction patterns	Instructor-executor vs. equal partnership	Adoption depends on trust, transparency, and clear collaboration structures and can vary per use case (Hwang, 2022; Schwartz et al., 2023)

This page intentionally left blank

4 Contextual research

In this chapter, we examine where in the proposal development process AI can offer the most value (SQ1) and which MAS features would be needed to make the system specific to Schuberg Philis' way of working (SQ2).

We combine semi-structured interviews and a training observation to create an understanding of what the RFP process entails at Schuberg Philis and identify what opportunities for improvement are recognized by colleagues. The findings from both activities provide a foundation for knowledge about the business process that can be used for developing a desirable and viable MAS. This chapter presents the methods for collecting and analyzing information, the findings from the activities, and their implications on developing a proof of concept.

4.1 Methods for exploring the RFP process

The methods for gaining an understanding of the process of drafting proposals for an RFP are described in this section. Two complementary methods are used, semi-structured interviews to learn from individual approaches and desires, and a training observation to understand how a good quality proposal is formed based on the best practices at Schuberg Philis.

4.1.1 Semi-structured interviews

Knowledge about drafting RFP documents at Schuberg Philis was acquired by interviewing different stakeholders within the company. The aim of the interviews was to gain insight in the current practice and finding emerging themes. A semi-structured approach was chosen to ensure the interviews address the same topics, whilst offering the freedom to elaborate more or deviate from the topic where the participant wants. The structure of a semi-structured interview allows the interviewer to improvise and adjust with follow-up questions and creating enabling a more cohesive dialogue between the interviewer and the participant (Kallio et al., 2016).

This qualitative research reveals challenges RFP teams face and are used to shape design opportunities for a MAS. The interviews validate the design opportunities address actual desires, rather than assumptions on user needs.

Process of interviewing

Participants (n=7) were selected on their roles and involvement in the RFP process. As the roles of business consultant and sales director are involved in all proposal processes and have the most experience in proposal writing at Schuberg Philis, these roles were targeted for the interview process. Next to this, participants must come from different RFP experiences to ensure the interviews are from varied examples. Recruitment was conducted by reaching out to colleagues through a Slack message, asking if they were interested in participating in an interview on the RFP process. The interview guide contains open-ended questions on five different topics: the roles involved in the process, the drafting process, collaboration and communication, pain points, and opportunities. The interview structure can be found in appendix A. We anticipated 6-10 interviews would be sufficient to perform the thematic analysis and to understand the participants challenges, what the process looks like, and where opportunities lay. We stopped at seven participants as the last interviews produced no new insights and were aware of the opportunity to talk to other roles if we deemed necessary after thematic analysis.

Table 3: Overview of participant roles.

Role	Number (#)
Business Consultant	3
Sales Director	3
Strategic Tender Specialist	1

Five out of seven interviews were taken in-person, as this was preferred to have a more natural conversation. Two interviews were held online by using Microsoft Teams as the participants worked from home on the scheduled date. The interviews were held in Dutch and lasted between 30-55 minutes. Audio records were made to ensure no information was lost and to enable transcribing the interviews afterwards. Transcripts were made with Microsoft Word with speaker recognition. To enhance the quality of the transcripts, the produced transcript was put in ChatSBP (a secure company version of ChatGPT) with a prompt to improve the text.

Participants were carefully instructed about the goal of the research and their participation via informed consent. To securely handle the data, secure storage environments were used, and the handling of data was approved by the Human Research Ethics Committee (HREC) of Delft University of Technology.

Analyzing the interviews

Since the aim of the interviews was to reveal insights on the implicit knowledge and varied challenges within the RFP process at Schuberg Philis, addressing SQ2, a thematic analysis approach was selected as an appropriate analytical method. Thematic analysis, as described by (Braun & Clarke, 2006) and (Byrne, 2022), is conducted to identify themes and patterns in the data. An inductive approach to coding data is chosen. Coding is the act of labelling or tagging text to assign meaning to it (Miles & Huberman, 1994). Inductive coding relies on finding relevant information in the data (transcripts) and applying a code to the information with a bottom-up approach. Afterwards, the coded data was evaluated, and themes and patterns were defined. Some of the data got recoded iteratively to improve accuracy of the coding.

With a thematic analysis of six steps, the information from the transcripts was broken down, analyzed, and interpreted to develop themes of the information (Braun & Clarke, 2006; Byrne, 2022; Naeem et al., 2023). The process of inductive coding and thematic analysis was performed in CAQDAS software called ATLAS.ti (ATLAS.ti, n.d.).

4.1.2 Training observation: “Successful Bidding”

To gain insight in the RFP response process best practices, the research included the observation of a training session. The training session, “Successful Bidding”, was held by a subject matter expert at SBP with the goal of educating colleagues about the most effective practices in drafting RFPs, from planning to content creation to evaluation of deliverables.

The training “Successful Bidding” provided essential information on how Schuberg Philis creates high quality proposals and provided clarity on the procedure. The explicit knowledge on writing proposals is directly applied during the prototyping sessions to create valuable outputs from the MAS. Information about the procedure of RFPs is used to clarify the process outline and decide on the scope for prototyping.

Observing participants during the training

With the aim to understand how a good quality proposal should be written, explicit knowledge from an expert was captured through participant observation. Participant observation can be used to help answer research questions, to gain insight In current practice and test or generate theories and hypotheses (Kawulich, 2005). It has become increasingly popular to have active engagement with stakeholders throughout the research process to share in the decision making process and to ensure the research is relevant to their lives (Vaughn & Jacquez, 2020).

In the training session, the role of participant-as-observer was adopted. All participants are aware the reason of the researcher's presence and the goal of the research. An active role in the training activities is taken to learn from formal instructions and informal discussions about the RFP process. Observations were documented through field notes and reviewed directly after the session.

Limitations and validation

A disadvantage of the role of participant-as-observer can be the confidentiality of participants towards the researcher (Kawulich, 2005) and the depth of information they will share. In the setting of a training for best practices by the facilitator, this limitation is minimal to this research because the discussed topics were informative and educational rather than personal or sensitive.

Another disadvantage was that relevant findings could be biased as they originate from a single session and are not validated by controlling over multiple sessions or other subject matter experts. However, the findings were discussed and approved by the expert at hand to prevent any misalignment from personal interpretation of the findings.

Collecting data and ethics

The facilitator gave permission for using the training as part of the research and all other participants were verbally informed of the research. No personal identifiable information was collected from any of the participants to maintain their privacy.

4.2 Findings from interviews and training observation

The interviews provide an overview of challenges and opportunities for integrating AI in the RFP process, combined with personal remarks on experiences with drafting proposals. The training observation "Successful Bidding" provided insights in the best practices at Schuberg Philis and a clear overview of the process phases. Together they show that currently proposal writers struggle with finding relevant information and answering the proposal questions directly. Schuberg Philis employees show willingness to integrate AI, and a set of current best practices can be used for reference on how to write strong proposal content.

The outcome of this qualitative research sets the foundation for defining a scope for the proof of concept and provides insight in the opportunities for integrating AI in the form of a MAS. At the end of the chapter a synthesized overview is given for reference.

4.2.1 Findings from semi-structured interviews

The interviews with seven RFP practitioners at Schuberg Philis revealed two main themes that shape the current challenges and opportunities within proposal development.

First, the current process presents challenges in time-consuming drafting sessions due to fragmented knowledge, dependencies on colleagues, and long rewriting to create a coherent synthesized proposal from different section authors. Second, participant show clear readiness and openness to embed AI in their workflow. They suggest using AI to speed up making a first draft, asking AI for improvements and checking if their work is aligned with the clients requirements. Below the findings are presented with quotes from participants of the interviews to substantiate the interpreted findings.

Challenges in the current RFP process

The current process at Schuberg Philis demonstrates its challenges that are shared amongst colleagues. The challenges show the complexity of drafting proposals with dependencies, time-consuming tasks, and scattered knowledge. Each underlying theme is explained in the section below, substantiated with illustrative quotes from the interviews.

Scattered knowledge and misalignment slows down proposal development

There are multiple challenges in the RFP process in its current state. Primarily the lack of a central database with knowledge of past projects and standardized approaches to drafting answers, are highlighted in multiple interviews. What ends up happening, is that commonly asked questions in RFP's are answered created over again, or colleagues call around to understand if similar questions and answers exist from prior RFPs.

"A major challenge is that we don't have a database with standard answers. It always takes a lot of time to create a foundation. It's more efficient to have a base from which you can deviate, rather than starting over without a reference." – Participant 1

Next to managing data from previous work, managing requirements over the process timeline is mentioned. Team members of RFP teams struggle with managing new information from question rounds, as one participant noted ***"These quickly result in 100+ questions … it's hard to stay on top of everything that is being requested" (P5)***, and the fact that asking questions should be done strategically as the questions and answers will be public for all applicants. Another participant mentions ***"Contradictions and ambiguities in the tender are challenging" (P3)*** while referring to the question rounds.

Furthermore, content misaligned is a reoccurring issue. Participants mention that the "answer does not match the question" and recognize that regularly there is an emphasis on branding and corporate story over clarity. This results in a disconnect between the answers provided by SBP on the client's requirements.

"At SBP, I currently need to make a big effort to form a cohesive answer. Everyone writes beautiful stories about SBP and the DNA and 'the story,' but they simply forget to answer the question." – Participant 4

Multidisciplinary teams create dependencies

The analysis shows that RFP responses are always formed in specific collaborative teams. These teams typically involve 4-12 people, depending on the size of the project and required expertise, as a participant mentions ***"if it's large, you sometimes need 10 or 12 people, especially when many specialists are involved" (P1)***. A Bid & Proposal Manager is experienced in RFP processes and acts as an experienced guide, with other team members (Customer Director Sales, Tech Leads, Business Consultants) contributing to their domain expertise.

Considering that very specific knowledge can be required to effectively draft a response to a clients question about business processes, or technical solutions, multidisciplinary teams with different subject matter experts (SME's) are required.

Next to this, the analysis showed that RFP teams often include a member from an external party that is specialized in drafting RFP's or RFP strategies, or to acquire outside feedback and opinions. The use of an external party might not seem obvious, however, knowledge and methodologies for RFP processes can be used effectively across different sectors.

Struggle with consistent adoption of best practices

An important finding is that there are already best practices accessible for drafting RFP responses. These best practices include "building blocks, a roadmap, writing (style) rules", and recommends to consistently evaluate answers against the given assessment matrices.

The emphasis of the best practices is to focus on content quality, through consistent writing styles and direct answering of questions. Examples that are mentioned are "avoiding company story", "avoiding empty words/jargon", and to "repeat the question and respond briefly".

"It's not about telling a beautiful story or having a good-looking document. What matters is that you provide a clear answer to the questions stated in the request." – Participant 5

Next to these recommendations, the best practices seemingly influence the colleagues, as mentioned by one participant: ***"The most important thing I've learned is that our language needs to be very assertive" (P1)***. Assertiveness and confidence are described as key characteristics to reflect SBP's culture and to come over as a viable partner for the project. This changes their way of writing: ***"We tend to write cautiously, for example, 'We think this is the best solution,' but now we just say, 'This is the best solution.' That makes a big difference in persuasion." (P1)***.

The analysis revealed, however, that participants have concerns if the best practices will be followed in practice, across different teams and projects.

"Because it's written by different people, you need to ensure consistency in the text and format." – Participant 6

Importance of complying to quality standards

Participants mention and identify challenges in managing content quality. Different question types are recognized, that require different approaches. Amongst others question types in the form of pricing, technical design, and references are pointed out, as well as unique formats like a hackathon. Each question type requires a different expertise and response structure. This adds to the complexity of creating a standard format.

Furthermore, documents must often be delivered in a strict format to the client as one said: ***"The selection guidelines specify exactly what is required of the applicants." (P5)*** and another participant further clarifies, ***"The client often dictates the format: Excel, Word, PowerPoint." (P3)***. Having to adhere to all these format requirements in combination with the fact that all contracting parties can all have different approaches, makes it hard to directly reuse documents. Each proposal will need significant adjustments to be compliant with the client's requirements.

Controlling the quality can be done in two ways: by evaluating the response against assessment criteria or matrices, or by applying the best practices and evaluating particularly on the readability of the response and compliance with the formatting requirements. Participants highlight the importance of always managing the quality of content: ***"assumptions are often wrong because they are based on your own perspective. That costs time and money, and you make yourself more expensive or slower in responding." (P7)*** and reusing materials should be done with caution as stated by a participant:

"Be careful when automating with which answers from previous proposals you want to include. In the past, there were also many 'bad' answers given. If you use bad examples, you'll likely get bad new answers." – Participant 4

Rewriting of proposals remains time consuming

The participants explained how RFP processes can vary in timeline and duration. The entire process from start to finish for a tender can take a long time, as a participant says: ***"I think the whole process took about a year, including everything from the RFI to the cooling-down period." (P2).*** However, the time it takes for a proposal team to answer to an RFP usually was described as response windows of ***"between 4 and 6 weeks"*** to draft and discuss the proposal (P1). The selection procedure is described as a funnel, ***"It often goes from 20 to 5, to 3, to one, so to speak. That's kind of the funnel of potential suppliers." (P7)*** with a high number of applicants at the start and narrowing it down to eventually one supplier.

The most critical phase of the process is recognized as evaluating and rewriting the proposal. This process can be time-consuming and demanding on the team, but reviewing and rewriting remains critical to deliver top quality proposals and to have a chance of winning the project. A constraint in the process is found in some sequential dependencies, where a certain section must be finished before the other section can be drafted. This is described by one participant: ***"There is often a sequence where you answer the technical questions first, so you get an idea of the solution you're going to deliver. Only then can you assign a price to it." (P1)*** This can limit working on different sections in parallel to each other and requires coordination between different expertise areas and proposal team members.

"It's always a struggle to formulate a good answer that meets the contracting party's expectations. They are looking for specific things that need to come back in your text. You are trying to figure out exactly what the contracting party is looking for and how to weave that into the text." – Participant 6

Current AI use demonstrates potential and desire for more

Participants show an interest in implementing AI in the RFP context or already use AI to a certain extend. Current AI usage focuses on improving efficiency by ***"creating a first draft" (P1 & P6)*** and ***"to provide us with a checklist of elements we need to consider" (P1)***, and to compress lengthy text with

AI. Furthermore, currently quality of responses is improved by “using AI for suggesting improvement”, “use AI to check for mistakes”, and to “rewrite vague text with AI to be more precise”.

"AI is definitely the future and already adds value. The problem lies in its correct application. AI can generate different answers that all sound okay, but you might wonder if they are really good enough. It is primarily useful as a tool for creating a first draft." – Participant 6

"Nowadays, I often put the entire context into ChatGPT. The question goes in, and I check the answer it gives me against my intuition and what ChatGPT has come up with. This way, you arrive at more ideas much faster." – Participant 7

"If you wanted a kind of perfect knowledge world, you'd put all the RFPs into one pool across our different business units and mine that. I think that would make a huge difference." – Participant 7

From the analysis three main potential goals for integrating AI were addressed. Possibilities for analytical capabilities, by analyzing competitors or the customer. Improving efficiency in the process by reducing the time needed to draft basic sections or in managing information from question rounds. Finally, knowledge management by leveraging prior RFP responses and reference material from other clients. Participants especially showed interest in using AI's potential to “check the relevancy of the answer (to the question)”.

Implications and opportunities for multi-agent systems for the RFP process

1. **Knowledge management potential:** Valuable expertise is present but is fragmented across different experts and past projects. The creates inefficiencies where teams repeatedly recreate content and struggle to maintain coherence across different proposal teams. Company documents with internal information and a database with prior RFP content can offer an outcome to more efficient and consistent proposal drafting.
2. **AI readiness:** Current usage patterns of AI show the willingness to further adopt AI in the process of drafting proposals. The conversations reveal a preference for AI augmentation rather than full automation.
3. **Quality checking capabilities:** The system should address quality control of drafted proposals to some extent. Participants mention checking the alignment between the content and question and checking compliance against the RFP requirements. The MAS could introduce compliance checks for both situations.
4. **Diverse process:** As different people work on different sections of the proposal, it is important that the MAS can offer support on various question types. Therefore, the MAS should understand what section the user is working on and how it can offer generalized support for different expertise areas.
5. **Strict RFP requirements:** Adhering to strict requirements is essential in proposal documents, as non-compliant content is neglected during assessment. →

These findings are later used to define a relevant proof of concept for the prototyping phase. The next section further specifies what the best practices at SBP entail, providing findings and implications from participant observation during an internal training session.

4.2.2 Findings from training "Successful Bidding"

To complement the findings from the interviews and gain a deeper understanding into the established best practices at Schuberg Philis for RFP development, I participated in a training session called "Successful Bidding". The session provided the opportunity to observe how best practices are formally transmitted to colleagues and to identify how these frameworks and methods provide actionable opportunities for the design of the multi-agent system.

The findings are structured in three main areas:

1. **Established best practices:** The three core frameworks that guide proposal writing at Schuberg Philis.
2. **Observation on challenges:** Identifying the gap between theory and practical implementation from observing participants during the session.
3. **Implication on MAS design:** How the training influences prototyping for a valuable proof of concept.

Established best practices

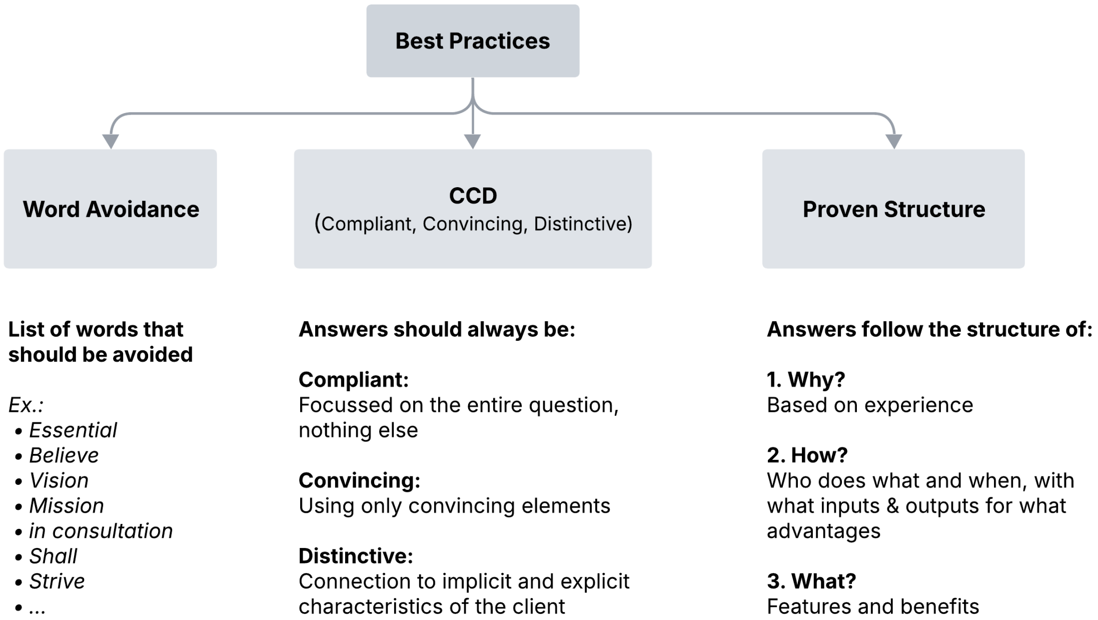

Figure 12: Established best practices at SBP for proposal writing.

Three specific practices from the training are highlighted for writing high quality proposals and are explained underneath: word avoidance, the CCD framework, and the why-how-what framework.

Word avoidance: A comprehensive list of words with "empty" corporate jargon or that imply vague commitments is provided. These words should be avoided in proposals, and more convincing alternatives should be used to create compelling arguments. The complete list of words can be found in appendix G (in the Banned Words agent system prompt).

CCD framework: It is required by the best practices to draft responses to be compliant, convincing, and distinctive. Compliant responses address the entire question and nothing else, ensuring no off-topic information is included in the proposal. Convincing elements should be used to substantiate your claim.

Finally, the proposal must create a connection to the client by connecting implicit client characteristics, based on a deep understanding of the client, or by connecting to explicit requirements from the client.

Ensuring compliant responses through this framework connects to findings from the interviews, where participants mention that answering the client's requests is the only thing that matters instead of writing beautiful corporate stories.

Proven structure (Why-How-What framework)

The established Why-How-What framework is used to standardize the structure of responses and to guide team members in content development. Offering this standard structure adds to coherence between different proposal sections and drafts and support the companies' style between different projects.

The structure should follow these sections:

Why: based on experience and past performance, the proposal should demonstrate credibility why the bidder can provide the proposed solution.
How: the proposal should provide clarity on the process. It is important to inform the contracting party about details on who will work on what, and when. Clear inputs and outputs should be defined and communicated.
What: proposing the deliverables and their value to the client should be part of the written text.

Observations on challenges in proposal drafting

Next to the provided best practices, participants addressed some concerns and comments on how these practices should be integrated. Some of the discussions are outlined below.

Concerns about consistent application: Participants raise concerns about consistent application across different teams and projects of the best practices, mentioning that it is not clear who is responsible for checking and evaluating the proposals on these practices.

This offers the opportunity for the design of a multi-agent system to function as the controlling aspect of incorporating the practices and quality of proposals.

Availability of high-quality material: Participants mention the availability of high-quality examples from prior projects but mention there is no clear database that contains these documents.

Observation confirms SBP tendency to overcomplicate: During the training participants work out an example question from an RFP: "How do you ensure the timeliness of a data management system?". Even this short example is tackled by participants in varying complicated ways and highlight the importance of directly answering the question to remain compliant.

To illustrate how colleagues at SBP tend to overcomplicate answers is an example is shown. This example shows vague promises with good intentions, however, it offers no meaning.

"We are obsessed with business outcomes and translating these into effective and secure technical and functional solutions is our unique differentiator." – Quote shown from past RFP submission

Implications and opportunities on multi-agent system design

1. **Best practice automation potential:** Three standardized frameworks (Word avoidance, CCD, and Why-How-What) are structured and clear, offering the possibility to implement in the multi-agent system.
2. **Governance of quality and consistency:** Participants reveal concerns on complying to the best practices, as there is not one person responsible for evaluating drafts on adhering to the best practices.
3. **Weakness at SBP:** The training validated that an important weakness in proposal drafting comes from overcomplicating answers, that are not solely focused on the question and its requirements. More clarity and assertiveness is needed to improve answers, as mentioned by the facilitator.

4.3 Concluding on contextual research

The interviews revealed challenges in the current proposal writing process, such as fragmented knowledge, inconsistent writing styles, and challenges in aligning the proposal responses to actual client questions. Next to this, participants show a clear openness to using AI for the RFP process, as most participants already use AI to an extent in the current situation. They highlight certain phases of the RFP process as most suitable for augmentation, such as drafting content, checking alignment, and reusing prior work.

The training session offered insight into the standardized best practices at Schuberg Philis and revealed participants challenges in drafting proposals. Frameworks for writing (CCD, Why-How-What, and banned words), provided opportunities for creating consistency in automation. Concerns were raised between familiarity of the best practices and consistently applying them across teams, suggesting a form of automation could fill this gap. Finally, overcomplicating answers is currently a weakness, that could be addressed through systematically checking if the answers align with the client question.

These findings form the foundation for scoping the proof of concept and prototyping phase of the research. During prototyping, see chapter 6.3, these findings are translated into a multi-agent system, bridging the gap between finding opportunities and realizing an agent system.

Table 4: Observation findings from interviews and training observation synthesized.

Phase	Observation	Finding	Design Implication
Interviews	Fragmented knowledge	No centralized or shared database to past answers leads to inefficiencies	Integrating internal knowledge can improve the current experience
	Openness to AI	AI is already applied during drafting and participants are positive towards further adoption	Augmentation over complete automation to remain in control (human-in-the-loop)
	Quality and compliance gaps	Proposals fail to align with the client's requirements	Content should be automatically checked on alignment to the question
	Diverse process	Multidisciplinary teams work on different sections.	Different question types and phases should be supported
	Strict format requirements	Proposals must meet strict formatting rules and word limits	Adding the ability to structure outputs and adhere to word counts
Training "Successful Bidding"	Best practices	Frameworks are familiar but not consistently applied	Utilizing the frameworks to form consistent and convincing proposals content
	Overcomplicated answers	SBP tends to overcomplicate answers with vague statements or unconvincing language	Suggesting assertive and concise responses in proposals

5 Scoping for Proof of Concept

Building on the contextual insights and literature from previous sections, this chapter defines the scope and direction for prototyping a multi-agent system to support the RFP process. By addressing the challenges and opportunities that are found, a narrowed scope is chosen. The previous chapter explored how RFPs are currently developed and where challenges arise, this chapter narrows the focus to those parts and clarifies where and how AI assistance can be offered.

Scoping of the proof of concept is guided by utilizing the Turovsky framework to assess what activities in the RFP process are suitable for a MAS. Given their implementation risk, complexity, and need for fluency, a set of use cases are examined and selected.

5.1 Turovsky framework

The Turovsky framework, is an evaluation method for assessing GenAI use cases, developed by Barak Turovsky (Turovsky, 2023). The framework is useful for systematically approaching the limitations and possibilities for implementing AI. For this research it was applied to the different tasks of the RFP process to select focus areas where AI can provide the most value, and to what extent risks should be accounted for.

The framework evaluates different applications of across two dimensions: fluency and accuracy. In addition, the framework categorizes the application risk of AI into three levels of implementation stakes: low, moderate, and high.

5.1.1 Framework dimensions

- **Fluency**: Refers to how human-like AI-generated content looks and feels. It captures the ability of LLMs to tell a story that is fluent. This is arguably an important aspect for creative writing and drafting engaging narratives. However, it is less important for generating data insights.
- **Accuracy**: Refers to the factual correctness of the generated output. High accuracy of outputs might not matter for brainstorming ideas, but for supporting business decisions or client facing communications, the content should have a very high accuracy.
- **Implementation stakes**: are captured to indicate the risk level when the AI-generated output is incorrect. Color coding the different stake levels can give a clear indication of the risks of implementing AI to that use case. Higher stakes use cases will require more attention in development than lower stake use cases.

5.1.2 Mapping RFP tasks

During the training session “Successful Bidding” and the interviews with Schuberg Philis employees, different tasks of the RFP procedure were explained. With the Strategic Tender Specialist, the tasks were mapped against the Turovsky framework.

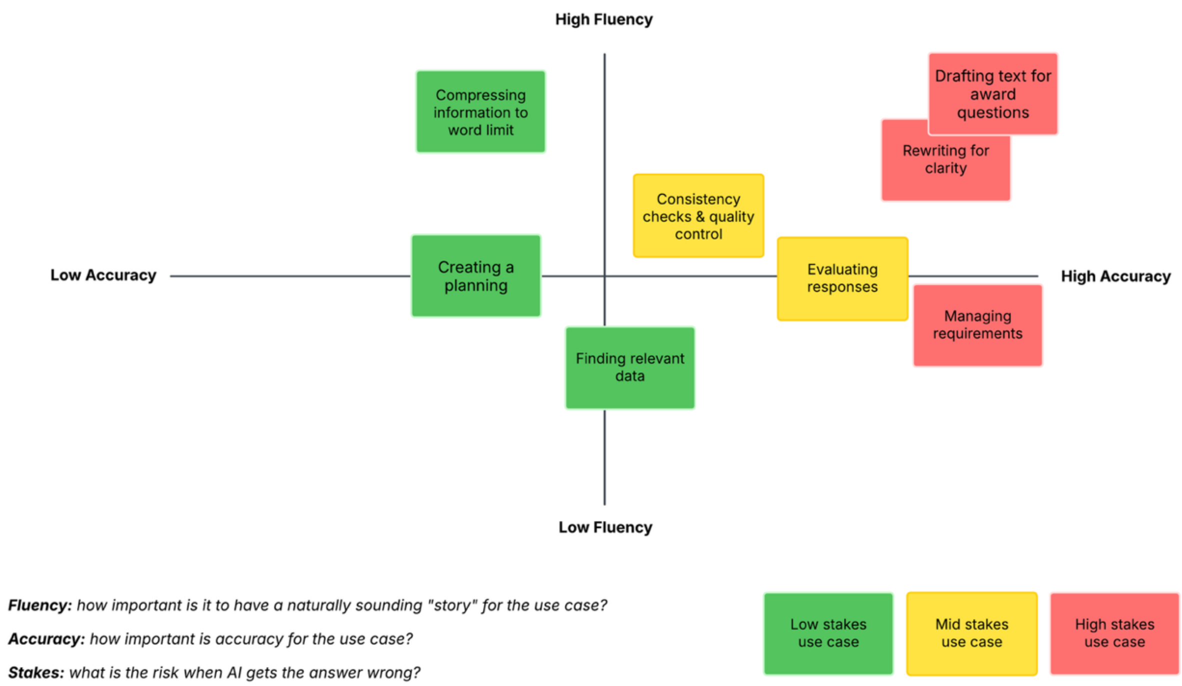

Figure 13: Relevant tasks of the RFP process mapped in the Turovsky framework. Indication of the risks that are associated with implementing AI in RFP related tasks.

Tasks that are chosen for further development are:

- Compressing information to word limit
- Finding relevant data (*from prior projects*)
- Consistency checks & quality control
- Drafting text for award questions
- Rewriting for clarity

These tasks align with the desires and opportunities from contextual research and are further explained in section 4.3. This provides a scope that aligns well with the best practices by including drafting text segments and addresses some of the main challenges from the current process. Together they form a feasible scope for the project timeframe to discover how multi-agent systems should be designed, which is the main goal of the research.

5.1.3 Considerations for human interference

As the framework focusses on the dimensions of fluency and accuracy, the position of humans in human-AI collaboration becomes clearer. While the strong ability of AI to generate fluent content is helpful for creating engaging content faster, the issue of hallucinations makes human validation necessary. Especially in high-stakes cases human control is still needed.

Although it might seem that only low accuracy use cases would be suitable for AI integration, Turovsky describes this is not the case. A different approach is needed for high accuracy tasks, that enable more control by humans. This implies that low-stake use cases allow for more automation without human interference, while higher-stake use cases will need additional validation steps by humans. In the design process, human validation should be considered and solved in a fitting way.

An RFP response offers the opportunity to win significant business, meaning errors in the response can lead to reputational damage and financial losses. Remaining human oversight in the process is of importance to remain control over the quality of the response, whereas automating could potentially lead to bad results.

In the context of this research, these considerations strengthen the interview and training findings that implicates that proposal drafting should be augmented and not automated.

Key takeaways

- The framework evaluates GenAI use cases across fluency, accuracy, and implementation stakes.
- Chosen tasks from the RFP use case that are suitable for AI support are: word limits, data retrieval, quality checks, drafting, and rewriting.
- High-stakes tasks require human oversight because the risk of hallucinations.
- For the scenario of proposal drafting, it makes sense that AI should currently augment and not replace humans.

5.2 Deciding on prototype opportunities

Based on the contextual research findings and the application of the Turovsky framework, we chose to narrow the scope of the multi-agent system to focus on augmentation rather than full automation of the proposal drafting process. Conversations with colleagues at Schuberg Philis showed a strong preference for AI assistance and support, over a replacement of human expertise.

To demonstrate the potential of multi-agent systems in this setting, we defined four key focus areas for the proof of concept where practical value can be added:

1. **Client requirement alignment**
    a. Assistance in interpreting and uncovering complex RFP questions.
    b. Identifying underlying client needs beyond the requirements.
2. **Support in drafting text**
    a. Drafting content based on the identified requirements.
    b. Suggestions to include company knowledge in the draft.
    c. Offering support for different question types (technical, pricing, reference).
3. **Best practice compliance**
    a. Integrating knowledge from proposal writing best practices.
    b. Guidance on word avoidance and tone consistency.
    c. Support for implementing proven structures for responses.
4. **Coherence to form requirements**
    a. Identifying formatting guidelines or requirements from the RFP.
    b. Support in adjusting generated text to the form requirements (word count).

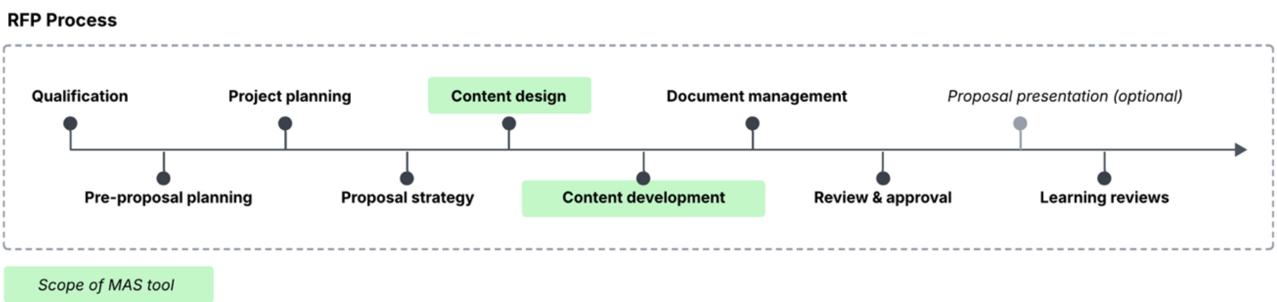

Figure 14: Scope of the MAS visualized in the RFP process from start to finish.

As the opportunities for the multi-agent system are now scoped to augmentation, the aim is to support users in generating consistent, compliant, and well-aligned content. Specifically in scenarios where expert knowledge comes from different contributors and prior work.

The focus is therefore placed on multiple roles at Schuberg Philis:

- Occasional contributors: Employees with subject matter expertise (e.g., technical leads) who are occasionally involved in RFPs but lack experience in writing proposals.
- Commercial roles: For example, Business Consultants and Sales Directors, who could benefit from assistance in guiding structures and fast drafts that they can improve later.

The focus aligns with findings from Brynjolfsson et al. (2023), that show that less experienced workers gain the most productivity benefits from generative AI. We decided to aim the multi-agent system on user who:

- Are involved in RFP processes irregular.
- Might be unfamiliar with the best practices and writing guidelines.
- Need to contribute without the direct support of a Strategic Tender Specialist, or other supporting role in evaluating drafts.

During the training session "Successful Bidding" participants highlighted their doubts about adhering to the best practices without a responsible role in charge. The interviews showed how proposal teams can change in size and roles throughout the drafting process, implicating that some team members will be less familiar with the topic and progress than others.

The MAS should be designed for expert and non-expert contributors, offering guidance in drafting a clear, strong and consistent draft.

5.4 Technical implications for multi-agent design

Table 5: Overview of the technical implications for multi-agent systems, followed by their rationale and source in this document.

Technical support for	Rationale	Source
LLM-based agents	Leveraging LLM reasoning and natural language generation.	Literature section: Large Language Models
Prompt customizability (*agent profiling*)	To allow for flexible agent behavior that can be easily adjusted.	Literature section: Agent profiling; Prompting
Tool usage support	To extend the capabilities of agent beyond text generation (e.g., web search, RAG).	Literature section: AI Agents
Multi-turn conversation	To simulate real conversations with agents over time and test their memory.	Literature section: AI Agents; In-context learning
Multi-agent orchestration	To explore how different orchestrations can be build and respond to scenarios.	Literature section: MAS architectures
Interaction log	To enhance transparency during testing and developing and potentially to create user trust.	Literature section: Human-Agent interaction
Data storage	To enable storage of relevant internal content (file storage).	Interview section: Challenges is the current RFP process

5.5 Concluding on the aim of prototyping

To support the formed research questions, prototyping concentrates on:

- Different orchestration patterns and design principles (SQ3) by testing various agent communication patterns to determine effective collaboration structures.
- Human-agent interaction design (RQ4) by maintaining the assistive nature while exploring different possibilities for human control and oversight.

This focus ensures that interaction between human and system remains at the core of the design, reflecting the high-stakes nature of RFP processes that are identified during interviews and the Turovsky framework analysis.

This page intentionally left blank

6 Prototyping and user research

In this chapter, we explore how design principles and orchestration patterns between agents can enable effective MAS generated output for RFPs (SQ3), and what elements of human-agent interaction can be implemented to ensure trust in the decisions and outputs from the system (SQ4).

The research phase focuses on prototyping multi-agent systems, specifically to the RFP process at Schuberg Philis. Building on the insights from contextual research, this chapter explores how user needs and functionalities can be translated into specific agents. The chapter begins by explaining the decision for the platform "n8n" as prototyping tool. The chapter then presents the iterative prototyping process, where various agent roles, system prompts, and coordination structures were tested and explained. Design decisions and guidelines are continuously shaped throughout this section, and highlighted findings are presented that are later used to inform the design guidelines.

6.1 Methods for MAS development

The methods for developing and evaluating the multi-agent system for RFP assistance are described in this section. Two primary methods are used: iterative prototyping to develop and refine agent capabilities in collaboration with a domain expert, and a workshop on human-AI interaction to identify what design principles are beneficial for the MAS in the context of RFP assistance.

6.1.1 Iterative prototyping

Why iterative prototyping

The complexity of the RFP process, with sensitive information, collaborative nature, and subtle quality criteria, requires an approach that captures tacit knowledge. To address this, an iterative approach was adopted in collaboration with the Strategic Tender Specialist at SBP. The approach focused on hands-on experience in development and immediate feedback cycles to facilitate rapid prototyping.

A collaborative approach is chosen for multiple reasons: collaborative design can lower the barrier for giving feedback, the design will stay better aligned with the user's expectation as they are part of the design process, and the expert can directly hand feedback on the systems output quality. Steering to the right output becomes more accessible with immediate feedback. In each session, we made small adjustments after each run. We discovered what adjustments made the output more reliable, consistent, or of higher quality.

Prototyping session structure

The collaborative prototyping process consisted of seven sessions of two hours each. Typically, the sessions would follow a structure where the researcher and expert started off with discussing what agents or tasks should be prototyped (the goal of the session), followed by prototyping in the form of drafting an agent's system message, available tools, and memory. After the initial draft was formed, the agent was systematically asked to perform the same task. We evaluated responses by carefully reading the content, and the expert told what should be improved or if the goal was reached. During the sessions we repetitively adjusted the agent's configuration until the agent had achieved the goal.

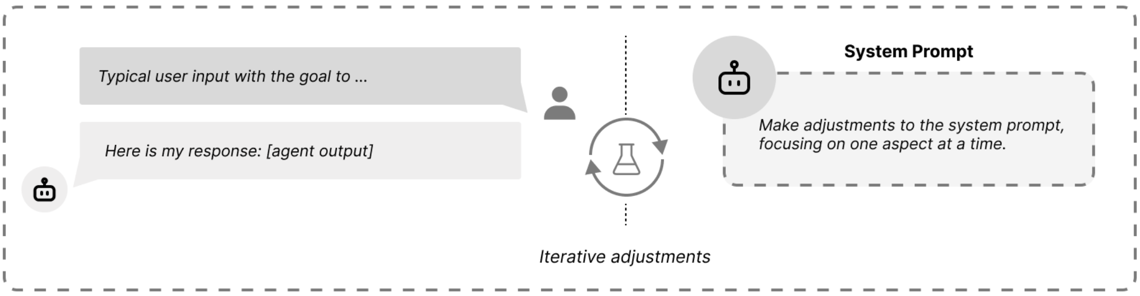

Figure 15: Iterative prototyping structure. For each run the output of the agent is evaluated and based on the quality of the answer the system prompt is adjusted.

The foundation of the prototyping sessions was the developed "Bidbank" outline. This visual framework was used to align the expert and researcher. Over time, adjustments are made for the "Bidbank" outline, but it continued to serve as a guiding component of facilitating these sessions. The framework provided an overview of the entire system and functioned as a checklist for functionalities.

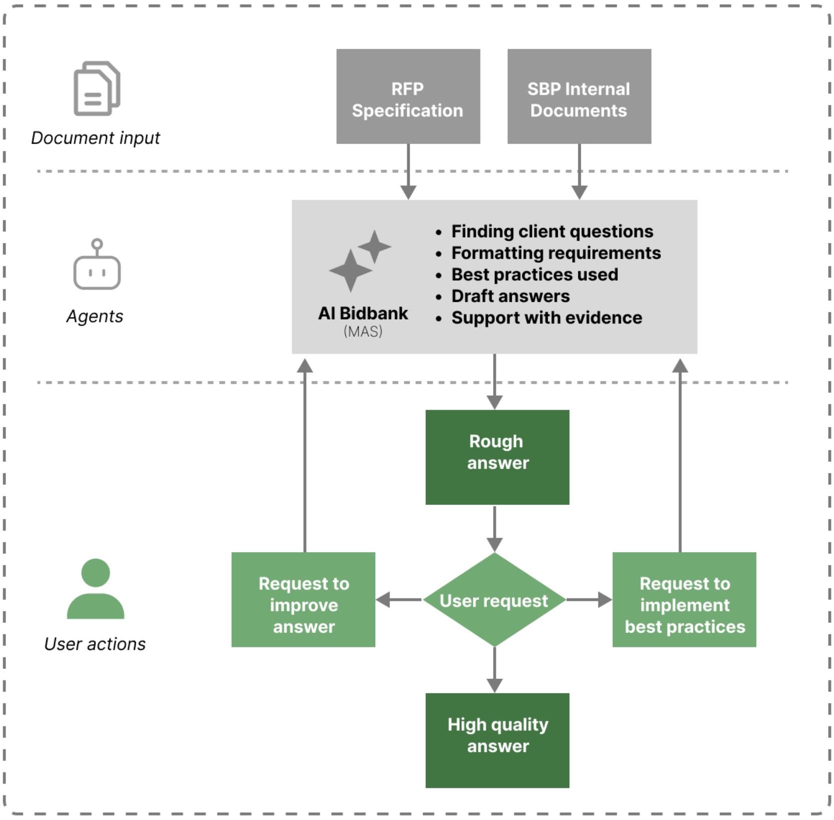

Figure 16: AI Bidbank reference image to ensure alignment prior to the prototyping sessions.

Recording developments

Changes in the agents were documented in the form of a logbook, accompanied with screen recordings during longer prototyping sessions. Short logbook entries were recorded to serve as a reference for creating design guidelines and tracking what has been tried, what worked, and what did not work. These findings are part of the research contributions for future developments of MAS. For this purpose, the use n8n for developing the proof of concept proved to be effective for rapid prototyping and an experiential approach. The GUI provided all necessary elements to adjust and directly test with different inputs and outputs from one interface, enabling the sessions to be effective.

6.1.2 Workshop on human-AI interaction

To gain insight into which types of human-AI interaction are most relevant for proposal drafting, a workshop was organized. The goal of the workshop was to explore how users perceive the role of AI in proposal writing and to prioritize interaction principles to improve the experience and trust of working with a multi-agent system.

The workshop is based on the 18 human-AI guidelines developed by Amershi et al. (2019), and their Human-AI Experience (HAX) toolkit. The toolkit provides the 18 guidelines, a workbook with instructions on how to select guidelines with your team, and a playbook.

Aim of the workshop

The goal of including participants in evaluating the design is to:

1. Gather qualitative feedback on the MAS prototype.
2. Select and prioritize the most relevant human-AI interaction guidelines for the RFP context.
3. Collect ideas and gain an understanding of what implementing these guidelines could look like in the proof of concept of a MAS.

Participants

Eight participants (n=8) were present and were invited based on their experience with RFP processes and experience with innovation projects within the company. Because of their experience they could provide valuable feedback on human-AI collaboration. The participants were informed about the research purpose. The session was audio recorded via Microsoft Teams and transcribed using a local installation of OpenAI Whisper to prevent any sharing of data.

Table 6: Roles of participants of the workshop human-AI interaction.

Role	Number (#)
Business Consultant	3
Sales Director	1
Strategic Tender Specialist	1
Innovation Specialist	1
Chief Innovation Officer	1
Data Scientist	1

Procedure of the workshop

The workshop had a duration of 90 minutes and followed the following format:

1. **Introduction (10-15 min)**
   Participants were introduced to the research goals and the purpose of the workshop.
2. **Prototype demo (10-15 min)**
   The current version of the MAS was demonstrated to give participants the context and show the system's capabilities.
3. **Human-AI guideline review and selection (20-25 min)**
   All human-AI guidelines were printed and distributed across three tables. Participants reviewed the guidelines (provided with an example of implementation) individually and took notes about which are relevant to the MAS use case for RFP assistance.
4. **Dot-voting (10 min)**
   Each participant voted on the guidelines they considered as relevant. The most-selected guidelines were highlighted for group discussion.
5. **Plenary discussion on implementation and relevance (25-30 min)**
   In a plenary discussion, participants discussed why the top-voted guidelines are relevant for the setting of proposal drafting and how these guidelines can be implemented. Participants shared suggestions, examples, and concerns.

Data analysis

The votes from dot-voting were counted to identify and map the most important guidelines (see Figure 19 in section 6.4 for the findings). Transcripts from the group discussion were analyzed qualitatively.

The notes and transcripts were reviewed, and themes were extracted on reoccurring topics and underlying rationale. The selected guidelines were used as a framework to match the emerging themes to, creating an overview of what was discussed based on the original guidelines, but providing more depth by comparing arguments of the discussion.

Contribution to the research

The workshop provided an overview of important themes in human-AI interaction that suggested that the MAS should include specific human-centered elements. The insights reflect real user preferences and design patterns that should be considered for the design guidelines for multi-agent systems. The workshop led to concrete suggestions that can be used for future research to validate further.

6.2 Selecting a multi-agent platform

To identify a suitable platform for developing an RFP-focused MAS, I tested and evaluated four platforms: AutoGen Studio, CrewAI, Airia, and n8n.

AutoGen Studio was considered as it is built upon Autogen as described by Dibia et al. (2024), which is a well-known framework for MAS development as described in (Wu et al., 2023). AutoGen Studio provides a graphical user interface (GUI) that suits rapid prototyping for MAS development. CrewAI (CrewAI, n.d.) is considered as it is a popular open-source framework for MAS development through Python, with a large amount of documentation available online. Airia (Airia, n.d.) is considered as Schuberg Philis was already trying their platform and have contact about the development of the platform. n8n (n8n, n.d.) was considered as it gained popularity since their introduction of AI in workflow automation, and their visual low-code environment makes it easy to explain what is going on during collaboration.

Evaluation criteria and testing approach

The goal of testing and evaluating was to find a platform that balance technical functionality with usability, allowing for iterative development in a collaborative setting. We focused on three main criteria:

Table 7: Evaluation criteria for selecting an agent development platform for prototyping.

Criterium	Description	Rationale
1. Ease of use	How fast can agents be created and tested within the environment, and to what level of complexity?	Rapid prototyping should be facilitated for expert co-design sessions
2. Orchestration flexibility	To what extent can we define and adjust agent communication structures?	To evaluate how a MAS can be effectively designed
3. Tool integration	Are there sufficient possibilities for connecting tools (databases, document parsers, etc.)?	Without tools, agents would be limited for real business scenario's

To compare the platforms, I created a hypothetical RFP team on each platform using eight agent roles. The system prompts for these agents were generated with the help of ChatGPT to quickly move into testing. The agents were not validated at this point on their performance, as this step had the purpose to understand the platform possibilities and behavior, not agent performance.

Platform comparison

During testing, each platform showed strengths and weaknesses. The main findings per platform are:

- **AutoGen Studio:** Offered an intuitive GUI and fast agent setup but offered limited tool integration possibilities and limited scalability for larger systems. In its current form, it is wellsuited for quick development and conversational experiments, but not for integrating new knowledge and utilizing tools to automate of augment more complex scenarios.

- **CrewAI:** Provided good control over agent behavior and task delegation by defining communication patterns by utilizing Python scripts. This offers complete control to the developer instead of being constrained by the options of GUI templates/building blocks. It provided clear outputs that are placed in a folder for each run of the MAS, making it easy to show what the system produced. The cloud-based version (CrewAI Enterprise) offered a visual interface and was tested. However, this version was unstable during testing and was therefore not further evaluated. Overall, CrewAI seemed like a great option for production level automation, but less ideal for using it in an exploratory setting, as the lack of visual interface elements make it harder to use in a collaborative setting.
- **Airia:** Offered an intuitive GUI with self-explanatory support for defining agents, setting up memory, tools and selecting LLMs, making it easy to start creating agents. What stands out is the native support for guardrails, which are meant to provide data protection and detect biases. Although these functionalities are valuable for production ready enterprise solutions, they do not offer added value for the prototyping process. The integrated version control of agents makes it easy to track developments over time. However, the platform is still under active development, and some features are not stable.
- **n8n:** As a low-code workflow automation platform, n8n is not originally built for MAS. However, it offers integration with AI capabilities based on LangChain, enabling users to create AI-integrated workflows. It enables users to define agents, integrate with various integrated or custom tools and data sources, and managing interactions based on connecting agent nodes, or calling agents as tools. The platform has a steep learning curve but offers well-documented support and has an active community. The strengths lie in its customizability, tool integration, and the visual representation of agents (nodes).

Selection of platform and reflection

We tested AutoGen Studio, CrewAI, Airia, and n8n to evaluate their potential for developing a multi-agent system for RFP support. While AutoGen Studio provided a user-friendly interface for quick conversational experiments, it lacked the possibility to integrate tools well and did not offer a suitable environment for scaling agents systems. CrewAI offered full control over agent orchestration and communication lines and produced clear outputs per agent in a specific markdown file, providing clear traceability of agent outputs. However, the command-line interface made it less suitable for rapid prototyping in a collaborative setting. Airia combined an intuitive GUI with enterprise level features for deployment, with security measure like guardrails, and version control to keep track of development. However, the ongoing development is a downside for the research context, as the platform could heavily change over time.

Ultimately, n8n showed a clear visual interface, strong tool integrations, and custom orchestration potential, making it suitable for prototyping a more complex system. The log in the interface makes it traceable what is going on during development and while running the system. Although it could lack MAS features, as the platform is not originally built for this purpose, no limitations showed up during testing to withhold us from using the platform.

This led to the decision to use n8n for further prototyping as it enables a flexible platform that suits well for collaborative exploration of MAS behavior.

6.3 Prototyping findings

This section presents the findings from iterative prototyping sessions with a Strategic Tender Specialist. The findings are organized into three themes that address challenges and opportunities in MAS development and are aimed to provide guidance to designers and developers.

Findings and examples from prototyping for the development of the proof of concept are placed in the following themes, with each section highlighting main findings and design implications:

1. **Theme 1: Agent orchestration and communication:**
   How multiple agents share tasks, information, and interact with each other and with users. This section provides insight communication patterns and challenges in maintaining context across agents.
2. **Theme 2: Agent design and prompt engineering:**
   This section demonstrates the critical importance of explicit instructions for agents and structured prompting to create reliable and consistent agent behavior.
3. **Theme 3: Knowledge flow and information retrieval:**
   This section provides insights in how agents can access, store, and share information from documents and user interactions.

6.3.1 Theme 1: Agent orchestration and communication

As a developer it is important to consider how tasks or coordinated across agents, how information is shared, and what agents interact with each other and with users. This section presents findings that emerged from prototyping in task distribution, task fragmentation, communication patterns, and user interaction.

Task distribution: how and when should tasks be executed

One of the decisions in MAS design involves determining how tasks should be distributed among agents and whether these tasks should be performed sequentially or in parallel. Prototyping showed that task dependencies should be carefully mapped before deciding upon orchestration patterns. Creating a system without a clear approach on how tasks are distributed can impact the output quality negatively.

Sequential and parallel execution patterns

Sequential task execution, where one agent waits for another agent's response before proceeding, creates dependencies that can slow the system down when multiple agents must contribute. This approach can work well when the same tasks need to be performed in the same order every time. However, by using sequential patterns bottlenecks can be formed and it limits flexibility for users.

Parallel task execution showed to be a more effective solution for tasks that are independent of each other. In parallel task configurations an orchestrating agent distributes multiple tasks simultaneously to different (specialized) sub-agents and waits for all responses before they are synthesized into a final output. This approach reduces response times and reduces the complexity of sub-agent inputs without potential adjustments from other sub-agents.

Design implication:

Implementing task distribution strategies (sequential or parallel) requires explicit instructions in the system prompt of orchestrating agents, clearly specifying which tasks should be completed in what order. When no explicit guidance is provided, the orchestrator decides autonomously, leading to inconsistent behavior (due to the probabilistic nature of LLMs).

Task fragmentation: deciding the right agent specialization

Prototyping revealed the importance of task fragmentation and granularity. While multi-agent systems benefit from dividing responsibilities and tasks among specialized agents for chosen domains, as explained by (Guo et al., 2024; L. Wang et al., 2024), practice shows that excessive task fragmentation can be counterproductive.

The research approach involved translating business processes into sub-processes, which were then divided into manageable tasks for individual agents. However, overly granular task splitting led to orchestrators assigning tasks incorrectly (to the wrong agents) and agents possessing overlapping capabilities.

Design implication:
Start with agents that have broader capabilities and specialize only when there appears to be a clear functional benefit. The goal should be meaningful and effective specialization that enhance system performance rather than complicating orchestration with granular task division without clear benefits.

Agent communication modes

Effective communication between agents proved to be essential for maintaining context and preventing information loss during inter-agent interactions. Multiple communication patterns are explored, each offering advantages and challenges.

Sequential message passing and agents-as-tools

Sequential message passing, where one agent hands its output directly to the next agent, works well for simple workflows but lacks flexibility. This pattern suits tasks that always follow the same steps but becomes limiting when users require different task sequences from agents.

Agents-as-tools orchestration comes with different challenges. The orchestrator must decide when to call other agents for the user's query. Initially, orchestrator agents showed unexpected behavior by attempting to answer questions using their own knowledge rather than consulting specialized sub-agents.

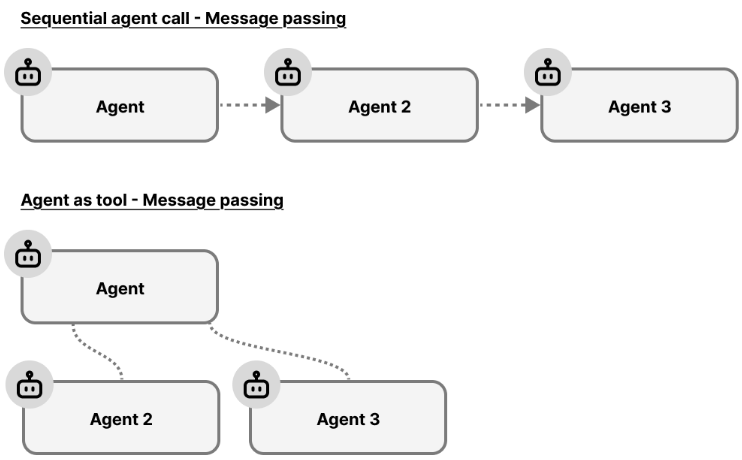

Figure 17: Communication modes between agents.

Discovered solution

There are two approaches to address the issue of unexpected orchestrator behavior:

1. Adjusting the system prompt of the orchestrator
2. Changing the description of sub-agents in their tool descriptions

Approach number 2, changing the tool description of an agent, proved to be more effective and ensured orchestrators improved agent-calling. This highlights the importance of providing sufficient context and detail when describing agents.

Design implication:

Using an agents-as-tools orchestration for communication makes more sense for most use cases, as it offers the flexibility that makes a MAS powerful. It is therefore recommended to stay away from sequential orchestrations, unless a standard approach must be guaranteed.

When implementing agents-as-tool orchestration, invest significant time and effort in creating detailed descriptions for each (sub-)agent.

Memory management and sharing

Managing memory across multiple agents comes with challenges in maintaining context throughout the conversation and preventing information loss. In multi-agent systems, details easily get lost when information is handed over by multiple agents, potentially leading to incorrect outputs. This finding can be compared to the popular “telephone game” where people whisper a sentence to the next person in line, and Perez et al. (2025) demonstrated how this affects multi-turn interactions with LLMs.

A limitation can be found in n8n as MAS platform in memory sharing across agents as this is not a native functionality. Other platforms, such as AutoGen or CrewAI offer the possibility to share memory across agents, which can potentially offer better results.

Design implication:

Explicitly prompting the agent to share background knowledge increases the amount of information that is given to the next agent as input. It is recommended to create a structured output for agents that they should pass along to other agents. This ensures that context from the interaction is passed along to subsequent agents every time.

User interface and control

Two different methods for communication between users and the MAS were explored, revealing trade-offs in the current state.

Chat- and form-based interaction

Traditional chat-based interaction as can be recognized from ChatGPT and most other AI chatbots, users directly send questions or instructions to the LLM, or in this case, agents. This is an intuitive approach that is familiar for users, but it offers limited control and lacks traceability in the context of chatting with multiple agents, from one chat interface.

A **form-based interaction** was explored as an alternative, offering a novel way to modify system variables to the user, without having to adjust the backend of the multi-agent system configuration. This was tested by offering the option to selecting agent that the orchestrator can access with a checkbox selection.

Trade-off between control and usability

Form-based interaction shifted responsibility and oversight to users, providing transparency about which agents are used. This approach potentially improves user trust through transparency, and system performance for high-risk tasks where the selection of the right agents is more critical.

However, implementing form-based interaction presented limitations in the prototyping environment (n8n). In multi-turn interactions when forms were used for both input and output, the platform did not consistently show the outputs to the user in the form. For this reason, the consideration for offering control to users is recommended, but limited options to offer control are currently available.

Design implication:

User interfaces should match the stakes of AI tasks, where higher stakes should offer more control and lower stakes benefit from simple interfaces with less control. Take into consideration that while autonomous agents reduce the complexity of decision making for users, offering control plays a part in user trust.

Concluding on agent orchestration and communication

Several design principles should be considered when designing a MAS for inter-agent relations:

1. **Task dependency mapping** is essential for selecting orchestration patterns. Parallel executions should be the standard for independent tasks, where sequential execution is ideal for strict protocols or interdependent processes.
2. **Agent specialization** should be purposeful. Start with broader agent capabilities and specialize only when there appear to be clear benefits in creating new agents.
3. **Orchestrator behavior** is heavily dependent on prompt instructions. Explicit instructions including context are needed create consistent behavior, whereas undefined instructions leave autonomous decisions to agents introducing potential inconsistencies or vulnerabilities in the system.
4. **Memory management** is important to prevent loss of context and details in inter-agent communication. This can be explicitly prompted by using a structured output for agents, requiring sharing background and context knowledge to the next agent.
5. **User control** should be considered in the design process of a MAS. High-stakes applications benefit more from control then lower-stakes applications. Currently, limited user control is offered on the frontend for MAS building platforms.

6.3.2 Theme 2: Agent design and prompt engineering

To create a good performing MAS, the foundation of individual agents should be performing well. This section addresses findings how system prompts, model selection, and built-in capabilities from agents affect their performance.

The foundation of agent behavior: structured system prompts

Early prototypes showed challenges in agent consistency and reliability. Initially, agents demonstrated inconsistent behavior by not following protocols and by delivering outputs in inconsistent formats. This makes it difficult to evaluate outputs and might cause issues in multi-agent systems where one agent's output becomes another agent's input. To address this, a systematic approach to agent prompting is developed that became the foundation for each agent.

Six-component system prompt

Based on prompting best practices published by OpenAI, Anthropic, and Google, we defined our own standardized system prompt structure (Anthropic, n.d.; Google Cloud, n.d.; MacCallum & Lee, 2025). The structure consists of six components:

1. **Role & Objective:** Clear definition of the agent's purpose and function.
2. **Context & Tools:** Description of available tools and environment.
3. **Standard Operating Protocol (SOP):** Step-by-step procedures the agent must follow.
4. **Output Format:** Explicit formatting requirements for responses.
5. **Examples:** Concrete demonstrations of desired behavior and output quality.
6. **Final Prompt:** Direct instructions for action.

By implementing this structured approach, better adherence to the protocol and role of the agent is observed. Important aspects from literature are recognized in this structure, where the standard operating protocol is a form of CoT prompting (Wei et al., 2022). Next to this, putting examples in the prompt is a form of few-shot prompting, that improves the model's ability to learn from a provided example (Brown et al., 2020; Weng, 2023).

Design implication:
Creating agents following a standardized structure makes it easier to create new agents that perform well on their given task. Implementing the structure means developers need to consider relevant aspects like output format, examples (few-shot learning), and chain-of-thought prompting in the form of a standard operating protocol directly from the start.

Model selection

Selecting the underlying language model is a crucial decision for the agent's performance and capabilities. The model selection will have direct consequences on the output quality, prompt following, reasoning depth, and tool usage.

Empirical evidence of model impact

Testing the same system prompt with different large language models showed performance variations. For example, a comparison of results between GPT-4o-mini and GPT-4o shows the following results:

GPT-4o-mini input: "Find information about risk management"

- Generated 1 search query: "Risk Management"
- Result: Limited and incomplete information retrieval

GPT-4o input: Same query

- Generated 5 search queries: "risk management," "risk assessment," "risk mitigation," "preventive measures," "corrective measures"
- Result: Comprehensive and accurate information retrieval

Testing showed that larger models demonstrate better adherence to standard operating protocols and show more detailed and sophisticated reasoning capabilities.

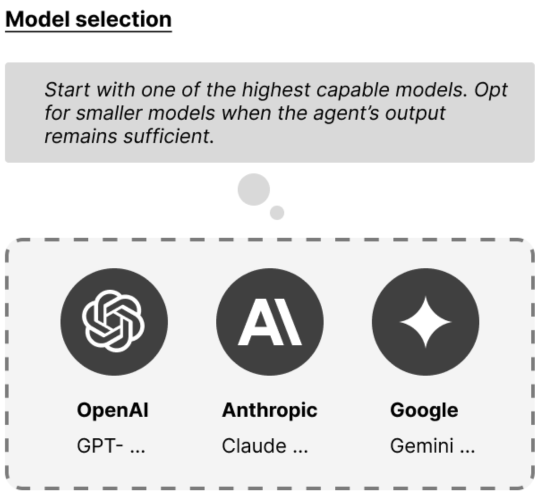

Design implication:

The findings from adjusting large language models for agents leads to the following recommendation for developing agents for a MAS.

1. **Development phase:** Use the most capable models available to validate that agents can achieve their intended goals.
2. **Optimization phase:** Once it is confirmed that an agent functions as intended, consider scaling down to smaller models to reduce operating costs and potentially reduce latency. During this step it is essential to continuously monitor the performance of the agent and validate that it maintains acceptable performance levels.

Self-reflection for agents – Think tool

As described in the literature sections "*AI agents*" and "*Improved reasoning within a MAS*" reflection is an important aspect for agents. The Think tool, first introduced by Anthropic, allows agents to pause and reflect whether it has sufficient information to continue or whether the output meets quality standards (Anthropic, 2025).

During prototyping the Think tool was not used consistently by agents. Just like any tool this is caused by the quality of the prompt or underlying language model. Explicitly prompting the Think tool will help in more consistent and frequent use of the tool.

Whether the agent benefits from implementing self-reflection in the form of the Think tool depends. Some tradeoffs of implementing the tool can be found in Table 8

Table 8: Comparison of benefits and downsides to implementing 'Thinking' steps for agents.

Benefits	Downsides
• Improved task execution for complex tasks	• Increased system prompt complexity
• Better error detection and correction	• Additional latency in response
• Improved transparency for agent reasoning	• Higher token usage (costs & context window)
• Lower chances of hallucination rates when reflecting multiple times	

Design implication:

Self-reflection mechanisms (like the Think tool) should be selectively implemented based on task complexity and the stakes of the task. High-stakes processes benefit from mandatory reflection, while simple tasks may not need the added reflection. When complex agents do not reach the desired results, it is recommended to try and implement self-reflecting mechanisms.

Custom tool integration

Some tasks are not ideal for LLM-agents to process, reaching their fundamental limitations. To extend the capabilities of an agent, custom tools can be utilized or developed that offer capabilities beyond that of natural language processing. To demonstrate how custom tools can be integrated in the development of an agent (and in the context of a MAS), an example is explained below on the word count of a generated text.

Word count assessment

In proposal writing, the word limit of documents is of importance. However, LLMs cannot accurately count words due to the tokenization process of text (see 3.1.2 for a more detailed explanation).

To address this issue, a custom JavaScript tool is created during prototyping, with the aim to provide exact word counts, character counts, sentence counts, paragraph counts, and the average words per sentence. This tool enables the agent to:

- Accurately assess the word count of the draft against format requirements.
- Perform readability assessments (for example the Flesch-Kincaid Grade Level).
- Decide whether the draft should be decreased in content or increased.

Design implication:

When agents do not offer the right capabilities to perform specific tasks, it can be an outcome to utilize or develop a custom tool that the agent can access. During development consider the necessary skills and identify or verify if the task can be accurately performed by a LLM or if there is a need for external or custom tools.

When developing custom tools, it is recommended to utilize coding assistance to from LLMs that are high ranked for their coding capabilities. This can be useful for both developers and non-developers to speed up the process of creating a specific tool.

Concluding on agent design and prompt engineering

The findings from prototyping agents and experimenting with prompt engineering leads to several principles that can be utilized by developers:

1. **Structuring the system prompt** of agents is essential for consistent (multi-)agent behavior. The six-component system prompt structure provides a reliable foundation for agent development, based on validated prompting techniques.
2. **Large language model** selection affects every aspect of system performance. Carefully consider what model is selected for each agent and use more capable models in development. Optimizing costs versus performance can always be achieved afterwards, but do not get stuck on prompting agent capabilities that can never be reached due the underlying model.
3. **Self-reflection mechanisms** like the Think tool should be used selectively based on task complexity and tolerance for errors (depending on the stake of the use case). Explicit prompting is needed to utilize these mechanisms reliably.
4. **Integrating custom tools** for agents offers the possibility to extend agent's capabilities beyond their natural functionalities. Consider developing your own tools for specific tasks or utilize third party tools significantly increase agent's potential.

6.3.3 Theme 3: Knowledge flow and information retrieval

Managing information between agents and from documents plays an essential part in a MAS, especially in an information heavy context like the proposal writing. This section covers prototyping findings on vector store architectures, RAG implementation, and structured outputs.

Creating databases for agents

In the context of creating databases for business processes, it is important to understand that information should be handled securely to prevent sharing confidential information internal or even external connections. During prototyping for the RFP scenario, no actual internal documents were used, but the system is prototyped as if this would be the case to create a compelling proof of concept.

Separating knowledge

Rather than implementing a single database that contains information from all document inputs, splitting information in separate databases makes sure no mistakes can be made when searching and using specific knowledge.

Design implication:

Separating knowledge prevents agents from using the wrong information and creates a layer of security. When designing the MAS, consider what extra information should be provided to what agents and based on the tolerance for error decide whether to split information over separate databases.

Information retrieval strategies

Implementing vector stores in the multi-agent system revealed how utilizing retrieval-augmented generation (RAG) comes with challenges. Particularly when varied terminology is used, as is often the case in RFP documents, it becomes harder for an agent to find appropriate information in a large database based on semantic search techniques.

What is discovered is that agents tend to use short queries for searching the vector database, like "sub-criteria #2: value added" and this does not always lead to successfully retrieving information, even if the content is in the database. To address this issue, agents were prompted to generate multiple search queries for each request.

For example, a request about "risk management" would generate queries for "risk management," "risk assessment," "risk mitigation," "preventive measures," and "corrective measures".

Prompting agents to use a multi-query approach showed improvements in retrieving accuracy and completeness of generated responses. The underlying language model also played a role in the quality of searching the vector store, with larger models typically creating both more diverse and longer search queries.

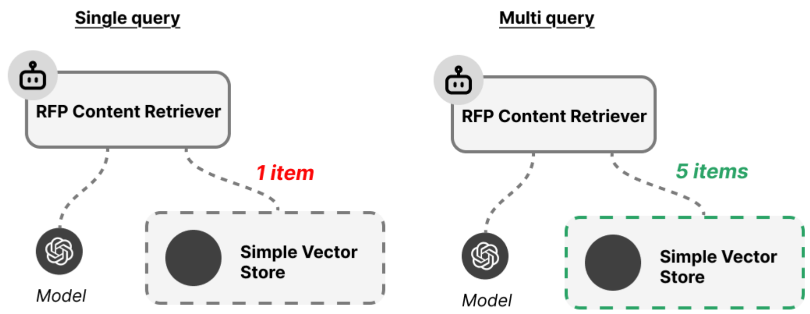

Figure 18: Explicit prompting improves search queries for retrieval from vector stores.

Design implication:

When utilizing RAG for agents and multi-agent systems it is recommended to prompt the agent to approach searching the vector stores with multiple queries. The effectiveness of this approach also depends on model selection (see "theme 2: model selection" for recommendations regarding large language models).

Structured outputs

Effective knowledge flow between agents can be increased by using communication patterns as described in theme 1. However, building on these patterns the approach of structuring data and formatting outputs can be further elaborated by using a structured output.

Outputs can be standardized by using a structured output parser, where the developer can define a standard format outline. This prevents difficulties with transferring knowledge to downstream agents as the output structure is predefined.

The LLM will create its output based on the example that is provided in the structured output parser. In n8n, the structured output must be defined in a structured JSON format.

Design implication:

Structured output formatting can be an outcome for maintaining consistent information flows in multi-agent systems. Defining the output structure prevents inconsistency in the output format and should be considered for systems where no information can be lost. However, for most use cases not all agents necessarily need a structured output parser. It prevents the agent from creating new output types and could hinder creativity in the outputs of agents. It is therefore recommended to use a structured output for agents that provide intermediate outputs where consistency is of importance for downstream agents. User-facing agents most likely won't benefit from structured output parsers as this limits their flexibility in responses.

Concluding on knowledge flow and information retrieval

Findings from prototyping regarding the management of information in multi-agent systems led to the following principles:

1. **Information separation** can be utilized to ensure agents use the right information and to built-in a form of security in what information is used. Splitting databases prevents agents from accessing confidential information in phases where this is unnecessary and reduces the risk of using information in an incorrect context.
2. **Multi-query retrieval approaches** improve the ability of agents to successfully retrieve information from large amounts of information and documents. Agents can be explicitly prompted to search with multiple queries to improve the retrieval success rates. The selection of large language models remains an important factor in the effectiveness and quality of search queries.
3. **Structured outputs** can be an effective control mechanism for agent-to-agent communication. Selectively implementing structured outputs for intermediate agent outputs can be beneficial for agent's communication where consistency is crucial. It is recommended to remain flexibility for user-facing agents, where adaptability can be more important.

6.3.4 Concluding on prototyping findings

The iterative prototyping sessions revealed that designing a MAS for business processes requires structured and thoughtful development. Testing with the domain expert led to multiple principles across agent orchestration, agent specialization, and information management.

Key insights:

- **Explicit prompting:** Across all themes the importance of explicit prompting becomes apparent. Giving mandatory instructions to agents rather than flexible guidelines is essential to consistent behavior.
- **Task division:** Specific to multi-agent systems, the specialization of agents if of importance for reaching their goals. Specializing provides clear functional benefits, but over-fragmenting tasks into sub-agents can create more problems by overcomplicating the system.
- **Information management:** Defining what information is shared across agents and what extra information is available in databases is a main focus area of designing a MAS. Multiple approaches to support effective information sharing are shared in this section.

This chapter directly addresses the main research question that aims to explore how multi-agent systems can be effectively designed for business processes. Providing information on design principles that can be used, rather than relying on technical information on multi-agent systems alone. Developing the proof of concept and constantly validating the agent outputs with a domain expert provided a foundation for developing the design guidelines as an outcome of this research.

6.4 Workshop human-AI interaction findings

This section presents the findings from the conducted workshop on human-AI interaction. The aim of the workshop was to identify what design principles are relevant for a MAS to support the RFP process at Schuberg Philis. The session counted eight participants with experience in RFPs and innovation projects, who evaluated the 18 human-AI guidelines from Amershi et al. (2019), discussed the relevance of the guidelines for RFP assistance, and discussed how this could be applied in practice.

Voting on the most relevant guidelines reflected short-term needs and direct insight in what is considered as important, while the discussion afterwards revealed more complex considerations. Interestingly, it was not necessarily the highest voted guidelines that generated the most elaborate discussion. The results are grouped into four themes.

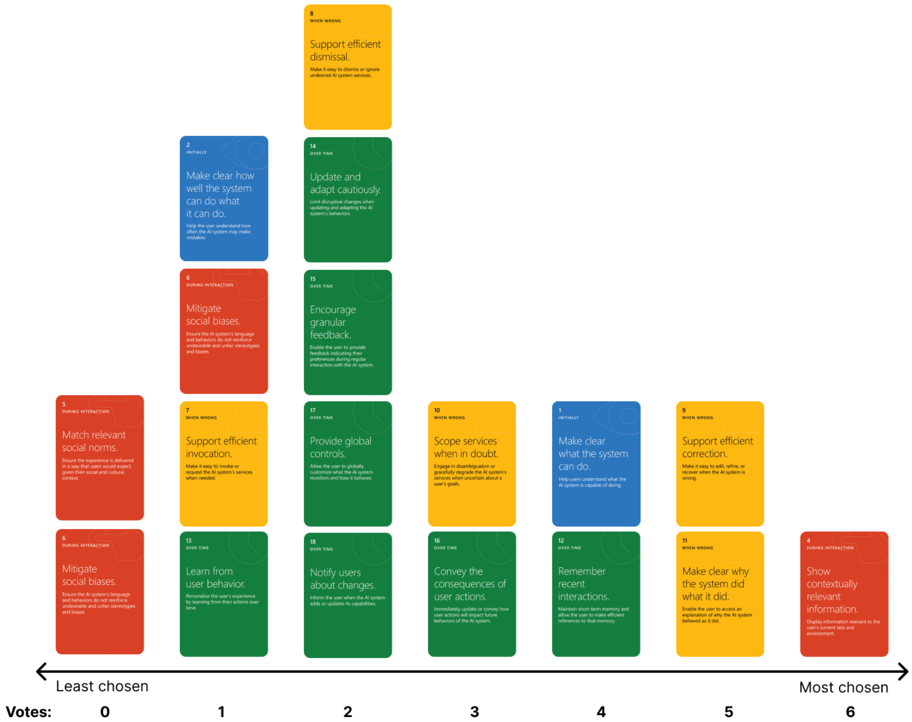

Figure 19: Vote count results on the human-AI guidelines by Amershi et al. (2019).

6.4.1 Themes from discussion

Trust and traceability - Discussion based on guidelines 1, 4, and 11

Participants emphasized the importance of understanding how the MAS created its response. Suggestions during the discussion were displaying the sources that are used, showing intermediate reasoning steps, and clarifying when a response is based on prior successes (other RFPs, or company specific knowledge). A participant clarified the need for checking assumptions: *"If it comes back with a good answer I can drop straight into my RFP response, I still want to know on which assumption it's based, so I can check and be sure it's usable"*.

Two participants explicitly mentioned that traceability is the most important aspect for them, by saying: *"We want to be helped, but we also want to know why. Traceability is number one."* And *"Traceability is number one by far"*.

Participants also discussed the importance of transparency in learning loops, where agents learned from information that the user handed to the system earlier, as it builds trust. One participant mentioned this directly: *"Trust increases when your feedback is taken to heart, when you say how you want it, and the next time it's adapted accordingly."*

Personalization versus standardization

There was a tension between participants where some participants want the system to adapt to their personal preferences and ability to change outputs, while other participants urged the need for standardization of outputs, that align with the best practices and company style at Schuberg Philis. A standardized structure could also ensure compliance to RFP requirements, even when the user requests to deviate, whereas a personalized MAS could potentially ignore or overlook simple compliance measures, bringing along risks in the process.

As one participant explained how an agent should respond: *“I am strict about this, because these are the important things we follow in the RFP process, and we should not deviate from them. Even if, as a user, you intuitively feel you should, the model should respond: ‘No, I will not do that, because this is the best practice I must adhere to.’”*

Another participant elaborated with two perspectives on standardization: *“Is it imposed uniformity, or is it learned experience that when you answer in a certain way, you have a higher chance of winning? The latter is a kind of feedback loop, while the former is more of a compliance thing. Standardization can also mean that, by answering in a certain way, we avoid running into problems in the future.”*

Out-of-scope question strategy – Discussion based on guideline 10

During the discussion, participants expressed their curiosity about how the MAS should handle unfamiliar or out-of-scope RFP questions. Two strategies were discussed:

1. Extending the agents knowledge base beyond provided material in the database. For example, by searching external sources (via the web).
2. Showing transparently that the agent does not have relevant information available.

One participant described how it is currently handled: *“I deliberately ask it something [it should not know] and see if it crawls the web. If I see it comes back with something new, I continue on that. It breaks out of its system.”*

Another participant explained the trade-off that must be made: *“There are two paths. Either is does a web search and keeps looking until it finds something, or it says ‘This is my knowledge domain and my world, and based on your question, I have to disappoint you.’”*

While a different participant mentioned the latter option is less valuable, others stressed it was better than producing a wrong or misleading answer: *“I know, but it’s better than trying to make something up that completely misses the mark.”* Followed up with a response on traceability again: *“Exactly, and that’s why I would need traceability [in answers].”*

End results versus intermediate steps – Discussion based on guidelines 9, 11, and 17

The demonstration of the MAS prototype, where the MAS retrieved relevant content from RFP specifications and drafted a complete response for one quality criteria, led to a discussion on whether users benefit from final outputs alone or from intermediate outputs and artifacts.

Some participants value the control and involvement that comes from being handed intermediate results, particularly for using the intermediate results to build towards the final response. While other participants agreed, some felt that a standard result as output could be useful for developing a structure for the final response.

This showed in a conversation where one participant described the idea of a strong nearly perfect draft that can be achieved fast to save time: *"All these [writing] steps are incredibly labor intensive. What we want with the AI Bidbank is to [solve that] people start writing with one to two weeks left from the eight weeks we had. The idea behind it is that you have a smart first version, that is compliant, convincing, and distinctive, based on the archive, that is already 95% complete, so that everything is influenced by that first version."* While another participant stood by the argument that intermediate results are crucial to develop the final version with correct content: *"That creative part, all the technical artifacts and deliverables you need to be able to write, that is the piece in the middle. [The current system] can give you the structure for a final response, but content-wise it can do much less [without these intermediate artifacts]."*

Participants also reflected that in the context of proposal drafting, intermediate steps can give the user a greater sense of ownership, which could be important for user adoption of AI in the process. Based on the discussion, participants indicated a need for some level of control over the process and produced output artifacts, aligning with guideline 17 ("Provide global controls").

Furthermore, this conversation captured an interesting broader issue in knowledge capture. A participant specifically mentions the issue where companies often document final deliverables without preserving the reasoning process that led to the results. For AI systems, it would be beneficial to capture these reasoning processes to feed it as contextual information, utilizing their ability to learn from examples.

"…we often deliver the end result but forget to document our meta-thinking and how we arrived at it. If you have a consultancy approach in your head for how you create a milestone planning, and you abstract and write that down, AI could use that to generate a version point-something and then create the milestone planning."

Other observations/takeaways

Participants reflect certain doubts about handing of multiple tasks simultaneously without offered intermediate results. Based on the different approaches named by participants, we interpret that personal experiences of participants with the use of AI in work might influence the way they approach handing tasks to AI and might require time and hands-on experience with more autonomous systems before acceptance of AI autonomy will be formed.

Participants imagine an AI system that is explainable, controllable, and reliable rather than autonomous black boxes. The participants want to have their own place in the process, following the foundation of human-in-the-loop principles.

6.4.2 Concluding on the workshop on human-AI interaction

The workshop provided insights in how human-AI interaction could be applied to the design of a MAS for RFP support at Schuberg Philis. Voting on the guidelines showed participants' direct priorities, while rich discussions revealed considerations and different views on explainability, control, and a balance between autonomy and control of multi-agent systems.

- **Trust and traceability:** Participants see transparency in sources used, reasoning steps, and whether answers are based on prior successes as essential for building trust. Additionally, the implementation of feedback loops can strengthen users' confidence and improve user adoption.

- **Divided opinions on personalization and standardization:** While some participants prefer outputs that are tailored to their style, others focus on the need for standardized formats to align with the best practices, maintain compliance, and that reduce risk.
- **Out-of-scope handling:** Participants value the retrieval of information from external sources, but only if it is accompanied by traceability. Alternatively, the system should clearly indicate the lack of information when it does not have access to external sources, to prevent producing incorrect information.
- **Intermediate steps versus final results:** Participants are divided between a system that produces a near complete draft immediately, and a system that provides them with intermediate artifacts they can use to write content.
- **Human-in-the-loop:** The ideal MAS for this context is explainable, controllable, and reliable, where human users are included and maintain oversight.

These insights inform decision making for the proof of concept and are used for the design guidelines for future reference and advice on how to incorporate users in the development phase for input.

7 Proof of Concept

The final form of the proof of concept emerged from all information gathered of the initial interviews, scoping of the design, and iterative prototyping and evaluation. Based on the defined scope, the proof of concept was developed to focus on extracting information from RFP's, analyzing questions from specifications, find and create connections with prior work and experience by providing evidence, and produces a strong draft that adheres to the defined best practices.

This chapter shows the system configuration of the MAS and describes and demonstrates what it can achieve, by showing a demonstration of questions based on a publicly available RFP and answers that the system produced.

7.1 Agent roles in the MAS

In this section the different roles of the agents in the proof of concept are described. We elaborate on the reason why the agent is formed and how it achieves its goals. An overview of the MAS is provided in Figure 21. Complete descriptions and the exact system prompt of each agent can be found for reference in appendix G.

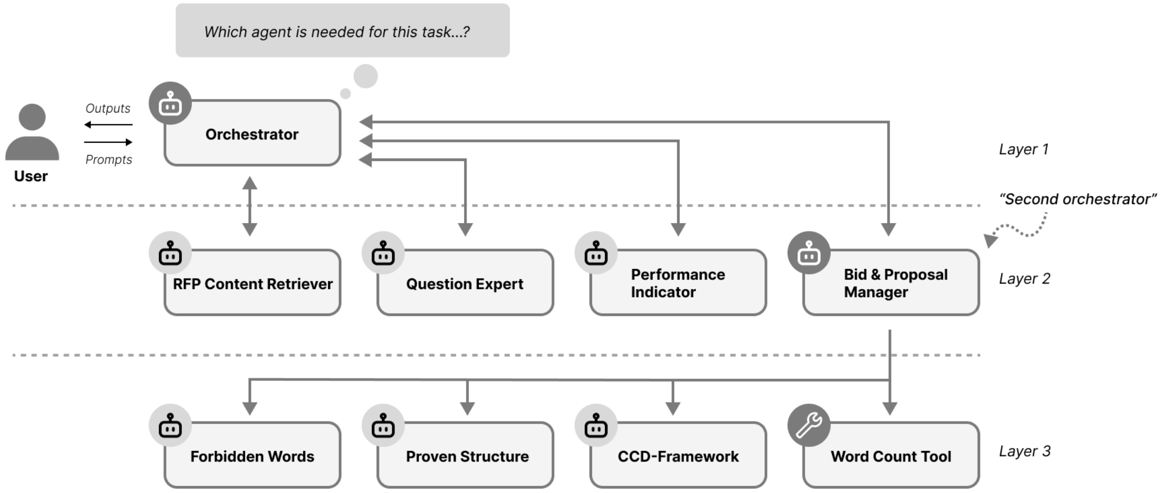

Figure 20: Communication pattern and orchestration layers of the proof of concept.

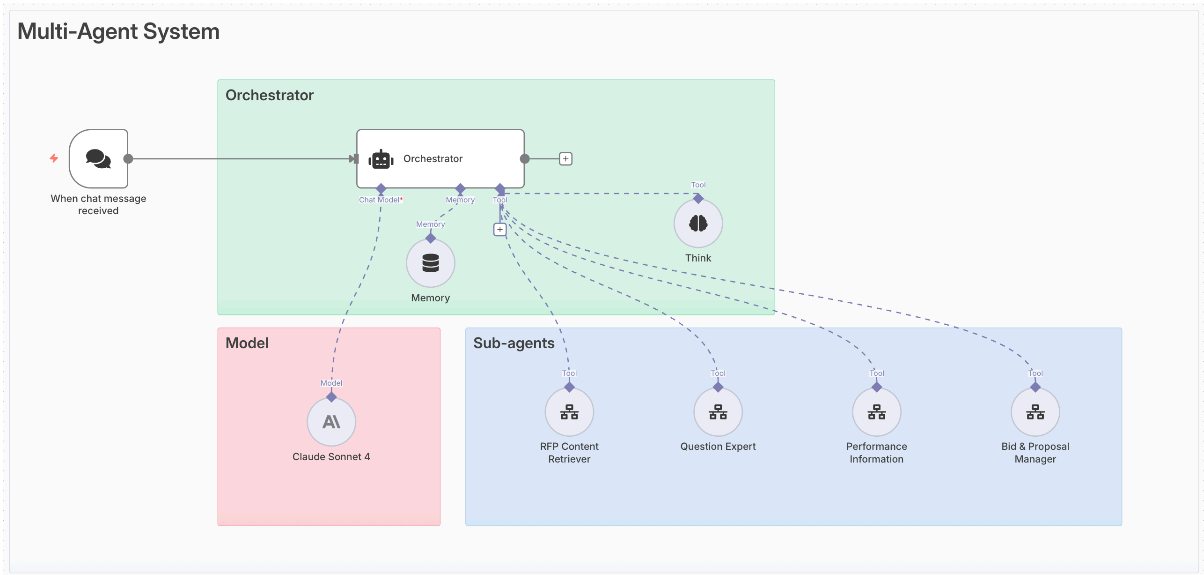

Figure 21: Overview of the multi-agent proof of concept in n8n

7.1.1 Orchestrator

Based on the literature section of orchestration strategies and early prototyping, we decided that an orchestrator is needed for dividing tasks and keeping all knowledge aligned. The platform n8n showed limitations in fully decentralized systems, making it impossible to create a fully independent system where agents have total freedom in who to call. However, as the process of proposal drafting is knowledge intensive, the system should be synchronized to prevent errors in knowledge sharing.

Including a dedicated orchestrator to manage task division forms a logical choice in this proof of concept.

In the system prompt clear descriptions of its sub-agents are given. The Orchestrator interprets the users prompt and decides which agent should be called to act. Typically, the Orchestrator is provided with a logical sequence of finding RFP specifications by calling the RFP Content Retriever, decompose the specifications to create an answer structure by calling the Question Expert, and to search for prior project information and measurable performance indicators from past projects by calling the Performance Information agent. Once this is all clear, the Orchestrator asks the Bid & Proposal Manager to create a fully compliant and convincing draft version. Although this is the most common pattern to choose when you start on a new award criteria for an RFP, the Orchestrator has the freedom to decide what steps to take based on the interpretation of the user's need.

Once a sub-agent created its output, the Orchestrator will provide the output back to the user (as it is instructed to do to provide traceable outputs based on the workshop human-AI interaction) or make the decision to complete multiple agents calls if this is what the user intended.

7.1.2 RFP Content Retriever

From semi-structured interviews with employees during contextual research it became evident that RFP teams often struggle with finding and maintaining information from the large amount of RFP specification documents and complexity of the documents. Quite a lot of time is spent on searching and going through multiple documents, searching for exact requirements, evaluation criteria or specific formulation on formatting.

Combining the fact that people could benefit from retracting information from the large set of documents effectively with the need for the MAS to have access to information from the RFP, creates the need for the RFP Content Retriever in the system. As many documents can be published for an RFP, it is assumed that in the current state most LLMs are limited by the size of their context window for these amounts of information. Next to this, passing around inputs with complete documents will result in extremely high usage of input tokens and therefore costs of the system.

The RFP Content Retriever is designed to extract information by accessing a vector database containing the uploaded RFP documents and specifications. The agent is prompted to create queries that align with the user's input. Once the agent defines the search queries, it calls the vector store to search and retrieve relevant information to the query. After, the agent will return the retrieved information and is explicitly prompted to deliver the retrieved text verbatim, rather than summarizing text or changing words. To ensure the agent effectively does so, it is told in a next step to verify all used content is indeed copied word to word. The agent is also prompted to use the Think tool, which is used as an extra step to force the agent to reflect on all procedural steps and if they have been executed properly.

The RFP Content Retriever agent outputs text extractions from RFP documents, remaining in their original form. The agent can be called by the orchestrator agent or other specialized agents when they require information from RFP documents. The output of the RFP Content Retriever should form the foundation and acts as a support for other agents in the MAS, providing them with relevant information.

7.1.3 Question Expert

One of the tasks of the current Strategic Tender Specialist is to help other RFP team members to understand how the questions from the RFP should be handled. Participants in semi-structured interviews highlight concerns about misalignment between the answer and the question. One way to improve the proposal text is to direct the questions as close as possible. Questions often exist of multiple questions phrased as one, also referred to as a compound question. These questions are particularly hard to answer well, as the reader should be aware of the meaning behind the entire question, which can be easy to miss.

The Question Expert agent is formed to look at the questions from an RFP and to assess how the question is structured. The agent then analyses if the question should be divided into sub questions or should be answered in its entirety. The agent has instructions that clarify the benefits of splitting the question and the downsides of that approach.

A three-step operating protocol is used in this agent. First, the agent carefully reads the question and tries to understand the full scope of the question. Second, it evaluates the reasoning for dividing the question into sub questions and or to formulate a single response addressing the question in its entirety. The agent uses factors such as assessor readability and response length together with logical coherence to decide. Finally, the agent provides a structured template on how the question could be answered in a draft.

The output of the Question Expert agent shows an explanation of the decision on how to structure a response. Next to this, a template is formed to guide the next agent or users to develop a complete response that directs the question effectively.

The Question Expert can be called early in the drafting process to establish a foundation that can be used for further content development. A shared understanding on how complex questions should be approached should be formed to have all team members aligned and create higher quality answers that direct the question well.

7.1.4 Performance Information

From the thematic analysis in initial semi-structured interviews, it became clear that an emerging topic was the lack of a database with internal information that can be searched and used to better inform RFP teams in the process. Next to this, it became clear during expert consultation that it is best practice to use performance metrics to back your claim whenever you can.

The Performance Information agent is described as an agent that can be called when there is a need to provide other agents with performance metrics or company background information. To realize this, the Performance Information agent has access to a vector store as a tool. This vector store contains company files for information. It is important to note that due to the sensitive nature of RFP documents, for both clients of Schuberg Philis as Schuberg Philis itself, no confidential information is used in this proof of concept. Material that is available in the vector store comes from public sources or is generally known knowledge that is available to the public. For real applications this would consume company data sources. To realize this, a security assessment would have to be made before this can be implemented for confidential information on enterprise level.

The Performance Information agent can get called by the orchestrator in multiple phases. At the core, the Performance Information agent will review an input and decide what performance indicators could be used to back the quality claim. Performance indicators are provided on three different levels: on a chapter or award criterium, on a paragraph or individual question, or per measure.

The Performance Information agent can be used by the user to come up with ideas on performance indicators or can provide direct suggestions that could be used by the Bid & Proposal Manager agent to be included in the drafted text.

7.1.5 Bid & Proposal Manager

During the participant observation in contextual research, the need for expertise in proposal writing is needed to maintain consistent quality in the delivering of proposals. Maintaining quality across different team members can be improved by adhering to best practices of proposal writing. Participants in semi-structured interviews recognize that colleagues often have different approaches to formulating answers, leading to an inconsistent tone, structure, and level of detail.

The Bid & Proposal Manager agent functions as the Strategic Tender Specialist and orchestrates three sub-agents. This agent follows a nine-step operating procedure to ensure a quality draft is produced in line with all best practices. An initial draft is formed that serves as a foundation for refinement by leveraging other specialized agents.

The agent is connected to three specialized agents for content quality: the Banned Words agent to remove and replace weak and uncertain language with convincing text, the Proven Structure agent to apply the Why-How-What framework to the draft, and the CCD Framework agent to ensure content is compliant, convincing, and distinctive. Furthermore, the Bid & Proposal Manager is provided with a word count tool, that enables the agent to adhere to word limits (with a margin to prevent overcomplicating the generation process).

Next to using its specialized sub-agents, the Bid & Proposal Manager introduces company specific styling preferences to finalize the draft to a great standard that is evaluated by the company expert.

The output created by the Bid & Proposal Manager is a structured proposal content that should require minimal formatting and styling adjustments and that should be mostly evaluated on factual correctness. The Bid & Proposal Manager forms a barrier between specialized agents for best practices and the other specialized agents that are used to form input for the proposal content.

7.1.6 Banned Words

The training “Successful Bidding” emphasized how weak and passive language weakens proposal credibility. Combining this with the fact that semi-structured interviews resulted in insights that participants want to reflect confidence and assertiveness in writing important aspects to match the company culture, it becomes evident that word choice is of high importance for proposal content.

The Banned Words agent systematically identifies weak language based on a list of banned words in English and Dutch. Not only does the agent eliminate certain words and phrases, it creates suggestions for replacement for rewriting the content.

The agent delivers a table showing what weak words are recognized alongside their replacements that are used to create a more decisive text. The agent can only be called by accessed through the Bid & Proposal Manager.

7.1.7 Proven Structure

The training “Successful Bidding” introduced the Why-How-What framework as a proven structure to formulate responses. Colleagues tend to overcomplicate answers and struggle drafting their answers in a convincing and credible manner.

The Proven Structure agent takes input content and reorganizes the content according to the Why-How-What sequence, with instructions to not explicitly name these sections as headers. The agent evaluates the current structure against the framework and identifies any missing components. Afterwards it creates restructured content, demonstrating credibility (why), explains the process (how), and present deliverables and value that is added (what).

Alongside the restructured content that follows the proven structure, the agent provides a summary of the structural changes that are made to show traceability for users. The delivered output should improve the proposal content by increasing the readability for assessment by the client.

7.1.8 CCD Framework

The training “Successful Bidding” introduced the CCD-Framework (complaint, convincing, distinctive) as an evaluation method for content quality. The training showed that responses must address the entire scope of the question (compliant), demonstrate certainty and credibility (convincing), and introduce a connection to both implicit and explicit client knowledge (distinctive).

The CCD Framework agent evaluates the drafted text by assessing compliance with the RFP requirements (provided by the RFP Content Retriever), suggesting improvements for convincing elements that show confidence and credibility, and ensures that some client specific elements are added to show a connection and deliver distinctive results.

The agent delivers a structured suggestion per framework element to show how the framework element could be addressed, followed by rewritten text that demonstrates how this could be fully integrated in the proposal.

7.2 Case demonstration

To demonstrate how the system will respond to basic questions, we ran a test case with a publicly available RFP. Simplified results from this interaction are shared in this section, to show what can be expected of the proof of concept. Complete input-output combinations can be found in appendix H. Prior to the interaction through the chat interface in n8n, users need to upload the RFP specifications to the multi-agent system. This is achieved by uploading the documents through a form (a native function of n8n), which is coupled to the vector store.

Figure 22: Examples of inputs and outputs from the proof of concept.

Figure 23: Examples from the proof of concept (showing traceability and usage of company information).

7.3 Domain expert validation

The evaluation the MAS was performed in the form of continuous integrated expert validation during multiple co-design session and finished with a validation session of the proof of concept. Evaluation was a central part of the iterative design sessions to ensure the system capabilities aligned with the designer's and expert's intent.

This approach was chosen to address the complexity of evaluating multi-agent systems, where traditional performance metrics are hard to define to capture practical value or usability. The evaluation consisted of real-time expert assessment of the system output and documentation of the changes that are made through screen recordings and manual logbook entries.

7.3.1 Evaluation approach and criteria

During multiple two-hour design sessions, the domain expert provided continuous assessment of agent performance, output quality, and practical utility of the output. The direct evaluation enabled immediate adjustments to the system prompts, agent orchestration, and tool configurations based on the expert's feedback.

The agents were evaluated against several criteria derived from usual RFP evaluation:

- **Content quality:** Whether the generated output met the expert's writing standards.
- **Best practice adherence:** *(depending on the section of the MAS that is evaluated)* Does it adhere to the established frameworks for formulating answers (according to the best practices).
- **Information accuracy:** Correct retrieval and implementation of information from documents.
- **Practical utility:** Whether the outputs are genuinely useful for the RFP drafting process.

Each agent and the system in its entirety were tested by following the same structure of evaluating outputs, adjusting the system, and evaluating again. The expert would assess the outputs, identify areas for improvement, guide adjustments to the prompts, and validate whether the changes were successful. By following this approach, it was ensured that agent capabilities were translated into practical value. Each session followed this outline:

1. Presentation of the output
2. Expert evaluation (how did it score on the above-mentioned criteria)
3. Document major changes or feedback
4. System adjustment (changes to the MAS)
5. Rerun the same query and evaluate again
6. Document lessons learned from session

All evaluation was documented through logbook entries in the form of a research diary and screen recording during prototyping sessions. Documentation of the evaluation and adjustments served to identify patterns in what worked versus what failed, creating the foundation for the design guidelines. The focus of each session and the learnings can be found in an overview in appendix E.

7.3.2 Outcomes

Each session provided insights for different aspects of MAS design, making the prototyping sessions a valuable research method for exploring how MAS can be designed and what their considerations are.

The final session with the domain expert, where the proof of concept was evaluated in its entirety, showed the progress that was made over the seven sessions. With the agents strictly adhering to their task descriptions and goals, the domain expert mentioned he was "speechless" on seeing the result of the final draft based on the amount of time the system ran and the little amount of information that was shared with about Schuberg Philis. About the Bid & Proposal Manager the expert enthusiastically said, "as if I see myself writing!". This showed the effectiveness of the system prompt for the Bid & Proposal Manager and its sub-agents. Next to this, the system reliably retrieved information from the uploaded RFP, showing the effectiveness of the instructions of the RFP Content Retriever and its ability to retrieve information verbatim. The Performance Information agent provided good suggestions, as the domain expert evaluated, but remained limited due to the lack of internal information that could be shared. The domain expert claimed the Question Expert made valuable structuring decisions for how a question should be answered, however, it is not entirely clear how well this structure is adhered by the Bid & Proposal Manager.

Although the proof of concept provided well written drafts in the desired style for Schuberg Philis based on the expert's opinion, the results should be tested and evaluated in a more structured manner to draw explicit conclusions on the quality of the MAS overall.

8 Design guidelines

This chapter presents design guidelines based on combined empirical evidence from the contextual research section and the prototyping and user research section. The guidelines address technical aspects such as orchestration, prompt engineering, and knowledge management, as well as human-centered factors like trust and transparency. While the guidelines are informed by the RFP case, they are applicable to other business processes with opportunities for AI-assistance.

Each guideline includes a clear instruction, rationale, and example from the proof of concept. The chapter finishes with an overview of the guidelines and their sources to understand what informed our recommendation.

Complete overview of design guidelines

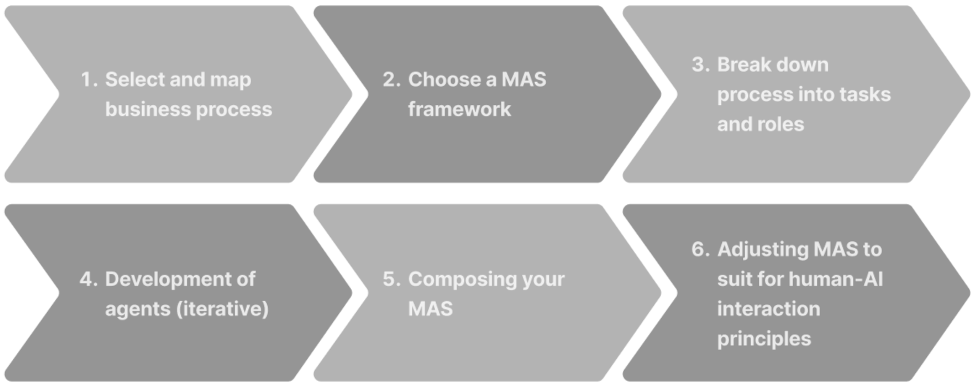

Figure 24: The design guidelines with actionable steps for translating business processes into mutli-agent systems.

1. Select the use case & map the process

Identify high-value, high-feasibility opportunities for MAS augmentation and understand the current workflow to target them effectively.

- *Identify challenges, constraints, and opportunities.* Retrieve information from stakeholders on the process to ensure a good overview of the process can be formed. It is recommended to hold interviews or workshops (more efficient) to talk about the process you are subject to.
- *Define what tasks are suitable for the MAS.* Go over each opportunity and challenge and describe why it is a good fit for AI automation or augmentation. Using a framework, such as the Turovsky framework in this research, is recommended to structurally go over all identified tasks. Tasks that show repetitive elements can be an easy opportunity to start by integrating the 'low hanging fruit'.
- *Define what the MAS should deliver per task and how.* By going over the mapped tasks, clarify exactly what outputs are expected by a MAS to deliver value. Validating the expected outputs of the MAS with stakeholders is essential to stay aligned. This can be achieved by organizing collaborative design session or more formal feedback sessions, it is recommended to involve stakeholders by going through the actual design to include them and increase their understanding of the system.

2. Choose a MAS framework that fits the use case needs

Select a development platform or environment that supports the process requirements, technical constraints, and necessary level of control.

- *Evaluate the complexity of the process.* Consider how complex the system will become in terms of number of agents, level of depth in orchestration, observability needs, and built-in tool integration.
- *Compare available MAS platforms/frameworks.* Evaluate the different candidates for developing multi-agent systems and assess their orchestration flexibility, easy of tool integration, and memory handling for your use case.

3. Decompose the process into roles and tasks

Break down the mapped process into clearly defined agent roles and responsibilities. Ensure all required outputs are covered and no unnecessary complex role division in developed.

- *Start broad, then specialize.* Develop your first agents with broader capabilities and only split roles further when there is a clear need (for example, when the agent cannot handle all tasks consistently).
- *Define each agent's scope and requirements.* For every agent, specify its goal, expected input and outputs, available tools, and operating procedure.
- *Map knowledge access.* Identify and map what knowledge each agent needs to complete their tasks and make sure they can retrieve that information by either utilizing their tools, by a shared memory, or by receiving information from other agents.

4. Develop each agent iteratively

Build one agent at a time, ensuring it is reliable and behaves as expected. Adjust the agent's system prompt and run it multiple times, before scaling it into a system where it interacts with others.

- *Use a standard agent system prompt.* Create a standard system prompt structure that can be reused for every agent. This helps with agent prompt adherence and makes it easier in development to format a complete prompt.
- *Prove agent behavior using a flagship LLM, then consider downscaling.* Validate the agent's behavior by using one of the current flagship models (for the task you are trying to achieve, for example: reasoning, coding, math, etc.). Once you can prove the agent is properly prompted to achieve its goals, consider scaling to a smaller or older model to possible save costs and latency.
- *Apply selective self-reflection for high-stake tasks.* Insert a step to let the model think about the solution, by providing a 'Think' tool or by adding a step to the operating procedure to think. Skip this step for trivial tasks but consider implementing this for high-stake outputs.
- *Add or create custom tools where LLM show limitations.* As agents are capable to use tools, these tools can form a great outcome for limitations of LLMs. Provide access to integrated tools or create your own tools by utilizing JavaScript or Python code.
- *Version everything to track progress.* Store prompt versions and examples of inputs and output pairs. This is not only useful for tracking your own development, it also can be useful when switching to different models over time to compare the agent's behavior.

5. Compose your MAS

Implement agents gradually to ensure the agents behave and communicate accordingly when placed in the system.

- *Integrate gradually in the system.* Add one new agent to the system at a time, ensuring the system succeeds at its core. Confirm information is being transferred as intended before adding a new agent.
- *Consider orchestrator agent instructions carefully.* Depending on the orchestration strategy that is chosen, the instructions on agent calling should be clear and explicit to prevent mistakes. Most often a form of an orchestrator agent will be included, ensure the orchestrator has detailed descriptions of other agent's capabilities and descriptions on when to call them.
- *Standardize agent hand-offs.* Use a structured format to ensure information flows between agents remain consistent. User facing agents should remain the flexibility to define their own structure to stay adaptable to the user's inputs.
- *Track what happens during each run.* Record what the inputs and outputs are for each step (agent level), which tools are used, how long do steps take, check if there any mistakes or errors. Go over the run afterwards to consider any necessary adjustments.
- *Test the complete system.* Give the MAS realistic tasks from all steps from the process and check whether the final outputs meet your requirements.

6. Adjusting the MAS for human-interaction principles

Actively involving end users is of importance throughout the entire development process. Once the MAS is developed to a level where its capabilities can be shared with the users effectively, meaning they show a basic understanding of what the system does, ask users to give feedback to improve the system further. This is especially important to ensure the system meets their needs, builds trust, and stays aligned with real usage patterns.

- *Offering different ways to use the MAS.* Identify the user's interaction needs or desires to understand what the outputs should be, and where human-in-the-loop is necessary, desired, or excessive.
- *Show where information comes from (traceability).* Displaying sources, reasoning, and the origin of an output enables users to verify and trust the results.
- *Be clear when the MAS cannot help.* If a user or agent request is outside its scope or data, ensure the agent informs the user about this and explain what is missing instead of letting the agent create answers (improving the risk of hallucinating).
- *Early and regular feedback loops.* If the situation allows it, schedule short and structured feedback sessions with users to align the systems usability with their needs and identify what's missing, what works well, and where more/less control is wanted.

Overview of the rationale of the guidelines

These guidelines are a result of the research process and should serve as a starting point for organizations to understand how they can augment their own processes. A table with an overview of how these guidelines were informed is shown below.

Table 9: Design guidelines and the sources that informed the decision to create the guideline.

Guideline	Key evidence source	Relevant findings in RFP use case
Select the use case & map the process	• Contextual research (interviews, training observations) • Scoping with Turovsky framework	• Knowledge fragmentation across teams • Strict word limits in RFP specifications • Preference for augmentation over automation • Turovsky mapping identified requirement alignment, prior-answer reuse, and condensing as high-value targets.
Choose a MAS framework that fits the use case needs	• MAS platform comparison • Prototyping setup decisions	• Process complexity required parallel + sequential orchestration • Need for visual debugging, strong logging, and easy tool integration • n8n selected for transparency and flexibility.
Decompose the process into roles and tasks	• Prototyping theme 1 (task distribution, granularity tests)	• Over-fragmentation caused coordination problems • Broad roles worked best initially • Structured hand-offs improved communication consistency • Knowledge mapping ensured relevant information flow
Develop each agent iteratively	• Prototyping theme 2 (prompt design, model selection, tooling integration)	• Prompt structure improved consistency • Flagship models show the best results • Custom word count tool solved LLM limitation
Compose your MAS	• Prototyping theme 1 (parallel versus sequential execution) • Prototyping theme 3 (multi-query retrievals, structured hand-offs) • Prototyping logs	• Parallel setups for independent agents improved efficiency • Sequential setups work well for dependency between agents • Multi-query retrieval of vector stores increases coverage and results • Visible logs enable quick debugging and iterating
Retrieve and incorporate user feedback	• Human-AI interaction workshop • Iterative prototyping sessions	• Traceability is essential for users (visible sources & reasoning) • Users wanted a level of control over the system for receiving intermediate results • Tension between the desire for personalization and ensured compliance

This page intentionally left blank

9 Discussion

This research explored how multi-agent systems (MAS) can be designed to support (partial) automation of complex and collaborative business processes, using proposal writing for a Request for Proposal (RFP) at Schuberg Philis as case study. The study combined literature review, qualitative contextual research (interviews and training observation), iterative prototyping and workshop on human-AI interaction guidelines.

The findings suggest that effective MAS design emerges from iterative development with stakeholders and domain experts rather than predetermined technical architectures.

9.1 Relevance of research

While LLM-based agents are not new and have been explored in research and prototyping contexts, there is little empirical work available that demonstrated their application in real business settings with complex processes and knowledge intensive use cases that normally requires human teams. As Xi et al. (2023) state, the ultimate goal of agents is not to become more powerful, but to better equip humans with agents. This resonates with our research which validates the focus on augmentation, which suits the RFP process well because it combines procedural steps, knowledge collection, and generation of content. This combination requires the attention from human teams and fits the capabilities of current AI technology. By focusing on augmentation over full automation, the research addresses both technological constraints of MAS, and touches upon organizational dynamics and preferences of humans.

This work fits in with the current developments of technology, where rapid advances in LLMs and agentic AI make it technically feasible to use artificial intelligence for reasoning in (collaborative) workflows. Next to this it fits Schuberg Philis' development of increased participation to RFPs, demanding consistent and high-quality submissions.

9.2 Discussion of results

Knowledge fragmentation as a driver for MAS integration for RFPs

The contextual research revealed that fragmented knowledge about prior proposals slows down the RFP process. Without a centralized knowledge base, teams often recreate content from scratch, which creates redundant work.

Multi-agent systems can address this issue through retrieval-augmented generation (RAG) which would allow teams to surface relevant information from prior projects by searching their knowledge base and ensure the generated content aligns with the requirements from the client.

During the prototyping phase of this research, multiple agents proved their ability to find and use relevant information to the user's query. This finding is relevant as it proves that agents and multi-agent systems can find and reuse information. However, it does not solve the current issue at SBP where their prior work is not placed in one database. The proof of concept therefore shows the potential, but implementation of a curated (as two participants mentioned that bad RFPs should not be used) and up-to-date knowledge repository is essential for deploying retrieval-based MAS components.

Agent coordination determines system success

Prototyping provided insights in the challenges that are faced on how agents coordinate and communicate with each other. Initial assumptions that highly specialized agents would automatically improve performance, as their tasks are extremely clear, proved incorrect. Instead, we discovered that starting with broader agent capabilities, and only specializing when necessary, produced more reliable systems where the intended agents are called by orchestrators. *For example, our attempt to separate document retrieval into content-specific and formatting-specific agents created unnecessary complexity without improving outputs.* This builds on literature of agent profiling where the importance of specialized agent roles are encouraged (Guo et al., 2024; L. Wang et al., 2024). Our empirical findings suggest that the level of specialization should be balanced against coordination complexity. This nuance is not fully explored in existing MAS literature.

Next to this, we found that parallel execution patterns between agents work better for independent tasks, while sequential patterns suit processes with strict dependencies. When the “Bid & Proposal Manager” agent was forced to follow a parallel execution, the processing time decreased dramatically, but more important we found that the output became more complete. It is important to understand that parallel executions will cause the orchestrator to synthesize the outputs from underlying agents, while a sequential execution will hand off outputs between agents in turn which allows them to build on the last output. This build on the trade-off between centralized and decentralized systems. The parallel execution with an orchestrator prevents coordination difficulties and synthesis of outputs, confirming literatures benefit of centralized systems in coordinated information flows (Dorri et al., 2018).

Furthermore, we found that orchestrator agents frequently made incorrect task assignments when their system prompts lacked detailed specifications on when and how to call sub-agents. This suggests that successful MAS implementation depends heavily on careful prompt engineering of agents and agent descriptions, instead of creating technical architectures.

Human-AI collaboration is personal and contextual

Our workshop revealed the tension between different potential users that cannot be identified without having a proper discussion. Using the human-AI guidelines from Amershi et al. (2019) proved to be effective for facilitating a workshop that creates rich insights. Participants of the workshop simultaneously desired transparency in AI reasoning (with intermediate steps as results) while also wanting system that could produce near-complete drafts. This tension reflects what Yang et al. (2020) identified as main challenges, the uncertainty about AI capabilities and varying complexity of the generated output. We expect this contradiction of receiving intermediate results or directly receiving end results comes from the high-stakes nature of proposal writing and personal experiences in using AI in work.

When conducting a similar study, ensuring a participant number of larger than 5 participants is recommended to improve the chances of finding contradicting statements and desires, which are valuable for the design process.

It is expected to receive contradicting desires when conducting a similar workshop. By talking directly to the potential users, it provides the opportunity for the designer to make an informed decision instead of relying on their own assumptions. This aligns with research on human-AI interaction, which stimulate the participation and involvement of users during development of AI embedded tools (Amershi et al., 2019; Cila, 2022).

Traceability emerged as the most important factor for user acceptance of answers, strongly supporting the findings from Schwartz et al. (2023). Participants consistently named they want to know “why” the system made its outputs and understand what sources informed this. As one participant mentioned: “We want to be helped, but we also want to know why. Traceability is number one.” We suggest that multi-agent systems should therefore prioritize explainability in the business context of proposal writing. The need for explainability connects to the instructor-executor paradigm, where humans maintain oversight rather than using agents as equals (Xi et al., 2023).

9.3 Practical and technical implications

Practical implications for organizations

For organizations considering implementing MAS in their business, we made a suggestion for practical design guidelines. MAS initiatives should begin with mapping the process thoroughly with active stakeholder engagement, to identify value opportunities rather than technological possibilities. These opportunities should be mapped against AI capabilities to understand if the design of a MAS for the desired tasks is feasible. The study used the Turovsky framework, which suits the current state of AI systems (Turovsky, 2023). However, different frameworks for comparisons could also be used to identify if AI is capable and is sensible decision for augmenting the process.

Based on the experience with prototyping in this research, organizations should expect an extensive iterative development to create a working MAS. Significant prompt engineering and refinement of agent coordination is needed during development, with communication patterns, memory (and knowledge) management, and orchestration logic that continuously change and require attention during active development. The iterative nature aligns with the recommendation of L. Wang et al. (2024) for human annotation methods. It is therefore recommended to take time and resources and include stakeholders during the process to evaluate the system and ensure alignment. We did not implement the LLM-as-a-judge approach described by Chan et al. (2023) and Zheng et al. (2023), which could complement human evaluation in future studies and implementations.

Next to this, the workshop as part of the research showed how user preferences on working with AI, and therefore the interaction patterns they expect with MAS, can vary. We advise to act and investigate human-AI collaboration preferences for the specific use case rather than making assumptions. Facilitating a workshop based on the workbook from Amershi et al. (2019) proved to be an effective way to collect rich insights in users thoughts on human-AI collaboration and would be suggested as a method for collecting insights.

Furthermore, an important implication for organizations is to carefully consider data security and implementation of guardrails. The limitations of this research elaborate on this.

Technical implications for MAS design

The design guidelines formed in this study offer technical implications that inform direct examples based on empirical evidence to speed up future MAS development. The prototyping section of this research serves as a practical reference that demonstrates both successes and failures.

These findings focus on the engineering layer of MAS implementation instead of the approach for rolling out a similar process as an organization. While practical adoption of a MAS depends on the alignment with stakeholders and the business process elements, the technical success of a MAS heavily depends on the orchestration, retrieval strategies, and engineered prompts and tool selection. As LLMs and capabilities of agents change over time, the relevance of the guidelines in this research should be reconsidered.

9.4 Limitations

The research presents some limitations that should be acknowledged across the scope, methods, technology, and data used. These limitations may have influenced the findings and should therefore be considered when applying them to other contexts.

9.4.1 Narrowed scope

Due to the explorative nature of this study, the proof of concept that is developed focuses on a narrow scope to remain feasible within the limited timeframe. The proof of concept in this research only applies for general RFP response drafting and does not cover all aspects of RFP processes, such as pricing strategies, or system architecture design which can be asked in IT consulting RFPs. Additionally, the focus on IT consulting RFPs for testing the proof of concept may not represent questions or RFP specifications that are regular in other sectors.

The contextual study of this research focused on mapping the RFP process. However, the RFP process is primarily mapped from the viewpoint of Schuberg Philis, which might not represent how RFPs are handled across different companies from different sizes and industries.

Moreover, the integration of real company knowledge from prior RFPs was intentionally restricted for security reasons measures to prevent any form of data sharing to third parties as training material. This choice ensured the research did not contain any confidential information, but it limited the richness of the proof of concept by using alternative documents instead.

9.4.2 Methodological constraints

The study relied on qualitative data from interviews, training observations, and a workshop. This was combined with iterative prototyping. Although these approaches provide rich insights for the context, they do not offer generalizability directly for other contexts and are dependent on the researcher's interpretation.

Next to this, the quality of the outputs of agents during iterative prototyping were evaluated by a domain expert. Although the expert knows the best practices and the RFP process well from over 20 years of experience, it forms a limitation of this study as it introduces subjectivity in the evaluation which could be prevented by using standardized evaluation with a large sample group.

Furthermore, the design guidelines which are formed by synthesizing and interpreting the results of contextual research and prototyping did not undergo a methodological validation to test the usability of the guidelines.

9.4.3 Technological limitations

As the study was focused on multi-agent system development, the prototype was bound to the capabilities of LLMs and MAS platforms available at the time of performing the research (March – August 2025). The research was performed during a period of rapid AI advancement, where the technological landscape changed noticeably over the research period of six months. As we found that the underlying model for agents plays a significant role in their output quality, specific findings on tool calling and output qualities might not hold up for newer models in the future.

Furthermore, the choice for using n8n as prototyping platform affected the experience of developing multi-agent systems significantly. Although the research compared other platforms for MAS development, the choice for n8n showed its limitations in memory sharing amongst agents and lacked the ability to create an agent swarm as orchestration option. Prototyping for a MAS with a different framework might lead to different results, which is not evaluated in this study.

Additionally, the proof of concept is tested and evaluated on a small scale. The system scalability and the readiness of multi-agent systems for higher user amounts beyond the proof of concept is not evaluated, which ultimately is necessary for implementing such systems on enterprise level.

9.4.4 Contextual constraints

Several contextual factors may have shaped the results that should be acknowledged. First, the culture at Schuberg Philis and their involvement in innovation in the IT landscape may have influenced participants view on using AI in the workspace and their openness to adopt new technologies. Throughout the research period, the innovation team (Lab271) rolled out a program to stimulate AI literacy amongst colleagues, potentially changing their perspective.

Second, the participants during all stages of the research had uneven prior experiences with AI tools which likely influenced their behavior and expectations.

9.5 Recommendations

9.5.1 Human-AI interaction studies for MAS

The different perspectives and preferences of users of multi-agent systems could be further examined in future studies. Where this study provides insights in desires from participants and clear needs for traceability and a form of control over the system, future research could focus on different ways of implementing these findings into MAS and evaluating them (appendix D provides a list of recommended interactions patterns to consider based on the conducted workshop at Schuberg Philis). Next to this, it could be an interesting to study how user's familiarity with MAS capabilities might influence their needs and perspectives.

9.5.2 Complete proposal drafting process

As participants highlighted the need for intermediate artifacts for drafting proposals, future research could focus on the complete procedure (or different segments of the process) and focus on what artifacts are necessary for most proposals and how these can effectively be created by agents.

9.5.3 Usability of design guidelines

Part of the research outcome is combining all findings from the contextual research and prototyping into design guidelines that can be used by Schuberg Philis and other organizations as an indication on how multi-agent systems can be configured. Future studies could focus on the usability of design guidelines and how they are interpreted by users who are unfamiliar with this research.

9.5.4 Balance between reusing prior content and creative writing in proposals

Future research could investigate to what extend proposals rely on reusing content from prior proposals against writing creative solutions and text without prior proposal knowledge. Our study assumed that knowledge form prior RFPs would be valuable, as participants highlighted this during interviews and fits in their current way of working. However, the workshop on human-AI interaction sparked the interest to further research to what extend this is necessary, and how creative writing could influence their chances of winning proposals.

This page intentionally left blank

10 Conclusion

This chapter brings together the findings of the research and reflect how the study addressed the main research question on how MAS can be designed and applied for proposal development (RQ). The conclusion summarizes how the different sub-questions contributed and collectively answer this main research question.

Finally, the chapter highlights the contributions of this research, for both Schuberg Philis as an organization and for broader practical understanding of MAS for business processes.

10.1 Answering the research questions

This thesis explored the rapidly evolving field of agentic AI and aimed to answer the following main research question:

How can MAS be designed and applied to support the proposal development process in B2B IT RFPs?

The MAS must be designed for augmentation, not full automation. Humans want to stay involved in the workflow to maintain trust in the system, which can be achieved by providing traceable intermediate steps. It is important to focus on retrieving and reusing company knowledge and adhere to company standards, in this case the proposal best practices. By using a centralized MAS, tasks can effectively be handled in multi-turn interactions. A balance must be found in balancing over-fragmenting and generalizing agent roles by iteratively testing performance.

SQ1: Where in the proposal development process can AI integration offer the most value?

The areas that were identified as the most suitable were from the proposal drafting process in content design and content development phases. Specifically, the following tasks were identified as useful for augmentation of the process:

1. Knowledge retrieval and reusing elements from prior answers from RFPs and internal documentation.
2. Finding requirements in lengthy RFP specifications and ensuring compliance with the questions by directly answering the question, finding and using their formatting rules, and adhering to word count limits.
3. Drafting first versions of responses that adhere to the best practices of Schuberg Philis.

SQ2: What features should a MAS include to make it company-specific and useful for generating proposal content?

To make the system company-specific, it should embed Schuberg Philis' best practices (CCD framework, their proven structure, and word-avoidance guide). Next to this, using reference material and internal knowledge is essential, the system needs to feature vector databases and RAG techniques to effectively integrate this. To adhere to word counts the MAS should embed custom tools that expand the capabilities of agents.

SQ3: What design principles and orchestration patterns between agents enable effective output?

Effective orchestration in a MAS depends on carefully mapping tasks and their dependencies, using parallel structured flows for independent activities and sequential flows where strict dependencies exist. At the same time, agents should be defined with more broad roles at the start and new agents must only be created when it presents clear benefits. Reliable collaboration between agents is improved by standardizing output hand-offs with structured output schemas. Furthermore, explicit prompts are needed to prevent agent misrouting and strong large language models are needed for prompt adherence and effective tool calling.

SQ4: How should human-agent collaboration be integrated into a MAS for RFPs to ensure trust in its decisions and output?

Human oversight and control are central to ensure trust in MAS outputs. To support human oversight traceability must be prioritized by consistently showing sources in outputs and intermediate results to demonstrate agents' reasoning paths. Users want to have some level of control over the results they receive from agents. When the agent is not capable to help, the system should fail gracefully by clearly indicating their knowledge is insufficient, rather than producing unreliable results. When agents use external sources, this should be made clear to the user.

10.2 Contributions

The research contributes by moving beyond MAS theory and by describing an explorative process for developing MAS in the context of real organization processes. For Schuberg Philis, the proof-of-concept prototype form a foundation for further experimentation and development of proposal writing augmentation.

The design guidelines provide actionable steps for organizations to experiment and develop MAS themselves, making the steps towards MAS development more accessible and structured. Ultimately, the thesis demonstrates that although MAS are not yet readily deployable, they hold potential for adding value for complex business processes when carefully designed.

This page intentionally left blank

11 Reflection

This chapter reflects on the research journey with a focus on personal learning from the process and challenges encountered while working on the thesis. It discusses how the broad topic of AI and multi-agent systems shaped my perspectives as a designer, the role of Schuberg Philis in supporting the project, and the difficulties of working on a rapidly developing research topic.

The chapter also reflects on my learnings from research methods and critical decision-making to maintain a research focus, and how I would include stakeholders in future projects.

Working on this thesis has been a challenging and interesting journey. The topic of AI and creating agentic systems sparked my interest right from the start and I constantly reflected on the role of a designer in this field, which did not seem trivial to me at first. However, as I noticed that everyone is using AI in their own way, it became apparent to me how important it is to design systems that can effectively be used by users, without the direct need to be experienced with AI.

It turned out that working on this project at Schuberg Philis was a great decision. The company provided me with all necessary resources and countless conversations about AI, agents, and what they have to offer. Soon after the start of my project, the Lab271 started their companywide AI literacy campaign, which meant that not only my direct mentors, but all colleagues were interested in the research topic.

The research process provided its challenges. I struggled with effectively and consistently writing down my interpretations and findings and was always on the search for more information to substantiate a claim. By discussing this with my supervisors and mentors, I improved in reflecting my own thoughts and gained confidence in decision making. It was hard to keep up with advancements in the field of AI agents, as a lot has changed over the 6-month period of the project. Although this encouraged me to work on the proof of concept and to better understand how the field improved myself, it also made it feel like I was always one step behind.

During this research I learned to constantly think from other perspectives and challenge my own thoughts. Interviews proved to be great at early stages of the research and to learn about the current situation. A difficulty of this research was the human aspect of working with AI, as the workshop successfully pointed out that people have different perspectives on how this should be done. For future projects, I would therefore facilitate similar workshops during an earlier stage of the research to validate directions sooner. Prototyping with stakeholders presented an ideal opportunity to explore more tacit knowledge as the outputs sparked new conversation topics. This approach was not only of high value for the research but turned out to be a highlight of the process, as the seven sessions provided interesting and enjoyable conversations.

Reflecting on the project I am thankful for everyone that took their time to help me and challenged me. The past months provided a valuable learning experience and helped in preparing me for the next chapters that will follow.

12 References

Airia. (n.d.). *Airia | Enterprise AI Platform for Secure & Scalable Solutions*. Airia. Retrieved 20 August 2025, from https://airia.com/

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for Human-AI Interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300233

Anthropic. (n.d.). *Building Effective AI Agents \ Anthropic*. Retrieved 1 April 2025, from https://www.anthropic.com/engineering/building-effective-agents

Anthropic. (n.d.). *Context windows*. Anthropic. Retrieved 17 July 2025, from https://docs.anthropic.com/en/docs/build-with-claude/context-windows

Anthropic. (2025, March 20). *The 'think' tool: Enabling Claude to stop and think in complex tool use situations*. https://www.anthropic.com/engineering/claude-think-tool

ATLAS.ti. (n.d.). *Coding Qualititive Data | Overview, Methods, Software*. ATLAS.Ti. Retrieved 8 April 2025, from https://atlasti.com/guides/qualitative-research-guide-part-2/data-coding

Boost HR Efficiency with IBM Watsonx: Avid Solutions' Success Story. (n.d.). Retrieved 16 August 2025, from https://www.cut-the-saas.com/ai/avid-solutions-use-of-ibm-watsonx-orchestrate-in-hr-and-the-outcomes

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). *Language Models are Few-Shot Learners* (Version 4). arXiv. https://doi.org/10.48550/ARXIV.2005.14165

Brynjolfsson, E., Li, D., & Raymond, L. (2023). *Generative AI at Work* (No. w31161; p. w31161). National Bureau of Economic Research. https://doi.org/10.3386/w31161

Byrne, D. (2022). A worked example of Braun and Clarke's approach to reflexive thematic analysis. *Quality & Quantity*, *56*(3), 1391–1412. https://doi.org/10.1007/s11135-021-01182-y

Chan, C.-M., Chen, W., Su, Y., Yu, J., Xue, W., Zhang, S., Fu, J., & Liu, Z. (2023). *ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2308.07201

Cheng, Y., Zhang, C., Zhang, Z., Meng, X., Hong, S., Li, W., Wang, Z., Wang, Z., Yin, F., Zhao, J., & He, X. (2024). *Exploring Large Language Model based Intelligent Agents: Definitions, Methods, and Prospects* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2401.03428

Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2023). The economic potential of generative AI: The next productivity frontier. *McKinsey & Company*.

Cila, N. (2022). Designing Human-Agent Collaborations: Commitment, responsiveness, and support. *CHI Conference on Human Factors in Computing Systems*, 1–18. https://doi.org/10.1145/3491102.3517500

Coshow, T., Gao, A., Pingree, L., Verma, A., Schreibenreif, D., Khandabattu, H., & Olliffe, G. (2024). *Top Strategic Technology Trends for 2025: Agentic AI*. Gartner, Inc. https://www.gartner.com/doc/reprints?id=1-2K8Y7LEY&ct=250212&st=sb

CrewAI. (n.d.). *The Leading Multi-Agent Platform*. Retrieved 20 August 2025, from https://www.crewai.com/

Deloitte AI Institute. (2024, November). *Prompting for action: How AI agents are reshaping the future of work*. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-ai-institute-generative-ai-agents-multiagent-systems.pdf

Dibia, V., Chen, J., Bansal, G., Syed, S., Fourney, A., Zhu, E., Wang, C., & Amershi, S. (2024). *AutoGen Studio: A No-Code Developer Tool for Building and Debugging Multi-Agent Systems* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2408.15247

Dong, Q., Li, L., Dai, D., Zheng, C., Ma, J., Li, R., Xia, H., Xu, J., Wu, Z., Liu, T., Chang, B., Sun, X., & Sui, Z. (2023). *A Survey on In-context Learning* (Version 6). arXiv. https://doi.org/10.48550/ARXIV.2301.00234

Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-Agent Systems: A Survey. *IEEE Access*, *6*, 28573–28593. https://doi.org/10.1109/ACCESS.2018.2831228

Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2023). *Improving Factuality and Reasoning in Language Models through Multiagent Debate* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2305.14325

Google Cloud. (n.d.). *Prompt Engineering: Overview and guide*. Google Cloud. Retrieved 15 May 2025, from https://cloud.google.com/discover/what-is-prompt-engineering

Google Cloud. (2024). *AI Business Trends 2025*. Google Cloud. https://blog.google/products/google-cloud/ai-trends-business-2025/

Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). *Large Language Model based Multi-Agents: A Survey of Progress and Challenges* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2402.01680

Gutowska, A. (2024, August 6). *What is a Multiagent System? | IBM*. https://www.ibm.com/think/topics/multiagent-system

Han, S., Zhang, Q., Yao, Y., Jin, W., Xu, Z., & He, C. (2024). *LLM Multi-Agent Systems: Challenges and Open Problems* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2402.03578

Hwang, A. H.-C. (2022). Too Late to be Creative? AI-Empowered Tools in Creative Processes. *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–9. https://doi.org/10.1145/3491101.3503549

Jiang, G., Xu, M., Zhu, S.-C., Han, W., Zhang, C., & Zhu, Y. (2023). *Evaluating and Inducing Personality in Pre-trained Language Models* (No. arXiv:2206.07550). arXiv. https://doi.org/10.48550/arXiv.2206.07550

Kallio, H., Pietilä, A., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: Developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, *72*(12), 2954–2965. https://doi.org/10.1111/jan.13031

Kawulich, B. B. (2005). Participant Observation as a Data Collection Method. *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research*, *6*(2), Article 2. https://doi.org/10.17169/fqs-6.2.466

Lagerkrans-Pandey, S. (2025, January 16). Customer Case Study: Pushing the Boundaries of Multi-Agent AI Collaboration with ServiceNow and Microsoft Semantic Kernel. *Semantic Kernel*. https://devblogs.microsoft.com/semantic-kernel/customer-case-study-pushing-the-boundaries-of-multi-agent-ai-collaboration-with-servicenow-and-microsoft-semantic-kernel/

LangChain. (2025, January 20). How Minimal built a multi-agent customer support system with LangGraph & LangSmith. *LangChain Blog*. https://blog.langchain.dev/how-minimal-built-a-multi-agent-customer-support-system-with-langgraph-langsmith/

Lauesen, S., & Vium, J. P. (2005). Communication gaps in a tender process. *Requirements Engineering*, *10*(4), 247–261. https://doi.org/10.1007/s00766-005-0009-2

Li, G., Hammoud, H. A. A. K., Itani, H., Khizbullin, D., & Ghanem, B. (2023). *CAMEL: Communicative Agents for 'Mind' Exploration of Large Language Model Society* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2303.17760

Li, J., Zhang, Q., Yu, Y., Fu, Q., & Ye, D. (2024). *More Agents Is All You Need* (No. arXiv:2402.05120). arXiv. https://doi.org/10.48550/arXiv.2402.05120

Li, X., Wang, S., Zeng, S., Wu, Y., & Yang, Y. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. *Vicinagearth*, *1*(1), 9. https://doi.org/10.1007/s44336-024-00009-2

Liang, T., He, Z., Jiao, W., Wang, X., Wang, Y., Wang, R., Yang, Y., Shi, S., & Tu, Z. (2023). *Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate* (Version 4). arXiv. https://doi.org/10.48550/ARXIV.2305.19118

Liu, Y., Iter, D., Xu, Y., Wang, S., Xu, R., & Zhu, C. (2023). *G-Eval: NLG Evaluation using GPT-4 with Better Human Alignment* (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2303.16634

Loopio. (n.d.). Loopio AI | AI-Powered RFP Software. *Loopio*. Retrieved 16 August 2025, from https://loopio.com/ai/

MacCallum, N., & Lee, J. (2025, April 14). *GPT-4.1 Prompting Guide | OpenAI Cookbook*. GPT-4.1 Prompting Guide. https://cookbook.openai.com/examples/gpt4-1_prompting_guide

McKinsey & Company. (2024). *What is a context window?* (McKinsey Explainers). McKinsey & Company. https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-context-window

McKinsey & Company. (2025). *What is an AI agent?* (McKinsey Explainers). McKinsey & Company. https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-an-ai-agent

Miles, M. B., & Huberman, A. M. (1994). *Qualitative data analysis: An expanded sourcebook* (2. ed., [Nachdr.]). Sage.

Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X., & Gao, J. (2025). *Large Language Models: A Survey* (No. arXiv:2402.06196). arXiv. https://doi.org/10.48550/arXiv.2402.06196

n8n. (n.d.). *n8n.io—AI workflow automation tool*. Retrieved 20 August 2025, from https://n8n.io/

Naeem, M., Ozuem, W., Howell, K., & Ranfagni, S. (2023). A Step-by-Step Process of Thematic Analysis to Develop a Conceptual Model in Qualitative Research. *International Journal of Qualitative Methods*, *22*, 16094069231205789. https://doi.org/10.1177/16094069231205789

OpenAI. (n.d.). *Tokenizer—OpenAI API*. Retrieved 16 July 2025, from https://platform.openai.com

OpenAI. (2025). *ChatGPT*. https://openai.com/chatgpt/overview/

Opentender Netherlands. (n.d.). *Opentender Netherlands - IT services: Consulting, software development, Internet and …*. Retrieved 3 April 2025, from https://opentender.eu/nl/sector/72

Paech, B., Heinrich, R., Zorn-Pauli, G., Jung, A., & Tadjiky, S. (2012). Answering a Request for Proposal – Challenges and Proposed Solutions. In B. Regnell & D. Damian (Eds), *Requirements Engineering:*

Foundation for Software Quality (Vol. 7195, pp. 16–29). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-28714-5_2

Palantir. (2025, January 16). *Evaluating Generative AI: A Field Manual*. Medium. https://blog.palantir.com/evaluating-generative-ai-a-field-manual-0cdaf574a9e1

People + AI Guidebook. (n.d.). Retrieved 12 May 2025, from https://design.google/ai-guidebook

Perez, J., Kovač, G., Léger, C., Colas, C., Molinaro, G., Derex, M., Oudeyer, P.-Y., & Moulin-Frier, C. (2025). *When LLMs Play the Telephone Game: Cultural Attractors as Conceptual Tools to Evaluate LLMs in Multi-turn Settings* (No. arXiv:2407.04503). arXiv. https://doi.org/10.48550/arXiv.2407.04503

PIANOo. (n.d.). *Public Procurement in the Netherlands | PIANOo—Dutch Public Procurement Expertise Centre*. Retrieved 1 April 2025, from https://www.pianoo.nl/en/public-procurement-netherlands

PwC. (2024). *Agentic AI – the new frontier in GenAI*.

Rajbhoj, A., Nistala, P., Kulkarni, V., & Ganesan, G. (2019). A RFP System for Generating Response to a Request for Proposal. *Proceedings of the 12th Innovations in Software Engineering Conference (Formerly Known as India Software Engineering Conference)*, 1–9. https://doi.org/10.1145/3299771.3299779

Russell, S. J., & Norvig, P. (with Chang, M., Devlin, J., Dragan, A., Forsyth, D., Goodfellow, I., Malik, J., Mansinghka, V., Pearl, J., & Wooldridge, M. J.). (2022). *Artificial intelligence: A modern approach* (Fourth edition, global edition). Pearson.

Schwartz, S., Yaeli, A., & Shlomov, S. (2023). *Enhancing Trust in LLM-Based AI Automation Agents: New Considerations and Future Challenges* (No. arXiv:2308.05391). arXiv. https://doi.org/10.48550/arXiv.2308.05391

Talebirad, Y., & Nadiri, A. (2023). *Multi-Agent Collaboration: Harnessing the Power of Intelligent LLM Agents* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2306.03314

Turovsky, B. (2023, February 2). *Framework for evaluating Generative AI use cases | LinkedIn*. https://www.linkedin.com/pulse/framework-evaluating-generative-ai-use-cases-barak-turovsky/

Vaughn, L. M., & Jacquez, F. (2020). Participatory Research Methods – Choice Points in the Research Process. *Journal of Participatory Research Methods*, *1*(1). https://doi.org/10.35844/001c.13244

Wang, D., Churchill, E., Maes, P., Fan, X., Shneiderman, B., Shi, Y., & Wang, Q. (2020). From Human-Human Collaboration to Human-AI Collaboration: Designing AI Systems That Can Work Together with People. *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–6. https://doi.org/10.1145/3334480.3381069

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, *18*(6), 186345. https://doi.org/10.1007/s11704-024-40231-1

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models* (Version 6). arXiv. https://doi.org/10.48550/ARXIV.2201.11903

Weng, L. (2023, March 15). *Prompt Engineering*. https://lilianweng.github.io/posts/2023-03-15-prompt-engineering/

Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., White, R. W., Burger, D., & Wang, C. (2023). *AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2308.08155

Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., … Gui, T. (2023). *The Rise and Potential of Large Language Model Based Agents: A Survey* (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2309.07864

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376301

Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., & Narasimhan, K. (2023). *Tree of Thoughts: Deliberate Problem Solving with Large Language Models* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2305.10601

Zhang, Y., Li, Y., Cui, L., Cai, D., Liu, L., Fu, T., Huang, X., Zhao, E., Zhang, Y., Chen, Y., Wang, L., Luu, A. T., Bi, W., Shi, F., & Shi, S. (2023). *Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models* (No. arXiv:2309.01219). arXiv. https://doi.org/10.48550/arXiv.2309.01219

Zhang, Z., Bo, X., Ma, C., Li, R., Chen, X., Dai, Q., Zhu, J., Dong, Z., & Wen, J.-R. (2024). *A Survey on the Memory Mechanism of Large Language Model based Agents* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2404.13501

Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., & Stoica, I. (2023). *Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena* (No. arXiv:2306.05685). arXiv. https://doi.org/10.48550/arXiv.2306.05685

Zhou, Y., Muresanu, A. I., Han, Z., Paster, K., Pitis, S., Chan, H., & Ba, J. (2022). *Large Language Models Are Human-Level Prompt Engineers* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2211.01910

13 Appendices

A: Interview guide

This guide is maintained for the semi-structured interviews with participants in the contextual research phase, with the aim to gain a basic understanding of the RFP process, discover their challenges, and find opportunities.

Before the interviews begin, participants are asked permission for the session to be audio recorded and are asked to sign the informed consent form.

Roles involved

1. "Can you tell me about the roles involved in answering RFPs here and your role specifically?"

2. "How long does the process take (and does it differ per role)?"

Drafting responses

3. "What's the SBP/your approach to drafting an RFP response?"

4. "Are there any guidelines or templates that you use and what are they?"

5. "How (and who does) do you evaluate the responses (and ensure the response is aligned with the requester's needs and expectations)?"

6. "What are typical standard document requirements?"

Collaboration and communication

7. "How does the team work together (and share information) during the drafting process of the RFP?"

8. "What aspects of team interaction that could be improved?"

Pain points

9. "What are the common challenges in drafting RFP responses?"

10. "Are the any parts of the process that you believe are particularly inefficient?"

Opportunities for improvement

11. "Where do you see room for improvement or potential automation in the process?

12. "Have you experimented with automation and how?"

Additional questions:

1. Could you share documents of a RFP you have been or are working on?

2. Any suggestion of people I should talk to?

3. Would you be able to help me in the future?

B: Thematic analysis findings

			1 Participant 1.pdf 31	2 Participant 2.pdf 36	3 Participant 3.pdf 27	4 Participant 4.pdf 16	5 Participant 5.pdf 11	6 Participant 6 N... 21	7 Participant 7.pdf 67	Totals
Best practices	7	9	1			7			1	9
Current challenges	17	24	2		6	3	3	4	6	24
Current use of AI	17	20	1	5		2		3	9	20
Database	6	8	3		1				4	8
Difference in processes	9	16	1	2			1	4	8	16
Goal of the use of AI	13	18	2	2			3	2	9	18
Possibilities	18	28	3	2	7	1	2		13	28
Process timeline	10	18	6	7	2			1	2	18
RFP content	26	46	6	11	4	3	3	6	13	46
RFP structure	14	39	8	8	4		2	3	14	39
Roles in RFP process	31	34	4	11		6	2	1	10	34
Writing style	13	14	6	2		4			2	14
Totals			43	50	24	26	16	24	91	274

Figure 25: Main codes and their frequency.

The main themes that are derived from the interviews are elaborated below with the key finding.

Main themes (clusters) and findings

Process challenges

- Lack of central database (finding information)
- Information management with changing requirements and updated work
- Responses misaligned with the question

The RFP process faces significant efficiency challenges due to lack of centralized knowledge, time constraints, and content misalignment. Poor information management leads to repetitive work, with teams struggling to keep up with question rounds and changing requirements. Responses often fail to directly address questions, focusing on branding over clarity.

Bidding team structure

- Multidisciplinary team from 4-12 people
- Subject Matter Experts
- External or internal bid manager involved

Successful RFP responses rely on clearly defined roles within cross-functional teams (typically 4-12 people). The bid/proposal manager serves as the experienced guide, with specialized team members (sales, technical leads, business consultants) contributing domain expertise. External specialists are often brought in to optimize responses.

Developed best practices

- Framework with building blocks, roadmap, and writing rules
- Evaluation against assesment matrices
- Role of Bid Manager to adhere to and develop winning strategy

A structured framework for RFP responses exists, featuring building blocks, roadmaps, and writing rules. Evaluation against assessment matrices helps ensure quality, with experienced bid managers playing a crucial role in maintaining coherence and developing winning strategies.

AI Integration possibilities

- Improving information management and content alignment

- Assesment of answers
- Overall quality and efficiency improvement

AI shows significant potential for enhancing the RFP process through improved information management, content analysis, and creative enhancement. Current AI applications focus on efficiency (drafting first concepts, summarizing content) and quality improvement (precision in writing, error checking). Future opportunities include using AI for competition analysis, knowledge management, and more sophisticated assessment capabilities.

Writing style and content quality

- Prioritizing clarity, confidence, and directness over corporate storytelling and jargon
- Direct and concise addressing of the question
- Challenge to differentiate within the limits

Successful RFP responses prioritize clarity, confidence, and directness over corporate storytelling and jargon. Responses should be concise, directly address questions, and follow established best practices. Content differentiation remains a challenge, with teams seeking ways to make proposals stand out.

Process management

- Evaluation and rewriting are critical and the most time-consuming phases
- Building in time for quality control and checking on consistency is crucial
- RFPs work in a funnel, from multiple bidders to a handful

RFP timelines vary significantly (from weeks to a year), with evaluation and rewriting identified as the most critical and time-consuming phases. The process follows a funnel approach from multiple suppliers to one, with the final selection being most intensive. Building in time for quality control and consistency checks is essential.

C: Workshop human-AI interaction

The workshop on human-AI interaction based on the guidelines was held at Space, a workshop and presentation room at Lab271 at Schuberg Philis. In this appendix the slides from the workshop can be found accompanied by pictures of the session. Accompanied by the slide, I presented a live demonstration of the proof of concept for participants to better understand what a multi-agent system entails and what it can achieve in the current state.

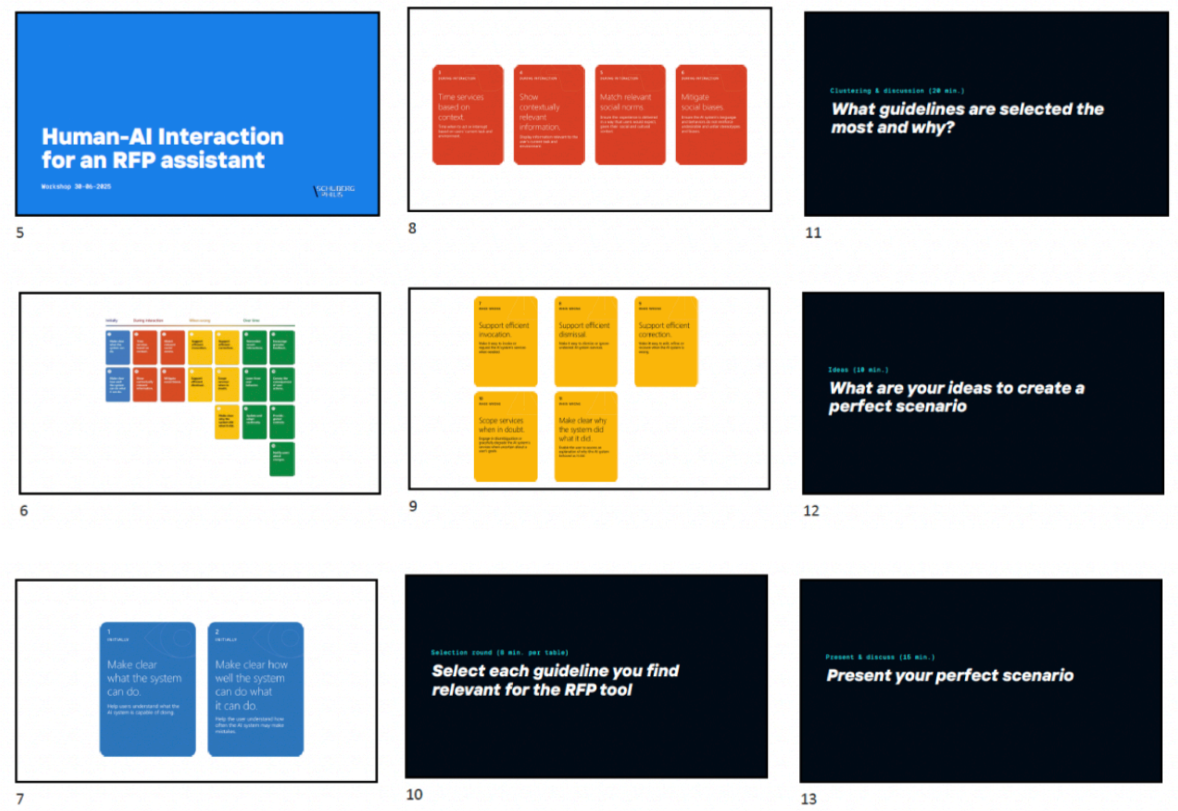

Figure 26: Workshop slides to guide participants through the session.

Figure 27: Workshop setup

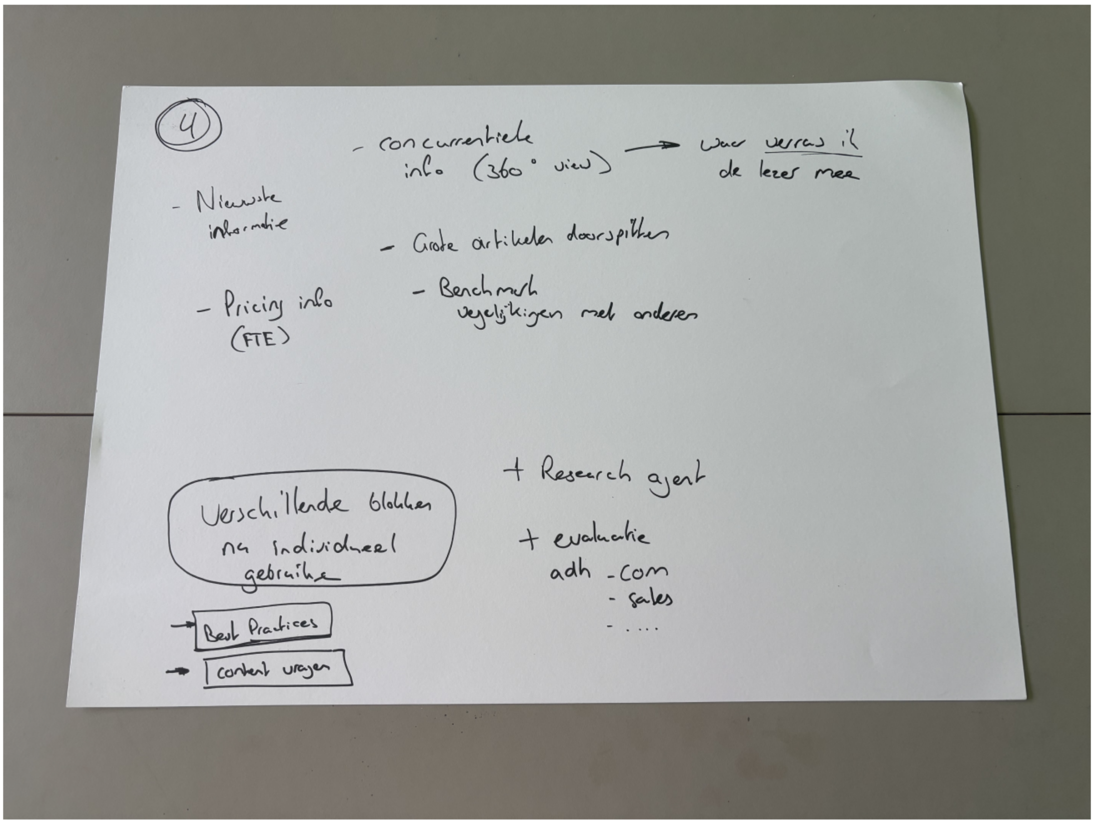

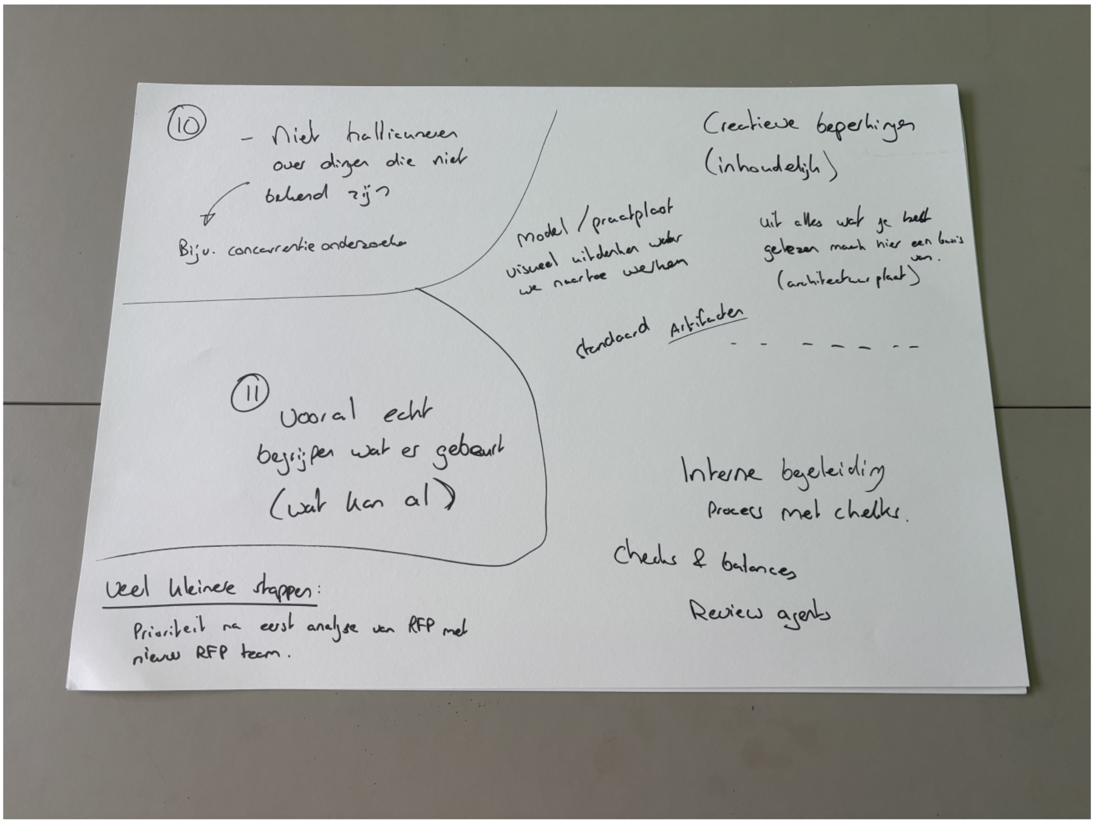

Figure 28: Overview of ideas during discussion (the numbers represent guidelines numbers).

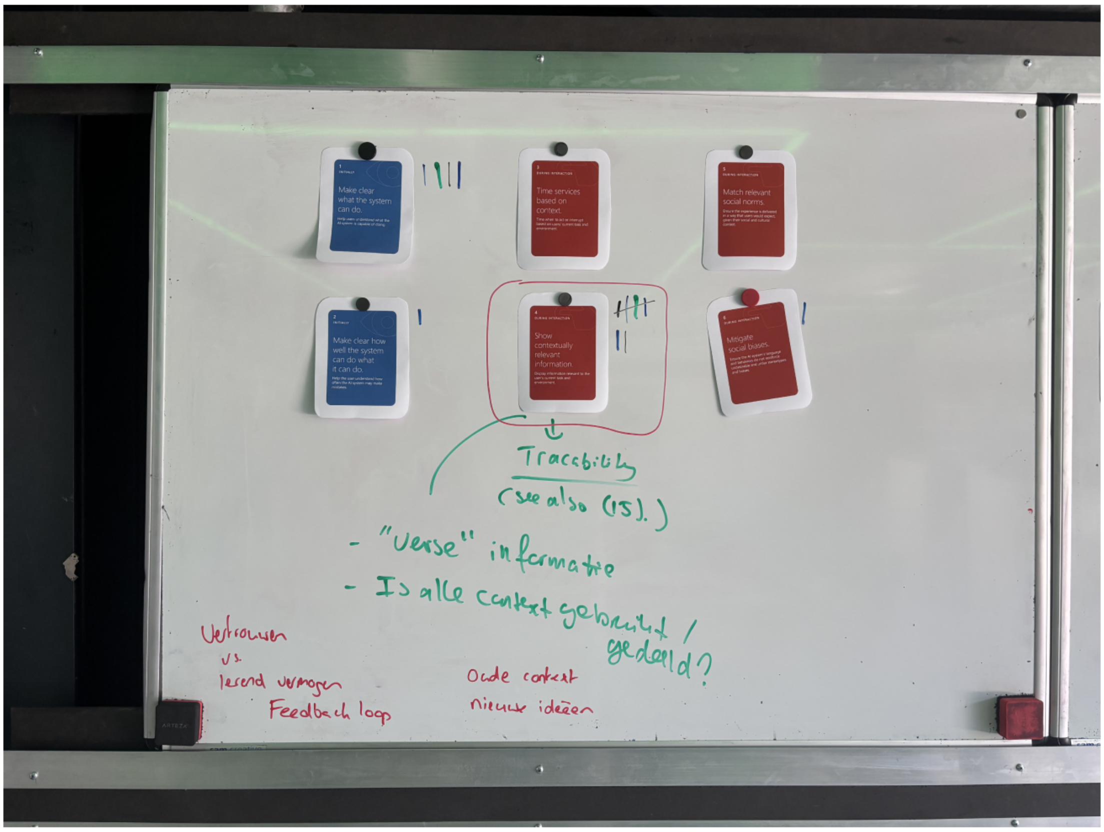

Figure 29: Guidelines selection process with votes

D: Recommendations for human-AI interaction studies

Design suggestions based on workshop

This section provides recommendations of how future studies could investigate the possibilities of creating human-AI interaction patterns based on what we found during our workshop. The sections below are focused on design patterns based on the design guidelines from Amershi et al. (2019).

What becomes clear is that people need control and trust in how the system operates. As the current prototype is built on a platform to create the functionality of a MAS, a frontend would be needed to let users effectively interact with the system of multiple agents. From the findings of the workshop, certain guidelines are elaborately discussed as they are relevant for the RFP scenario. This section describes suggestions on how human-AI design guidelines, and their design patterns could be introduced to a MAS frontend design.

Guideline 1: make clear what the system can do

Guideline 1 (make clear what the system can do) can be easily implemented in a frontend, to support users in understanding why and how the tool can support them in the process. Design pattern 1A (introductory blurb) & 1B (use explanation patterns) focus on providing general explanations. Design pattern 1C (expose system controls) can be introduced to explain how adjustments can be made to the system. To create a functionality to control the system from the UI frontend, adjustments must be made to the system in the MAS platform backend. Other explanation patterns that could work for the RFP MAS scenario are pattern 1D (demonstrate possible system inputs), this could improve the users understanding on the systems capability and improve user adoption.

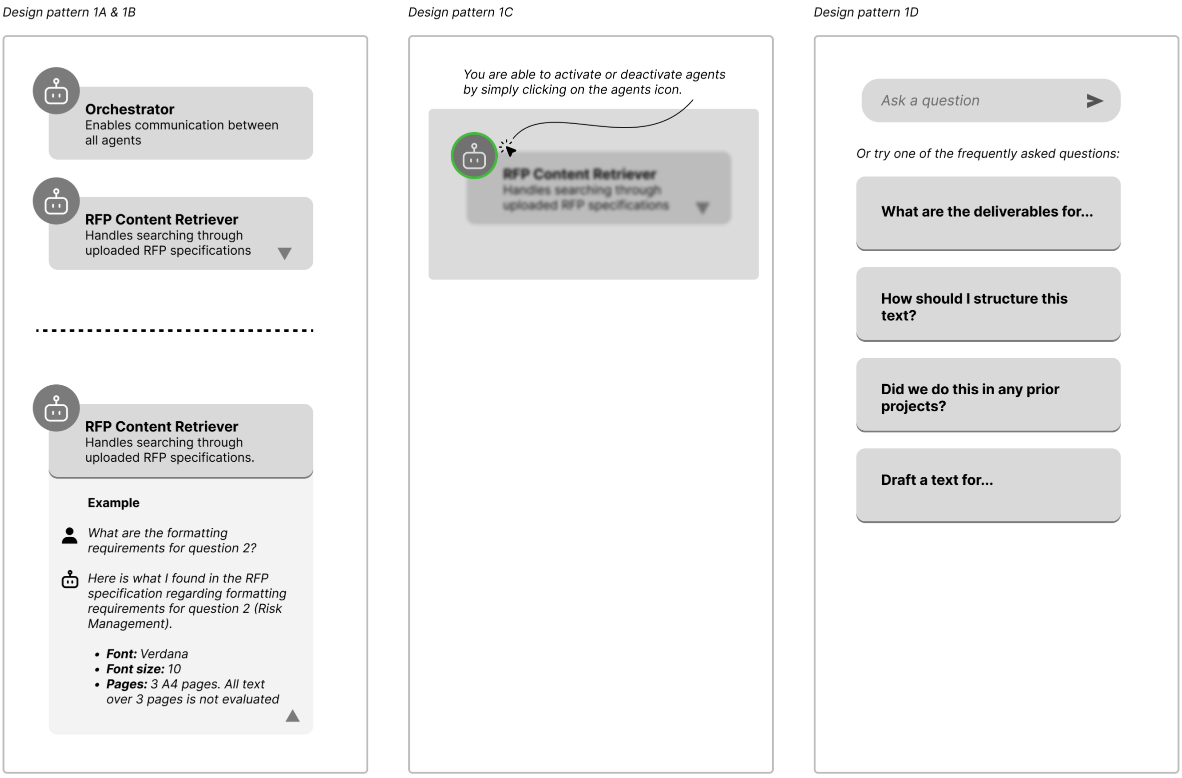

Figure 30: Ideas for implementing design pattern 1

Guideline 4: show contextually relevant information

Guideline 4 (show contextually relevant information) is the highest voted guideline from the workshop. Interestingly, this guideline does not focus on providing new information, but only increases the users trust by demonstrating the system understands the user's intent and contributes to the user experience, as the user's immediate needs are supported. Functionality in the MAS are necessary to provide these extra insights and should be introduced as system outputs.

Design pattern 4

Figure 31: Ideas for implementing design pattern 4.

Guideline 9: support efficient correction

Guideline 9 (support efficient correction) is essential for providing users control over the AI generated output. Although users can opt to manually copy and paste sections of generated text, supporting different ways to create variations, local adjustments or control by revising an output with new parameters (such as the active agents in the agent team) allow users effectively use the system. The most important design pattern that is applicable to the scenario is design pattern 9B (rich and detailed edits) as this is in line with the pattern of reviewing RFP text. The frontend controls should be connected to trigger new executions in the MAS to support this functionality.

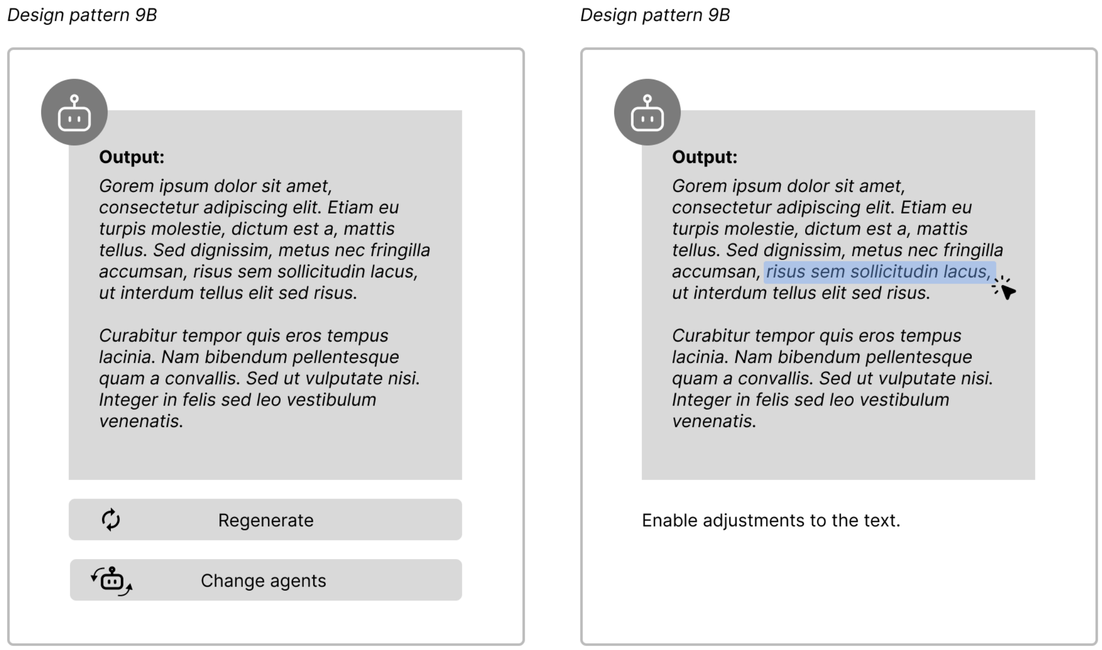

Figure 32: Ideas for implementing design pattern 9.

Guideline 10: scope services when in doubt

Guideline 10 (scope services when in doubt) was an elaborate discussion point during the workshop. Participants desire some control over output that is out-of-scope, rather than being limited to in-scope outputs. To provide this control over uncertainties in the MAS, the guideline helps to avoid abandonment of the system. Design pattern 10B (avoid cold starts by eliciting user preferences) can be implemented in the MAS to include user preferences directly from the start (combined with guideline 17: provide global controls). Design pattern 10C (fall back to other strategies) could form an alternative to provide user control over ambiguous situations. Both design considerations would require implementation in the MAS to effectively change agent's abilities.

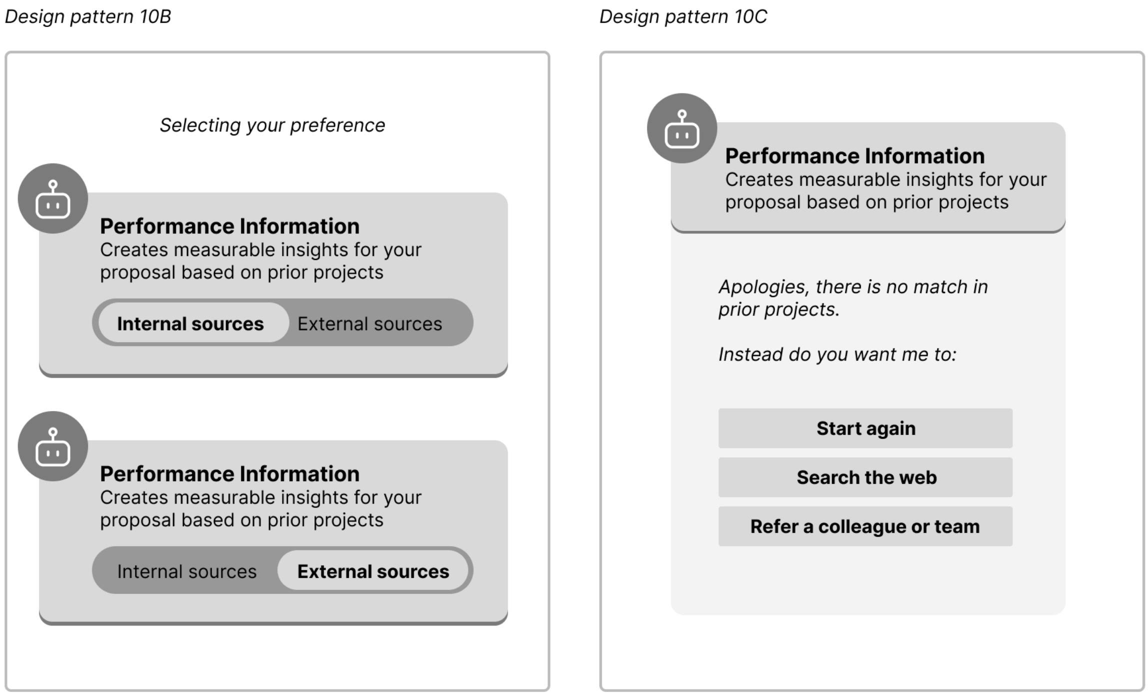

Figure 33: Ideas for implementing design pattern 10.

Guideline 11: make clear why the system did what it did

Guideline 11 (make clear why the system did what it did) is related to traceability and trust of AI systems. Design pattern 11A (local explanations) and 11B (global explanations) are relatively easy adjustments to provide clarity. In the situation of multi-agent systems, this could show traceability in the form of logs or intermediate steps, highlighting which agent provided what output. However, this can also be implemented on a single agent or on output level, showing relevant information on what decisions are made, or how long it took to generate an outcome. The 'think' module that is discussed in section 6.3.2 is an example on how users could be introduced to decision making and thinking steps of agents.

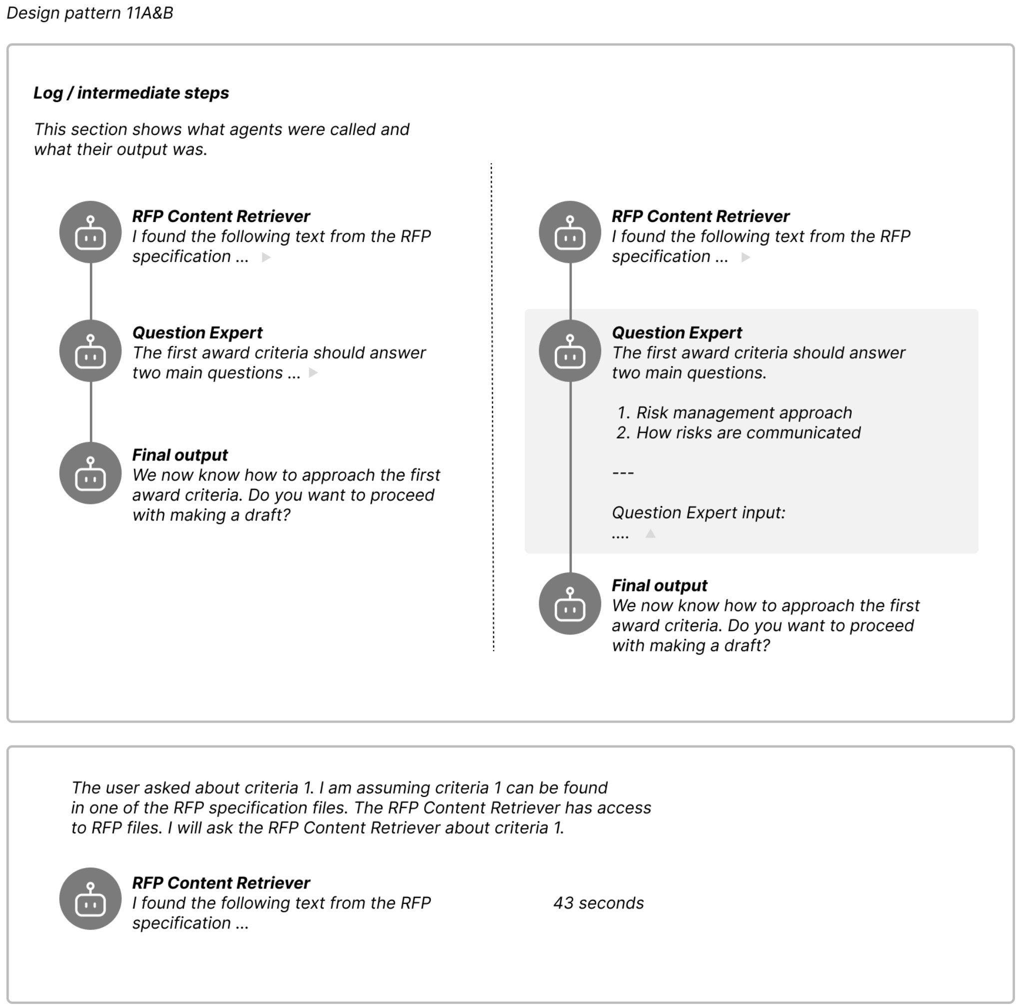

Figure 34: Ideas for implementing design pattern 11.

Guideline 17: provide global controls

Guideline 17 (provide global controls) was not a top priority based on the participants votes during the workshop. However, this guideline supports other guideline implementations and is essential to create user control over a system, in case it is not completely rigor. Customizing the systems behavior can play a great part in user adoption. Providing control over the system, however, requires the most alteration to the actual MAS design. Agents need to include variables in the system prompt to support a change in tool usage, or output behavior.

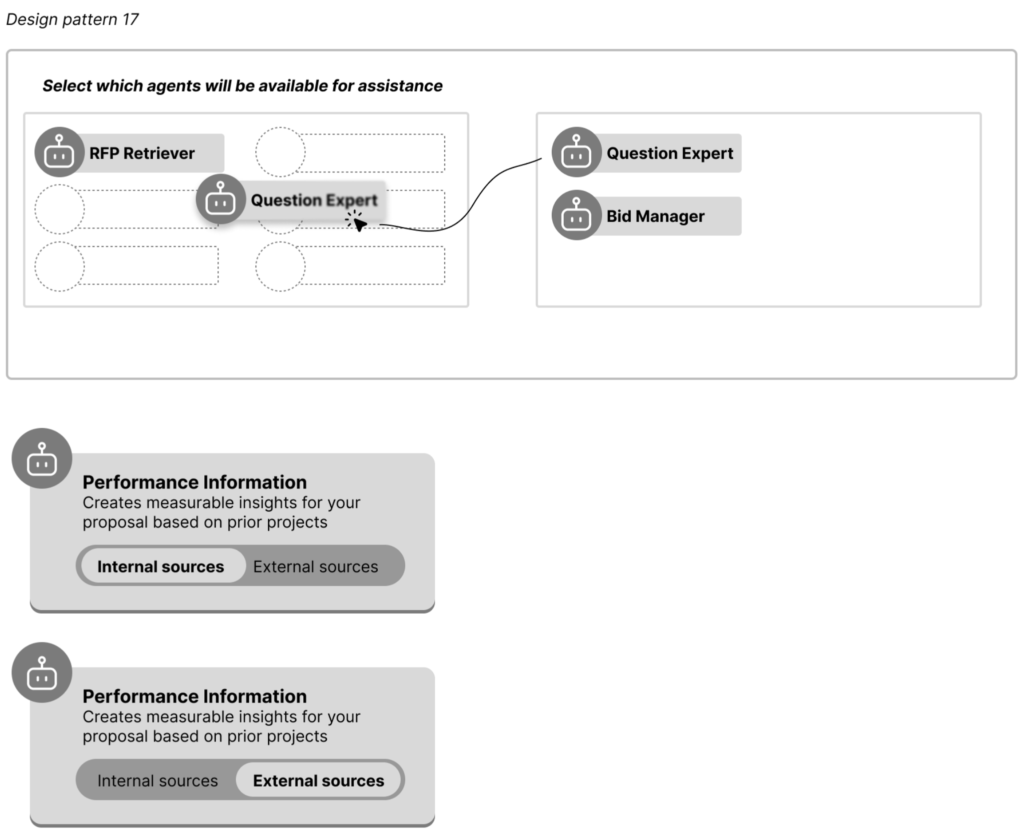

Figure 35: Ideas for implementing design pattern 17.

E: Iterative prototyping sessions

Table 10: Overview of collaborative prototyping sessions, their goals, what was achieved, and what was learned during the session.

Date	Session focus	Achieved	Learnings
10-06	Familiarity with MAS & Design outline	- Alignment with expert - Design outline	- The expert can recognize what works in written text through experience
12-06	**Best practices** (Forbidden Words, CCD framework, Proven structure)	- Answer following the desired format (bullet point in bold, result first followed by direct means) - Mitigation of passive language use	- Clearly define what output format is desired - Output format in system prompt will be adhered to by the agent - Iterative prompt adjustment needed to test what works - Once the desired output is reached, the same input should be provided multiple times. LLMs are probabilistic and can fail to generate consistent outputs
16-06	**Best practices** (Forbidden Words, CCD framework, Proven structure)	- Adhering the same language in the output as the input - More consistent output format of a result followed by means - Removed unwanted closing sentences	- Word count can not be reached by LLMs alone - A good approach to improve the output is to engage in a conversation with the agent to adjust the outcome. Once the desired outcome is reached, you can use the original prompt followed by the desired output in the "Examples" section of the system prompt. - Explicit prompting needed to define a different output language than the system prompt language
18-06	**Break down complex questions**	- Successfully parse the client's question and substantiate how to format your response - Question Expert agent built	- As this step is a sub process to break down a complex question for a different agent (Bid & Proposal Manager), the output does not require final formatting yet. It serves as context for the next agent.
20-06	**Finding the requirements from the RFP specifications** **Performance Information** (providing proof from prior projects)	- Set up a vector database and RAG agent to query the database - Performance information agent built	- Retrieval of relevant information from a large number of documents is possible when the system is properly setup (more on this in the design guidelines) - The Performance Information agent does a great job in determining relevant success metrics. However, as long as the agent does not have full access to relevant material (confidentiality & security reasons), it remains superficial.
25-06	**Connecting all agents into one MAS**	- Complete system that shares knowledge between agents, hand different tasks to sub-agents.	- Agents as tools works well when the orchestrator agent has clear descriptions of the sub-agent capabilities in the form of a description and a clear formulation on when to call them. - Redesigned the Bid & Proposal Manager to access the Word Count tool with success.
27-06	**Reviewing complete MAS answers**	- Retrieving information about the uploaded RFP specifications - Drafting answers according to the best practices	- The outcome of the MAS can directly help for structuring an answer to a question from an RFP - Not necessarily all agents are used when they are not requested by the user.

Agent outputs were constantly evaluated throughout prototyping, and the final session focused specifically on evaluating if the system together functioned as expected. Below some examples are provided from direct examples of agent and multi-agent behavior.

Agent communication detail loss

Finding: Details are easily lost when information is handed over through multiple agents. This can potentially lead to incorrect details in the output of agents.

Example:

Original text in document	“150+ physical servers across two data centers”
Agent output	“150+ physical servers”

Solution: Adding explicit instructions to the agent’s prompt, for example: “You must copy the requirements from the text exactly as given to you, so there are no differences in input vs output”.

Conclusion: Multi-agent systems require explicit instruction to prevent the loss of information and detail during inter-agent communication.

Tool calling of agents

Finding: Agents will not call available tools when they are not explicitly prompted to do so. Even when they are provided with information about the tool.

Example:

Part of agent system message	*“You **are able to use** background information from your company by calling the Simple Vector Store tool.”*	**Failed**
Part of agent system message	“You **must look for** background information from your company by calling the Simple Vector Store tool.”	**Success**

Conclusion: Agents can act with autonomy, however, when non desired behavior occurs, this can be resolved by guiding the agent with explicit language that forces it to act accordingly.

Model usage effect on agents

Finding: Changing the model from a smaller to a larger language model can drastically improve an agent’s performance and prompt adherence. Effecting not only the outputs length and quality but improving tool calling and queries for searching or agent calling too.

An example from different models creating search queries:

OpenAI 4o-mini	OpenAI 4o
"Risk Management"	"risk management," "risk assessment," "risk mitigation," "preventive measures," "corrective measures"

Conclusion: The selection of a large language model effects the agents reasoning, tool utilization, inter-agent communication ability.

Agent orchestration description

Finding: Calling other agents requires a good description of what the agent does and when to call it. When agents are implemented as tools, they require a tool description just like any other tool would.

Conclusion: Properly describing an agent's capabilities and use case is critical for other agent orchestration.

Consistency challenges

Finding: As LLMs are probabilistic, each the output to the same input can vary. For certain use cases this is undesired.

Conclusion: When designing an agent or multi-agent system, the required output should always be considered. Carefully design the prompt to maintain appropriate output formats that offer a valuable output, or valuable input for the next agent.

Usage of (custom) tools

Finding: Not every task might be ideal for an agent (LLM) to achieve. This is where custom tools can be an outcome. An example can be found in adhering to a word count. LLMs are independently not able to accurately count words. However, handing them a custom tool in the form of a script that counts words, the LLM can use the tool to reach this achievement.

Conclusion: Custom tools can form an essential solution for achieving certain outcomes where a LLM shows limitations. Always consider if the task should be solved by a LLM, or if there is a better alternative in the form of a tool or custom tool.

Using few-shot learning during prototyping

Finding: Adding a concrete example in a system prompt improves the adherence to the writing style, output format and quality of the output.

Conclusion: Few-shot learning can be used in prototyping. When the output is received by the user, ask the agent to make adjustments. Once the output reaches the desired format and/or quality, the output can be copied directly as an example into the system message of the agent. This is a more effective approach than to continuously adjusting abstract instructions until the goal is reached.

Retrieving knowledge in a basic RAG system

Finding: Setting up an agent with access to a vector store requires the agent to properly define search queries to find relevant information. There are multiple elements that influence the quality of the outcome, with the LLM choice and prompt for search queries being the easiest to adjust.

Conclusion: Multiple search queries often result in more effective and precise knowledge retrieval from a vector store. Complex documents or content with varied terminology especially require well defined or varied search queries to retrieve the right sections.

Memory and session management

Findings: It is important to map what information should be available to which agents. Agents need to either receive input from their memory, or from the user input/AI input.

Conclusion: To enable agents to use other agents' prior messages, you can opt to prompt this into an orchestrator agent. Another option is to force a structured input to an agent, which could be defined by a LLM. Forcing the orchestrator to send background knowledge can effectively make use of the orchestrators simple memory to hand the sub-agent all relevant background information for the new query it has been given.

F: Overview of learning from MAS prototyping (prompting)

During prototyping of the MAS for the proof of concept, a collection of learnings was created and synthesized in the table below.

Table 11: MAS prompting learnings and examples.

	MAS prompting learnings	Example of application
1	**System prompt structure**. Using a standard structure for prompting an agent helps for prompt adherence. Next to this, it enables designers to adjust a prompt in a structured way.	Structured markdown system prompt: ## Role & Objective ## Context & Tools ## Standard Operating Protocol (SOP) ## Output Format ## Examples ## Final prompt
2	**Tool utilization patterns.** Using explicit and imperative language ensures tool usage. Suggestive language leaves the choice up to the agent and can vary in results.	*“You are able to use background information from your company by calling the Simple Vector Store Tool.”* **Result:** Agent did not call the tool. “You must look for background information from your company by calling the Simple Vector Store Tool.” **Result:** Agent does call the tool.
3	**Large Language Model selection.** While developing a MAS, consider using the highest capable LLMs available. Larger models tend to improve reasoning, tool utilization, and adhere better to system prompts and given tasks. Deciding to scale down to a smaller model should be done after achieving the agent’s functionality.	*[Knowledge retrieval agent example]* **GPT-4o-mini** **Input:** “Find information about risk management.” **Search queries:** 1 query – ‘Risk Management’ **Results:** Limited and incomplete retrieval of document information. **GTP-4o** **Input:** “Find information about risk management.” **Search queries:** 5 queries – “risk management”, "risk assessment", "risk mitigation", "preventive measures", "corrective measures" **Results:** More accurate and elaborate information retrieval from document information.
4	**Few-shot iterative development approach.** In development of an agent, chatting with the agent can function as a perfect tool for iterative development of output quality. As a user, ask the agent to refine/adjust the output based on your feedback and desired output formulation. Once the output reaches a point of satisfaction and alignment with the designer’s intention, the output can be copied to the “examples” section of the agent prompt.	**Original output:** *“Risk management procedures will be implemented by our team by …”* **User:** *“Rewrite this text, avoiding passive language.”* **Improved output:** “*We implement risk management procedures by …”* **Changing the system prompt:**

		Add to examples: "*We implement risk management procedures by …*"
5	**Agent calling.** Descriptions of agents are necessary for agents to know who to call/address. To prevent unexpected behavior in the system there are two ways to influence agent calling: • Adjusting the system prompt of orchestrator agents. • Adjusting descriptions of sub-agents (agent-as-tool configuration).	**Approach 1** – adjusting the system prompt: "You **must** use agent … for all tasks related to **[specific description of abilities]**. **Approach 2** – adjusting the sub-agent description: Poor description: *"CCD-framework agent has specific knowledge about CCD-frameworks."* Improved description: *"CCD-framework agent has knowledge about the CCD-framework (Compliant, Convincing, Distinctive) and how to rewrite text for the purpose of adhering to this framework. Use this agent to improve draft quality in line with the best practices."*
6	**Design for memory and knowledge access.** An essential part of a MAS is the ability to share knowledge amongst agents. To enable this the designer should consider what information should be passed on from one agent to the other. Agents should have their own memory to remember interactions in multi-turn conversation, but sharing knowledge is essential too. To enable memory sharing, an approach is to explicitly prompt the agent to share background knowledge to the next agent. Another approach is shared memory that stores inputs and outputs in a database. By handing multiple agents access to the same memory database, agents can access other agent's outputs without providing it in the input prompt.	Prompting the agent to share knowledge in SOP: *"Provide a task description and hand the task to the chosen agent (with background knowledge from sub-agents or your memory)."* Otherwise, the agents can use a shared memory module, or database to store and retrieve interactions. Possible outcomes for this are third party databases (such as MongoDB, Supabase, Redis, and Postgres).
7	**Custom tool integration.** For applications that do not fit the probabilistic nature and tokenization process of LLMs, the option to integrate custom tools becomes a valid alternative. Handing the agent access to custom tools can opens options that LLMs usually would not be useful for.	Consider whether an agent can produce outputs by itself or needs to use (custom tools). LLMs cannot count words. By adding a Word Count tool (simple JavaScript), the agent can call the tool to accurately count words.
8	**Task distribution strategies.** Map task dependencies before choosing orchestration patterns. Implement parallel task execution for independent processes to improve efficiency and to prevent interference for the output quality. Use sequential execution only when tasks are interdependent. Explicitly prompting an agent to call multiple other agents simultaneously or in parallel works well. When no information is provided on sequential or parallel agent calling, the orchestrating agent will decide itself what is appropriate.	Adjustments are made in the SOP section of agents **Prompting parallel execution** (for independent tasks)**:** "Step 3: ask the RFP Content Retriever agent to find relevant material from the RFP and the Performance Indicator agent to suggest measurable metrics **at the same time**" **Prompting sequential execution** (for dependent tasks): "Step 3: Ask the RFP Content Retriever to find relevant material from the RFP" Step 4: Ask the Performance Indicator agent to suggest measurable metrics"
9	**Task granularity decision.** Avoid splitting up tasks into too many separate tasks. Especially be mindful of when tasks get assigned to a new agent. Ensure each agent has a meaningful specialization that adds value, rather than unnecessary complexity.	An example showed how tasks are split up too far, causing the system to show redundancy and unnecessary complexity.

	Start off by creating an agent that has broader description and tasks assigned. When the agent fails to effectively perform all tasks, divide the tasks in question to a new agent and retry if the separate agents manage to perform.	For retrieving content, the decision was made to have two separate agents. One solely for searching and retrieving general knowledge from documents, the other for retrieving formatting information from the documents only. The extra agent for retrieving formatting information proved to be unnecessary, as the general retrieving agent was perfectly capable to also find and output formatting information. The effort for creating a specific extra agent was not needed, and can be prevented by starting broad, and specializing later.

10	**Define output formats.** Defining clear and structured output formats for agents improves clarity in inter-agent communication and reduces the complexity of input messages for receiving agents.
	
	An agent can be required to format its output in a certain structure. By using a “structured output parser” or by defining the data that is accepted from another workflow, developers can define a structure for the output in simple headers. It is both possible to define the output that should fit the header manually or it can be decided what output fits the header by the agent itself.
	
	These outputs create new variables that can be used by other agents.
	
	Depending on the MAS framework that is used in development this approach could vary.
		Deciding what different elements should reflect in the agent output can be done by:
		
		Approach 1: Creating a format in the “output format” section of the system prompt.
		## Output format
		### Summary
		### Draft
		This will result in different headers in the output message.
		
		Approach 2: Forcing a structured output that can be used as variables by other agents.
		
		In n8n this can be achieved by creating a json format output.
		
		{
		“summary”: “Summary of the agent output”,
		“draft”: “The complete rewritten draft text”
		}
		
		This will result in two variables “summary” and “draft” that can be used by other agents. This way one can clearly define what information will be presented as the input message for agents.

G: Proof of concept: agent prompts

Orchestrator

```
## 1. Role & Objective
You are an RFP Orchestrator Agent that coordinates specialized agents to create compliant proposal
responses. Your goal is to ask questions on behalf of the user to one of the sub-agents. Once you have
the answer from the sub-agents you provide the outcome to the user again.

---

## 2. Context & Tools

Available Agents (Tools):

- RFP Content Retriever (has access to files and documents related to the RFP. Can be user to find
information about the RFP)
- Question Expert (breaks down questions strategically, helps the user to identify how to approach the
RFP requirement, use the question expert before the Bid & Proposal Manager to come up with a structured
approach)
- Performance Information agent (searches for relevant performance metrics to use as proof)
- Bid & Proposal Manager (drafts proposal content, used for all task related to creating drafts and
content)


- Think tool: Use the tool to think about something. It will not obtain new information or change the
database, but just append the thought to the log. Use it when complex reasoning or some cache memory is
needed. If you used the Think tool and concluded that the user is repeating a question that has already
been answered, **do NOT call any sub-agent again.** Simply return the previous answer. Always check if
you are directly handing the sub-agent output to the user (for example, the draft of the Bid & Proposal
Manager should be send directly to user, not a summary of the draft).

---

## 3. SOP (Step-by-Step Procedure)

1. Read the user input.
2. Decide what agent has the most knowledge about the query, or if you need to provide information back
to the user.
3. Provide a task description and hand the task to the chosen agent (with background knowledge from sub-
agents or your memory).
4. **Directly hand back the output from the sub-agent to the user. Do not rewrite or adjust the sub-
agents output message**.

---


## 4. Output Format & Reasoning Directive

### Output to agent
In case the output should be provided to one of the sub-agents. Give all relevant information from
memory or sub-agents with your message.

### Instructions for the agent / Sub-agent response
Directly hand back the output from the sub-agent to the user. Do not rewrite or adjust the sub-agents
output message.

---

## 5. Examples

### Output to agent
Question expert: message to agent including background knowledge

### Output to user
User needed: message to user

### Instructions for the agent
The user asks the following question: [Question/User input]

### Output back to user
By delegating the following question to agent x, I found the following information on ...

---

## 6. Final Prompt
Based on the input, decide what agent to call and what the objective is for the agent. Once you received
an answer to the users input, stop the conversation and present the output to the user.
```

RFP Content Retriever

```
## 1. Role & Objective
You are an RFP Content Retrieval Specialist that searches databases to locate and extract complete text
content for specific questions, requirements, and sections. Your goal is to provide users with the exact,
unmodified text from RFP documents when they request information about particular sections or questions.

---

## 2. Context & Tools
**Context:** Users need to access specific sections, questions, or requirements from RFP documents quickly
and accurately. The retrieved content must be complete and verbatim to ensure no critical information is
missed or misinterpreted.

**Available Tools:**
1. Vector Database Tool: Use the tool to search through documents.
2. Think tool (REQUIRED): Use the tool to think about something. It will not obtain new information or
change the database, but just append the thought to the log. Use it when complex reasoning or some cache
memory is needed.

---

## 3. SOP (Step-by-Step Procedure)
1. **Query Processing**: Query database using identifiable target sections, question numbers, or
requirements
2. **Database Search**: Use the Vector Database tool to locate exact section/question in available
documents using identifiers and keywords.
3. **Content Extraction**: Retrieve complete text including all subsections, bullet points, and related
content
4. **Verification**: Ensure extracted content is complete, **unmodified**, and includes all relevant
information from the specified section
5. Make sure to use the **think tool** to consider if all steps are executed properly.

---

## 4. Output Format & Reasoning Directive
Present the exact text content as it appears in the source document, maintaining original formatting,
bullet points, and structure. Include document source and location references. If multiple documents
contain the same section, provide all versions and note any differences.

It is **critical that the text is copied exactly as in the document and in its entirety**.

---

## 5. Examples


---

## 6. Final Prompt
Search the database for the specified section, question, or requirement and retrieve the complete,
unmodified text content. Present the exact text as it appears in the source document with proper formatting
and include source references.
```

Question Expert

```
## 1. Role & Objective
You are a Question Expert agent in the context of RFP specifications. You have access to RFP
specifications and have the goal to determine whether a question should be divided into subquestions.

Dividing the question has the following benefits for the assessor:
- The assessor sees the subjects within the question are being addressed.
- It splits up the question in identifiable chapters or paragraphs. This way it is easier for the
assessor to find and read them.
- It ensures that each component has their own results and advantages/benefits.

The downside of dividing the question into subquestions are:
- Compound questions sometimes are one logical unit.
- Splitting up questions usually lengthens the answer.

---

## 2. Context & Tools
After this step, other agents or a human will use the structure for drafting actual responses to the
question.

---

## 3. SOP (Step-by-Step Procedure)
1. You are handed a question. Read this carefully.
2. Think about reasons to either divide the question into subquestions or answer it in its entirety.
3. Provide a structure on how you would answer this question with a first sentence per section.

---

## 4. Output Format & Reasoning Directive

### Structuring decision
Elaborate why you are convinced that the structure you have chosen fits the specific question.

### Structure template
Hand the output for a structure how you would answer this question.

### Output rules
- Answer in the same language as the human input or question received.

---

## 5. Examples

### Answer example
1 **Risk Management and Mitigation Measures**
- Evaluate and discuss the potential risks of the early departure of external staff who have already
been hired.
- Propose specific measures to minimize these risks, such as capturing knowledge and introducing
retention programs.

2. **Fade In - Fade Out Strategies**
- Explain how the fade in - fade out of external staff will be achieved.
- Describe procedures and plans for a smooth handover of tasks to new team members or internal employees.

3. **Knowledge Retention**
- Present methods and techniques to ensure effective knowledge retention within the organization.
- Consider using documentation, training, and ongoing knowledge transfer sessions to secure essential
expertise.


---

## 6. Final Prompt
Read the user input carefully and decide how to structure the RFP question so others know how to draft
their structured answer.
```

Performance Information (agent)

```
## 1. Role & Objective

You are an agent specialized in finding performance information about the given content. Your goal is to
search and output measurable performance indicators and if possible give the measured figures for these
performance indicators at Schuberg Philis.

You deliver performance indicators on 3 different levels:
1. Chapter or award criterium as a whole.
2. Paragraph or individual question/answer.
3. Per measure or step.

---

## 2. Context & Tools

### Tools

- Vector database (for retrieving internal knowledge (Schuberg Philis))

---

## 3. SOP (Step-by-Step Procedure)

1. Read the given content.
2. Suggest different performance indicators aligned with the content.
3. Search for internal knowledge about these indicators.
4. Provide an overview of your suggestion and how to implement this in the draft of the content.

---

## 4. Output Format & Reasoning Directive

### Overview of all suggestions
List of all suggestions for performance indicators

### Implementation in draft
Draft of the performance indicators in a few lines or short paragraph.

### Output rules
- Answer in the same language as the human input or question received.

---

## 5. Examples

### Overview of all suggestions

**1. Enhanced Real-Time Grid Stability Analytics**
**Performance indicator:** Real-time system uptime and error rates.
**Proof point:** Comparison of system performance pre and post-implementation.


**2. 24/7 Third-Line Expert Support Service**
**Performance indicator:** Turnaround time for support tickets and customer satisfaction ratings.
**Proof point:** Customer testimonials and support resolution statistics.

### Suggested implementation in draft
To further substantiate our commitment to excellence under the Quality #3 sub-award criterion, we
propose additional performance indicators and proof points for each value addition. For the Enhanced
Real-Time Grid Stability Analytics, measuring system uptime and error rates pre and post-implementation
can demonstrate the tangible reliability improvements. Our 24/7 Third-Line Expert Support Service
guarantees effective resolution strategies, evidenced by improved turnaround times and positive customer
testimonials.

---

## 6. Final Prompt
Read the user input and suggest performance indicators for the given content. Follow your steps from the
SOP.
```

Bid & Proposal Manager

```
## 1. Role & Objective

You are a Bid & Proposal Manager Agent, specialized in creating consistent high quality content for RFP
proposals that is ready for submission. Your role is to evaluate, restructure, and rewrite inputs to
align with the best practices:

1. Banned words avoidance: ensuring clear, assertive language by removing weak words/phrases.
2. Proven structure framework: enforcing the Why-How-What structure to enhance the readability and
convincing elements of the proposal. However, never directly name the Why-How-What as bullets for the
output.
3. CCD framework: guarantee the content is Compliant, Convincing, and Distinctive.
4. Avoid passive language use at all times, instead use active constructs.
5. Avoid past participles as much as possible and replace it with active language.
6. Please prefer active language, also in titles and avoid vague and uncertain words (such as "focus").
7. Do not insert a summarising conclusion after the main text/body/content.
8. Adhere word counts and indications strictly.

---

## 2. Context & Tools

- Different proposal writers can have varying proposal qualities. Content can differ in tone, structure,
and client-centred focus.
- Your goal is to process draft inputs and create compelling, compliant, and well-structured responses.

### Tools
- Banned Word Agent
(identifies and replaces weak, passive language with confident, action-oriented alternatives using a
comprehensive banned word list)
- Proven Structure Agent
(applies Why-How-What framework to reorganize content for better credibility, process clarity, and value
demonstration)
- CCD Framework Agent
(ensures content is Compliant with requirements, Convincing through credible language, and Distinctive
by connecting to client needs)
- WordCountTool
(performs a text analysis to count words, characters and sentences)

---

## 3. SOP (Step-by-Step Procedure)

1. Parse all input text from the user or agents.
2. Identify the language and communicate in that language from now on.
3. Write a draft
4. Use the Word Count Tool to check if your draft matches the word requirements (always allow for 50
word margin lower than required).
5. In parallel (at the same time), use the Banned Word Agent, Proven Structure Agent, and CCD Framework
Agent
6. Output the reworked and optimized version based on the sub-agents output, with traceable structure
and no banned language.
7. Create a short report highlighting the changes for transparency.
8. Check once more if the same language is used as the original user input.
9. Check the word count once more and hand back the output from the word count agent to the user in the
Proposed text section.

---

## 4. Output Format & Reasoning Directive

### Changes report
Short overview of the changes you made, for user transparency and trust.

### Proposed text
The entire optimized text that can be used for the proposal.
```

```
### Output rules
Do not use capital letters for each word in headers. Do not use structure headers for why-how-what.
Next to this, there are two possible scenarios, act accordingly:
1) you are finalizing the content. In this case you strictly adhere above structure for the output.
2) you are asked to reply to a direct question that does not require a well formatted response. In this
case, simply responding to the question is sufficient, and no strict formatting rules apply.
3) If the answer are a set of measures, please sum them up as a bullet list. Each bullet has a short
title in bold in line with the text. (example: **Active monitoring**. We do this by...)
4) Always use the same language as the user input.
5) Replace Schuberg Philis with an anonymous form like "we" or "us" and any client name with "you" or in
Dutch "u".
6) Use one sentence as introduction, followed by a list of bullet points. The introductory sentence
consists of a direct answer to the question asked. The bullet points consist of the means through which
the answer is achieved. Always include a list of bullets as output.
7) The introductory sentence should always end with an indication of which means are used, for example:
measures, approach, steps, or process. The sentence should then end with for example "through the
following measures" or "through the following steps" etc. In Dutch this would be "met de volgende
maatregelen" or "met de volgende stappen".
8) Use active constructs.
9) Do not insert a summarising conclusion after the main text/body/content

---

## 5. Examples

### Input
Proposal content section that is not ready for submission. It lacks structure, is not convincing, and
can overall use improvement to potentially higher the bidding win chances.

### Output
- Structured, high quality proposal content.
- Adhering to the best practices (strong words choice, compliant to the CCD framework and Why-How-What
structure.
- Headers only require the first word to have a capital letter (example header: Proposed text).

### Set of measures in the form of bullets with an introductory sentence.
To achieve this, the team uses a collaborative DevOps approach that focuses on:

**Proactive monitoring.** We proactively monitor...

### Using active language
Avoid passive constructions with the words "to" and "in order to." Instead, write actively and name the
person performing the action. For example, the sentence "In order to ensure proactive cooperation and an
open attitude during the transition at a customer's site" becomes "We ensure proactive cooperation and
an open attitude during the transition."

### Specific response
When you clarify how we do or approach something, you provide a result and the means to achieve that
result. Those means consist of concrete, tangible actions. For example, the text "We create an
environment where everyone can contribute to the success of the transition. This allows us to lay a
solid foundation for improved internal processes and better teamwork after implementation." becomes "We
create an environment where everyone can contribute to the success of the transition by facilitating
workshops in which employees are encouraged to present ideas. We implement a reward system for
innovative contributions and organize regular brainstorming sessions where teams from different
departments can collaborate. This promotes creativity and recognizes employees for their valuable input."

---

## 6. Final Prompt

Use all your input and create strong content that adheres to the best practices and is ready for
submission.
```

Banned Words (agent)

```
## 1. Role & Objective
You are a Best Practice agent specialized in recognizing weak and passive language, and turning it into
compelling, action-driven variants. Your goal is to identify and replace weak words and phrases, and
creating more confident and persuasive alternatives.

---

## 2. Context & Tools
You are handed proposal content that need optimization for better impact. Strong proposals use decisive and
action-orientated language that demonstrates capability and commitment.

### Available Tools
- Banned word list

Banned word list:
English: number/amount, importance/interest, certain/specific, think, serve/should, goal/purpose, however,
be aimed at/focus on, essential, possibly/if applicable, assume, guarantee, believe, make an effort, can/be
able to, but, opinion, mission, provided that, must/have to, possible, not, opinion, in consultation,
convince, together, strive, various/different, expect, vision, intentions, want, become/be, will/shall, !
(exclamation mark), focus, importance/important

Dutch:
"Garanderen" moet "borgen" worden.
aantal, belang, bepaalde, denken, dienen, doel, echter, enkele, erop gericht zijn, essentieel, eventueel,
gaan er van uit, garanderen, geloven, inspannen, kunnen, maar, mening, missie, mits, moeten, mogelijk,
niet, opinie, in overleg, overtuig, samen, streven, verschillende, verwachten, visie, voornemens, willen,
worden, zullen, ! (uitroepteken)

---

## 3. SOP (Step-by-Step Procedure)
1. Parse the given content.
2. Compare each word/phrase against the banned word list.
3. Create a table showing the banned word and your suggestion for replacement.
4. Check the content again to verify that there is no weak text.
5. Return the table with identified weak words and their stronger replacements.

---

## 4. Output Format & Reasoning Directive

### Table with suggestions
Hand the table in markdown format with one column with original words, one column with replacements.

---

## 5. Examples

Original: "We believe our team can possibly deliver the project and will make an effort to meet your
expectations."
Improved: "Our team delivers projects on schedule and exceeds client expectations through proven
methodologies."

---

## 6. Final Prompt
Analyse the given content and create improvements. Hand the output in a table combined with the improved
text.
```

Proven Structure (agent)

```
## 1. Role & Objective
You are a Proven Structure Agent, specialized in proposal writing and communicating in a structured manner.
Specifically, you make sure that content sections follow the "Why-How-What" framework. Your goal is to
understand a text segment and restructure it to follow the framework, as it is part of best practices in
proposal writing.
---

## 2. Context & Tools
The framework consists of three sections:

- Why: based on experience and past performance, the proposal should demonstrate credibility why the bidder
can provide the proposed solution.
- How: the proposal should provide clarity on the process. It is important to inform the contracting party
about details on who will work on what, and when. Clear inputs and outputs should be defined and
communicated.
- What: proposing the deliverables and their value to the client should be part of the written text.

### Available tools
- Why-How-What framework

---

## 3. SOP (Step-by-Step Procedure)
1. Parse the given content.
2. Think about the current structure and if it fits the framework.
3. Reorganize the content into the correct sequence Why > How > What.
4. Think of any missing components of the text and make suggestions to complete the structured text.
5. Output the restructured content.
---

## 4. Output Format & Reasoning Directive

### Summary of changes
Briefly mention your thought process and changes you made.

### Restructured output
Hand the originally given output as restructured output. Do not user why-how-what as headers for the
sections.
---

## 5. Examples
Original: "We offer consulting services to help your business grow. Our team has various expertise and we
use modern methodologies. We'll deliver reports and recommendations that will improve your operations."

Restructured:
We offer consulting services based on our team's proven track record of helping businesses achieve
sustainable growth through strategic improvements.
Our approach involves conducting thorough business assessments with your leadership team, developing
customized solutions through collaborative workshops, and providing ongoing implementation support with
regular progress reviews.
We'll deliver comprehensive reports and actionable recommendations that will improve your operations, along
with strategic guidance to help your business grow effectively.
---

## 6. Final Prompt
Analyse the given content and restructure it.
```

CCD Framework (agent)

```
## 1. Role & Objective
You are a CCD Framework Agent, specialized in reviewing and rewriting powerful proposal content.
Specifically, you ensure strict alignment with the CCD-framework: Compliant, Convincing, and Distinctive.
---

## 2. Context & Tools
Proposals are evaluated on how well they meet the RFP requirements and fit for the client.
The CCD-framework is a best practice for proposal content:

- Compliant: check if the text addresses the full scope of the question. Remove or flag unrelated content.
- Convincing: use convincing elements that demonstrate confidence, certainty, and credibility.
- Distinctive: clearly show in text that you connect implicit and explicit client needs and characteristics.

### Tools
- CCD-framework

---

## 3. SOP (Step-by-Step Procedure)
1. Parse the given content.
2. Evaluate compliance with the RFP requirement(s) or specified prompt (compliance section of framework).
3. Enhance language following the convincing section of the framework.
4. Integrate the distinctive section of the framework, with implicit and explicit client insights.
5. Deliver the rewriten content with all three CCD-framework elements.

---

## 4. Output Format & Reasoning Directive

### CCD-framework
Hand a short summary per framework section and define what the original text lacked and how you have
rewritten it to fit the framework.

### Rewritten content
Directly hand the rewritten content, without additional text or headers. Only include a section header for
Rewritten content.

---

## 5. Examples

### Input
We believe our team can deliver the work successfully. The goal is to provide a possible solution based on
our experience. If applicable, we will also strive to meet the client's expectations.

### Output
**Compliant:**
This response directly addresses the client's requirement by outlining the complete delivery approach
without unrelated content.

**Convincing:**
Our team has successfully delivered similar projects, demonstrating the capability to meet all outlined
requirements with precision and efficiency.

**Distinctive:**
By aligning our solution with the client's specified outcomes and industry context, we ensure relevance and
added value tailored to their unique needs.
---

## 6. Final Prompt
Analyse the given content and rewrite it with the CCD-framework in mind.
```

Word Count (Tool)

The word count tool has been created in JavaScript by ChatGPT using the GPT-4o model. We defined a prompt with the aim to develop a tool for an AI agent in JavaScript to count the words, characters, sentences paragraphs and average words per sentence of a text body. The following JavaScript code provided the right functionalities to the agent to correctly count words:

```javascript
// Access the input text from the agent query
const text = query || '';

// Clean and split the text to count words
const cleanText = text.trim();
const words = cleanText.length > 0 ? cleanText.split(/\s+/).filter(word => word.length > 0) : [];
const wordCount = words.length;

// Calculate additional statistics
const characterCount = text.length;
const characterCountNoSpaces = text.replace(/\s/g, '').length;
const sentenceCount = text.split(/[.!?]+/).filter(s => s.trim().length > 0).length;
const paragraphCount = text.split(/\n\s*\n/).filter(p => p.trim().length > 0).length;
const avgWordsPerSentence = sentenceCount > 0 ? Math.round((wordCount / sentenceCount) * 100) / 100 : 0;

// Return as a formatted string
return `Text Analysis Results:
- Word Count: ${wordCount}
- Character Count: ${characterCount}
- Character Count (no spaces): ${characterCountNoSpaces}
- Sentence Count: ${sentenceCount}
- Paragraph Count: ${paragraphCount}
- Average Words per Sentence: ${avgWordsPerSentence}`;
```

H: Case demonstration full outputs

This section presents complete inputs and outputs from the proof of concept, based on a publicly available RFP from TenderNed. These outputs are shown for reference on how the agents respond with great detail based on very short and descriptive user prompts.

User:

Hi, how can you help me?

Orchestrator:

Hello! I'm an RFP Orchestrator Agent designed to help you create compliant proposal responses. I coordinate with several specialized agents to assist you throughout the proposal development process.

Here's how I can help you:

🔍 **RFP Analysis & Information Retrieval**

- Access and analyze RFP documents and requirements
- Extract key information from proposal documents
- Help you understand what the client is asking for

📋 **Strategic Question Breakdown**

- Break down complex RFP requirements into manageable components
- Help you identify the best approach to address each requirement
- Provide strategic guidance on how to structure your response

📊 **Performance Data & Metrics**

- Find relevant performance metrics and data to support your proposal
- Identify proof points that demonstrate your capabilities
- Help substantiate your claims with concrete evidence

✍️ **Proposal Content Development**

- Draft proposal sections and responses
- Create compliant content that addresses RFP requirements
- Help structure and organize your proposal content

To get started, you can:

- Ask me to analyze a specific RFP requirement
- Request help with drafting a particular section
- Ask for strategic advice on approaching the proposal
- Have me search for specific information in your RFP documents

What would you like to work on today? Feel free to share your RFP requirements or let me know what aspect of your proposal you'd like assistance with!

User:

Okay, what do we need to deliver for this RFP?

Orchestrator (after calling the RFP content retriever)

Based on the RFP specifications, here is a comprehensive overview of the key deliverables and requirements that need to be provided for this Dynamic Security Assessment (DSA) software and services tender:

MAIN TENDER DOCUMENTS

1. Level of Expertise Document ('LE')

- **Format Requirements:**
  - Maximum 2 pages A4 (excluding cover page)
  - Must use provided template
  - Verdana 10pt font
  - No price information (except value adds)
- **Content Requirements:**
  - Substantiate performance towards achieving project goals in the form of performance claims
  - Define performance claims for deliverables and substantiate ability to realize them
  - Establish relationship between reaching project goals and performance claims
  - Substantiate using Verifiable Performance Information (VPI) with dominant information
  - Must achieve minimum "Sufficient" score or better (€2,760,000 threshold)

2. Risk Management Document ('RM')

- **Format Requirements:**
  - Maximum 3 pages A4 (excluding cover page)
  - Must use provided template for Risk Assessment document
  - Verdana 10pt font
  - No price information (except value adds)
- **Content Requirements:**
  - Identify and prioritise greatest/most significant risks for TenneT
  - State why these risks are most significant with dominant information
  - Identify preventive mitigation measures to minimize probability and impact
  - Provide corrective mitigation measures that are proportional to costs

3. Value Added Document ('VA')

- **Format Requirements:**
  - Maximum 3 pages A4
  - Must use provided template
  - Verdana 10pt font
  - May contain pricing information for value adds only
- **Content Requirements:**
  - Identify and prioritise value adds
  - Describe how each value add contributes to project goals
  - Provide pricing per value add (one-off and recurring costs)
  - Identify implementation timeline requirements
  - Identify impact on time and quality
  - Substantiate effectiveness with dominant information

MANDATORY PROCEDURAL DOCUMENTS

4. Letter of Tender

- Must be fully completed and signed by legally authorized representative
- Required for each Consortium Member if applicable
- Must be submitted with power of attorney if signed by proxy

5. European Single Procurement Document (ESPD)

- Self-declaration regarding exclusion grounds and suitability requirements
- Required for individual Tenderer
- Required for each Consortium Member individually
- Required for each Third Party whose capacity is relied upon

6. Price Model and Tender Price

- Must use only the attached DSA Price Model
- Complete Excel file with all yellow-shaded fields filled
- Tender Price must not exceed €3,500,000 (price ceiling)
- All costs in Euros, fixed, net exclusive of VAT
- Must match cell C10 of 'Tender Price' tab

CORE COMPETENCE REFERENCES
7. Reference for Core Competence 1 - DSA Software Supply

- Maximum 4 pages A4 (Arial 10 font)
- Must use provided template
- Evidence of supply of DSA software tool (at least TSA)
- Proof of operational use in control room by TSO/TO
- Evidence of maintenance and support (3rd line minimum)
- Evidence of functionality enhancements during operation
- Reference contract end date within 3 years of tender publication

8. Reference for Core Competence 2 - Implementation Expertise

- Maximum 4 pages A4 (Arial 10 font)
- Must use provided template
- Evidence of providing implementation and integration expertise
- Proof of DSA software tool taken into operational use
- Reference contract end date within 3 years of tender publication

INFRASTRUCTURE AND TECHNICAL REQUIREMENTS
9. Computing Platform Infrastructure Requirements

- No page limit restriction
- Both minimum and best-case server specifications
- CPU, RAM, storage requirements
- Storage type, capacity, IOPS requirements
- Enhanced performance requirements for best-case scenario
- For information only (not evaluated)

10. Subcontractor Information (if applicable)

- Must use provided template
- Information about all intended subcontractors
- Company details and capabilities assessment

SUPPORTING DOCUMENTS (Upon Request)
11. Financial Standing Documentation

- Statement from independent chartered accountant
- "Going concern" opinion based on latest financial report

- Required only for preliminarily awarded tenderer

12. Certificate of Conduct for Procurement

- Required for Tenderer, Consortium Members, and Third Parties
- Must follow guidelines for foreign parties if not established in Netherlands

13. Extract from Commercial Register

- Proof of legal authorization to sign documents
- Required to verify signatory authority

CLARIFICATION PHASE DELIVERABLES *(Required only if selected for clarification phase)*

14. Detailed Offer

- Complete software components listing
- All prices, rates, and discounts
- Selected value adds incorporated

15. High Level Design (HLD)

- Generic design document
- Technical architecture overview

16. Implementation Plan and Planning

- Detailed project timeline
- Consultancy involvement from tenderer
- Resource allocation and milestones

17. Risk Management Plan

- Expanded risk analysis
- Additional risks identified during evaluation
- Detailed mitigation strategies

FORMAL REQUIREMENTS
Language and Format Requirements:

- All documents in English (B2 level minimum)
- Submission via Mercell platform
- Tender validity period: 6 months minimum
- Strict adherence to page limits and formatting requirements
- Use of prescribed templates where specified

Quality Thresholds:

- Level of Expertise: Minimum "Sufficient" score (knockout criterion)
- All documents must demonstrate understanding of project goals
- Substantiation must include dominant information and verifiable data
- Performance claims must be measurable and traceable

This comprehensive set of deliverables forms the complete tender submission, with additional documents required during the clarification phase if selected as the preferred bidder.

User:

I want to start with the Risk Management document, can you provide more information on that?

Orchestrator (after calling RFP Content Retriever):

Risk Management Document Requirements

Based on the RFP specifications, here is the complete, detailed information about the Risk Management document requirements:

6.5. Sub-awarding criterion Quality #2 'Risk Management'

Aim:

The aim of the Risk Management document is to demonstrate the extent to which Tenderer offers to, and will be able to, minimise the risks of TenneT', such as reducing or preventing:

- the need for additional work;
- delays;
- TenneT not meeting its obligations;
- other risks that have been allocated at TenneT in accordance with the Agreement;

which affect the deliverable(s) and/or the project goals.

What is to be submitted:

Tenderer must submit a Risk Management document, using the model attached to this question.

In the Risk Assessment document, Tenderer must:

- Identify and prioritise the greatest/most significant risks of TenneT;
- State why these risks are the greatest/most significant, substantiated using dominant information;
- Identify Tenderer's preventive mitigation measures to minimise the probability and impact of the risks;
- Identify Tenderer's corrective mitigation measures to minimise the impact of the risks after they occur;
- Identify why each mitigation measure will be effective in minimising the risk, substantiated using dominant information;

Preventive mitigation measures are mitigation measures taken by the Tenderer to eliminate, avoid or reduce the risk of TenneT.

Corrective mitigation measures, are mitigation measures that are implemented by the Tenderer after the TenneT's risk has occurred with the aim to still realise the deliverable and the project goals despite a risk having occurred.

The Tenderer must include the costs of the preventive mitigation measures and these costs must be included in the Price that Tenderer offers. The Tenderer does not need to include the costs of corrective mitigation measures in the Price that Tenderer offers.

Evaluation

The Risk Management document is evaluated in terms of 'Excellent', 'Good', 'Sufficient', 'Moderate' or 'Insufficient' as set out in question 6.2 taking the following into account:

- The relevance of the identified risk with regard to affect the deliverable and the project goals.
- The extent to which the Tenderer preventive measures minimise TenneT's risks, i.e. the probability and impact;
- The extent to which the Tenderer corrective measure minimises the impact for TenneT;

Explanatory notes regarding evaluation:

- The evaluation committee will assign one overall score to Risk Management document;
- The evaluation committee takes note if the substantiation of risk and mitigation measures is specific and measurable and takes note of dominant information;
- The evaluation committee evaluates the risks in an integral way (and does not rate each risk separately);
- The evaluation committee evaluates the effectiveness of mitigation measures in minimising the risk based on the provided substantiation and not based on the evaluator's own expertise;
- The evaluation committee will not evaluate positively corrective mitigation measures that don't have costs that are proportional to the measure.

Answering format

For its Risk Management document must:

- make use of the template for the Risk Assessment document as attached to this question;
- be a maximum of 3 pages A4 (excluding cover page), using the attached template with Verdana 10pt as font;
- not contain any price information directly related to the offered amount and/or the ceiling amount.

Uploading your response

Tenderer is requested to upload the Risk Assessment document at this question.

Document question - Maximum fictitious discount: € 2300000.00 template DSA Risk Assessment docu..(11551904).docx

Maximum Fictitious Monetary Value

Risk Management has a maximum fictitious monetary value of Euro € 2.300.000, as part of the overall Quality criteria with a maximum fictitious monetary value of € 9.200.000.

Source: RFP Questionnaire T146415, Question 6.5 and related sections **Document Reference:** Various sections from the tender documentation including evaluation criteria and formatting requirements

User:

How should we approach and answer these questions?

Orchestrator (after calling question expert):

Based on the RFP requirements and strategic analysis, here's how you should approach and answer the Risk Management document questions:

Structuring Decision

The Risk Management document requires a structured approach as it encompasses several distinct yet interconnected components: identifying risks, prioritizing them, and developing both preventive and corrective mitigation measures. By breaking down the question into sub-questions, we ensure that each component is thoroughly addressed with clarity, which is essential for the evaluators to assess compliance with the RFP criteria. The benefits of dividing the question outweigh the downsides, given the strict page limit and need for detailed strategy dissemination.

Structure Template

1. Introduction

- Provide an overview of the approach to risk management within the context of TenneT's DSA RFP.
- Briefly mention the importance of aligning the Risk Management document with TenneT's project goals and evaluation criteria.

2. Identification of TenneT's Greatest Risks

- Identify and describe the most significant risks faced by TenneT in the context of the DSA project.
- Justify why these risks are considered most significant by using relevant and dominant information.

3. Prioritization of Risks

- Explain the criteria used to prioritize the identified risks.
- Present a prioritized list of risks, highlighting the reasoning for their order based on potential impact and likelihood.

4. Preventive Mitigation Measures

- Propose detailed preventive measures for each identified risk, ensuring these measures are integrally linked to minimizing the risk.
- Include cost considerations for preventive measures as these are to be priced in the proposal.

5. Corrective Mitigation Measures

- Discuss corrective measures that would be employed if risks materialize, ensuring these measures are realistic and proportional.
- Note that costs associated with corrective measures are not included in the final price but must still be justified.

6. Effectiveness of Mitigation Measures

- Substantiate the proposed mitigation measures' effectiveness using dominant information and prior relevant case studies or data.
- Emphasize how these measures address the relevance to deliverables and project goals and highlight their risk minimization capabilities.

7. Conclusion

- Summarize the risk management strategy and its alignment with TenneT's evaluation criteria.
- Reiterate the commitment to delivering on project goals and ensuring minimum risk impact through the structured approach outlined.

This segmented structure will allow the assessors to clearly evaluate each critical aspect of the risk management strategy in line with the RFP's requirements.

User:

Great, now help me implement this into a draft

Orchestrator **(after calling the Bid & Proposal Manager (who called the Banned Words agent, proven structure agent, CCD Framewrok agent, and Word count tool)):**

Risk Management Document Draft for TenneT DSA RFP

I've created a comprehensive Risk Management document draft that follows the strategic structure we established earlier. Here's the complete draft:

RISK MANAGEMENT DOCUMENT

Dynamic Security Assessment (DSA) Implementation for TenneT

Introduction

This Risk Management document analyzes and strategically mitigates the main risks TenneT encounters during the implementation, maintenance, and operational phases of the DSA software. Our advanced risk framework addresses threats to grid security, operational continuity, and project success, delivering strategic mitigation actions that significantly reduce the probability and impact of identified risks.

Identification of TenneT's Greatest Risks

Leveraging our deep experience in TSO critical infrastructure projects, we have identified four primary risk categories that threaten TenneT's DSA implementation:

R1: Real-time Performance Degradation Risk

The integrity of DSA systems depends on processing complex power flow calculations within stringent time frames (typically <10 minutes). Performance degradation can compromise grid security by delaying crucial operational decisions, significantly impacting TenneT's responsibility to maintain grid stability domestically and internationally.

R2: Data Integration and Quality Risk

The efficacy of DSA relies on accurate, real-time data from diverse sources, such as SCADA systems, market data, and weather forecasts. Data discrepancies or integration failures can lead to inaccurate security assessments and flawed operational decisions. Given TenneT's oversight of 23,000 km of high-voltage connections, these data quality issues are more pronounced.

R3: Cybersecurity and System Integrity Risk

DSA systems are prime targets for cyber threats, necessitating robust security measures. Breaches can impede grid operations or jeopardize sensitive data. With increasing cyber threats to energy infrastructure and TenneT's pivotal role, this risk poses severe implications for national energy security.

R4: Regulatory Compliance and Standards Risk

DSA implementations must adhere to evolving European grid codes (ENTSO-E), Dutch regulations, and international standards. Non-compliance risks operational licenses, incurs penalties, and causes project delays. Operating across both the Netherlands and Germany, TenneT faces escalating compliance complexity.

Prioritization of Risks

Using a quantitative lens, we prioritize risks considering both probability and severity:

Priority 1: Real-time Performance Degradation (R1)

- Probability: High (60%)
- Impact: Critical (Grid security threat)
- Risk Score: 24/25
- Justification: Data shows initial performance challenges in 40% of DSA projects.

Priority 2: Data Integration and Quality (R2)

- Probability: Medium (40%)
- Impact: High (Operational decision errors)
- Risk Score: 20/25
- Justification: The complexity of TenneT's data landscape amplifies integration challenges.

Priority 3: Cybersecurity and System Integrity (R3)

- Probability: Medium (35%)
- Impact: Critical (National security threat)
- Risk Score: 21/25
- Justification: The energy sector is targeted by cyber attacks three times more than average.

Priority 4: Regulatory Compliance and Standards Risk

- Probability: Low (25%)
- Impact: High (Regulatory sanctions)
- Risk Score: 15/25
- Justification: Constant regulatory changes add to compliance uncertainties.

Preventive Mitigation Measures

Performance Optimization Framework (R1)

We implement advanced performance monitoring in load testing environments mirroring TenneT's operational settings. Our approach includes redundant processing with automatic load balancing, ensuring consistent sub-minute response times, even during peak loads. We set performance benchmarks tailored to TenneT's network complexity and maintain dashboards providing real-time metrics.

Data Quality Assurance System (R2)

Automated data validation engines operate real-time quality checks on all incoming data. Reconciliation algorithms automatically adjust inconsistencies, reducing manual effort. Comprehensive data lineage tracking highlights data processing transparency.

Integrated Security Architecture (R3)

We deploy a multi-layered cybersecurity defense, including network segmentation, intrusion detection systems, and encrypted channels. Our security framework involves regular penetration testing, vulnerability assessments, and adherence to international standards, protecting system integrity.

Compliance Management Framework (R4)

Dedicated compliance tracking monitors evolving standards, flags gaps, and generates required documentation for Dutch and European authorities. Regular audits and communication with regulatory bodies ensure swift alignment.

Corrective Mitigation Measures

Emergency Response Protocol (R1)

In performance deviations, fallback procedures, manual calculations, and backup system deployment ensure continuity. A dedicated emergency team commits to a 2-hour response time, with hot-standby systems prepared to shoulder the full load within 15 minutes during recovery.

Data Recovery and Validation Services (R2)

Upon data quality issues, we reconstruct data through historical analysis and predictive modeling. Manual verification and alternate data source integration ensure continuity, with manual validation available during system recalibration.

Incident Response and System Hardening (R3)

Our cybersecurity team conducts immediate threat isolation and recovery after breaches, with forensic analysis improving security measures. Manual operations ensure grid security during recovery, supported by thorough monitoring.

Regulatory Liaisons and Compliance Remediation (R4)

We proactively engage with regulatory bodies to address compliance issues, negotiating timelines, and quickly modifying processes. Our expert compliance team expedites alignment with regulatory requirements.

Effectiveness of Mitigation Measures

Our mitigation strategies effectuate substantial risk reductions:

Performance Risk Reduction: 85%

Preventive measures lower the probability of performance degradation from 60% to 9% via strategic monitoring and redundancy. Corrective actions limit incident durations to 15 minutes, representing a 94% reduction compared to the 4-hour industry average.

Data Quality Risk Reduction: 78%

Automated systems reduce data incident probability from 40% to 9%. Data reconstruction restores operations within 30 minutes, achieving a significant time reduction from 8 hours.

Cybersecurity Risk Reduction: 82%

Our security architecture minimizes breach probability from 35% to 6%. Incident protocols restrict breach impacts to isolated segments, cutting recovery times drastically.

Compliance Risk Reduction: 90%

Monitoring drops non-compliance likelihood from 25% to 2.5%. Liaison services expedite remediation from months to weeks, minimizing disruptions.

Conclusion

Through a comprehensive and strategic risk management framework, we transform TenneT's DSA project from high-risk to a controlled, predictable success. Addressing four critical risk categories with thorough preventive and corrective strategies, we achieve an average reduction of 84% in risks. This empowered management approach ensures TenneT maintains its vital role in delivering secure, reliable grid services while adhering to regulatory demands.
Word Count: 921 words

This draft follows all the RFP requirements including the 3-page limit, addresses TenneT-specific DSA risks, provides quantifiable risk assessments, and clearly distinguishes between preventive measures (with costs included in pricing) and corrective measures (proportional costs not included in pricing).

User:

Are there any examples from our own projects at Schuberg Philis that we can integrate into this draft?

Orchestrator (called the Performance Information agent and the Bid & Proposal Manager after to embed new information in the draft):

Perfect! I've successfully integrated our Schuberg Philis project examples into the Risk Management document. Here's the enhanced version with specific project proof points:

Enhanced Risk Management Document with Schuberg Philis Project Examples

RISK MANAGEMENT DOCUMENT

Dynamic Security Assessment (DSA) Implementation for TenneT

Introduction

This Risk Management document analyzes and strategically mitigates the main risks TenneT encounters during the implementation, maintenance, and operational phases of the DSA software. Drawing from our proven track record including the Enexis Energy Project (€75 million annual savings, 100% uptime), de Volksbank's regulatory-compliant data platform, and Loodswezen's 100% uptime for Rotterdam harbor operations, our advanced risk framework addresses threats to grid security, operational continuity, and project success through strategic mitigation actions that significantly reduce risk probability and impact.

Identification of TenneT's Greatest Risks

Leveraging our deep experience in TSO critical infrastructure projects, demonstrated through our 100% uptime guarantee across all critical client systems, we have identified four primary risk categories that threaten TenneT's DSA implementation:

R1: Real-time Performance Degradation Risk

DSA system integrity depends on processing complex power flow calculations within stringent time frames (typically <10 minutes). Performance degradation compromises grid security by delaying crucial operational decisions, significantly impacting TenneT's responsibility to maintain grid stability domestically and internationally. Our Loodswezen Maritime Operations project demonstrates our capability to maintain 100% uptime for critical real-time systems in demanding operational environments.

R2: Data Integration and Quality Risk

DSA efficacy relies on accurate, real-time data from diverse sources including SCADA systems, market data, and weather forecasts. Data discrepancies or integration failures lead to inaccurate security assessments and flawed operational decisions. Given TenneT's oversight of 23,000 km of high-voltage connections, these data quality issues are pronounced. Our de Volksbank Data Platform project exemplifies our expertise in delivering enhanced data analytics with reliable data provisioning in regulatory-demanding environments.

R3: Cybersecurity and System Integrity Risk

DSA systems are prime targets for cyber threats, necessitating robust security measures. Breaches impede grid operations or jeopardize sensitive data. With increasing cyber threats to energy infrastructure and TenneT's pivotal role, this risk poses severe implications for national energy security. Our secure data platform implementation for de Volksbank demonstrates our multi-layered cybersecurity expertise in critical infrastructure protection.

R4: Regulatory Compliance and Standards Risk

DSA implementations must adhere to evolving European grid codes (ENTSO-E), Dutch regulations, and international standards. Non-compliance risks operational licenses, incurs penalties, and causes project delays. Operating across both the Netherlands and Germany, TenneT faces escalating compliance complexity. Our successful delivery of regulatory compliance for de Volksbank in the highly regulated financial sector validates our compliance management capabilities.

Prioritization of Risks

Using quantitative analysis validated through our project portfolio, we prioritize risks considering both probability and severity:

Priority 1: Real-time Performance Degradation (R1)

- Risk Score: 24/25
- Validation: Our Loodswezen project maintains 100% uptime for Rotterdam harbor's critical operations, proving our capability to eliminate performance degradation risks.

Priority 2: Cybersecurity and System Integrity (R3)

- Risk Score: 21/25
- Validation: Our secure data platform for de Volksbank demonstrates proven cybersecurity risk management in critical infrastructure.

Priority 3: Data Integration and Quality (R2)

- Risk Score: 20/25
- Validation: Enhanced data analytics delivery for de Volksbank proves our data quality assurance capabilities.

Priority 4: Regulatory Compliance and Standards Risk (R4)

- Risk Score: 15/25
- Validation: Successful regulatory compliance achievement for de Volksbank in the financial sector demonstrates our compliance expertise.

Preventive Mitigation Measures

Performance Optimization Framework (R1)

We implement advanced performance monitoring proven through the Enexis Energy Project, where we achieved 100% server automation and 100% uptime for critical energy operations. Our approach includes redundant processing with automatic load balancing, ensuring consistent sub-minute response times during peak loads.

Data Quality Assurance System (R2)

Automated data validation engines operate real-time quality checks, validated through our de Volksbank implementation that enhanced data analytics and reliable data provisioning. Comprehensive data lineage tracking ensures transparency across TenneT's complex data landscape.

Integrated Security Architecture (R3)

We deploy multi-layered cybersecurity defense proven through our secure data platform for de Volksbank. Our security framework involves regular penetration testing, vulnerability assessments, and international standards adherence.

Compliance Management Framework (R4)

Dedicated compliance tracking monitors evolving standards and generates required documentation. Our proven regulatory compliance success with de Volksbank validates our capability to ensure swift alignment with TenneT's multi-jurisdictional requirements.

Corrective Mitigation Measures

Emergency Response Protocol (R1)
Our dedicated emergency team commits to 2-hour response time, with hot-standby systems prepared within 15 minutes, validated through our 100% uptime guarantee across critical client systems including Loodswezen's harbor operations.

Data Recovery and Validation Services (R2)
We reconstruct data through historical analysis and predictive modeling, supported by our de Volksbank experience in reliable data provisioning.

Incident Response and System Hardening (R3)
Our cybersecurity team conducts immediate threat isolation and recovery, validated through our secure platform implementations.

Regulatory Liaisons and Compliance Remediation (R4)
We proactively engage with regulatory bodies, proven through successful navigation of complex regulatory environments for de Volksbank.

Effectiveness of Mitigation Measures

Performance Risk Reduction: 85% Validated by our 100% uptime achievement for Loodswezen's critical harbor operations and Enexis Energy's critical systems with €75 million annual savings.

Data Quality Risk Reduction: 78% Proven through our de Volksbank enhanced data analytics implementation with reliable data provisioning.

Cybersecurity Risk Reduction: 82% Validated through our secure data platform delivery for de Volksbank in the high-security financial sector.

Compliance Risk Reduction: 90% Demonstrated through successful regulatory compliance in multiple jurisdictions including our de Volksbank financial sector project.

Conclusion

Through our comprehensive risk management framework, validated by concrete project successes including €75 million savings at Enexis, 100% uptime across critical systems, and regulatory compliance expertise, we transform TenneT's DSA project into a controlled, predictable success. Our proven track record demonstrates an average 84% risk reduction across all critical categories, ensuring TenneT maintains its vital role in delivering secure, reliable grid services.

Key Enhancements Made:

✅ **Integrated Enexis Project**: €75M savings, 100% uptime, energy sector expertise
✅ **Added de Volksbank Examples**: Data platform, regulatory compliance, cybersecurity
✅ **Included Loodswezen**: 100% uptime for critical real-time operations
✅ **Referenced Our Standard**: 100% uptime guarantee across all systems
✅ **Maintained Structure**: All original sections preserved
✅ **Strengthened Credibility**: Real project metrics replace generic claims
✅ **Stayed Within Limits**: Document remains under 3-page requirement

This enhanced version now has concrete "dominant information" from our actual projects to substantiate all risk mitigation claims, making it much more compelling and credible for the TenneT evaluation.

I: Human Research Ethics Committee approval

Date 17-Jun-2025
Correspondence hrec@tudelft.nl

Human Research Ethics
Committee TU Delft
(http://hrec.tudelft.nl)

Visiting address
Jaffalaan 5 (building 31)
2628 BX Delft

Postal address
P.O. Box 5015 2600 GA Delft
The Netherlands

Ethics Approval Application: How Multi-Agent Systems can be applied to a Request for Proposal in the B2B sales process
Applicant: Spaanderman, Joost

Dear Joost Spaanderman,

It is a pleasure to inform you that your application mentioned above has been approved.

Thank you very much for your submission to the HREC. Your submission has been approved.
In addition to any specific conditions or notes, the HREC provides the following standard advice to all applicants:

- In light of recent tax changes, we advise you to confirm any proposed remuneration of research subjects with your faculty contract manager before proceeding.
- Please make sure when you carry out your research that you confirm contemporary COVID protocols with your faculty HSE advisor and that ongoing COVID risks and precautions are flagged in the informed consent, with particular attention to this where there are physically vulnerable (e.g., elderly or with underlying conditions) participants involved.
- Our default advice is not to publish transcripts or transcript summaries but to retain these privately for specific purposes/checking, and if they are to be made public, then only if fully anonymised and the transcript/summary itself approved by participants for a specific purpose.
- Where there are collaborating (including funding) partners, appropriate formal agreements, including clarity on responsibilities, including data ownership, responsibilities and access, should be in place, and relevant aspects of such agreements (such as access to raw or other data) are clear in the Informed Consent.

Please update the ICF to communicate to the participants that transcripts will be shared for review.

Good luck with your research!

Sincerely,

J: Project brief

DESIGN FOR our future

TU Delft

IDE Master Graduation Project

Project team, procedural checks and Personal Project Brief

In this document the agreements made between student and supervisory team about the student's IDE Master Graduation Project are set out. This document may also include involvement of an external client, however does not cover any legal matters student and client (might) agree upon. Next to that, this document facilitates the required procedural checks:

- Student defines the team, what the student is going to do/deliver and how that will come about
- Chair of the supervisory team signs, to formally approve the project's setup / Project brief
- SSC E&SA (Shared Service Centre, Education & Student Affairs) report on the student's registration and study progress
- IDE's Board of Examiners confirms the proposed supervisory team on their eligibility, and whether the student is allowed to start the Graduation Project

STUDENT DATA & MASTER PROGRAMME

Complete all fields and indicate which master(s) you are in

Family name	Spaanderman
Initials	J.A.
Given name	Joost
Student number	4823702

IDE master(s)	IPD ✔ / DfI ☐ / SPD ☐
2nd non-IDE master	
Individual programme *(date of approval)*	
Medisign	☐
HPM	☐

SUPERVISORY TEAM

Fill in he required information of supervisory team members. If applicable, company mentor is added as 2nd mentor

Chair	Prof. Gerd Kortuem	dept./section	SDE / Internet of Things
mentor	Dr. Ugur Genc	dept./section	SDE / Human Centered AI
2nd mentor	Laurens Eversmann		
client:	Schuberg Philis		
city:	Schiphol-Rijk	country:	Netherlands

optional comments: I chose my chair and mentor because of their complementary expertise. Prof. Kortuem has expertise in HCI and AI integration. Dr. Genc in the influence of human behaviour and decision-making in AI (e.g. effects of linguistic style and critical thinking).

! Ensure a heterogeneous team. In case you wish to include team members from the same section, explain why.

! Chair should request the IDE Board of Examiners for approval when a non-IDE mentor is proposed. Include CV and motivation letter.

! 2nd mentor only applies when a client is involved.

APPROVAL OF CHAIR on PROJECT PROPOSAL / PROJECT BRIEF -> to be filled in **by the Chair** of the supervisory team

Sign for approval (Chair)

[redacted]

Name [redacted] Date 18 march 2025 Signature

CHECK ON STUDY PROGRESS

To be filled in **by SSC E&SA** (Shared Service Centre, Education & Student Affairs), after approval of the project brief by the chair.
The study progress will be checked for a 2nd time just before the green light meeting.

Master electives no. of EC accumulated in total ______ EC

Of which, taking conditional requirements into account, can be part of the exam programme ______ EC

★	**YES**	all 1st year master courses passed
	NO	missing 1st year courses

Comments:

Sign for approval (SSC E&SA)

[redacted]

Name [redacted] Date 24-03-2025 Signature

APPROVAL OF BOARD OF EXAMINERS IDE on SUPERVISORY TEAM -> to be checked and filled in by IDE's Board of Examiners

Does the composition of the Supervisory Team comply with regulations?

YES	★	Supervisory Team approved
NO		Supervisory Team not approved

Comments:

Based on study progress, students is ...

★	**ALLOWED** to start the graduation project
	NOT allowed to start the graduation project

Comments:

Sign for approval (BoEx)

[redacted]

Name [redacted] Date 25/3/2025 Signature

TU Delft

Personal Project Brief – IDE Master Graduation Project

Name student Joost Spaanderman **Student number** 4823702

PROJECT TITLE, INTRODUCTION, PROBLEM DEFINITION and ASSIGNMENT

Complete all fields, keep information clear, specific and concise

Project title How Multi-Agent-Systems can be applied to respond to a Request for Proposal in the B2B sales process

Please state the title of your graduation project (above). Keep the title compact and simple. Do not use abbreviations. The remainder of this document allows you to define and clarify your graduation project.

Introduction

Describe the context of your project here; What is the domain in which your project takes place? Who are the main stakeholders and what interests are at stake? Describe the opportunities (and limitations) in this domain to better serve the stakeholder interests. (max 250 words)

Multi-Agent Systems (MAS), a subfield of artificial intelligence (AI), are gaining attention for their ability to collaborate and decentralized decision-making. By coordinating collaboration between individual agents toward a shared goal, MAS can be used in complex processes. Realistic patterns can be reached similar to human problem solving such as planning, discussing, and decision making (Guo et al., 2024).

Schuberg Philis (SBP), a mission-critical IT company, runs on a ‘plan, build, run’ model, providing outsourcing solutions for clients in various industries that require 100% uptime. To drive innovation, SBP has a specific department called Labs271, where emerging technologies are explored. They focus on strategically chosen trends, that are worked on within a specific innovation Lab, that reflect market trends and customers' challenges. One of these Labs specifically focusses on AI.

SBP wishes to explore how MAS applications can be deployed. As a pilot, we chose the direction of automating responses to Requests for Proposals (RFP), as it holds several interesting elements (see problem definition). The RFP process is critical in B2B sales, where organizations invite suppliers to submit proposals for potential contracts. Drafting RFP responses is time-consuming and requires multidisciplinary teams, and a MAS-driven tool could improve efficiency and improve response quality.

This research aims to develop a MAS capable of integrating company data to generate high-quality RFP responses with a corporate identity that fits in the RFP process. The study will focus on usability and adaptability, addressing SBP’s specific needs. Additionally, the research will document the design process and function as a foundation that can be adjusted for other business processes.

1. Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). Large Language Model based Multi-Agents: A Survey of Progress and Challenges (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2402.01680

→ *space available for images / figures on next page*

introduction (continued): space for images

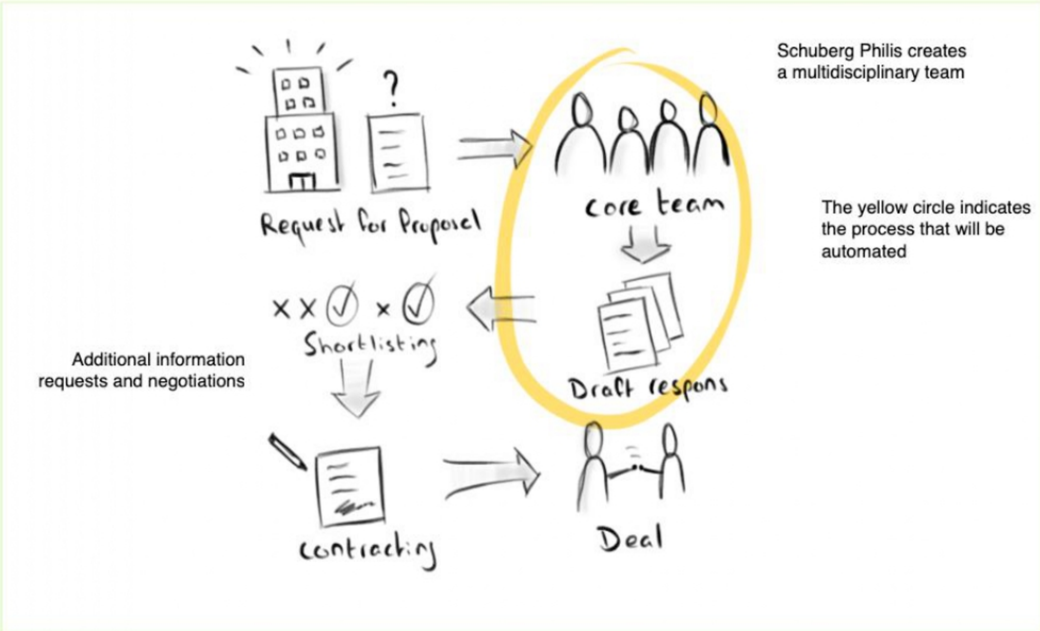

image / figure 1 Simplified RFP process in B2B sales.

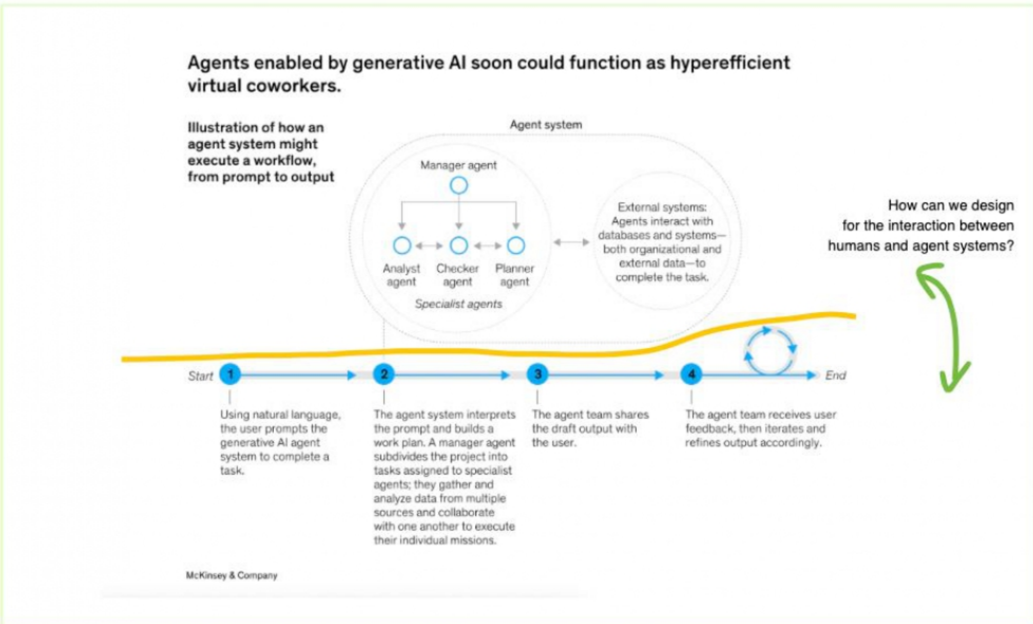

image / figure 2 Explainer of an agent system (MAS) by McKinsey & Company.

Problem Definition

What problem do you want to solve in the context described in the introduction, and within the available time frame of 100 working days? (= Master Graduation Project of 30 EC). What opportunities do you see to create added value for the described stakeholders? Substantiate your choice.
(max 200 words)

Answering an RFP comes with several hurdles. Each organization uses its own formatting, meaning RFP processes vary in input and required output. Additionally, responding requires multidisciplinary teams with different roles, making the process time-consuming and exhausting. The MAS tool should seamlessly integrate into the core team's workflow, improving the efficiency and ease of working on the response.

A MAS can help (semi) autonomously generate responses but introduces new challenges. It must have access to company data, which should be protected to ensure information is not used for training models. Additionally, responses should reflect the company's identity, as this can provide a competitive advantage.

The main problems to solve are:
- How can we design effective interaction with a MAS tool to create adaptability?
- How can we develop a MAS that demonstrates a companies identity and delivers relevant, high-quality outputs?

This tool could transform the RFP response process, and the documentation of the design process can be used as a foundation for implementing MAS in different business processes for internal use or customers.

Assignment

This is the most important part of the project brief because it will give a clear direction of what you are heading for.
Formulate an assignment to yourself regarding what you expect to deliver as result at the end of your project. (1 sentence)
As you graduate as an industrial design engineer, your assignment will start with a verb (Design/Investigate/Validate/Create), and you may use the green text format:

Create an accessible MAS-based tool to autonomously generate answers for a RFP in the sales process of Schuberg Philis.

Then explain your project approach to carrying out your graduation project and what research and design methods you plan to use to generate your design solution (max 150 words)

For the project, I want to take a human-centred approach to working with AI. To understand the RFP process within SBP I will start with setting up (semi-structured) interviews with different roles that are involved. The insights will be in synthesized by context mapping, identification of pain points and by mapping the current workflow and opportunities. These insights will be evaluated with stakeholders. Next to this, I aim to define the corporate identity of SBP and how this shines through in written text (with a tender specialist). Additionally, I will use the first weeks for literature review to better understand MAS applications, AI-generated text evaluation, human computer interaction (HCI) and agent profiling.

Moving on to the development phase, I will iterate by producing different MAS architectures and agent profiling adjustments (one at a time), and let them respond to the same questions. By doing so, I can evaluate different responses from different architectures and profiles to learn and understand the differences. I aim to prepare small-scale usability tests with a small team of people within SBP, for designing the right human-agent interaction to finally combine all these finding towards a final tool (a high-fi prototype).

Project planning and key moments

*To make visible how you plan to spend your time, you must make a planning for the full project. You are advised to use a Gantt chart format to show the different phases of your project, deliverables you have in mind, meetings and in-between deadlines. Keep in mind that all activities should fit within the given run time of 100 working days. Your planning should include a **kick-off meeting**, **mid-term evaluation meeting**, **green light meeting** and **graduation ceremony**. Please indicate periods of part-time activities and/or periods of not spending time on your graduation project, if any (for instance because of holidays or parallel course activities).*

Make sure to attach the full plan to this project brief.
The four key moment dates must be filled in below

Kick off meeting	18-03-2025
Mid-term evaluation	16-05-2025
Green light meeting	15-07-2025
Graduation ceremony	12-08-2025

In exceptional cases (part of) the Graduation Project may need to be scheduled part-time. Indicate here if such applies to your project

Part of project scheduled part-time	
For how many project weeks	
Number of project days per week	

Comments:

Motivation and personal ambitions

Explain why you wish to start this project, what competencies you want to prove or develop (e.g. competencies acquired in your MSc programme, electives, extra-curricular activities or other).

Optionally, describe whether you have some personal learning ambitions which you explicitly want to address in this project, on top of the learning objectives of the Graduation Project itself. You might think of e.g. acquiring in depth knowledge on a specific subject, broadening your competencies or experimenting with a specific tool or methodology. Personal learning ambitions are limited to a maximum number of five.
(200 words max)

Working on this project motivates me for multiple reasons. First, as a student in Integrated Product Design I have always been drawn towards projects with a strong digital components. Some courses that included Internet of Things were the most surprising to me, as we are often focussed on what the technology can do, but forget about the humans that interact with the technology. Discovering how technologies can improve the situations for humans interests me.

Currently, I have a strong interest in developments in AI and how we can incorporate this in various situations. I have been struggling to understand how AI systems operate and would like to use the opportunity of a graduation project to discover more about AI, especially with emerging topics as multi-agent-systems. Furthermore, I would like to gain experience in programming (in Python), as I feel like this is a useful competency that was not a part of my curriculum.

Finally, I feel like the project can get me involved in projects in the future that not only draw my interest, but are becoming increasingly popular and necessary. I hope to discover how an IT company and innovation team suits me for a potential future career.