Journal of Software Engineering and Applications, 2025, 18(9), 373-388
https://www.scirp.org/journal/jsea
ISSN Online: 1945-3124
ISSN Print: 1945-3116

Research Report on Multi-Agent Collaboration in Shipping: Expert Agent Construction and Application

Chong Pan

Cosco Shipping Technology Co., Ltd., Shanghai, China
Email: pan.chong@coscoshipping.com

How to cite this paper: Pan, C. (2025) Research Report on Multi-Agent Collaboration in Shipping: Expert Agent Construction and Application. *Journal of Software Engineering and Applications*, **18**, 373-388.
https://doi.org/10.4236/jsea.2025.189022

Received: August 29, 2025
Accepted: September 21, 2025
Published: September 24, 2025

Abstract

As a vital pillar of global trade, the shipping industry possesses vast and complex expertise, data, and scenarios. Compiling shipping research reports demands high professional skills from data scientists and industry analysts, proving both time-consuming and labor-intensive. With the rapid advancement of artificial intelligence technology, the integration of large language models, RAG systems, and Multi-Agent systems with specialized shipping scenarios has emerged as a key direction for assisting shipping research and enhancing data analysis efficiency. Traditional RAG single-agent systems integrated with large models typically follow a linear "retrieval-generation" workflow when generating research reports. This approach suffers from disconnects between retrieval and generation, neglects conceptual connections, amplifies content hallucinations, and lacks the ability to generate comprehensive narratives and insights. This paper addresses the limitations of RAG-based single-agent systems in handling complex tasks requiring deep reasoning, multi-step planning, and tool collaboration. Drawing on multi-agent system principles, it designs a hierarchical collaborative expert agent architecture comprising management and execution layers. Detailed functional specifications for each module are provided. It proposes a four-stage framework for generating shipping report content and conducts a case study analyzing the impact of the Red Sea crisis on the global shipping landscape. Results demonstrate that this expert agent achieves significant improvements in completeness, accuracy, and research depth compared to single-agent RAG systems when generating reports. The findings demonstrate that the constructed shipping report expert agent substantially enhances the completeness of shipping content retrieval and the quality of report generation, elevating the intelligence level of shipping information analysis. This holds significant practical implications for advancing the intelligent transformation of the shipping industry.

Keywords

Shipping Industry, Large Language Model, Multi-Agent System (MAS), Shipping Professional Knowledge Base, Online Knowledge Graph, AI Agent

1. Introduction

As the lifeline of global trade, the shipping industry possesses vast and complex specialized data and industry knowledge spanning multiple domains including vessel technology, route planning, port management, international trade regulations, and maritime laws. Simultaneously, it faces multiple complex challenges such as persistent supply chain volatility, heightened geopolitical risks, and accelerated green and low-carbon transformation. With the rapid advancement of artificial intelligence technology, developing intelligent systems to generate specialized research reports with deep industry insights—quickly, efficiently, and accurately—has become a crucial initiative for driving the industry's intelligent transformation.

The preparation of shipping research reports demands high professional expertise from data scientists and industry analysts, and is both time-consuming and labor-intensive. This is particularly true when integrating massive, heterogeneous, and dynamic shipping data and knowledge—such as Automatic Identification System (AIS) data, port operation data, global economic indices, and news sentiment—which requires significant time and effort to organize. Traditional RAG systems with integrated large language models typically follow a linear "retrieval-generation" workflow, often confined to single tasks. They lack comprehensive narrative and insight generation capabilities, resulting in issues such as disconnect between retrieval and generation, overlooked conceptual connections, loss of long-tail information, and amplified hallucinations in generated content. When confronted with complex tasks requiring deep reasoning, multi-step planning, and tool collaboration, these systems prove inadequate.

This study focuses on the application of multi-agent systems in generating shipping reports. By simulating human information processing, decision-making, and operational procedures, it aims to develop a modular, multi-agent collaborative agent solution with enhanced reasoning, planning, and tool utilization capabilities. This approach elevates AI applications in the shipping sector from "operational assistance" and "single-purpose analysis" to a new level of "comprehensive knowledge creation".

This study first analyzed several system architecture styles for multi-agent systems. By integrating the business processes and characteristics identified in shipping industry research reports, it constructed the system architecture for the shipping report expert agent. Detailed designs were developed for each agent module within the architecture, notably proposing a four-stage framework mechanism for generating shipping report content. A case study was conducted analyzing the impact of the Red Sea crisis on the global shipping landscape. Ultimately, an expert agent construction plan compliant with systems engineering principles was established.

2. Architectural Concepts of Research Report Expert Agents

2.1. Multi-Agent System Overview

Multi-Agent Systems (MAS) originate from Distributed Artificial Intelligence (DAI), with the core idea being to decompose complex problems into solutions achieved through collaboration among multiple relatively simple, autonomous agents, as shown in **Figure 1**. According to a seminal IEEE review [1], agents typically exhibit characteristics such as autonomy, responsiveness, proactivity, and sociality. Communication, coordination, negotiation, and cooperation among agents enable the handling of large-scale, distributed, and dynamically changing tasks that RAG systems and single agents cannot effectively manage.

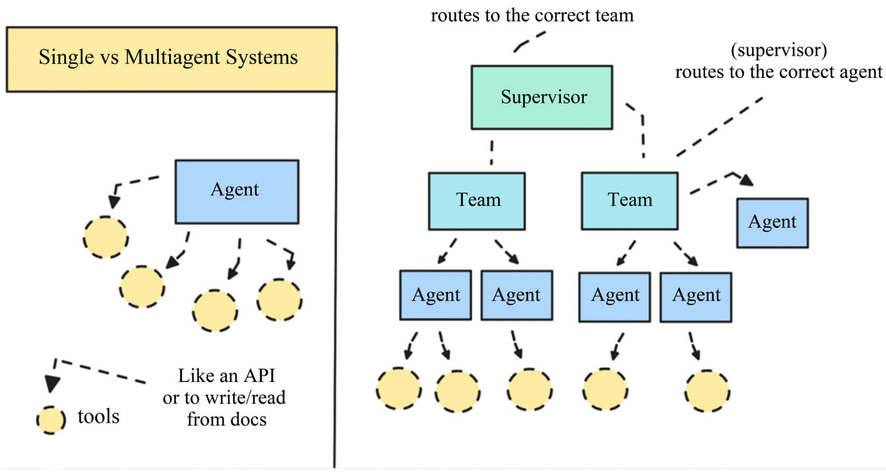

Figure 1. MAS system.

2.2. Architectural Analysis of the Shipping Report Expert Agent

The architecture of Multi-Agent Systems (MAS) is crucial for designing expert agents in shipping reports. Early research, such as the Carnegie Mellon University report [2], has already positioned MAS as a specialized software architecture style. Mainstream architectures include:

(1) Hierarchical Architecture

The hierarchical architecture is the framework most aligned with traditional human organizational management models, its design philosophy rooted in the concept of decomposing complex problems step by step. Agents are organized across distinct tiers, forming a clear tree-like or pyramid structure. As described in Milvus' research paper, this structure enables management at different levels of abstraction—from top-level strategic planning to bottom-level execution. Control flow operates with higher-level agents defining global strategies and decomposing objectives, then delegating specific subtasks to lower-level agents. These lower-level agents focus on task execution and tactical decision-making. Communication primarily follows a vertical pattern: superiors issue commands and tasks to subordinates, while subordinates report status and results to superiors. Horizontal communication between peers is typically restricted and requires coordination through superiors. This architecture effectively reduces complexity, facilitates structured scaling, and offers advantages in clear task control and unified objectives. However, it also introduces challenges such as limited flexibility, reduced adaptability, and coordination delays.

(2) Federated Architecture

The core of a federated architecture lies in "coordination" rather than "control", providing a means for multiple independent systems to work together. A federation comprises several autonomous, potentially heterogeneous MAS or organizational systems that collaborate externally through shared protocols and standards. Participants retain high autonomy in control flow, with collaboration grounded in voluntary participation, negotiation, and mutual benefit rather than coercive directives. This architecture fosters organizational cooperation while safeguarding data sovereignty and privacy. However, it also carries risks of complex coordination mechanisms and high trust costs.

(3) Decentralized Architecture

Decentralized architecture (often referred to as networked or swarm-based architecture) emphasizes the emergence of complex global behavior through simple local interaction rules. Structurally, it adopts a flat network topology where all agents function as peer-to-peer nodes. Any agent can directly communicate and interact with any other agent within the network. There is no central control unit; the system's overall behavior "emerges" through the collective local interactions and self-organizing processes of numerous agents, rather than being pre-planned. This architecture offers high flexibility and adaptability, strong robustness, but poses challenges for global coordination, with behavior exhibiting unpredictability.

The generation of shipping research reports is a highly structured, knowledge- and data-intensive, cross-organizational collaborative process encompassing the entire journey from knowledge and data collection to the final report release. Aligned with current industry business processes, it primarily consists of the following six steps: (1) Goal Definition and Requirements Analysis; (2) Knowledge and Data Collection and Integration; (3) Data Cleaning and Preprocessing; (4) Data Analysis and Modeling; (5) Report Generation and Visualization; (6) Review, Release, and Feedback. The entire task workflow exhibits a hierarchical, sequential decision-making and execution structure with clear command relationships between superiors and subordinates. In this scenario, adopting either a federated MAS architecture or a decentralized MAS architecture would incur additional costs in information flow between control modules. It would also struggle to handle out-of-order execution of sequential steps, while the advantages of these architectures in flexibility and randomness would be difficult to leverage. When employing a decentralized MAS architecture, high failure rates and prolonged execution times may also arise. In contrast, a hierarchical MAS architecture demonstrates superior adaptability to sequential execution scenarios, achieving the highest levels of stability and execution success rates.

Therefore, this study will draw upon the principles of a layered architecture to construct a system that combines top-level coordination with specialized division of labor at the lower levels.

3. Construction Plan for the Expert Agent in Shipping Research Reports

3.1. System Architecture Design

3.1.1. Layered Collaborative Architecture

The shipping research report expert agent draws upon the concept of a layered agent architecture, adopting a hybrid multi-agent framework based on the classic "Orchestrator-Worker" model. By simulating an efficient team of human research experts, the entire agent system is decomposed into multiple sub-agents with independent functionalities, while being hierarchically divided into management and execution layers [3] [4], as shown in Figure 2.

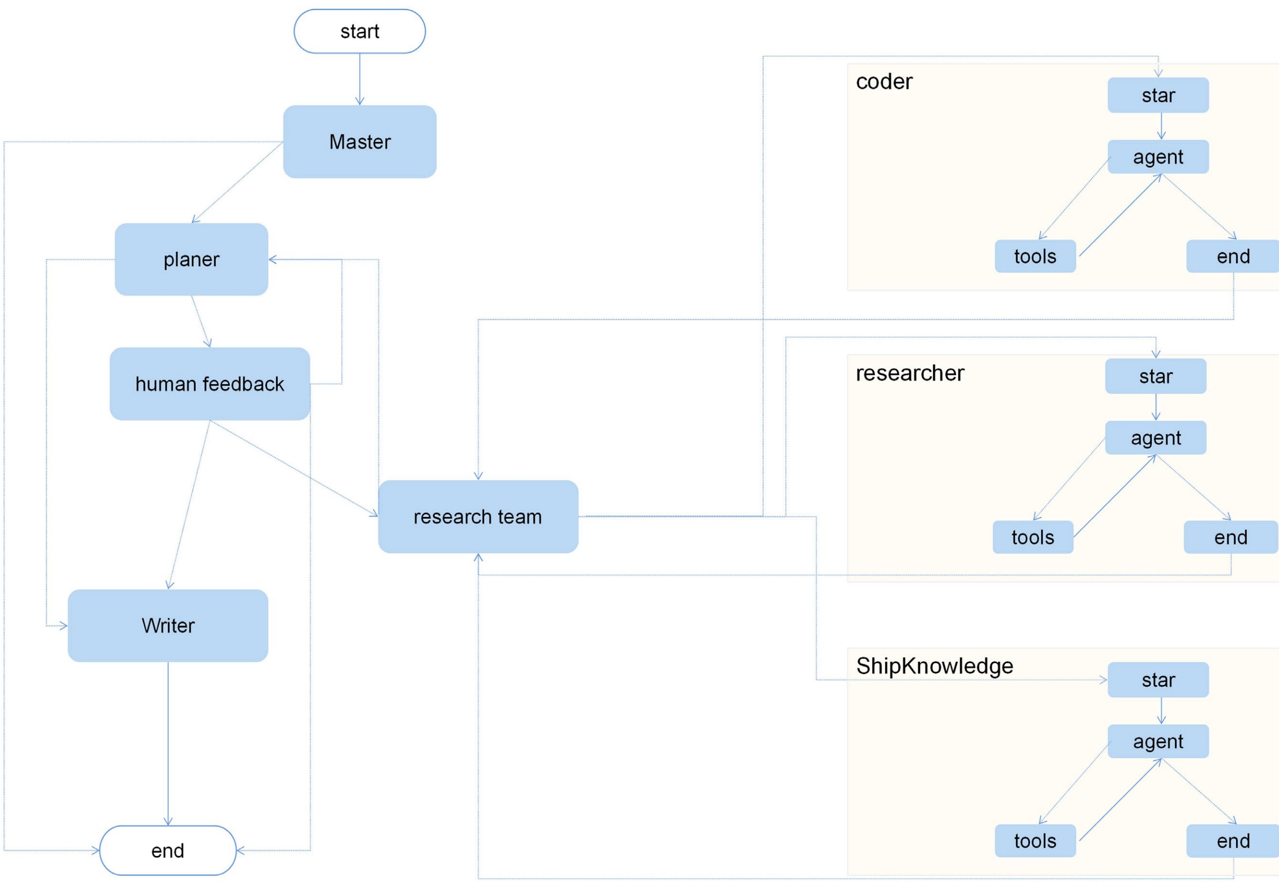

Figure 2. Over all architecture diagram.

Management Team: Responsible for requirements analysis, resource planning, task allocation, and quality control. Composed of the Master (Project Manager) AI, Planner (Architect) AI, and Human Feedback (Requester) AI, it assumes the roles of project manager, architect, and requester.

Execution Layer: Composed of a set of parallel "functional expert agents" playing roles across different domains. It consists of the Background Agent, Research Team Agent, Ship Knowledge Agent, Researcher Agent, Coder Agent, and Writer Agent, which execute specific tasks.

This research design facilitates the step-by-step decomposition of complex report generation problems into simple, concrete, and executable tasks. Each agent can focus on executing specific tasks, ensuring both logical task decomposition and professional execution depth. It effectively manages user requirements to achieve consistent objectives.

3.1.2. Role Responsibilities and Core Processes

Table 1. Expert agent role responsibilities.

Agent Type		Role	Responsibilities
Master AI Agent		Project Manager	The team's "brain" and chief strategist, responsible for the initial understanding of requirements and resource allocation.
Planner AI Agent		Architect	The team architect, responsible for solution design and task decomposition.
Human Feedback Agent		Demand Side	The client side of the team, responsible for reviewing proposals and ensuring quality control.
Background Agent		Background Research	Team investigator responsible for conducting "background research" or "preliminary research" on user requirements.
Research Team AI Agent	Ship Knowledge AI Agent	Shipping Knowledge Research	Content Team—Responsible for managing and developing the shipping knowledge base and building online knowledge graphs.
	Researcher Agent	Research on Report Content	Content Team—Responsible for researching, executing, and compiling tasks related to report content.
	Coder Agent	Algorithm Programming	Content Team—Responsible for data processing, data analysis, and code execution for reporting tasks.
Writer AI		Report Writer	The team's report writer is responsible for generating clear, professional research reports using shipping report templates.

As shown in Table 1 and Figure 3, The Project Manager (Master) guides the Planner in formulating plans based on user-input tasks, background research, and feedback from the requester (Human Feedback). It continuously evaluates the execution status and completeness of subtask results, reflecting and replanning when necessary. The Planner constructs a directed acyclic graph (DAG) for subtasks and dynamically selects appropriate execution-layer agents to achieve structured and adaptive multi-step execution. Execution-layer agents execute specific subtasks, with the Writer ultimately generating the final report.

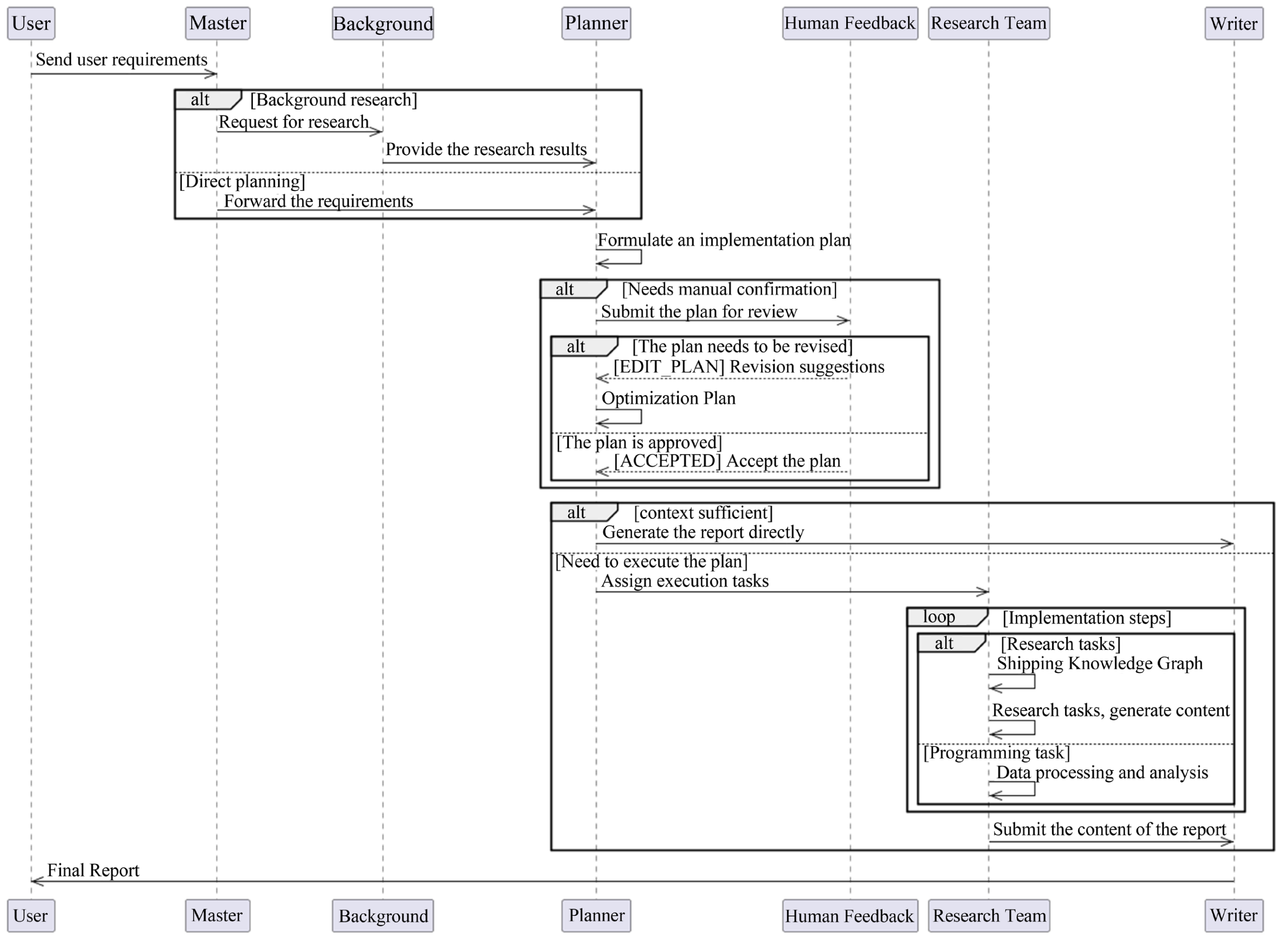

Figure 3. Multi-agent interaction diagram.

3.2. Key Functional Module Design

3.2.1. Shipping Knowledge Base

The construction of the shipping knowledge base forms the foundational capability core of the entire report generation system. The dataset used to construct this knowledge base originates from COSCO SHIPPING Technology Co., Ltd.'s shipping industry big data. It encompasses vessel basic information and operational data, port basic information and operational data, geographic information such as global satellite nautical charts, massive volumes of AIS global message data (including real-time and historical data), and extensive meteorological observation data (including real-time, historical, and 7-day forecast data). Drawing upon the experience gained from building the transportation knowledge graph at Chang'an University [5], this system adopts a combined bottom-up and top-down approach. It utilizes an ontology framework to guide data extraction while simultaneously enriching the ontology structure with data. This approach overcomes the limitations of single-method approaches, holistically enhancing the quality and usability of the knowledge graph.

Our maritime knowledge graph construction follows this path, as shown in Figure 4: first, we collect data from the shipping sector, leveraging the diverse professional shipping data and knowledge provided by our company's Ship Vision platform. Key data types in shipping include structured data (e.g., vessel registration information, port operational data, port facility data), typically stored in databases for easy extraction and processing; semi-structured data (e.g., maritime regulations, industry standard documents), requiring information extraction techniques to identify key insights; and unstructured data (e.g., shipping accident investigation reports, news articles, social media content), necessitating analysis and extraction through natural language processing (NLP) technologies. Finally, construction follows this technical pathway: leveraging natural language processing (NLP) techniques like named entity recognition (NER) and relation extraction (RE) to automatically identify core entities from text—such as "vessels", "ports", "routes", "companies", "cargo", and "events"—along with their complex relationships (e.g., "owned by", "calls at", "serves", "affected by"). A vast and intricate semantic network is then constructed using graph databases like Neo4j, providing rich contextual information for the agent's reasoning. This enables the agent to understand "why" rather than merely "what".

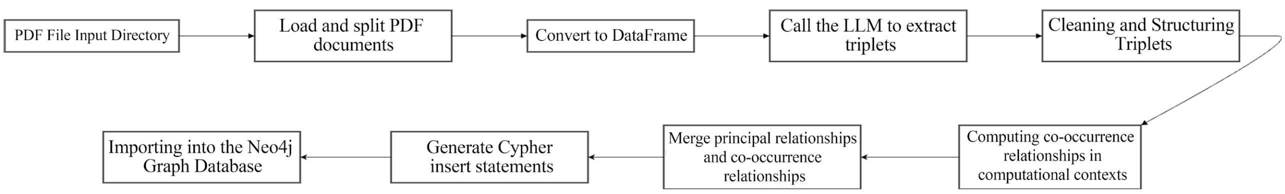

Figure 4. Workflow of the shipping knowledge base.

Through multiple rounds of expert feedback [6], key entities and relationships within the shipping domain were identified to construct a knowledge graph encompassing vessel types, port information, route planning, regulations, and collision avoidance rules, As shown in Figure 5. The knowledge base not only supports data queries and relationship analysis but also provides report generation support through semantic matching, including original text summaries and detailed descriptions.

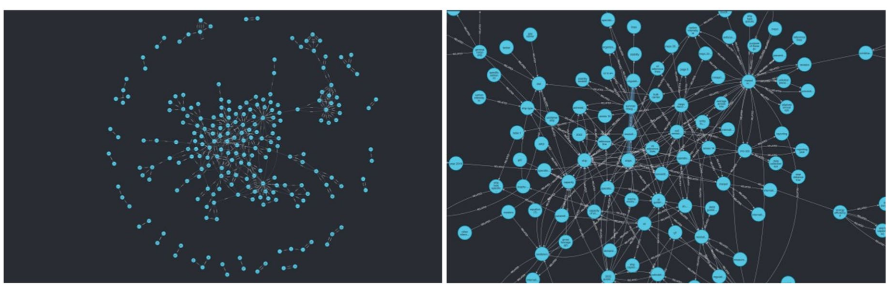

Figure 5. Shipping knowledge graph.

3.2.2. Management Agent Design

The management agent comprises the Master agent, Planner agent, and Human Feedback agent. The functions and objectives of each agent are as follows:

■ **Master Agent for Dynamic Task Allocation**

Master AI assumes the role of project manager, with core capabilities in dynamic task allocation and resource distribution. Based on the complexity of user requirements, it will initiate background investigations and collaborate with users as needed, flexibly selecting from three execution modes, as shown in **Table 2**:

Table 2. Task execution mode.

Mode	Applicable Scenarios	Execution Process
Simple report	Simple Report (Generation of COSCO Shipping Lotus's Voyage Report for the Past Six Months)	Directly generated by the report Agent based on the knowledge base (individual + expert), offering the highest efficiency; 1 AI agent, 3 - 10 tool invocations.
Combined report	Combined Report (e.g., Research Report on the Progress of Intelligent Ship Management)	The research team calls the API to retrieve information, integrates content from the knowledge base (individual + expert), and generates reports in a standardized format; 2 - 4 agents, each making 10 - 15 tool calls.
Complex research	Complex Research (e.g., Analysis of the Impact of the Red Sea Crisis on the Global Shipping Landscape)	Complex Multi-Step Reasoning: Personal Knowledge Base + Shipping Knowledge Base and Knowledge Graph + Web Retrieval + Other Tools, Activating All Agents with Clearly Defined Responsibilities.

This on-demand dynamic scheduling mechanism enables the system to achieve the optimal balance between efficiency and effectiveness across different reporting tasks.

■ **Planner Agent Based on DAG Task Planning**

The Planner AI Agent as the team's architect, responsible for designing solutions and decomposing tasks. By communicating and refining technical proposals with the requester (user), it provides the execution team with clear technical guidance and task breakdowns. The output of task planning is a directed acyclic graph (DAG) in JSON format, explicitly defining dependencies between each subtask and ensuring execution sequence. The plan is iteratively refined through interaction with the Human Feedback (requester) agent. This enables the system to handle complex research tasks while maintaining sufficient flexibility and controllability [7].

We employ a large-model prompt strategy combining"thought chains and structured sketches" for construction. In task planning, we integrate directed acyclic graphs (DAGs) to analyze task dependencies, decompose tasks into multiple subtasks, and construct DAGs to represent the execution order of these subtasks. Nodes in the DAG represent tasks, while edges denote dependencies between tasks. When generating the global task graph, we employ a structured sketch strategy to represent each task's reasoning process as concise, structured steps. This approach reduces redundancy in reasoning and enhances computational efficiency. By leveraging DAGs and generating structured sketches, the global task graph can be produced in a single forward pass. This method avoids the high-latency iterations characteristic of ReAct, thereby improving the efficiency and real-time performance of task planning, as shown in **Figure 6**.

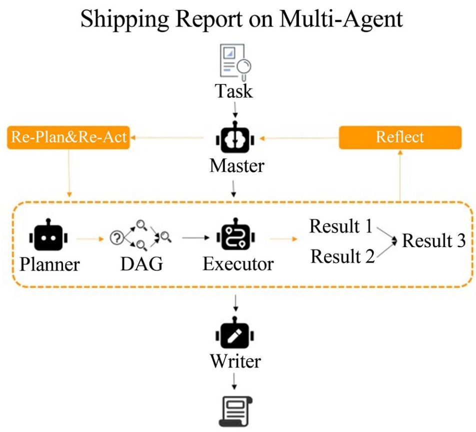

Figure 6. DAG-Based task planning.

- **Human Feedback Agent Based on Human-in-the-Loop (HITL)**

As the demand side of the team, the Human Feedback agent is responsible for reviewing proposals and ensuring report quality. After the Planner generates a research plan, it pauses the execution process to await user review. It collects user feedback and revision suggestions on the research plan, determines the next execution path, and ensures the AI-generated research plan aligns with user expectations [8].

This Human-in-the-Loop (HITL) design philosophy integrates human expertise—through user intervention at critical junctures—with the automation capabilities of intelligent agents. It achieves transparency, controllability, explainability, and trustworthiness in agent decision-making, aligning with responsible AI principles. This approach enhances user experience while ensuring system outputs are both efficient and reliable.

3.2.3. Design of Execution-Level Intelligent Agents

The Execution Layer Agent primarily handles specific tasks and consists of three major components: the Background Agent, the Research Team Content Team Agent (including the Ship Knowledge Agent, Researcher Agent, and Coder Agent), and the Writer Agent. The functions and objectives of each agent are as follows:

- **The Background AI that pre-emptively prepares all functionalities for you**

The Background Agent (also known as the Background Investigator) is a specialized agent within the Shipping Report Expert Agent System. Functioning like an experienced assistant, it proactively conducts "background investigations" or "preliminary research" for user queries before formulating the formal planning report generation strategy. After a user inputs a query, it automatically searches for and collects relevant background materials using tools like search engines and web crawlers. This provides the subsequent Planner AI with richer contextual information, enabling it to formulate more reasonable and comprehensive research plans.

This design, which proactively supplements information to reduce misjudgments or omissions caused by insufficient data in the Planner, automatically fills in missing details to make the system smarter, thereby enhancing research quality.

■ **Research Team AI: Turning Plans into Reality**

The Research Team agent is a collaborative team comprising three specialized agents: the Ship Knowledge agent, the Researcher agent, and the Coder agent. It is primarily responsible for executing research plans formulated by the Planner. The Ship Knowledge Agent retrieves and acquires the most relevant knowledge fragments along with sufficient contextual content from the shipping knowledge base, enabling the online generation of task knowledge graphs. The Coder AI provides a Python code execution environment, handling data and knowledge collected by the Researcher AI. It focuses on data processing, computation, and analysis. The Researcher AI invokes various tools to collaborate with the Ship Knowledge and Coder AIs in executing specific research plans, ultimately generating high-quality research content.

(1) Ship Knowledge AI Agent

The Shipping Report Expert Agent is an expert agent tailored for shipping scenarios and operations. Its core report content typically comprises a personal knowledge base, a shipping expert knowledge base, and relevant web search knowledge. Among these, the Ship Knowledge Agent primarily handles retrieval, recall, reordering of content from the personal and shipping expert knowledge bases, as well as the construction of an online knowledge graph. The weighting for integrating shipping knowledge content is as follows:

Shipping Report Content = Shipping Expert Knowledge Base*X + Personal Knowledge Base*Y + Latest Online Knowledge*Z;

Among these, X, Y, and Z represent weight configurations. We stipulate that $X \geq Y \geq Z$, meaning the shipping expert knowledge base carries the highest weight while online knowledge holds the lowest weight. Furthermore, we specify that $X \geq Y + Z$, ensuring the shipping expert knowledge base accounts for over 50% of the weight to guarantee the report's professionalism and accuracy. In summary, we first set the weight of the shipping expert knowledge base to 50% and ensure the weight of the personal knowledge base exceeds that of web knowledge, allowing the report to fully reflect personalized experience. The final weight parameter configuration is determined as: X = 50%, Y = 30%, Z = 20%.

(2) Researcher (Dynamic Researcher) Agent

The creation of the Researcher AI is a dynamic process. We do not view the AI as an "all-purpose technician" who must memorize how to use every tool, but rather shape it into a "wise researcher." Based on the research task, the agent first uses the tool retrieval module to select the most relevant tools and MCP services, forming a "small yet powerful" toolset. This toolset, combined with the reasoning capabilities of a large language model (LLM), collectively defines the "research boundary" for the current task. The specific workflow is as follows:

a) Further decompose sub-tasks to precisely identify the key capabilities required to complete the current task, such as "reading files".

b) Autonomously generate tool requirements by describing the tool's type and functionality in natural language.

c) Utilize the tool retrieval module to filter the most relevant tools and MCP services, forming a toolset. Combine this with the reasoning capabilities of large language models to construct a researcher agent for the current task.

■ **The Writer AI that turns research into results**

The Writer AI is the "writing specialist" within the Shipping Report Expert AI suite. Its primary function is to synthesize all knowledge, data, analytical findings, charts, and other information gathered by the Content Team AI. It then organizes this content into shipping reports structured and styled to present complex research in clear, accessible language with logical flow. Beyond generating text reports, we've also designed it to produce multiple content formats, including PowerPoint presentations and podcasts.

■ **Four-Stage Content Generation Framework**

Traditional RAG-based shipping report generation typically follows a linear "retrieval-generation" process, suffering from issues such as disjointed retrieval and generation, overlooked conceptual connections, amplified generative hallucinations, and a lack of comprehensive narrative and insight generation capabilities. The shipping domain is complex and far-reaching, requiring research reports to accomplish intricate professional tasks involving deep reasoning, multi-step planning, and collaboration among multiple tool agents. Such highly specialized scenarios demand exceptionally professional knowledge and data—not fabricated by large models [9]. Therefore, we developed a four-stage content generation framework that dynamically constructs knowledge graphs for user tasks online and uses these graphs to generate shipping report content.

We designed a four-stage content generation framework comprising preparation, real-time knowledge mapping, evaluation, and fusion stages. The preparation and real-time knowledge mapping stages utilize the Ship Knowledge agent, while the review evaluation and fusion stages employ the researcher agent and coder agent, as detailed below:

Preparation Phase: Conduct further background research and intent identification for the research task to expand relevant knowledge, extract core keywords, and construct multiple keyword queries for retrieving information from shipping expert knowledge bases, personal knowledge repositories, and conducting networked searches. Real-time Knowledge Graph Construction Phase: For retrieved knowledge fragments and search content, re-ranking is performed based on predefined content weighting ratios. Subsequently, llmgraph is utilized to construct a real-time knowledge graph (KG) centered around candidate shipping report content.

Evaluation Phase: Based on the established multidimensional relevance assessment framework, we comprehensively evaluate four dimensions: literal relevance (keyword matching), semantic relevance (deep semantic understanding), structural relevance (alignment between document type and query intent), and subject-specific relevance (coverage of specialized concepts). Utilizing PageRank analysis to assess graph structure, we calculate the importance score for each node. By combining node importance with entity evaluation results, we select the most relevant entities as primary external knowledge sources.

Integration Phase: Utilize LLM to consolidate its most relevant entities and associated content, combining code analysis to generate research content for the report task.

This mechanism ensures that the final research output is more task-relevant and incorporates more comprehensive industry expertise, significantly enhancing the depth, breadth, and professionalism of the report's content.

4. Case Study: Analysis of the Red Sea Crisis' Impact on the Global Shipping Landscape

4.1. Case Background and Research Task Definition

Background: Since late 2023, escalating geopolitical tensions in the Red Sea and Gulf of Aden have led to multiple attacks on merchant vessels. To mitigate risks, major global container shipping companies have successively announced suspensions of Red Sea routes, instead rerouting ships around the Cape of Good Hope at the southern tip of Africa. This sudden development has severely disrupted one of the world's busiest trade corridors.

Task: We issued an advanced directive to the shipping intelligence agent: "Conduct a comprehensive and in-depth analysis of the short- and medium-term impacts of the Red Sea crisis on the global shipping landscape, and generate a research report."

4.2. Report Generation and Results Evaluation

We break down the effects into four dimensions for observation, As shown in Table 3:

Table 3. Experimental results latitude observation.

Dimension	RAG-based Single Agent	Shipping Expert Agent (Multi-Agent System)	Experimental results
Accuracy rate	92.3%	98.1%	↑5.8 percentage points, primarily driven by multi-agent breadth-first search, industry knowledge bases, and manual review.
Structural integrity	37%	93%	Multi-agent systems: The management agent first locks the outline, then assigns tasks, while specialized agents write reports; Single-agent systems: Prone to missing modules.
In-Depth Insights (Number of Original Perspectives)	2 items	7 items	Multi-agent systems integrate the three-tiered outcomes of "breadth-first search, industry knowledge, and expert insights" for deep analysis, whereas single-agent systems can only provide superficial summaries.
End-to-end processing time (minutes, including manual review)	9 min	28 min	Multi-agent systems execute workflows through division of labor and collaboration modeled after human experts; single-agent systems generate workflows sequentially.

Figure 7. Shipping report expert agent (multi-agent system).

红海危机对全球航运格局的影响研究报告
（2025 年 8 月版）

目录
一、执行摘要
二、事件回顾与风险地图
三、短期影响（2023Q4–2024Q4）
3.1 航线结构突变
3.2 运力与运价冲击
3.3 港口与内陆瓶颈
3.4 供应链外溢效应
四、中期影响（2025–2027）
4.1 运力再平衡与价格战风险
4.2 联盟重组与竞争格局
4.3 贸易路线多元化
4.4 ESG 与监管新变量
五、情景模拟与财务测算
六、政策与企业应对建议
七、附录（数据表、图表索引）

一、执行摘要
2023 年 11 月–2025 年 5 月，胡塞武装对曼德海峡—苏伊士轴线的袭击造成全球 30% 集装箱量与 12% 总贸易量改道。短期看：
• 亚欧航线往返航距增加 8 500–10 000 海里，有效运力被抽走 15–20%；
• 即期运价最高上涨 260%（欧线），全球平均上涨 4 倍；
• 鹿特丹、新加坡等枢纽出现"空箱短缺+船舶集中到港"双重拥堵。

中期看：
• 2025–2026 年将有 350 万 TEU 新船交付，叠加复航后 6% 运力一次性释放，供需缺口或扩大至 7–8%，运价面临 60–70% 暴跌风险；
• 班轮公司由"航线竞争"转向"风险—效率—技术"三维竞争，联盟边界开始松动；
• 印度–中东–欧洲经济走廊（IMEC）、墨西哥近岸等替代通道获得资本加速注入。

二、事件回顾与风险地图
• 时间轴：2023–11 胡塞首次导弹攻击 → 2024–02 三大班轮公司全部停航苏伊士 → 2025–03 加沙停火协议 → 2025–05 胡塞威胁重启攻击。
• 风险区：曼德海峡 50 海里内为高风险，战争险保费由 0.07% 升至 2%。

三、短期影响（2023Q4–2024Q4）
3.1 航线结构突变
— 苏伊士通行量–55%，好望角航线运量 +130%。
— LNG、LPG 船几乎 100% 改道，导致欧洲冬季能源到岸延迟 12–18 天。

3.2 运力与运价冲击
— 全球有效运力损失 15–20%（以 TEU-英里计）。
— 欧线即期运价 2024–07 攀至 8 400 USD/FEU，仍比危机前高 4 倍。

3.3 港口与内陆瓶颈
— 2024Q2 鹿特丹平均等泊时间 72h→118h；新加坡空箱缺口 15%。
— 沙特陆海联运（吉达–达曼）运量环比 +200%，成为中东区域应急通道。

3.4 供应链外溢效应
— 欧洲汽车厂 VMI 库存由 3 天升至 12 天，仓储成本 +9%。
— 全球平均航程延长导致额外 CO_2 排放 1.3 亿吨，相当于德国全年 14%。

四、中期影响（2025–2027）
4.1 运力再平衡与价格战风险
基准情景（红海 2025Q4 全面复航）：
• 6% 全球运力立即释放，叠加 150 万 TEU 新船，2026 年供需差–7.2%；
• 欧线即期运价或回落至 1 500 USD/TEU，接近现金成本。

悲观情景（冲突反复）：
• 运力仍被锁死 10–12%，高运价再延续 18–24 个月，班轮公司现金流累计增厚 320 亿美元，但加剧全球通胀 0.3–0.4 个百分点。

4.2 联盟重组与竞争格局
• 双子星、海洋、Premier 三大联盟 2025 年起互换舱位比例由 22% 降至 12%，跨联盟舱位共享减少 → 价格战烈度上升。
• 马士基斥资 20 亿美元升级数字化调度，MSC 押注绿色甲醇船舶，技术差距拉大。

4.3 贸易路线多元化
• IMEC 铁路–海铁联运预计 2027 年可承担亚欧 4–6% 箱量，运输时间较传统苏伊士路线缩短 40%。
• 墨西哥、土耳其近岸产能 2025–2027 年新增 120 万 TEU 生成量，分流远东–北美东岸部分货物。

4.4 ESG 与监管新变量
• 欧盟 Fit-for-55 将 2026 起征收航运碳排放交易附加费，绕行好望角导致单位 CO_2 增排 35%，或触发额外 20–30 USD/TEU 碳成本。
• 美国《2025 海运改革法》草案要求班轮公司公开停航/加价理由，对联盟运力管控形成法律掣肘。

五、情景模拟与财务测算
（见附录表 A1–A3）
• 情景一：复航+运力过剩，2026 年行业 EBIT Margin 由 2024 年的 18% 降至 3%。
• 情景二：冲突长期化，2026 EBIT Margin 维持 14%，但 CAPEX 增加 25% 用于安全与替代燃料。

六、政策与企业应对建议
对政府/国际组织

1. 建立"红海安全走廊"多边机制，引入中国、海合会联合护航，降低保费与风险溢价。
2. 将 IMEC、北极航线等纳入 G7/G20 基础设施合作框架，给予出口信贷支持。

对班轮公司

1. 制定"动态运力池"——通过数字化平台在 7 天内可调配 5% 全球运力，对冲复航后的即期运力冲击。
2. 与货主签署"弹性运价协议"：当苏伊士复航后 60 天内，运价跌幅超过 40% 时触发运费分成机制，锁定双方现金流。

对货主/制造业

1. 建立"30% 货量走替代通道"的硬性 KPI，降低单一通道依赖。
2. 采用"提前 45 天锁舱 + 选择性空运"混合策略，应对 JIT 失效。

七、附录
• 表 A1：主要航线绕行后里程与耗时对比
• 表 A2：2024–2027 全球集装箱运力交付与拆解预测
• 图 B1：欧线即期运价 2022–2025 实际与情景走势
• 参考文献与数据源列表（略）

Figure 8. Single agent based on RAG.

Based on experimental data, as shown in Figure 7 and Figure 8: single-agent systems exhibit limited knowledge scope, simple structures, frequent hallucinations, and rapid generation speeds. In contrast, multi-agent systems demonstrate comprehensive superiority across three dimensions: factual accuracy, structural integrity, and depth of insight. This advantage stems from decomposing shipping reports into a highly specialized yet real-time collaborative micro-production line, yielding significantly better results.

5. Conclusions and Outlook

In this paper, we introduce an expert agent for generating shipping research reports. This multi-agent system-based agent is specialized for the shipping industry, designed to produce shipping research analysis reports driven primarily by individual knowledge bases, shipping expert knowledge bases, and networked information. By establishing a hierarchical agent architecture for management and execution layers, and through the collaborative work between the Master Agent, Planner Agent, Human Feedback Agent, Background Agent, Research Team Agent (including Ship Knowledge Agent, Researcher Agent, Coder Agent), and Writer Agent. Crucially, the implementation of the four-stage content generation framework designed in this study enables the agent to significantly outperform human experts in completing complex, multi-step reasoning tasks within shipping report generation scenarios. Compared to traditional RAG systems relying on single-agent generation, it substantially enhances the accuracy of critical data and the overall quality of shipping report content. In the applied case study analyzing the Red Sea crisis's impact on global shipping patterns, the Shipping Research Report Expert AI demonstrated robust practical potential. It effectively handled large-scale data analysis and the organization of vast shipping knowledge, delivering accurate, comprehensive, and concise research analysis reports. However, the generation of this research report relies on a shipping expert knowledge base and a four-stage content generation mechanism. Despite our commitment to rigor, certain limitations remain. First, data and knowledge within the shipping industry, particularly regarding policies and regulations, may contain gaps due to the influence of international circumstances. Second, conclusions derived from the model's knowledge-based reasoning may be subject to hallucination risks, although we have incorporated manual review mechanisms to mitigate such errors. Additionally, the rapid evolution of the shipping industry introduces the possibility of domain concept drift.

Future work will focus not only on continuously updating knowledge and data to ensure the timeliness of conclusions, but also on introducing a Multi-Agent Reinforcement Learning (MARL) framework [10]. This will enable the system to continuously self-optimize its analytical strategies and implementation paths through user feedback and the evolution of real market data. All AI agents will share a global reward, with penalty terms introduced to prevent redundant, erroneous, redundancy, and verbosity. This facilitates the collaborative evolution of the shipping report expert agent from "local optimization" to "global optimization". Additionally, efforts will concentrate on agent context engineering to establish a unified memory hub. This addresses issues such as context loss, information silos and resource wastage, difficulties in information integration, and decision-making paralysis. It allows multiple agents to collaborate based on a shared information foundation, synthesizing consistent conclusions.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

[1] Dorri, A., Kanhere, S.S. and Jurdak, R. (2018) Multi-Agent Systems: A Survey. *IEEE Access*, **6**, 28573-28593. https://doi.org/10.1109/access.2018.2831228

[2] Shehory, O.M. (1998) Architectural Properties of Multi-Agent Systems. Carnegie Mellon University.

[3] Xu, W., Mao, Y., Zhang, X., Zhang, C., Dong, X., Zhang, M. and Gao, Y. (2025) Dagent: A Relational Database-Driven Data Analysis Report Generation Agent. http://dx.doi.org/10.48550/arXiv.2503.13269

[4] Xiao, Z., Fu, X., Zhang, L., Ponnambalam, L. and Goh, R.S.M. (2017) Data-Driven Multi-Agent System for Maritime Traffic Safety Management. 2017 *IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Yokohama, 1-6 October 2017, 1-6. http://dx.doi.org/10.1109/ITSC.2017.8317707

[5] Li, B., Chen, L., He, Z., Liu, J. Li, H., Wang, N. and Tan, X. (2025) Construction of the Knowledge Graph for Undergraduate Transportation Engineering Major at Chang'an University. *Creative Education Studies*, **13**, 427-437.

[6] Li, Y., Liu, X., Wang, Z., Mei, Q., Xie, W., Yang, Y., *et al.* (2024) Construction of a Large-Scale Maritime Element Semantic Schema Based on Knowledge Graph Models for Unmanned Automated Decision-Making. *Frontiers in Marine Science*, **11**, Article 1390931. https://doi.org/10.3389/fmars.2024.1390931

[7] Li, Y.C., Cai, H.Y., Kong, R., Chen, X.R., *et al.* (2025) Towards AI Search Paradigm. http://dx.doi.org/10.48550/arXiv.2506.17188

[8] Saad, A. (2025) CH-MARL: Constrained Hierarchical Multiagent Reinforcement Learning for Sustainable Maritime Logistics. http://dx.doi.org/10.48550/arXiv.2502.02060

[9] Liu, J., Yi, H., Bi, S., Feng, J. and Qi, G. (2025) From Superficial to Deep: Integrating External Knowledge for Follow-Up Question Generation Using Knowledge Graph and LLM. *Proceedings of the 31st International Conference on Computational Linguistics*, Abu Dhabi, 19-24 January 2025, 828-840.

[10] Wang, Y. and Zhao, Y. (2025) Multiple Ships Cooperative Navigation and Collision Avoidance Using Multi-Agent Reinforcement Learning with Communication. *Ocean Engineering*, **320**, Article 120244. https://doi.org/10.1016/j.oceaneng.2024.120244