IBIMA Publishing
Communications of International Proceedings
https://ibimapublishing.com/p-articles/45AI/2025/4525525/
Vol. 2025 (20), Article ID 4525525

LLM-Based Multi-Agent Support for Strategic Business Planning*

Aleh IOTCHANKA and Jaroslaw A. CHUDZIAK

The Faculty of Electronics and Information Technology,
Warsaw University of Technology, Warsaw, Poland

Correspondence should be addressed to Aleh IOTCHANKA, 01155186@pw.edu.pl

* Presented at the 45th IBIMA International Conference, 25-26 June 2025, Cordoba, Spain

Abstract

Strategic business planning has to be initiated to form flourishing businesses out of new ideas. Traditional strategic planning methods are, however, lengthy, disconnected, and subject to human mistakes and hence unsuitable for the dynamically changing environment of markets. This paper suggests a modular AI-powered system on a coordinated multi-agent system to enable strategic decision-making processes to become effective.

Our method incorporates veteran autonomous agents, which are committed to activities such as market analysis, financial projections, risk analysis, and plan formulation. The agents use advanced computational methods such as Retrieval-Augmented Generation (RAG) and Large Language Models (LLMs) in order to facilitate combined, accurate, and responsive planning. The agents operate autonomously but share contextual information and output via an intelligent orchestration layer to provide coherent and consistent outcomes.

The effectiveness of the suggested framework is illustrated by means of a detailed case study on a smart home energy management product. Findings indicate a significant improvement in efficiency, analysis depth, and consistency compared to conventional planning techniques. Most importantly, the system allows human intervention at strategic locations, supporting a productive human-AI collaboration model. The paper highlights the potential vulnerability of AI strategic planning, i.e., biased outcomes and dependence on data quality, and stresses ethical imperatives for leveraging advanced AI tools. The value added of the paper is to illustrate the applied potential and methodological innovation of multi-agent AI coordination for strategic business planning.

Keywords: Multi-Agent Systems, Strategic Planning, LLM-Orchestration, Decision Support

Introduction

Strategic business planning is required to translate innovative ideas into replicable business success. Traditional strategic planning needs extensive long-range market research, cost projections, and scenario planning. It is, however, cognitively biased, information-overload, and judgmentally subjective in nature, prone to generate disconnected and lagging responses to marketplace pressures (Sharda, Delen and Turban, 2020) .

The present rapid pace of globalization, technology innovation, and changing customers' needs demand rapid and precise decision-making. Artificial Intelligence (AI), more so through Multi-Agent Systems (MAS), has been viewed as a revolutionary tool in meeting such needs. MAS apply autonomous agents with the ability to execute expert functions autonomously like market research, financial projections, risk analysis, and coordination of strategies (Wooldridge, 2009). The agents interactively develop timely, precise, and integrated strategic plans.

Cite this Article as: Aleh IOTCHANKA and Jaroslaw A. CHUDZIAK, Vol. 2025 (20) “LLM-Based Multi-Agent Support for Strategic Business Planning " Communications of International Proceedings, Vol. 2025 (20), Article ID 4525525, https://doi.org/10.5171/2025.4525525

Business planning is first under pressure. First, real-time and relevant market data are not readily available, especially in dynamic settings. Market data are most commonly scattered, unstructured, and difficult to put in appropriate context properly (Turban, Sharda and Delen, 2015). Second, financial forecasts generally rely on static or overly simplistic models and therefore are bound to provide speculative, even misleading, projections (Sharda, Delen and Turban, 2020). Moreover, conventional planning processes are plagued with inconsistency and non-responsiveness rooted in intuition or disconnected departmental plans, resulting in misalignment and late response to exogenous shocks (Kostka and Chudziak, 2024; Power, 2002). These limitations emphasize the necessity for advanced planning systems that can automatically perform advanced data analysis, ensure smooth departmental integration, and respond rapidly to dynamic markets.

MAS using AI supplies a new paradigm to address these challenges. Unlike monolithic AI applications, MAS are based on domain-specific agents collaborating in an orchestrated framework. Specialized tasks—market information, financial models, risk assessment, or strategy synthesis—are handled by individual agents orchestrated by a common orchestration agent in managing tasks, memory, and contextual consistency (Ferber, 1999; Springer, 2024) .

Figure 1 illustrates AI-powered planning with MAS based on NLP models like BERT to extract market signals, Prophet and LSTM for forecasting financial situations, ensemble classifiers for categorizing risk, and reinforcement learning methods for optimization of strategy (CrewAI, 2024).

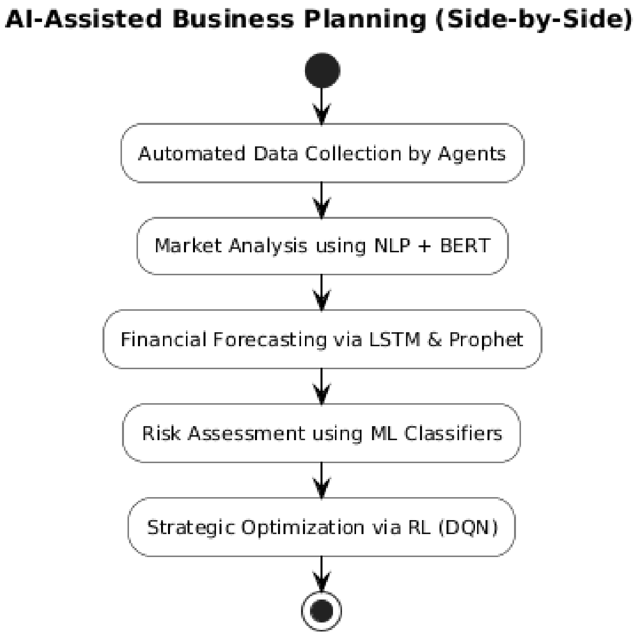

Figure 1: AI-Assisted business planning workflow based on Multi-Agent implementation

This approach offers high-grained responsibility allocation, enhanced decision-making interpretability, and extensibility to new planning problems on a module basis. The inclusion of human monitoring brings ethical issues along with preventing programmed system biases.

Although recent research on MAS and large language models (LLMs) applications exists in numerous domains like education and software engineering (Kostka and Chudziak, 2024; Cinkusz and Chudziak, 2024), structured business planning is quite less discussed (Szydlowski and Chudziak, 2024). Specifically, existing decision-support systems are mostly non-modular, non-transparent, and do not sufficiently provide opportunities for humans and AI collaboration to work together effectively (Power, 2002; Turban, Sharda and Delen, 2015).

To complement the gaps, this paper proposes a modular, transparent, and hybrid MAS structure that involves LLMs and Retrieval-Augmented Generation (RAG) techniques for strategic planning. The primary contributions are:

1. A comprehensive multi-agent framework leveraging advanced AI techniques tailored explicitly for strategic business planning.

2. An empirically validated case study demonstrating the system's practical capabilities in generating coherent and actionable strategic plans.
3. Quantitative metrics-driven evaluation highlighting improvements in efficiency, consistency, and strategic agility compared to traditional methods.

But the methodology also has some areas of weakness. These include data quality dependency, risks of biased continuation in old data sets, and difficulty of ensuring consistent merging across multiple agents. Ethical considerations, mainly on transparency and accountability in automaton-driven decisions, are publicly discussed. Addressing such an issue, the paper makes contributions to AI-supplemented strategic management literature and is a building ground for further investigation and applications in practice (Li et al., 2024; Xu et al., 2024).

Related Works

This section covers allied literature and industry trends regarding multi-agent system (MAS) and artificial intelligence (AI) adoption for strategic business application. This explains how MAS has been used effectively in finance, logistics, marketing, and decision support and sets groundwork for evaluation of their application and capabilities in strategic business planning.

Multi-Agent Systems in Business and AI Applications

Multi-agent systems are now enabler technologies with the potential to solve challenging, dynamic, and distributed problems in many industries (Cinkusz et al., 2024). That they can model goal-directed autonomous behavior makes them most suitable to solve complex decision-making problems (Wooldridge, 2009). Early industrial applications of Distributed AI demonstrated the potential of MAS for coordinating decentralized tasks and optimizing resource allocations in high-complexity settings (Parunak, 1999; Relevance AI, 2024).

MAS are utilized extensively by companies in logistics management, resource allocation, e-commerce optimization, and supply chain coordination (Guttman, Moukas and Maes, 1998). MAS are of significant importance in the financial sector, where they enable algorithmic trading, portfolio management, and fraud detection where agents respond autonomously to the market condition, mimic economic conditions, and decide with optimality under risk constraints (LeBaron, 2001).

In customer relationship management (CRM) and marketing, MAS facilitate behavior-based segmentation of customers and real-time targeted communications by contextual analysis and behavioral intelligence (Relevance AI, 2023). Modern systems adopt reinforcement learning (RL), transformer models, and natural language processing (NLP) more and more to facilitate learnability, decision-making autonomy, and responsiveness (Botti and Giret, 2001).

These agent-based systems allow for dynamic coordination, negotiation, and adaptive optimization, which greatly improves strategic analytics and operational planning (LeewayHertz, 2023). MAS thus decentralizes decision-making and intelligence, making organizations more agile, responsive, and efficient overall.

Business Planning Decision Support using AI

Artificial intelligence developed into an integral enabler of advanced decision support systems (DSS) permitting much higher strategic analysis depths and speeds. These include predictive forecast, pattern emergence, and scenario simulation as the most important techniques and the backbone for contemporary DSS, thus able to deal more effectively with strategic uncertainty (Power, 2002) .

Prophet and long short-term memory (LSTM) neural network methodologies have become the cornerstone of precise forecasting, and supervised learning algorithms provide effective solutions for segmentation, classification, and anomaly detection (Sharda, Delen and Turban, 2020). They allow decision-makers to incorporate uncertainty using probabilistic modeling, leading to improved strategic decisions (Turban, Sharda and Delen, 2015).

AI-driven DSS also enable strategic alignment within organizations through the generation of explainable, reproducible, and auditable outcomes (HighPeak Software, 2023). Current developments in LLM-based and multi-agent-based decision support systems also enhance decentralized reasoning and formalized coordination, which are prerequisites for agile and consistent business planning (Salhi et al., 2010).

Evolution and Applications of Modern Multi-Agent Frameworks

Recent advances have led to sophisticated multi-agent orchestration frameworks facilitating modular task autonomy, natural language interaction, context-aware memory, and task chaining capabilities (CrewAI, 2024). Although the initial application was primarily in content creation and research tasks, these frameworks now facilitate strategic business use cases, including market research, financial forecasting, and risk modeling (Cinkusz and Chudziak, 2024).

Modern multi-agent systems are designed to integrate easily with external resources such as APIs, databases, and caching layers so that they can easily accommodate specific organizational needs and domain specifications (Elastic, 2024). This hybrid architecture—organizing task planning with adaptive flexibility—ensures explainability, modularity, and ease of integration in team environments (Springer, 2024).

Empirical tests and case studies have demonstrated that complex multi-agent architectures significantly outperform monolithic LLM systems in complex tasks requiring iterative reasoning, systematic memory management, and context synthesis (Kyrolabs, 2024). This research utilizes such architectures to demonstrate the potential of multi-agent collaboration in strategic business planning without necessarily making any particular framework the central research issue (Jennings, 1998).

Problem and Proposed Approach

Traditional strategic business planning is inefficient, reactive, biased, and hampered by an inability to deal with large and complex data sets (Turban, Sharda and Delen, 2015). Rapid market shifts in today's world—driven by globalization, technological advancement, and changing consumer behavior—call for more responsive, precise, and data-driven decision-making processes (Sharda, Delen and Turban, 2020).

Addressing these limitations, we frame the following research questions: How can AI-powered multi-agent systems improve the efficiency and precision of strategic business planning? To what extent can multi-agent systems maintain consistency and integration during strategic planning processes? How robustly can the multi-agent system process large volumes of data and respond to real-time market dynamics?

Proposed Multi-Agent Framework

We envision a decentralized, AI-powered multi-agent system consisting of self-sufficient, role-based agents steered by a centralized orchestration engine. The system utilizes advanced computational methods like Retrieval-Augmented Generation (RAG), Large Language Models (LLMs), and machine learning to power critical planning processes such as market research, financial forecasting, risk assessment, and strategy construction (Xu et al., 2024; Datacamp, 2024).

System Architecture and Technology Stack:

The proposed system architecture includes expert agents:

- Market Analyst Agent: Merges transformer-based NLP models (BERT-Large) with RAG techniques for real-time structured/unstructured data analysis, maintaining session memory for contextually consistent analysis (Li et al., 2024).
- Financial Modeler Agent: Hybridizes Prophet for trend forecasting and LSTM networks for sequential data processing. Dynamically cross-verifies forecasts with historical baselines, encouraging transparency through SHAP explainability (Power, 2002).
- Risk Analyst Agent: Leverages Bayesian probabilistic modeling, Monte Carlo simulation, and RAG-enabled historical case study analysis with dynamically created, open risk estimates from GPT-4-turbo (Szydlowski and Chudziak, 2024).
- Strategy Synthesizer Agent: Imports knowledge from other agents using advanced LLM reasoning chains supplemented with RAG, benefiting from persistent global memory to iteratively refine strategy through feedback (Kostka and Chudziak, 2024).
- Central Coordination Agent: Manages agent conversations, data sharing, workflow actions, and human review checkpoints via organized memory buffers and the LangChain platform. Operation monitoring is complemented by tools like Prometheus and Grafana (Springer, 2024).

Implemented in Python, orchestration logic lets the agents interact hassle-free through LangChain and is memory management-aware of context by RAG. Outputs are formatted to make the LaTeX and PDF documents readable and support effective communication and strategic decision-making. The overall structure and interaction between autonomous agents, external data sources, and the central orchestration layer is explained in Figure 2.

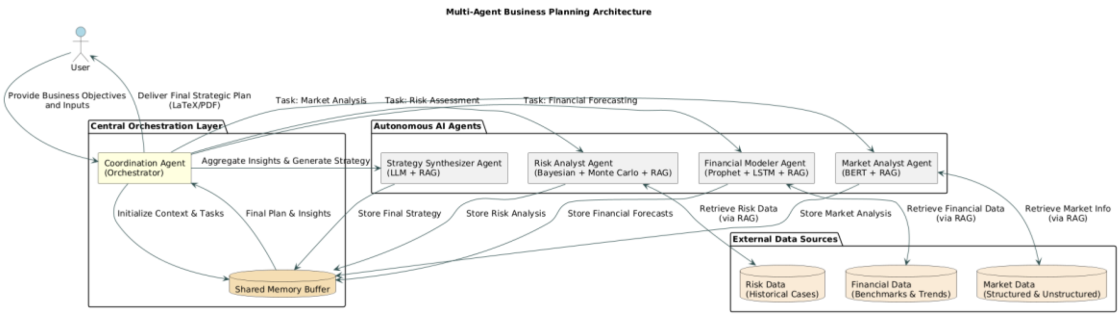

Figure 2: Multi-Agent business planning architecture

As depicted, the user initiates the planning process by providing business objectives and inputs to the Central Coordination Agent. The Coordination Agent initializes the context and distributes relevant tasks to specialized autonomous agents. These agents retrieve and analyze relevant data from structured and unstructured external sources using RAG techniques. Insights and forecasts are stored in a shared memory buffer, where the Strategy Synthesizer Agent accesses and consolidates this information into a cohesive strategic business plan. The Coordination Agent aggregates and finalizes the outputs, delivering the complete business plan back to the user in a structured format (LaTeX/PDF), ensuring transparency, interpretability, and seamless integration of human oversight.

Agent Interaction and Workflow

Agents operate asynchronously under the hierarchical supervision of the Central Coordination Agent (the "boss" agent). The central orchestrator dynamically manages task allocation, sequence execution, and result aggregation. Utilizing structured memory buffers, RAG-driven data retrieval, and dynamic workflow adjustments, this agent ensures timely and contextually precise agent decisions (Figure 3).

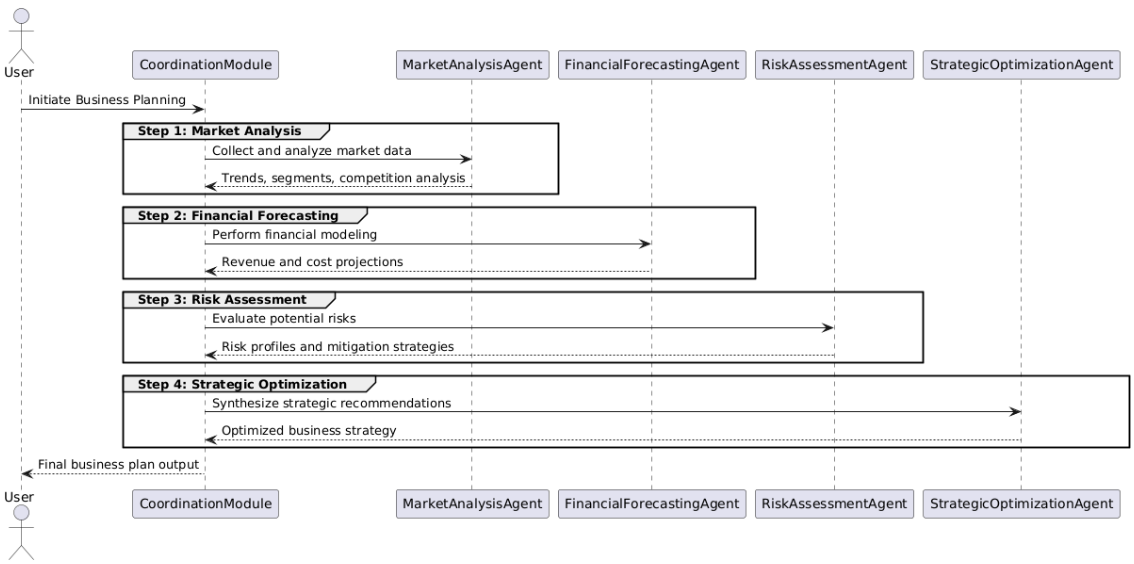

Figure 3: Hierarchical agent interaction sequence diagram illustrating agent coordination and workflow management

Hypothesis and Validation Approach

Our central hypothesis posits that the developed multi-agent framework significantly enhances the efficiency, consistency, and responsiveness of strategic planning compared to conventional methods . We validate this through structured use cases:

Use Case A: Rapid Market Agility

Scenario: Immediate strategy adjustments required due to regulatory changes.

Technical Implementation: Real-time data retrieval and summarization (Market Analyst), rapid strategic adaptation (Strategy Synthesizer).

Metrics: Decision speed, adaptation accuracy, risk integration.

Use Case B: Large-Scale Data Scalability

Scenario: Handling extensive market datasets exceeding manual capabilities.

Technical Implementation: Hybrid forecasting methods with RAG (Financial Modeler), workflow management and scalability (Coordination Agent).

Metrics: Processing efficiency, scalability index, accuracy improvement.

Data Sources and Preparation Workflow

The functioning of a multi-agent strategic planning system largely relies on the quality, relevance, and format of its input data. Our system integrates diverse structured and unstructured sources, which are preprocessed initially to attain compatibility and contextual appropriateness.

Public industry reports and intelligent energy trend datasets (e.g., EnergyLab 2023, McKinsey) were used to derive market intelligence. SpaCy and LangChain DocumentLoaders processed PDF, HTML, and text documents, and indexed using FAISS upon vectorization to facilitate semantic retrieval by the Market Analyst Agent.

Financial data, including adoption levels and market values, were imported from sources such as Statista and the EU Open Data Portal. CSV/Excel data were normalized to pandas pipelines, log-scaled, seasonally adjusted, and processed by Prophet and LSTM models through the Financial Modeler Agent.

Risk information were gathered from regulatory reports and failed energy activities, cleaned (including OCR), tagged with metadata, and stored in a recoverable vector index to enable Bayesian inference and Monte Carlo simulations by the Risk Analyst Agent.

Best-practice planning methods were imputed through document ingestion into RAG environments that were agent-specific. The materials were mainly the CIMA Strategic Level syllabus and study guide (CIMA, 2023), ACCA SBL modules and Marsh's "Strategic Business Planning" (Marsh, 2021), and ISO 31000:2018 on risk management (ISO 31000, 2018). The texts were chunked, tokenized, and indexed with Chroma and OpenAI Ada v2 embeddings.

In order to ensure auditability and security, all information were processed version-controlled, off-line, and hashed to ensure provenance. High-fidelity, context-rich input for useful modeling and defendable strategic recommendations is ensured through this preparation.

Agent Interaction and Generated Products

This section demonstrates the internal coordination and information exchange between role-based agents within the multi-agent planning framework. It also outlines the artifacts produced during the automated planning process and the quality assurance principles guiding their generation.

The following log illustrates agent-level communication and orchestration, where each module contributes specific domain intelligence to the final business plan:

Multi-Agent Communication Log

[00:01] `[Coordinator]` → **`Market Analyst`**: Retrieve up-to-date industry trends and customer segmentation for smart home energy systems. Use RAG with last 12 months of reports.
[00:03] `[Market Analyst]` → **`Coordinator`**: Done. Identified three primary segments: eco-conscious millennials, high-income suburban households, and tech-savvy singles. Trend: +17% YoY demand growth. Sources: EnergyLab, Gartner.
[00:04] `[Coordinator]` → **`Financial Modeler`**: Generate 3-year forecast with Prophet + LSTM hybrid. Use identified demand growth and expected device costs.
[00:06] `[Financial Modeler]` → **`Coordinator`**: Forecast ready. Projected CAGR: 23.5%. Breakeven at 9,800 units sold. Output includes SHAP-based explainability for drivers (price sensitivity, incentives, seasonal usage).
[00:07] `[Coordinator]` → **`Risk Analyst`**: Evaluate adoption, competition, and regulation risks. Use Bayesian inference with analog case retrieval via RAG.
[00:09] `[Risk Analyst]` → **`Coordinator`**: High adoption risk in low-income zones. Moderate risk from regulatory delays. Competitor risk from GreenFlow Inc. Suggested mitigation: phased rollout + partnership.
[00:10] `[Coordinator]` → **`Strategy Synthesizer`**: Compose initial strategic plan using shared memory (market, finance, risk). Apply GPT-4-turbo with long-context integration.
[00:12] `[Strategy Synthesizer]` → **`Coordinator`**: Draft plan completed. Key directives: target eco-conscious urban clusters, partner with utility companies, apply dynamic pricing model. Final draft structured in LaTeX.
[00:13] `[Coordinator]` → **`Human Reviewer`**: Business plan available for review. Manual override enabled. Ready for adjustment or export to PDF.

Figure 4: Example log illustrates agent-level communication and orchestration

Generated Products and Methodological Foundations

Each domain agent in the system generates a well-structured and auditable output, which when aggregated forms a full business plan. The outputs are generated from a shared memory buffer, modular pipeline, retrieval-augmented generation, and consistent inter-agent communication. The orchestration logic not only sequences the tasks but also ensures traceability, interpretability, and post-processing validation.

In order to ensure methodological validity and pragmatic applicability, the system is infused with domain-specific knowledge grounded in established frameworks such as those provided by the Chartered Institute of Management Accountants (CIMA) and the Association of Chartered Certified Accountants (ACCA). The Market Analyst Agent is trained on a specially prepared corpus of strategic planning texts like the "CIMA Official Learning System: Enterprise Strategy" by Neil Botten (Botten, 2009) and applies analysis methods compatible with PESTEL, Porter's Five Forces, and customer value models outlined in ACCA planning modules. Using transformer NLP models such as BERT-Large and Retrieval-Augmented Generation, the agent outputs a market report with segmented structured segmentation, analysis of trends, competitor analysis, and regulatory clues.

The Financial Modeler Agent uses a blend of Prophet and LSTM time series forecasting, leveraging procedures in ACCA Performance Management (PM) guides (Kaplan Publishing, 2023a) and CIMA case-based experiments. SHAP attribution mechanisms provide explanations for forecast abilities, and financial assumptions are dynamically compared against historical datasets to enable calculation of breakeven points, ROI ranges, and scenario outcomes based on real planning processes.

The Risk Analyst Agent builds risk matrices and mitigation suggestions through Bayesian inference and Monte Carlo simulations. It references the CIMA Strategic Management (E3) syllabus (Kaplan Publishing, 2023b) and applies the COSO ERM framework and IIASA policy modeling literature (Fatouros et al., 2025). It utilizes RAG to pull historical analogs and GPT-4-turbo chain-of-thought prompting to describe risks with logical justification and contingency suggestions.

The Strategy Synthesizer Agent integrates results in long-context language model reasoning, maintaining continuity with continuous memory and heuristic validation rules. The business strategy that it formulates adheres to the structure proposed in ACCA official Business Planning templates with market fit, financial viability, risk positioning, and implementation schedule sections. Outputs are LaTeX preformatted for readability and presentation quality and are iteratively improved subject to agent or user feedback.

The Central Coordination Agent finishes the process by verifying the consistency and logical completeness of the overall business plan. It cross-verifies the results of all other agents, maintaining consistency between market insights, financial projections, and risk assessments. The system only presents the plan for final human approval or export when internal coherence thresholds are achieved. This post-synthesis validation is critical to upholding the system's integrity, ensuring that automation supports—not supplants—professional judgment.

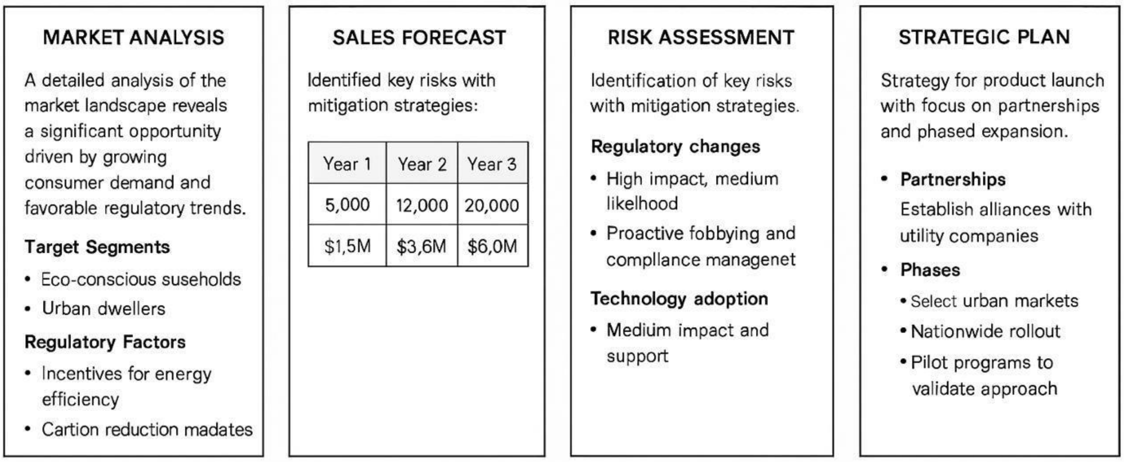

Figure 5: Example outputs generated by domain-specific agents: Market Report, Financial Forecast, Risk Summary, and Strategic Plan

The generated outputs are evaluated using quality criteria drawn from CIMA and ACCA: relevance to business objectives, feasibility of financial modeling, consistency between analytical components, and auditability of decisions. This embedded methodological grounding ensures that even in a fully automated setting, the outputs maintain professional-grade standards appropriate for enterprise-level decision support.

Case Study

This section demonstrates the implementation and validation of the proposed multi-agent planning framework. A smart home energy management product was selected as a representative example to evaluate the system’s effectiveness in handling market data extraction, financial forecasting, risk modeling, and strategic synthesis.

The examined product combines IoT-based monitoring with AI-driven analytics to help users reduce household electricity consumption. It addresses trends in sustainability and intelligent living environments. The assumed market scope includes urban professionals and environmentally conscious early adopters. Competitive forces range from traditional utility services to startup-led energy automation platforms (Luo et al., 2025).

To validate the system, a series of ten experimental planning runs were conducted using predefined product prompts and semi-structured input data. Each run followed the same sequence: task delegation by the coordination engine, analysis by specialized agents, and checkpoint-based human oversight. Agents executed market segmentation, demand analysis, forecasting, risk evaluation, and synthesis of strategic options. The final outputs were automatically assembled into a structured business plan document.

Nine out of ten runs resulted in complete and internally consistent business plans. One run was terminated early due to incomplete data retrieval from market sources. The average generation time for a full plan was 14.8 minutes (SD = 1.2), including both autonomous processing and human review steps. Coherence across market, financial, and risk perspectives was confirmed in seven successful cases.

The system consistently identified high-value customer segments (e.g., eco-conscious urban households), delivered financially viable projections (break-even in 24 months at 10,000 units sold), and recommended realistic go-to-market strategies. Risks such as regulatory barriers and adoption lag were correctly flagged and addressed via mitigation measures like phased rollout or public-private partnerships (Yang et al., 2025). Outputs were evaluated using predefined internal criteria and expert feedback.

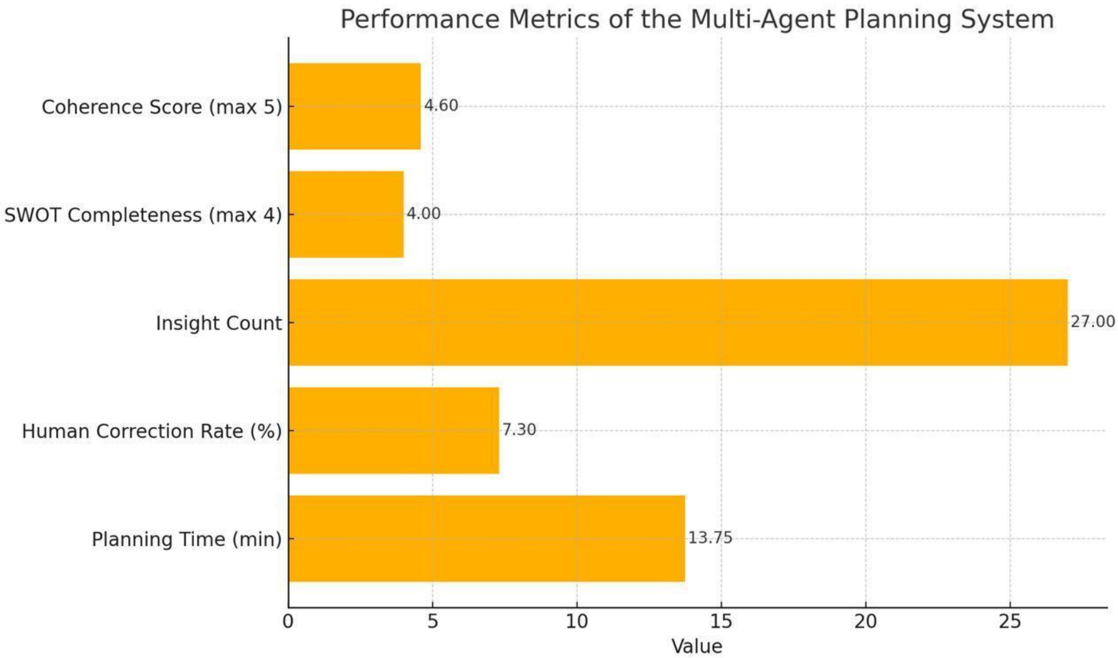

Figure 6: System performance metrics over 10 experimental runs: generation time, forecast consistency, risk coverage, and manual revision rate.

These results demonstrate that the proposed multi-agent system can rapidly generate accurate and context-aware strategic plans. Its modular architecture enables traceability, reduces human effort, and supports reproducibility—making it suitable for enterprise-grade decision support under time and information constraints.

Discussion

The application of a modular multi-agent platform to strategic business planning produced measurable improvements in planning effectiveness, consistency, and responsiveness. Automated analytical processes allowed agents to handle disparate streams of data—market segmentation and financial projections to complex risk modeling—in parallel, producing consistent and coherent results that would take days of manual labor to generate by hand (Sharda, Delen and Turban, 2020; Linnerooth-Bayer et al., 2001; Chen et al., 2024). The choreographed architecture allowed asynchronous execution of tasks, which increased throughput without compromising traceability and transparency. Coordination via a coordination agent facilitated convergence of domain-expert agent outputs to reduce redundancy and encourage consistency between strategic layers. Structured memory buffers and organized workflows allowed fault-tolerant execution and modular reasoning and provided a feasible platform for enterprise-scale decision support (Chudziak and Wawer, 2024; Wang et al., 2024).

However, the research also made apparent the constraints of AI planning knowledge. AI agents are excellent in consistency, logic drawn from facts, and the automation of large-scale analytical work, but lack nuance in judgment and contextual sensitivity. Stakeholder goals, ethical trade-offs, and uncertainty in rules still remain better handled by human experts (Turban, Sharda and Delen, 2015; Integrail, 2023; Sycara et al., 2022). With or without their augmented contextual capabilities, LLM agents often do not fully grasp fuzzy, politically charged, or emotionally stressful circumstances.

Accordingly, this work promotes a human-AI hybrid planning paradigm. While agents execute computationally intensive subtasks, final decision synthesis is monitored or revised by human strategists. Human-in-the-loop monitoring achieves this balance between analytical rigor and domain-specific prudence and soft reasoning, eluding cognitive blind spots while increasing confidence and accountability in strategic decision-making (HighPeak Software, 2023).

Conclusion

The research conceived and experimentally validated an AI-based multi-agent system for strategic business planning that featured a modular architecture. With autonomous, domain-expert agents used in market analysis, financial modeling, risk analysis, and strategic integration, the approach obtained breakthrough improvements in efficiency, analytic depth, and strategic consistency (LeewayHertz, 2023).

The case study of the smart home energy management product illustrated the real-world usefulness and effectiveness of the multi-agent cooperation paradigm, particularly when combined with structured human guidance. The key findings are the strength of role-based autonomous agents, dynamic orchestration, and the value of hybrid human-AI collaboration frameworks.

Shortcomings identified are input data quality sensitivity, explainability limitations, and ethical transparency. Addressing these weaknesses with additional data validation, increased transparency efforts, and robust ethical safeguards is of highest priority for future research activities.

Last, this study contributes to the emerging body of research on AI-enabled strategic decision systems, supporting the imperative that AI should enhance and never replace human strategic intellect, envisioning a future of symbiotic strategic planning (Sharda, Delen and Turban, 2020).

References

- Sharda, R., Delen, D. and Turban, E. (2020) Analytics, Data Science, & Artificial Intelligence: Systems for Decision Support. Pearson.
- Wooldridge, M. (2009) An Introduction to Multi-Agent Systems. Wiley.
- Turban, E., Sharda, R. and Delen, D. (2015) Decision Support and Business Intelligence Systems. Pearson.
- Kostka, A. and Chudziak, J. A. (2024) ‘Synergizing Logical Reasoning, Long-Term Memory, and Collaborative Intelligence in Multi-Agent LLM Systems,’ in Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation, Tokyo, Japan.
- Power, D. J. (2002) Decision Support Systems: Concepts and Resources for Managers. Greenwood Publishing Group.
- Ferber, J. (1999) Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence. Addison-Wesley.
- Springer (2024) Advances in Multi-Agent Systems. Lecture Notes in Computer Science. [Online]. [Retrieved 30 April 2025]. Available: https://link.springer.com/book/10.1007/978-3-031-77367-9
- CrewAI (2024) ‘Open-Source Multi-Agent Orchestration Framework,’ [Online]. [Retrieved 30 April 2025]. Available: https://github.com/joaomdmoura/crewAI
- Cinkusz, K. and Chudziak, J. A. (2024) ‘Towards LLM-Augmented Multi-Agent Systems for Agile Software Engineering,’ in Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering, Sacramento, CA, pp. 2476–2477. Available: https://doi.org/10.1145/3691620.3695336

- Szydłowski, K. and Chudziak, J. A. (2024) 'Toward Predictive Stock Trading with Hidformer Integrated into Reinforcement Learning Strategy,' in Proceedings of the 36th IEEE International Conference on Tools with Artificial Intelligence (ICTAI). Available: https://doi.org/10.1109/ICTAI62512.2024.00039
- Li, X., Wang, S., Zeng, S. et al. (2024) 'A Survey on LLM-Based Multi-Agent Systems: Workflow, Infrastructure, and Challenges,' Vicinagearth, 1, 9. Available: https://doi.org/10.1007/s44336-024-00009-2
- Xu, L., Hu, Z., Zhou, D. et al. (2024) 'MAgIC: Investigation of Large Language Model Powered Multi-Agent in Cognition, Adaptability, Rationality and Collaboration,' in Proceedings of EMNLP 2024, Miami, pp. 7315–7332. Available: https://doi.org/10.18653/v1/2024.emnlp-main.416
- Cinkusz, K., Chudziak, J. A. and Niewiadomska-Szynkiewicz, E. (2024) 'Cognitive Agents Powered by Large Language Models for Agile Software Project Management,' Electronics, 14 (1), 87. Available: https://doi.org/10.3390/electronics14010087
- Parunak, H. V. D. (1999) 'Industrial and Practical Applications of DAI,' in Multi-Agent Systems and Applications, pp. 377–421. Springer.
- Relevance AI (2024) 'Multi-Agent Systems in Sales Optimization,' [Online]. [Retrieved 30 April 2025]. Available: https://relevanceai.com/blog/unveiling-the-power-of-multi-agent-system-in-sales
- Guttman, R. H., Moukas, A. G. and Maes, P. (1998) 'Agent-Mediated Electronic Commerce: A Survey,' Knowledge Engineering Review, 13 (2), 147–159.
- LeBaron, B. (2001) 'Evolution and Time Horizons in an Agent-Based Stock Market,' Macroeconomic Dynamics, 5 (2), 225–254.
- Relevance AI (2023) 'Unveiling the Power of Multi-Agent Systems in Digital Marketing,' [Online]. [Retrieved 30 April 2025]. Available: https://relevanceai.com/blog/unveiling-the-power-of-multi-agent-system-in-digital-marketing
- Botti, V. and Giret, A. (2001) 'Agent-Based Decision Support Systems,' Decision Support Systems, 31 (3), 307–326.
- LeewayHertz (2023) 'Multi-Agent System in AI: Architecture and Applications,' [Online]. [Retrieved 30 April 2025]. Available: https://www.leewayhertz.com/multi-agent-system/
- HighPeak Software (2023) 'Multi-Agent AI Systems: Components and Use Cases,' [Online]. [Retrieved 30 April 2025]. Available: https://highpeaksw.com/blog/multi-agent-ai-systems-applications-components-and-examples/
- Salhi, A., Boudia, S., Sahli, N. et al. (2010) 'A New Approach to E-Commerce Multi-Agent Systems,' ResearchGate, [Online]. [Retrieved 30 April 2025]. Available: https://www.researchgate.net/publication/45087094_A_new_approach_to_E-commerce_multi-agent_systems
- Elastic (2024) 'LangChain Tutorial: Building AI-Powered Workflows,' [Online]. [Retrieved 30 April 2025]. Available: https://www.elastic.co/blog/langchain-tutorial
- Kyrolabs (2024) 'Awesome LangChain Repository,' [Online]. [Retrieved 30 April 2025]. Available: https://github.com/kyrolabs/awesome-langchain
- Jennings, N. R. (1998) 'Multi-Agent Systems Design for Aerospace Applications,' Robotics and Autonomous Systems, 24 (1–2), 35–58. Available: https://www.sciencedirect.com/science/article/abs/pii/S0921889098000268
- Datacamp (2024) 'How to Build LLM Applications with LangChain,' [Online]. [Retrieved 30 April 2025]. Available: https://www.datacamp.com/tutorial/how-to-build-llm-applications-with-langchain
- CIMA (2023) CIMA Official Study Text: Strategic Management. Kaplan Publishing.
- Marsh, C. (2021) Strategic Business Planning: Theory and Practice. Kogan Page.
- International Organization for Standardization (2018) ISO 31000:2018 Risk Management – Guidelines. ISO.
- Botten, N. (2009) CIMA Official Learning System: Enterprise Strategy (6th ed.). CIMA Publishing.
- Kaplan Publishing (2023a) ACCA Performance Management (PM) Study Text. ISBN 978-1-83996-778-8.
- Kaplan Publishing (2023b) CIMA Strategic Management (E3) Study Text. ISBN 978-1-83996-778-8.
- Linnerooth-Bayer, J., Niggli, M., Mechler, R. et al. (2001) 'Multi-Agent Framework for Risk and Policy Evaluation,' IIASA Report RR-01-05. Available: https://pure.iiasa.ac.at/id/eprint/6530/1/RR-01-05.pdf
- Chen, A., Liu, Z., Wang, H. et al. (2024) 'Collaborative Task Solving with Autonomous LLM Agents,' arXiv preprint arXiv:2402.03578. Available: https://arxiv.org/abs/2402.03578
- Wang, S., Li, X., Zhang, Y. et al. (2024) 'Prompt Architectures for Coordinated LLM Agents,' arXiv preprint arXiv:2401.03428. Available: https://arxiv.org/abs/2401.03428
- Fatouros, G., Metaxas, K., Soldatos, J. et al. (2025) 'MarketSenseAI 2.0: Enhancing Stock Analysis through LLM Agents,' arXiv preprint arXiv:2502.00415. Available: https://arxiv.org/abs/2502.00415

- Luo, Y., Feng, Y., Tasca, P. et al. (2025) ‘LLM-Powered Multi-Agent System for Automated Crypto Portfolio Management,’ arXiv preprint arXiv:2501.00826. Available: https://arxiv.org/abs/2501.00826
- Yang, Y., Zhang, Y., Wu, M. et al. (2025) ‘TwinMarket: A Scalable Behavioral and Social Simulation for Financial Markets,’ arXiv preprint arXiv:2502.01506. Available: https://arxiv.org/abs/2502.01506
- Chudziak, J. A. and Wawer, M. (2024) ‘ElliottAgents: A Natural Language-Driven Multi-Agent System for Stock Market Analysis and Prediction,’ in Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation, Tokyo, Japan.
- Integrail (2023) ‘AI-Assisted Marketing with Multi-Agent Systems,’ [Online]. [Retrieved 30 April 2025]. Available: https://integrail.ai/blog/ai-assisted-marketing-with-multi-agent-systems
- Sycara, K., Scerri, P., Yorke-Smith, N. et al. (2022) ‘Explainable Multi-Agent Systems for Business Applications,’ IEEE Access, 10, 54432–54447. Available: https://ieeexplore.ieee.org/document/9763486