A LLM - Based Agent Workflow for Financial Decision Support Using Financial Databases

Kejie Wang*

AIEN Institute, Shanghai Ocean University, Shanghai, China

*2392242@st.shou.edu.cn

Abstract

The rapid growth and increasing complexity of financial databases have made it challenging for users to efficiently retrieve and interpret relevant information for decision-making, particularly when queries are expressed in natural language. Traditional financial analysis processes often rely on human experts to translate user intentions into structured data queries, which introduces subjectivity, limited scalability, and operational inefficiencies. Existing automated financial systems, on the other hand, lack sufficient semantic understanding to effectively bridge natural language input and structured financial databases. This paper proposes an LLM-based agent workflow for financial decision support using financial databases. The proposed framework employs large language models to interpret unstructured user queries and transform them into structured prompts aligned with predefined financial data categories. By organizing financial databases into a hierarchical dictionary of semantic labels, the workflow enables accurate data mapping, automated retrieval, and integrated analytical reasoning within a unified system architecture. The agent-based design allows different functional components to collaborate sequentially, supporting flexible and scalable financial information processing. The proposed workflow is presented as an applied conceptual framework, emphasizing system design and workflow integration rather than algorithmic optimization. An illustrative example is provided to demonstrate the workflow execution and evidence-grounded output generation.

Keywords:

Master of law, financial decision-making, financial database.

1. Introduction

The rapid expansion of financial databases has significantly transformed the way financial information is generated, stored, and accessed. Modern financial databases now contain large volumes of heterogeneous data, including market prices, corporate fundamentals, macroeconomic indicators, and sector-level statistics [1]. While these databases provide unprecedented opportunities for data-driven financial decision-making, their increasing scale and structural complexity have also created substantial barriers for effective information access, particularly for users without professional financial or technical backgrounds [2].

In practical decision-making scenarios, users often express information needs in natural language rather than in structured query formats [3]. Translating such natural language queries into precise data retrieval and analytical tasks typically requires extensive domain knowledge and familiarity with database structures. As a result, financial decision-making processes continue to rely heavily on human experts to interpret user intent, select relevant data sources, and synthesize analytical results [4]. Although this expert-driven approach can be effective, it introduces several limitations, including subjectivity, limited scalability, high operational costs, and inconsistent decision quality across different analysts [2]. These challenges highlight a growing gap between user-oriented decision needs and the technical complexity of financial databases.

Recent advances in large language models (LLMs) offer new opportunities for addressing this gap [6,7]. LLMs have shown strong capabilities in natural language understanding, contextual reasoning, and task decomposition, enabling them to interpret complex user queries and generate structured representations of intent. These capabilities suggest that LLMs can serve as effective interfaces between human users and structured data systems. However, directly applying LLMs to financial databases poses several challenges. Financial data platforms are highly structured, rely on precise semantic labels, and often require strict alignment between user intent and predefined data categories. Without appropriate system-level design, LLM-generated outputs may remain disconnected from the underlying database architecture, limiting their practical usefulness for financial decision support [8].

Agent-based workflows provide a complementary perspective for addressing these challenges [9]. By decomposing complex tasks into multiple functional components, agent-based systems enable sequential processing, role specialization, and modular system design. Within a financial decision-support context, different agents can be responsible for intent interpretation, data mapping, retrieval coordination, and analytical reasoning. When combined with LLMs, agent-based workflows offer a structured way to operationalize natural language understanding within complex data environments. Rather than treating LLMs as standalone decision-makers, this approach embeds them as functional components within a broader workflow, emphasizing system logic, process integration, and interpretability.

In this study, we propose an LLM-based agent workflow for financial decision support using financial databases. The proposed framework is designed to transform unstructured natural language queries into structured prompts aligned with predefined financial data categories. A hierarchical financial data dictionary is introduced to organize database labels and semantic classes, enabling accurate intent mapping and automated data retrieval. The agent-based workflow coordinates multiple functional modules in a sequential manner, supporting flexible and scalable financial information processing within a unified system architecture.

This work is presented as an applied conceptual framework, with an emphasis on system design and workflow integration rather than algorithmic optimization or performance benchmark. An illustrative example is provided to demonstrate how the proposed workflow processes a natural language query, retrieves relevant financial data, and generates decision-support outputs. The contribution of this paper lies in offering a practical and extensible workflow that connects natural language understanding, agent-based automation, and financial database analysis, thereby supporting more accessible and consistent financial decision-making.

2. Literature Review

Research related to this study can be broadly categorized into three streams: (i) AI/NLP for financial decision support, (ii) agent-based and workflow-oriented systems, and (iii) financial information systems.

2.1. AI and Natural Language Processing in Financial Decision Support

AI has been widely used for financial decision-making tasks such as risk assessment, asset pricing, portfolio optimization, and market prediction [5,6]. With advances in NLP, textual sources (e.g., financial reports, news, and social media) have also been leveraged to support financial analysis [11,12]. More recently, LLMs have demonstrated strong capabilities in natural language understanding for tasks such as summarization, sentiment analysis, and question answering [6,7,11]; however, most studies remain task-level and provide limited guidance on integrating natural language understanding into end-to-end decision-support workflows over large-scale structured financial databases.

2.2. Agent-Based and Workflow-Oriented Intelligent Systems

Agent-based systems decompose complex tasks into multiple functional agents, enabling modularity, role specialization, and coordinated execution [9,13,14]. Workflow-oriented designs further emphasize orchestrating multi-stage processing (e.g., data collection, transformation, analysis, and output generation) to improve scalability and maintainability. Despite these advantages, existing work is often domain-specific and rarely focuses on operationalizing natural language understanding through agent workflows in structured financial data environments.

2.3. Financial Information Systems and Database-Driven Decision Making

Financial information systems support data-driven decisions by providing access to large-scale structured financial databases and decision-support functionalities [1,3,18]. Yet effective use typically requires knowledge of database schema, semantic labels, and analytical conventions, creating barriers for non-professional users [2,4]. While prior studies have explored interface and usability improvements, comprehensive system-level designs that align natural language user intent with predefined financial data categories and labels remain underdeveloped.

2.4. Research Gap and Motivation

Overall, the three research streams have largely evolved in parallel [6,10]. Existing studies seldom provide an end-to-end framework that systematically connects natural language input, agent-based workflow coordination, and structured financial databases, particularly with respect to semantic alignment between unstructured queries and predefined database labels [2]. This gap motivates the LLM-based agent workflow proposed in this study.

3. Problem Definition

The problem addressed in this study arises from the growing gap between user-oriented financial decision needs and the technical complexity of large-scale financial databases. Although modern financial databases provide extensive and high-quality structured data, effectively transforming user queries into actionable decision-support outputs remains a challenging task. This section formally defines the problem by analyzing the characteristics of user queries, the structure of financial databases, and the limitations of existing interaction mechanisms.

3.1. Characteristics of User-Oriented Financial Queries

In practical financial decision-making scenarios, users often express their information needs in natural language [17]. Such queries are typically informal, ambiguous, and context-dependent. Unlike structured database queries, natural language inputs may lack explicit references to specific data fields, financial instruments, or analytical methods. For example, a user may ask a question related to market trends, corporate performance, or investment risks without specifying the exact variables or datasets required for analysis.

From a system perspective, these queries present multiple challenges. First, user intent is often implicit and must be inferred from linguistic context rather than from predefined query structures. Second, a single query may involve multiple dimensions of financial analysis, such as time horizons, market segments, and performance indicators. Third, different users may express similar decision needs using diverse linguistic forms, making it difficult to rely on rule-based or template-driven query interpretation. These characteristics highlight the need for flexible semantic understanding mechanisms capable of translating unstructured input into structured analytical tasks.

3.2. Structural Complexity of Financial Databases

Financial databases are typically organized using highly structured schemas designed to ensure data accuracy, consistency, and efficient retrieval [1]. Data elements are categorized using predefined labels, hierarchical classifications, and domain-specific taxonomies. While such structures are essential for managing large-scale financial data, they also introduce barriers to intuitive user interaction.

Accessing relevant information from financial databases usually requires precise alignment between user intent and database labels. Users must understand how financial concepts are represented within the database, including naming conventions, categorical hierarchies, and data dependencies. This requirement creates a mismatch between user-centric decision thinking and system-centric data organization. As a result, effective database utilization often depends on financial professionals who possess both domain expertise and technical knowledge of database structures [2,4]. This reliance on expert intermediaries limits scalability and introduces subjectivity into the decision-making process.

3.3. Limitations of Existing Interaction and Automation Approaches

Existing approaches for interacting with financial databases can be broadly classified into manual expert-driven processes and automated or semi-automated systems [4]. Manual processes rely on financial analysts to interpret user queries, identify relevant datasets, and synthesize analytical results. Although effective in complex cases, this approach is time-consuming, costly, and difficult to standardize. Differences in analyst experience and judgment may also lead to inconsistent decision outcomes.

Automated systems aim to reduce human involvement by using predefined rules, query templates, or domain-specific interfaces [2]. While such systems improve efficiency for routine tasks, they often lack the flexibility required to handle diverse natural language queries. Rule-based automation struggles with linguistic ambiguity and fails to adapt to evolving user needs. More advanced AI-based tools address specific analytical tasks but are commonly designed as standalone modules rather than as integrated decision-support workflows. Consequently, these systems do not fully automate the end-to-end process from user query interpretation to database-driven decision output.

3.4. Problem Statement

Based on the above analysis, the core problem addressed in this study can be defined as follows: how to design a system-level workflow that effectively bridges unstructured natural language user queries and structured financial databases to support automated and consistent financial decision-making. Addressing this problem requires a mechanism that can (1) accurately interpret user intent expressed in natural language, (2) align inferred intent with predefined financial data categories and labels, and (3) coordinate data retrieval and analysis within a unified decision-support process.

Importantly, this problem is not solely a matter of improving individual algorithms or analytical models. Rather, it is a system design challenge that involves semantic alignment, workflow coordination, and functional integration across multiple processing stages. A successful solution must support flexibility in user interaction while maintaining compatibility with the structured nature of financial databases.

Motivated by this problem, the present study focuses on the design of an LLM-based agent workflow that integrates natural language understanding, data mapping, and decision-support processes. By emphasizing workflow logic and system architecture, the proposed approach aims to provide a scalable and interpretable solution for financial decision support using financial databases.

4. Framework Design

This section presents the design of the proposed LLM-based agent workflow for financial decision support using financial databases. The framework is developed from a system engineering perspective, emphasizing workflow coordination, semantic alignment, and functional modularity. Rather than optimizing individual algorithms, the design focuses on constructing a controllable and interpretable process that bridges natural language user queries and structured financial databases.

4.1. Design Objectives and System Principles

The primary objective of the proposed framework is to reduce the gap between user-oriented financial decision needs and the technical complexity of financial databases. In practical settings, users often seek decision-support insights without possessing detailed knowledge of database schema or analytical conventions. Therefore, the system is designed to support accessible and consistent decision-making while maintaining compatibility with existing financial data infrastructures.

To achieve this objective, the framework adopts an agent-based workflow architecture that avoids two common extremes. On one hand, fully manual expert-driven processes lack scalability and consistency. On the other hand, end-to-end black-box artificial intelligence systems may generate opaque outputs that are difficult to validate in financial contexts. The proposed approach positions itself between these extremes by embedding large language models within a structured workflow that constrains and guides system behavior.

The design of the framework follows several key principles. First, modularity ensures that each functional component has a clearly defined responsibility. Second, controllability is emphasized by constraining language model outputs through structured representations and predefined mappings. Third, interpretability is supported through explicit intermediate artifacts that allow system behavior to be traced and reviewed. Finally, compatibility with existing financial databases is maintained by decoupling semantic interpretation from database-specific implementations.

4.2. Overall Workflow Architecture

The proposed framework is organized as a sequential workflow consisting of five main stages: natural language input processing, intent interpretation and task structuring, semantic mapping via a financial data dictionary, database retrieval coordination, and analytical reasoning with output generation. Each stage produces structured intermediate outputs that serve as inputs to subsequent stages, forming an end-to-end decision-support process from user query to final response. as show in fig 1.

An agent-based design is adopted to implement this workflow. Each stage is handled by a dedicated agent responsible for a specific functional role. Agents operate sequentially rather than in parallel, reflecting the logical dependencies between tasks. By passing structured intermediate representations between agents, the system ensures that decisions made at earlier stages can be validated and constrained by later stages. This design improves transparency and reduces error propagation across the workflow.

Within this architecture, the large language model plays a central but constrained role. Rather than acting as an autonomous decision-maker, the LLM functions as a semantic processing component that supports intent interpretation, task formulation, and explanation generation. Its outputs are governed by predefined formats and workflow logic, ensuring alignment with system objectives and database constraints.

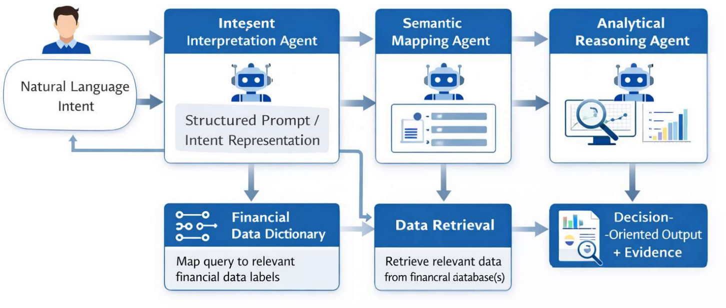

Fig. 1 Work flow chart

4.3. Natural Language Understanding and Intent Structuring

Natural language financial queries are inherently complex, often involving implicit assumptions, ambiguous terminology, and multiple analytical dimensions. Users may express decision needs using informal language without specifying exact data fields, time horizons, or evaluation criteria. This complexity presents a significant challenge for direct database interaction and motivates the need for an intermediate interpretation layer.

The intent interpretation agent is responsible for analyzing the semantic content of user queries and extracting key decision elements. These elements may include target entities, relevant financial metrics, temporal constraints, and the overall analytical intent, such as comparison, trend analysis, or risk assessment. Importantly, this agent does not directly interact with financial databases. Instead, it focuses on transforming unstructured input into a structured representation that captures inferred user intent.

Based on the extracted intent, the system generates a structured task representation that formalizes the analytical request. This representation follows predefined formatting rules and serves as a standardized description of the task to be performed. By introducing this intermediate structure, the framework reduces ambiguity and ensures that downstream agents operate on explicit and consistent task definitions rather than on raw natural language.

4.4. Semantic Mapping via Financial Data Dictionary

A core challenge in financial decision support lies in aligning user-oriented concepts with the structured labels used in financial databases. While users may refer to abstract notions such as performance or valuation, databases organize information through specific fields, identifiers, and hierarchical taxonomies. To address this semantic mismatch, the proposed framework introduces a financial data dictionary as an intermediary layer. as show in fig 2.

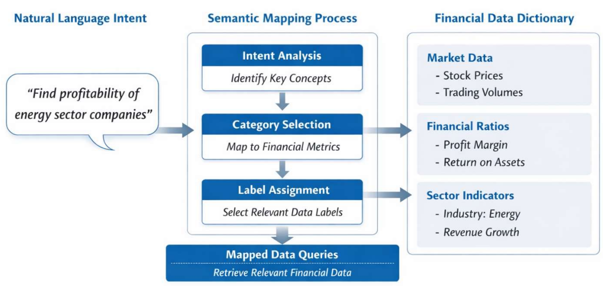

Fig. 2 Frame diagram

The financial data dictionary organizes database elements into hierarchical semantic categories that represent common financial concepts. High-level categories correspond to broad analytical domains, while lower-level labels map directly to specific database fields. This structure enables the system to interpret abstract decision intents in terms of concrete data requirements without exposing database complexity to users.

A dedicated mapping agent uses the financial data dictionary to align structured task representations with relevant data labels. By externalizing semantic alignment into a dictionary-based mechanism, the framework decouples intent interpretation from database-specific schema. This design improves maintainability and allows the system to adapt to changes in database structure by updating the dictionary rather than redesigning the entire workflow. This design is particularly important in large-scale financial databases where heterogeneous schema and evolving label structures significantly increase query formulation complexity.

4.5. Data Retrieval, Evidence Control, and Traceability

Once relevant data categories and labels have been identified, the retrieval agent coordinates database queries and data access. This agent is responsible for generating query requests based on mapped labels and retrieving the required data from financial databases. Its role is strictly limited to query execution and data aggregation, ensuring a clear separation between semantic interpretation and technical data access.

To support trustworthy decision-making, the framework incorporates mechanisms for evidence control and traceability. Each stage of the workflow produces structured artifacts that can be recorded and inspected, including intent representations, mapping results, and retrieved datasets. These artifacts enable the system to link analytical outputs to underlying data evidence and to support review or audit when necessary.

In addition, constraints are imposed on output generation to reduce the risk of unsupported statements. Analytical results and explanatory text are generated based on retrieved data and predefined analytical logic. By explicitly linking outputs to data sources, the framework enhances interpretability and supports responsible use in financial decision-support contexts.

4.6. Analytical Reasoning and Robustness Handling

The final stage of the workflow focuses on analytical reasoning and output generation. Based on retrieved financial data, the reasoning agent performs analysis aligned with the structured task definition. This may include summarizing trends, comparing entities, or organizing key indicators into decision-relevant narratives. The agent operates within predefined analytical boundaries and does not introduce external assumptions beyond available data.

Robustness is a critical consideration due to the inherent ambiguity of natural language queries. The framework incorporates mechanisms to handle common sources of uncertainty, such as ambiguous entity references, unspecified metrics, or unclear temporal constraints. When ambiguity is detected, the system may generate multiple candidate interpretations or apply conservative default assumptions that are explicitly reported in the output.

By explicitly managing uncertainty and limiting the scope of automated reasoning, the framework aims to prevent error amplification and unsupported conclusions. These robustness mechanisms ensure that the system provides reliable and transparent decision-support outputs, even when user input is incomplete or ambiguous.

5. Illustrative Example

To illustrate the operation of the proposed LLM-based agent workflow, this section presents a representative financial decision-support scenario. The example demonstrates how a natural language query is processed through the workflow stages, from intent interpretation to data retrieval and analytical output generation. The purpose of this illustrative example is not to evaluate system performance, but to clarify the internal logic and coordination of the framework in a realistic decision-making context. as show in fig3.

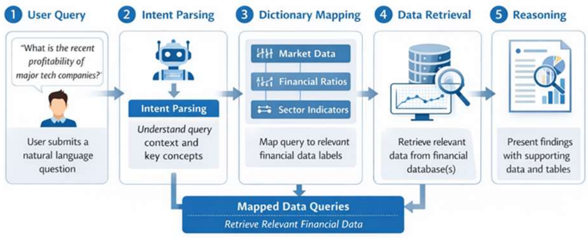

Fig. 3 Examples

5.1. Scenario Description and User Query

Consider a user seeking investment-related decision support under uncertain market conditions. The user does not possess professional knowledge of financial databases or analytical tools and expresses the decision need in natural language. A representative query may be formulated as follows:

"Given the current market environment, should exposure to technology-related equities be increased, and what evidence supports this decision?"

This query reflects common characteristics of real-world financial decision needs. It is high-level and decision-oriented, does not specify particular financial instruments or metrics, and implicitly requires comparative analysis, temporal context, and supporting evidence. Such a query cannot be directly translated into a structured database request without intermediate interpretation and semantic alignment.

5.2. Intent Interpretation and Task Structuring

Upon receiving the user query, the intent interpretation agent analyzes the semantic content to identify key decision elements. From the query, the agent infers the following components: the target asset category (technology-related equities), the decision objective (evaluating whether to increase exposure), and the need for supporting evidence. Temporal context is implicitly associated with recent or current market conditions, although no explicit time horizon is specified.

Based on this interpretation, the agent generates a structured task representation that formalizes the analytical request. This representation may include elements such as target sector, relevant performance indicators, comparative benchmarks, and an analytical intent categorized as evaluation and recommendation support. Importantly, the structured representation does not yet reference specific database fields or data sources. Instead, it serves as an intermediate abstraction that captures user intent in a form suitable for systematic processing.

The structured task representation is then passed to downstream agents. By explicitly documenting inferred assumptions, such as the interpretation of "current market environment," the workflow ensures that subsequent stages operate on transparent and reviewable task definitions.

5.3. Semantic Mapping and Data Retrieval

After the user intent is interpreted and transformed into a structured representation, the workflow proceeds to semantic mapping and data retrieval as a unified process. At this stage, the objective is to translate high-level user intents into machine-interpretable categories that can be directly aligned with financial database structures.

The system leverages a predefined financial data dictionary to support this mapping process. The dictionary organizes financial information into hierarchical categories, such as markets, industries, firms, and financial indicators, each associated with standardized semantic labels. Based on the structured prompt generated by the intent interpretation agent, the workflow identifies the most relevant categories and labels within the dictionary, thereby narrowing the search scope and reducing ambiguity in subsequent data access.

Once the semantic mapping is completed, the data retrieval agent queries the financial databases using the selected labels and category constraints. Rather than performing exhaustive searches across the entire database, the retrieval process is guided by the mapped semantic structure, enabling efficient and targeted access to relevant financial data. This design allows the workflow to remain scalable across heterogeneous data sources while maintaining consistency in query logic.

By integrating semantic mapping and data retrieval into a single coordinated step, the proposed workflow ensures a smooth transition from natural language intent to structured financial data. This integration reduces dependency on manually crafted query templates and supports flexible adaptation to diverse user queries in financial decision-support scenarios.

5.4. Analytical Reasoning and Output Generation

Based on the retrieved data, the analytical reasoning agent synthesizes decision-support insights aligned with the original user intent. Rather than generating unconstrained recommendations, the agent operates within predefined analytical boundaries. For example, it may summarize recent performance trends of technology equities relative to broader market benchmarks, assess volatility patterns, and highlight valuation indicators that inform risk and return considerations.

The final output is organized into a user-facing response that includes both analytical observations and supporting evidence. Key statements are grounded in retrieved data, and underlying assumptions—such as default time horizons or benchmark selections—are explicitly reported. This structured output enables users to understand not only the conclusion but also the basis on which the conclusion is formed.

By constraining analytical reasoning within the workflow and linking outputs to data evidence, the system reduces the risk of unsupported or speculative conclusions. The resulting response supports informed decision-making while maintaining interpretability and consistency. In cases where the generated output involves ambiguity or insufficient information, the workflow applies clarification or fallback mechanisms to ensure that the decision-support results remain coherent and reliable.

6. Discussion

This section discusses the implications, strengths, and limitations of the proposed LLM-based agent workflow for financial decision support. Rather than evaluating quantitative performance, the discussion focuses on system-level contributions, practical applicability, and directions for future research.

6.1. Practical Implications for Financial Decision Support

The proposed framework has several practical implications for financial decision-support systems [15,16]. First, by enabling natural language interaction with financial databases, the workflow lowers the technical barrier for accessing complex financial information. Users are not required to understand database schema, query languages, or analytical conventions, which broadens the accessibility of data-driven decision support beyond professional analysts.

Second, the agent-based workflow improves consistency in decision-support processes. Traditional expert-driven analysis often depends on individual judgment and experience, leading to variability in interpretation and outcomes. By formalizing intent interpretation, semantic mapping, and data retrieval within a structured workflow, the proposed framework supports more standardized and repeatable decision-support outputs. This consistency is particularly valuable in organizational settings where comparable analyses are required across different users or time periods.

Third, the emphasis on evidence linking and traceability enhances interpretability and trust. Financial decision-making often requires justifications that can be reviewed and audited. By explicitly linking analytical outputs to retrieved data and intermediate representations, the framework supports transparent decision support without relying on opaque model behavior. This property is especially relevant for applications involving risk management, compliance, or internal reporting. The benefits of this approach become more pronounced as database scale and schema complexity increase.

6.2. Advantages of Workflow-Based Integration over End-to-End Models

An important design choice of the proposed framework is the use of a workflow-based integration rather than an end-to-end black-box model. While end-to-end approaches may offer efficiency in narrowly defined tasks, they often obscure intermediate reasoning steps and limit system controllability [16]. In contrast, the workflow-based design decomposes the decision-support process into functional stages, each with clearly defined responsibilities.

This decomposition offers several advantages. It allows system designers to impose constraints on language model behavior, reducing the likelihood of unsupported or speculative outputs. It also facilitates modular updates: improvements to intent interpretation, dictionary mapping, or retrieval mechanisms can be introduced without redesigning the entire system. From an operational perspective, such modularity supports maintenance, debugging, and incremental system evolution.

Furthermore, the workflow-based approach aligns well with the structured nature of financial databases. By treating large language models as semantic processing components rather than autonomous decision-makers, the framework balances flexibility in user interaction with the rigor required for financial data processing. This balance addresses common concerns about deploying language models in sensitive decision-support contexts.

6.3. Limitations of the Proposed Framework

Despite its advantages, the proposed framework has several limitations that should be acknowledged. First, the study presents an applied conceptual design rather than an empirically validated system. While the illustrative example demonstrates logical feasibility, the framework's effectiveness in real-world deployments depends on implementation quality, data availability, and system integration.

Second, the framework assumes the existence of a well-maintained financial data dictionary that accurately reflects database semantics. Constructing and updating such a dictionary may require domain expertise and ongoing effort, particularly in environments where financial databases evolve rapidly. Without careful maintenance, semantic misalignment may still occur.

Third, the current design focuses on structured financial databases and does not explicitly address unstructured data sources such as textual disclosures or news streams. Integrating unstructured data would introduce additional challenges related to data preprocessing, relevance assessment, and evidence integration, which are beyond the scope of the present study.

Finally, although the framework includes mechanisms for handling ambiguity, fully resolving complex or highly subjective queries may still require human intervention. The system is intended to support decision-making rather than to replace expert judgment in all cases.

6.4. Future Research Directions

Several avenues for future research emerge from this study. One direction involves implementing and empirically evaluating the proposed workflow in real-world financial decision-support settings. Such evaluations could examine system usability, consistency of outputs, and user trust, providing quantitative and qualitative evidence of effectiveness.

Another direction concerns extending the framework to incorporate additional data modalities. Integrating unstructured textual data with structured financial databases may enhance decision support but would require careful design to maintain evidence traceability and interpretability. Future research may also explore adaptive dictionary mechanisms that learn semantic mappings over time while preserving system control.

Finally, the agent-based workflow could be extended to support interactive decision-support scenarios. By incorporating user feedback loops and clarification mechanisms, the system may better accommodate evolving decision needs and uncertainty. Exploring these extensions would further strengthen the applicability of workflow-based approaches for intelligent financial information systems.

7. Conclusion

This paper proposes an LLM-based agent workflow for financial decision support using financial databases. Motivated by the growing gap between user-oriented decision needs and the technical complexity of large-scale financial data systems, the study focuses on the design of a structured and interpretable workflow that bridges natural language queries and structured financial databases.

The proposed framework integrates large language models within an agent-based workflow architecture, emphasizing intent interpretation, semantic mapping via a financial data dictionary, controlled data retrieval, and evidence-grounded analytical reasoning. Rather than relying on end-to-end black-box models, the design highlights workflow coordination and functional modularity as key mechanisms for improving accessibility, consistency, and interpretability in financial decision support. An illustrative example demonstrates how the framework processes a natural language query through successive workflow stages to generate data-supported decision insights.

The contribution of this study lies in its system-level perspective. By formalizing the interaction between natural language understanding, agent-based automation, and financial database analysis, the framework offers an applied conceptual solution that complements existing task-specific and algorithm-focused approaches. The workflow-oriented design supports scalability and adaptability while remaining compatible with existing financial data infrastructures.

While the framework is presented at a conceptual level, it provides a foundation for future implementation and empirical evaluation. Further research may explore real-world deployments, integration with additional data modalities, and interactive extensions that incorporate user feedback. Overall, this study contributes to the development of intelligent financial information systems by demonstrating how workflow-based system design can operationalize large language models in structured financial decision-support environment.

References

[1] H. Chen, R. H. L. Chiang, and V. C. Storey, "Business intelligence and analytics: From big data to big impact," MIS Quarterly, vol. 36, no. 4, pp. 1165–1188, 2012.

[2] D. Arnott and G. Pervan, “A critical analysis of decision support systems research,” Journal of Information Technology, vol. 29, no. 4, pp. 293–306, 2014.

[3] J. Li, Z. Liu, S. Wang, and J. Zhang, “Natural language interfaces for data analytics: A survey,” IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 4, pp. 3571–3586, 2023.

[4] D. J. Power, Decision Support Systems: Concepts and Resources for Managers. Westport, CT, USA: Quorum Books, 2002.

[5] X. Yang, Y. Liu, and H. Zhang, “AI-enabled decision support systems in financial services: A review,” Decision Support Systems, vol. 153, p. 113668, 2022.

[6] Y. Zhang, X. Yang, and X. Chen, “Large language models in finance: A survey,” ACM Computing Surveys, vol. 55, no. 12, Art. no. 254, 2023.

[7] T. B. Brown, B. Mann, N. Ryder, et al., “Language models are few-shot learners,” in Advances in Neural Information Processing Systems, vol. 33, pp. 1877–1901, 2020.

[8] F. Doshi-Velez and B. Kim, “Towards a rigorous science of interpretable machine learning,” arXiv:1702.08608, 2017.

[9] N. R. Jennings, “On agent-based software engineering,” Artificial Intelligence, vol. 117, no. 2, pp. 277–296, 2000.

[10] X. Xu, X. Wang, Y. Li, and M. Haghighi, “Intelligent decision support systems: A survey,” Journal of Intelligent Information Systems, vol. 52, pp. 267–293, 2019.

[11] D. Araci, “FinBERT: Financial sentiment analysis using pre-trained language models,” arXiv:1908.10063, 2019.

[12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in Proc. NAACL-HLT, Minneapolis, MN, USA, 2019, pp. 4171–4186.

[13] S. Alter, “Nothing is more practical than a good theory: A work system theory perspective,” Journal of the Association for Information Systems, vol. 18, no. 9, pp. 671–690, 2017.

[14] M. Wooldridge, An Introduction to Multiagent Systems, 2nd ed. Chichester, U.K.: Wiley, 2009.

[15] V. Dignum, Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way. Cham, Switzerland: Springer, 2019.

[16] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, “A survey of methods for explaining black box models,” ACM Computing Surveys, vol. 51, no. 5, Art. no. 93, 2018.

[17] Y. R. Shrestha, S. M. Ben-Menahem, and G. von Krogh, “Organizational decision-making with AI,” Journal of Management Information Systems, vol. 36, no. 3, pp. 729–765, 2019.

[18] H. J. Watson, “Tutorial: Business intelligence and analytics,” Communications of the Association for Information Systems, vol. 34, pp. 1245–1248, 2014.