A unified architectural framework for retail AI agents based on tiered functional decomposition

Ioannis Patias
Faculty of Mathematics and Informatics, Sofia University "St. Kliment Ohridski", Sofia - Bulgaria[1]
patias@fmi.uni-sofia.bg

Abstract: *The rapid emergence of autonomous AI agents in commerce is reshaping how consumers interact with retailers, how decisions are made, and how value is exchanged across digital ecosystems. Yet despite accelerating innovation, the retail sector lacks a unified architectural model that defines how such agents should ingest information, reason about complex commercial contexts, and execute actions in a transparent, interoperable, and trustworthy manner. This paper introduces a unified architectural framework for retail AI agents based on tiered functional decomposition, offering a structured, three-tier model (Input, Model, Output) that standardizes agent behavior across diverse retail environments.*

The Input Tier formalizes how agents acquire and normalize information, including user intent, preference profiles, product catalogs, pricing signals, inventory data, logistics constraints, and retailer policies. By defining a consistent intake layer, the framework ensures that agents operate on coherent, comparable, and semantically aligned data structures.

The Model Tier provides the cognitive core of the agent. It encompasses goal decomposition, relevance scoring, multi-objective optimization, negotiation logic, risk assessment, and alignment safeguards. This tier transforms raw inputs into structured decisions, enabling agents to act as rational, preference-aligned economic participants capable of navigating complex retail ecosystems.

The Output Tier defines how agents act on decisions and how they learn from outcomes. It includes action execution (e.g., purchasing, reserving, scheduling), transparent justification, user-agent dialogue, cross-agent communication, feedback ingestion, and adaptive preference updating. This tier closes the loop between perception, reasoning, and action, enabling continuous improvement and long-term alignment with user goals.

Together, these tiers form a modular, interoperable architecture that can be adopted by retailers, marketplaces, and technology providers to ensure predictable, explainable, and scalable agent behavior. The framework supports progressive capability levels, enabling retailers to evolve from informational agents to fully autonomous purchasing systems. By establishing a shared architectural foundation, this work aims to accelerate the development of agentic commerce and foster a more efficient, transparent, and user-aligned retail ecosystem.

KEYWORDS: RETAIL AI AGENTS, TIERED FUNCTIONAL ARCHITECTURE, AGENTIC COMMERCE, INTEROPERABILITY FRAMEWORKS

1. Introduction

This section introduces the concept of retail AI agents, positioning them as emerging autonomous systems that mediate interactions between consumers, retailers, and digital commerce infrastructures. It outlines the motivations driving their development, the limitations of current retail technologies, and the need for a unified architectural framework capable of supporting scalable, interoperable, and trustworthy agent behavior. The section also situates retail AI agents within broader trends in agentic commerce and multi-agent systems, establishing the conceptual foundations for the tiered functional decomposition proposed in this paper. By clarifying the problem space and the research gap, it prepares the ground for the architectural model that follows.

Retail is undergoing a profound transformation as artificial intelligence shifts from passive recommendation engines to **autonomous agents capable of perceiving, reasoning, and acting within complex commercial environments** [1]. These retail AI agents represent a new class of intelligent systems designed to interpret user intent, evaluate product options, negotiate with retailer or marketplace agents, and execute transactions with minimal human intervention [2]. Their emergence reflects broader advances in agentic AI, multi-agent coordination, and preference-aligned decision systems, all converging to redefine how value is created and exchanged in digital commerce [3].

Despite rapid progress, the deployment of retail AI agents remains fragmented. Existing systems are typically proprietary, narrowly scoped, and optimized for isolated tasks such as product ranking, conversational assistance, or dynamic pricing [4]. They lack a **shared architectural foundation** that defines how agents should ingest information, structure reasoning, and produce actions in a transparent and interoperable manner. This absence of standardization limits scalability, inhibits cross-platform collaboration, and increases the risk of misalignment between user goals, retailer policies, and agent behavior.

The need for a unified framework is amplified by the increasing complexity of retail ecosystems. Modern commerce involves heterogeneous data sources, dynamic pricing mechanisms, intricate logistics networks, and diverse regulatory constraints. Retail AI agents must navigate these environments while maintaining alignment with user preferences, ensuring fairness, and preserving trust. Without a structured architectural model, agents risk becoming opaque, unpredictable, or incompatible across platforms.

This paper addresses this gap by introducing a **tiered functional decomposition** that organizes retail AI agent capabilities into three coherent layers: **Input, Model, and Output**. This decomposition provides conceptual clarity and operational modularity. The **Input Tier** defines how agents acquire and normalize information from users, retailers, and external ecosystems. The **Model Tier** captures the reasoning processes that transform inputs into decisions, including optimization, negotiation, and alignment safeguards. The **Output Tier** governs how agents act on decisions, communicate with stakeholders, and learn from feedback.

By formalizing these tiers, the proposed framework establishes a foundation for **interoperability, explainability, and progressive capability development**. It enables retailers, marketplaces, and technology providers to adopt agentic systems in a structured manner while ensuring consistent behavior across diverse environments. Ultimately, this architecture aims to accelerate the evolution of agentic commerce and support the creation of intelligent, trustworthy, and user-aligned retail ecosystems.

2. Background and related work

This section situates retail AI agents within the broader landscape of agentic commerce, multi-agent systems, and retail AI research. It reviews foundational concepts, highlights limitations of existing architectures, and identifies gaps that motivate the need for a unified tiered framework. By examining prior work in autonomous agents, retail decision systems, and interoperability standards, the section establishes the conceptual and technical context for the proposed three-tier functional architecture.

2.2. Formal definitions

We start with the foundational terminology necessary for understanding the conceptual and technical framing of agent-driven retail systems. A retail AI agent can be understood as an autonomous computational system designed to interpret user intent, analyze retail-specific information, and carry out commercial actions in a manner that remains aligned with both user preferences and the operational policies of the retailer [2]. Such an agent does more than simply process commands. It engages in contextual reasoning, evaluates available options, and selects actions that optimize outcomes within the constraints of the retail environment. This definition emphasizes autonomy, interpretive capability, and policy- aligned execution as core attributes.

Building on this idea, the notion of agentic commerce describes a broader economic paradigm in which autonomous agents mediate interactions, negotiations, and transactions on behalf of individuals or organizations [5]. In this paradigm, agents are not passive intermediaries but active participants capable of evaluating offers, comparing alternatives, and making decisions that reflect the goals of the entity they represent. Agentic commerce therefore shifts traditional commercial processes toward a model where software agents assume responsibilities that previously required direct human involvement, enabling more efficient, scalable, and personalized economic exchanges.

To support such sophisticated behavior, functional decomposition provides a methodological framework for structuring an agent's internal architecture. Functional decomposition involves dividing the agent's capabilities into modular layers that distinguish between information intake, reasoning processes, and action execution [6]. By separating these components, designers can create systems that are more interpretable, maintainable, and adaptable. The intake layer focuses on perception and data acquisition, the reasoning layer handles inference and decision- making, and the execution layer translates decisions into concrete actions. This structured approach ensures that complex agent behaviors emerge from well- defined, interoperable modules rather than monolithic designs, ultimately enabling more robust and scalable retail AI systems.

2.2. Related work in multi-agent systems

Multi- agent system research offers a rich foundation for understanding how autonomous entities can coordinate, negotiate, and make decisions within distributed environments. Over several decades, this field has produced a variety of theoretical models and interaction protocols that describe how agents form beliefs, pursue goals, and resolve conflicts in shared computational spaces. Prominent among these are BDI (Belief- Desire- Intention) architectures [7], which formalize how agents represent mental states and translate them into rational action. Also applicable are contract- net protocols [8], which provide structured mechanisms for task allocation and cooperative problem- solving. Finally, widely used, for intermodal behaviors, we see the utility- based agent models [9], which frame decision- making as the optimization of quantifiable preferences or outcomes. These frameworks supply essential conceptual tools for reasoning about autonomy, coordination, and strategic behavior.

However, despite their generality and expressive power, these classical models do not directly confront the domain- specific challenges inherent to modern retail ecosystems. Retail environments introduce constraints that extend beyond abstract coordination problems, including extreme catalog heterogeneity, rapidly changing product attributes, and complex fulfillment logistics that span inventory management, delivery routing, and real- time availability [4]. Additionally, retail agents must operate within strict regulatory and compliance boundaries, such as consumer protection rules, data- handling requirements, and marketplace fairness policies. These factors impose practical limitations and operational nuances that traditional multi- agent frameworks were not designed to capture. As a result, while multi- agent system research provides valuable building blocks, it requires significant adaptation and augmentation to support the demands of agentic commerce in large- scale, real- world retail settings.

2.3. Retail AI and decision systems

Retail AI has historically centered on predictive modeling techniques designed to optimize specific operational tasks such as product recommendations, demand forecasting, and dynamic pricing [10]. These systems have delivered substantial commercial value, but they are typically engineered as narrowly scoped, monolithic pipelines that excel at a single function while lacking the adaptability required for broader autonomous behavior [11]. Their architectures are often inflexible, with tightly coupled components that make it difficult to extend capabilities, incorporate new data modalities, or support interactive decision- making. As a result, traditional retail AI solutions are effective within predefined boundaries but are not well suited for the open- ended reasoning and multi- step action planning expected of autonomous agents.

The emergence of LLMs has introduced a new layer of flexibility, enabling systems that can interpret natural language, generalize across tasks, and perform higher- order reasoning. However, despite their impressive capabilities, LLMs alone do not provide the structural discipline needed for reliable agentic behavior. Without a well- defined architectural framework, their decision processes remain opaque, inconsistent, and difficult to audit or standardize across retail workflows. This gap highlights the need for modular, interpretable agent architectures that can harness LLM reasoning while ensuring predictable, policy- aligned execution in complex retail environments.

2.5. Interoperability standards

Efforts such as schema.org, GS1 standards [12], and OpenAPI specifications [13] have contributed important building blocks for interoperability across digital retail ecosystems. These initiatives establish shared vocabularies, product identifiers, and interface descriptions that allow systems to exchange structured information with greater consistency and reliability. They help reduce fragmentation in how product attributes are represented, how services expose their capabilities, and how platforms communicate basic operational details. Despite their value, these standards remain focused on data formats and service schemas rather than on the behavioral logic required for autonomous agents. None of them articulate how an agent should reason about goals, evaluate constraints, or select actions within a retail workflow. They also do not specify protocols for multi- step decision- making or coordinated execution. This absence creates a critical gap that the proposed three- tier architecture seeks to address by defining the behavioral and procedural foundations necessary for agentic retail systems.

2.5. Conceptual diagram: positioning retail AI agents

The schema (see figure 1) situates retail AI agents within the continuum of retail technologies, illustrating how they extend beyond traditional systems such as search engines, recommendation models, and pricing algorithms. While existing tools provide valuable but isolated capabilities, they lack autonomy, structured reasoning, and cross-system interoperability. Retail AI agents occupy a higher conceptual layer, integrating perception, decision-making, and action within a unified architecture.

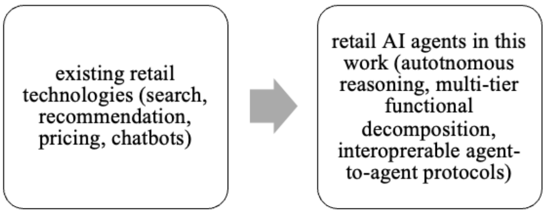

Fig. 1 *Retail AI agents within the continuum of retail technologies.*

This section emphasizes that agents are not incremental improvements but a categorical shift toward systems capable of negotiating, coordinating, and transacting on behalf of users. By positioning agents above legacy components, the diagram highlights the need for a tiered functional framework that can support their expanded role in agentic commerce.

3. Conceptual foundations

The conceptual foundations of the proposed framework establish the principles, requirements, and design goals that guide the development of retail AI agents. As retail ecosystems evolve toward agent-mediated interactions, it becomes essential to articulate the underlying assumptions and constraints that shape agent behavior. This section defines the theoretical basis for tiered functional decomposition, clarifies the operational environment in which retail AI agents must operate, and identifies the architectural qualities necessary for scalable, trustworthy, and interoperable agentic systems.

Retail AI agents operate at the intersection of user intent, commercial data, and autonomous decision-making. Their design must therefore reflect three core principles. First, **modularity**, which ensures that perception, reasoning, and action remain separable yet interoperable components. Modularity enables independent evolution of capabilities, simplifies debugging, and supports heterogeneous implementations across retailers. Second, **alignment**, which requires agents to act consistently with user preferences, retailer policies, and regulatory constraints. Alignment mechanisms must be embedded throughout the architecture, ensuring that autonomy does not compromise trust or compliance. Third, **interoperability**, which allows agents to communicate, negotiate, and coordinate across platforms, retailers, and service providers. Without interoperability, agentic commerce risks becoming fragmented and inefficient.

Dynamic pricing, variable inventory, complex logistics, and diverse product taxonomies characterize the operational environment of retail AI agents. Agents must therefore handle uncertainty, incomplete information, and conflicting objectives. This necessitates a structured approach to functional decomposition, where each tier (Input, Model, and Output) addresses a distinct aspect of the agent’s lifecycle.

4. Three-tier functional architecture

This section introduces the central contribution of the paper: a unified three- tier functional architecture designed specifically for retail AI agents (see figure 2). The framework articulates the structure, responsibilities, and coordinated interactions of the Input, Model, and Output tiers, establishing a clear separation of concerns that supports modularity and interpretability. The Input tier governs data acquisition and intent interpretation, the Model tier handles reasoning and decision- making, and the Output tier manages action execution across retail systems. Together, these components form a coherent architectural blueprint, illustrated through formal descriptions and conceptual diagrams that clarify their roles within agentic retail environments.

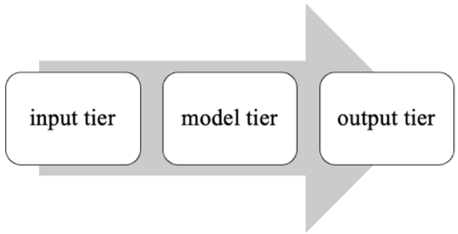

Fig. 2 *A unified three- tier functional architecture designed for retail AI agents.*

4.1. Input tier: perception and context acquisition

The **input tier** is responsible for acquiring, normalizing, and structuring all information required for downstream reasoning. It integrates signals from users, retailers, and external ecosystems, ensuring that the agent operates on coherent and semantically aligned data.

The key components of the input tier include:

- user intent and preferences: captures goals, constraints, and long-term behavioral patterns that guide decision-making.
- product catalog and attributes: normalizes retailer catalog data into a consistent schema for comparison and reasoning.
- pricing, promotions and loyalty: incorporates dynamic prices, discounts, and user-specific benefits.
- inventory and logistics data: provides stock levels, delivery windows, and fulfillment constraints.
- retailer policies: includes return rules, payment options, and compliance requirements.
- external ecosystem data: integrates reviews, sustainability scores, competitor prices, and third-party metadata.

The diagram (see figure 3) illustrates how diverse data sources flow into the input tier and are transformed through a normalization layer into a unified retail data model. The purpose is to ensure that all downstream reasoning operates on consistent, structured information, regardless of retailer systems, catalog formats, or external data variability.

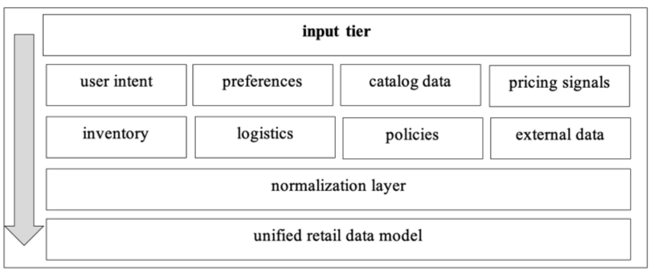

Fig. 3 *Diverse data sources flow into the input tier and are transformed through a normalization layer into a unified retail data model.*

4.2. Model tier: reasoning, optimization, and decisioning

The **model tier** forms the cognitive core of the retail AI agent. It transforms normalized inputs into actionable decisions through structured reasoning, optimization, and alignment.

The key components of the model tier include:

- goal decomposition and task planning: breaks user intent into actionable sub-tasks and defines the reasoning workflow.

- relevance scoring and ranking: evaluates product suitability based on attributes, preferences, and contextual signals.
- multi-objective optimization: balances competing factors such as price, quality, delivery time, and sustainability.
- negotiation and offer evaluation: assesses bundles, promotions, and dynamic pricing opportunities.
- risk and trustworthiness assessment: identifies unreliable sellers, low-quality products, or fulfillment risks.
- alignment and safety constraints: ensures decisions remain consistent with user preferences and retailer policies.
- explainability and traceability: generates interpretable reasoning paths for transparency and auditability.

The diagram (see figure 4) shows how the model tier integrates multiple reasoning modules (ranking, optimization, negotiation, and alignment) into a unified decision synthesis process. The output is a structured set of recommended or selected actions, ready for execution by the output tier.

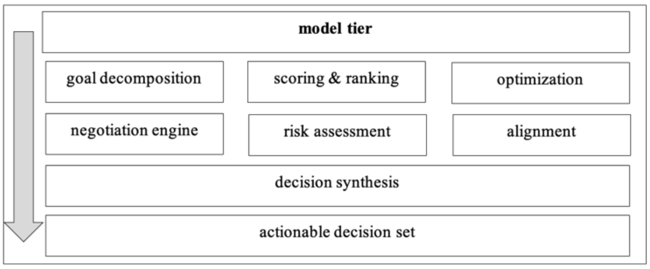

Fig. 4 *The model tier integrates multiple reasoning modules into a unified decision synthesis process.*

4.3. Output tier: action execution and adaptive learning

The **output tier** operationalizes decisions and closes the perception-reasoning-action loop. It executes commercial actions, communicates with users and other agents, and integrates feedback to refine future behavior.

The key components of the model tier include:

- action execution: performs tasks such as adding items to cart, purchasing, reserving, or scheduling delivery.
- transparent justification: provides human-readable explanations for decisions and trade-offs.
- cross-agent communication: interacts with retailer, logistics, and payment agents to complete workflows.
- feedback ingestion: captures user satisfaction, corrections, and post-purchase signals.
- adaptive preference updating: refines user models based on behavior and feedback loops.
- lifecycle automation: supports replenishment, warranty reminders, and subscription management.

The diagram (see figure 5) highlights how the output tier not only executes actions but also feeds insights back into the agent's preference model. This creates a continuous learning loop, enabling the agent to improve alignment and performance over time.

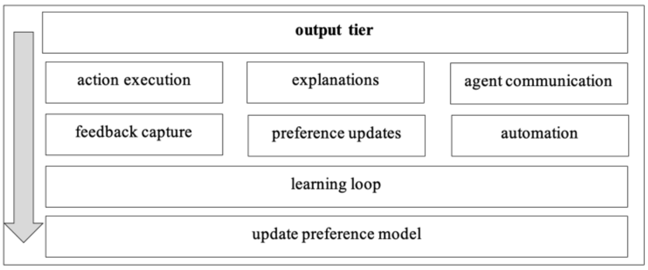

Fig. 5 *The output tier not only executes actions but also feeds insights back into the agent's preference model.*

The three- tier functional architecture outlined in this section establishes a clear and modular foundation for retail- focused agentic systems. By separating perception, reasoning, and action into distinct yet interoperable tiers, the framework supports scalable development, transparent behavior, and seamless integration across heterogeneous retail environments.

Conclusions

The transition toward agentic commerce marks one of the most significant paradigm shifts in the history of retail technology. As AI agents increasingly mediate interactions between consumers, retailers, and logistics networks, the need for a coherent architectural foundation becomes critical. This paper has introduced a unified framework based on tiered functional decomposition, offering a structured approach to designing, evaluating, and standardizing retail AI agents.

By decomposing agent functionality into input, model, and output tiers, the framework provides clarity, modularity, and interoperability. The input tier ensures that agents operate on consistent, semantically aligned data, enabling fair comparisons and reliable decision- making across heterogeneous retail environments. The model tier establishes a principled foundation for reasoning, optimization, and alignment, ensuring that agents behave as rational, preference- aligned economic actors. The output tier operationalizes decisions, enabling transparent action execution and continuous learning through feedback loops.

This architecture not only supports current retail use cases (such as product discovery, personalized recommendations, and dynamic pricing) but also anticipates future capabilities, including autonomous negotiation, multi- agent collaboration, and fully automated purchasing workflows. By defining clear interfaces and responsibilities for each tier, the framework enables retailers to adopt agentic technologies progressively, reducing integration complexity while maintaining alignment with user expectations and regulatory requirements.

Moreover, the proposed architecture lays the groundwork for industry-wide standards. As retailers, marketplaces, and technology providers converge on shared protocols and functional definitions, the ecosystem can evolve toward greater transparency, efficiency, and user empowerment. The framework thus serves as both a conceptual foundation and a practical blueprint for the next generation of retail AI systems.

Ultimately, this work aims to catalyze a shift toward more interoperable, trustworthy, and intelligent retail ecosystems, where autonomous agents enhance decision- making, streamline operations, and deliver meaningful value to consumers and businesses alike.

References

1. Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic ai: Autonomous intelligence for complex goals–a comprehensive survey. IEEe Access. Accessed Jan 2026, at: https://ieeexplore.ieee.org/abstract/document/10849561
2. Mishra, L. N., & Senapati, B. (2025). Retail Resilience Engine: An Agentic AI Framework for Building Reliable Retail

Systems With Test-Driven Development Approach. IEEE Access. Accessed Jan 2026, at: https://ieeexplore.ieee.org/abstract/document/10930951

3. Ntumba, C., Aguayo, S., & Maina, K. (2023). Revolutionizing retail: a mini review of e-commerce evolution. Journal of Digital Marketing and Communication, 3(2), 100-110. Accessed Jan 2026, at: https://tecnoscientifica.com/journal/jdmc/article/view/365/184
4. Hunt, W., & Rolf, S. (2022). Artificial intelligence and automation in retail. Friedrich Ebert Stiftung. Accessed Jan 2026, at: https://uniglobalunion.org/wp-content/uploads/Artificial-Intelligence-and-Automation-in-Retail_EN.pdf
5. Allouah, A., Besbes, O., Figueroa, J., Kanoria, Y., & Kumar, A. (2025). What is your ai agent buying? evaluation, implications, and emerging questions for agentic e-commerce. Evaluation, Implications, and Emerging Questions for Agentic E-Commerce (August 04, 2025).
6. Hernandez, I., Watson, B. C., Weissburg, M. J., & Bras, B. (2024). Using functional decomposition to bridge the design gap between desired emergent multi- agent- system resilience and individual agent design. Systems Engineering, 27(5), 911-930.
7. Agiollo, A., & Omicini, A. (2025). Integrating Machine Learning into Belief-Desire-Intention Agents: Current Advances and Open Challenges. arXiv preprint arXiv:2510.20641.
8. Kang, J., Ren, S., & Wang, C. (2022, August). A Situated Contract Net Protocol Realization in Command and Control. In Proceedings of the 5th International Conference on Information Science and Systems (pp. 49-56).
9. Diallo, A. O., Lozenguez, G., Doniec, A., & Mandiau, R. (2025). Utility-based agent model for intermodal behaviors: a case study for urban toll in Lille. Applied Intelligence, 55(4), 282.
10. Muthukalyani, A. R. (2023). Unlocking accurate demand forecasting in retail supply chains with AI-driven predictive analytics. Information Technology and Management, 14(2), 48-57.
11. Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges. arXiv preprint arXiv:2505.10468.
12. Alhava, O., Arola, T., Torro, O., Järvenpää, M., Järvinen, T., & Ruottinen, B. (2025). AI-Agent Application for Semantic Data Enrichment in Ventila-tion Systems Using National Nomenclature for IFC and GS1-Based Product Information.
13. Ocansey, J. T. (2024, September). Enhanced Interoperability and Consistency in Heterogeneous Systems with CorrLang and OpenAPI. In Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems (pp. 200-203).