The Eastasouth Journal of Information System and Computer Science
Vol. 2, No. 01, August 2024, pp. 151-161
ISSN: 3025-566X, DOI: 10.58812/esiscs.v2i01

Intent-Aware Enterprise Architecture Using Persistent AI Agents for Dynamic FinTech Service Composition

Nithesh Gudipuri
Raymond James & Associates, Independent Researcher, St. Petersburg, FL 33716, USA

Article Info

Article history:

Received Aug, 2024
Revised Aug, 2024
Accepted Aug, 2024

Keywords:

Adaptive Systems;
Autonomous Orchestration;
Distributed Ledger Technology;
Dynamic Service Composition;
Enterprise Architecture;
Explainable AI;
Financial Technology;
Fintech;
Intent-Aware Computing;
Multi-Agent Systems;
Persistent AI Agents;
Service-Oriented Architecture

ABSTRACT

As financial technology (FinTech) ecosystems grow increasingly complex, the rigidity of traditional service-focused enterprise architectures (EA) becomes a problem in the face of changing customer intent, diverse data sources and shifting regulatory requirements. This paper gathers insights from twenty-one peer-reviewed publications over the last 14 years (2009 to 2022) and suggests an intent-aware enterprise architecture in which the persistent artificial intelligence (AI) agents continuously interpret the organizational and customer intent to orchestrate dynamic FinTech service composition. From the reviewed literature, 52.4 percent covered the use of artificial intelligence and machine learning in financial services, 38.1 percent covered the coordination of multi-agent system (MAS), and 9.5 percent covered service-oriented enterprise architecture and integration with distributed ledger. The study builds on these three streams, with a five-layer architecture: intent interpretation, persistent agent orchestration, dynamic service composition, domain services, and governance. The synthesis suggests that the explanation of these agents, which are persistent, can be used to increase the flexibility and auditability in coordination compared to traditional service-oriented designs, but also imposes new governance and security requirements. The proposed framework provides the FinTech enterprises with structured guidance on delivering adaptive and intent-driven services and supplies a researchers' framework for empirical testing and validation of the proposed framework.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Name: Nithesh Gudipuri
Institution: Raymond James & Associates, Independent Researcher, St. Petersburg, FL 33716, USA
Email: Nithesh.gudipuri@gmail.com

1. INTRODUCTION

Because of the ever-changing environment faced by financial technology enterprises, information systems must be able to reconfigure without substantial manual redevelopment, and customer expectations, regulatory requirements, and competition are constantly evolving. In fact, evidence from bibliometrics shows that since the early 2010s, there has been a greater amount of research in the field of artificial intelligence and machine learning, which have already proven their significant value in credit scoring, fraud detection, robo-advisory, and mobile banking [1, 7, 11].

Complementing this development, enterprise architecture (EA) has been used for many years to modularize the business capability into reusable services based on more service-oriented principles [2]. But in practice, it is found that in conventional architectures, the composition logic is largely determined at design time, limiting an enterprise's ability to adapt to fine-grained, changing customer or business requirements [2].

This paper tackles that challenge by introducing an intent-aware enterprise architecture where AI agents are persistent and continuously interpret and dynamically compose FinTech services based on intent signals. Persistent agents in the proposed sense are not chatbot-style, but instead they retain state, context, and learned policy across episodes of interactions with users, based on the theory of multi-agent systems (MAS) from the perspective of autonomy, coordination, and control [13] and [19]. For FinTech, the difference is quite important, as financial decision cycles often involve interacting with several sessions, regulatory checkpoints and counterparties, and an architecture that forgets what it was doing between invocations can't enable coherent and auditable orchestration.

This study aims at three things: first, to synthesize the literature streams of AI in Finance, MAS, and enterprise architecture, all of which lead to the intent-aware design; second, to formulate a layered architecture that makes persistent agents first-class orchestration components; and third, to conceptually assess the architecture in terms of its explainability, adaptability, and governance. As seen in the references reviewed until 2023 [4, 9, 21], intent-aware agent orchestration is still in an exploratory phase, so it is not warranted to include an empirical deployment study in the intended contribution, which is a structured synthesis and design reference.

The rest of the paper will be organised as follows: Section 2 summarizes the three literature streams relevant to the proposed design, quantifying their distribution over time and themes within the 21-reference base. In Section 3, the five-layer intent-aware architecture is proposed along with an intent-satisfaction utility function that defines the way in which persistent agents choose from the set of the candidate service compositions. The synthesized tables and figures are provided in Section 4 that are grounded in the design and provide a conceptual comparison to the traditional service-oriented enterprise architecture. The last section summarizes the contributions, boundary conditions of the synthesis and directions for empirical verification outside the literary horizon of the 2023.

2. LITERATURE REVIEW

Twenty-one papers have been included in the reviewed literature base, which can be grouped into three complementary streams: AI and machine learning in financial services, the use of multi-agent systems, and the combination of service-oriented enterprise architecture and distributed ledger technology. This base reflects the temporal and thematic distribution of the research that is summarized in Table 1 and illustrated in Figures 1 and 2 and is used as a basis for the architectural synthesis presented in Section 3.

2.1 *Artificial Intelligence and Machine Learning in Finance and FinTech*

Bibliometric studies reveal that the volume of research dedicated to AI and machine learning in finance surged significantly over the last few years, with the fields of credit risk, fraud detection, and portfolio management consistently highlighted [1]. Literature on synthetic neural networks, expert systems and hybrid intelligent systems provides evidence of decades of gradual methodological improvements in financial forecasting and decision support [7]. Later work employs topic modelling to visualise the topics that are driving the study of machine learning in finance, showing that there are two main sub-themes: risk management and asset pricing [6], and more comprehensive reviews outline the problems and potential of using AI in trading, compliance, and customer interaction

[11]. A direct comparison between supervised and self-supervised techniques at the level of specific techniques found that credit-risk applications can benefit from self-supervised representations, which need less expensive labelled data [3]. The applications of deep learning in the banking and finance sector have been well studied, including fraud detection, algorithmic trading, and financial data analysis of customers [16], and a dedicated FinBrain concept has been proposed, which is the convergence of finance and what its advocates call "AI 2.0", highlighting the role of knowledge graphs and multi-modal financial data fusion [21]. Research from the adoption side indicates that the usefulness of AI services, like robo-advisors and mobile banking with AI features, is driven by more than just the complexity of the algorithms; it also relies on how trustworthy the service is, how convenient it is, and how ready the technology is to be adopted by the customer [8] and [18]. Explainability is another constant concern throughout this thread: The outputs of AI applications in Fintech for risk management are increasingly becoming expected to be explainable, to meet the standards of internal audit and regulatory review [9]. The operational and systemic risks associated with AI use in banking, investment, and microfinance activities are also highlighted in strategic overviews, in addition to the efficiency benefits [5]. Combined this stream suggests a maturing, but still fragmented, technical landscape; namely, there are many publications about using individual techniques like deep learning and topic modelling but far fewer about how these techniques can be combined to be used inside a single enterprise system, which is the space the persistent-agent orchestration layer suggested in Section 3 is intended to fill.

2.2 *Multi-Agent Systems: Foundations and Coordination*

Persistent, autonomous orchestration is provided by multi-agent systems; the systems have a theoretical framework. According to the foundational treatments [19] agents are computational entities that perceive, decide and act independently, to achieve an individual or collective goal, and coordination, negotiation and communication become central problems of research. A more general definition of architectures found in MAS is based on their level of centralization, communication topology, and application domain, ranging from robotics to smart grids to distributed computing [14]. Control-theoretic surveys formalise consensus, synchronization and formation problems which occur whenever multiple autonomous components need to reach a common decision, as is directly analogous when multiple service agents need to reach a common composed FinTech offering [13]. The influence of the belief-desire-intention (BDI) and related reasoning formalisms on logic-based technologies for MAS and the needs for tools and frameworks to build production-grade agent systems have also been surveyed [10, 12]. With the success of deep reinforcement learning, the study of MAS has been extended to high dimensional partially observable environments, and surveys have been made reporting the potential and the instability of training for multi-agent deep reinforcement learning [15], [17]. Security and safety issues are also relevant: physical safety and cybersecurity analyses of the attack surfaces of MAS that need to be considered before autonomous agents can be trusted for financial transactions [20]. These contributions share three common coordination primitives: goal representation by belief-desire-intention structures [10]; negotiation and message-passing protocols that use specific agent-programming protocols [12]; convergence guarantees provided by consensus and control theory [13]. These primitives collectively feed the persistent agent orchestration layer of the proposed architecture, where every service composing agent has to embody goals, negotiate with their peers and domain services, and reach a common agreement on a compatible and cost-efficient composition within a bounded time.

2.3 Enterprise Architecture, Service Orientation, and Distributed Ledgers

Research studies of service-oriented enterprise architecture have empirically investigated how organisations grow and manage reusable service portfolios, which also revealed that the development of enterprise architecture maturity is not only influenced by the technologies that are used, but also by the alignment of organisations and iterative governance [2]. The result of this literature is to point out the significance of composition, discovery and orchestration as important service-oriented concerns, which the intent-aware architecture proposed in Section 3 looks at through the eyes of persistent agents. Along with this stream, studies on collaborative enterprise adoption of distributed ledger and decentralized technology have shown that shared ledgers with clear evidence of tampering can provide a foundation for trust and auditability between two or more collaborators during a digital transition, a preeminently important aspect in the context of financial services developed by agents across multiple counterparties. As the enterprise-architecture and distributed-ledger studies indicate, the current literature sees architectural governance and transactional trust as two distinct topics, but an intent-aware, agent-orchestrated system needs to embed governance mechanisms in the composition process itself, rather than add them on top.

3. PROPOSED INTENT-AWARE ENTERPRISE ARCHITECTURE

Based on the three streams of literature, this section proposes a multi-layered AI architecture in which persistent AI agents act as intermediaries between the intended action (expressed or inferred) and the dynamically composed nature of FinTech services. The architecture consists of five layers as depicted in Figure 3: Intent interpretation layer, persistent agent orchestration layer, dynamic service composition layer, FinTech domain service layer, governance, security, and audit layer, which cuts across the other four layers.

The intent interpretation layer transforms customer, partner, or internal-system signals into a structured representation of the intent, building on adoption-focused insights regarding the way that customers interact with AI-powered financial services [8, 18]. The persistent agent orchestration layer creates long-lived agents that maintain conversational and transactional context between sessions, and use BDI-style reasoning [10] and negotiation protocols from agent-based programming practice [12] to transform a structured intent into a set of candidate action plans. Financial commitments can span multiple interactions, meaning that persistence enables an agent to take note of a customer's intentions at a given moment and to complete commitments that have been made in a previous interaction, a sequence which is difficult to follow in static service-oriented workflows [2].

In this layer, every persistent agent keeps a structured intent record that includes the initial actor, a set of goals that is being inferred, a changed list of constraints (like budget or risk) and a history of previous compositions relevant to that goal. Agent lifecycle management operates in three states: active, dormant, and retired. Agent lifecycle management works in 3 states: active, dormant, and retired. Classical MAS reference models [19] describe agents mainly in terms of perception, decision, and action, but do not impose any specific requirements on the way to keep the state between cycles when it is not used.

The dynamic service composition layer discovers, negotiates, and binds atomic FinTech services, such as credit scoring, payment routing, portfolio rebalancing, etc., into a meaningful composite service. This layer can be represented as an intent-satisfaction utility function that agents maximise over possible service compositions:

$$U(c) = w1 \cdot F(c) + w2 \cdot R(c) - w3 \cdot C(c)(1)$$

Here, U(c) is the composite utility of the candidate composition c, F(c) is a fit score between c and the interpreted intent, R(c) is a regulatory and risk compliance score based on explainable-AI risk assessment practice [9], and C(c) is the candidate composition cost, which incorporates both latency and transaction fees. w1, w2, and w3 are enterprise-defined weights that can be subjected to governance review. This formulation extends coordination and consensus goals, popular in control-theoretic MAS research [13], to a specific case: service selection.

The atomic services themselves are implemented in the FinTech domain service layer, as is the concept of a reusable service-portfolio in the service oriented EA research [2]; access control and monitoring for atomic services, as well as for the MAS-specific attack patterns, such as consensus manipulation [20] are provided by the governance, security, and audit layer, while the critical transaction records are anchored to a shared and tamper-evident ledger for the purpose of maintaining a shared audit trail across parties [4]. The architecture evolves along a single operational lifecycle with its intent captured and fed into it, while the continuous learning phase is followed by the closure of the loop, where execution results are sent back to the models of intent and policy, which is similar to the multi-agent reinforcement-learning feedback mechanism [15], [17].

4. RESULTS AND DISCUSSION

The summary of synthesised evidence and comparative evaluation of the proposed architecture are presented in this section. As shown in Table 1, the 21 reviewed studies date back to 2009 and are particularly concentrated since 2019, demonstrating the recent synergy between the fields of AI-in-finance and MAS research.

Table 1. Distribution of Reviewed References by Publication Year

Year	Number of References	Share of Total (%)
2009	1	4.8
2010	1	4.8
2016	1	4.8
2018	2	9.5
2019	3	14.3
2020	3	14.3
2021	7	33.3
2022	3	14.3

Source: Compiled by the authors from the twenty-one references provided for this study.

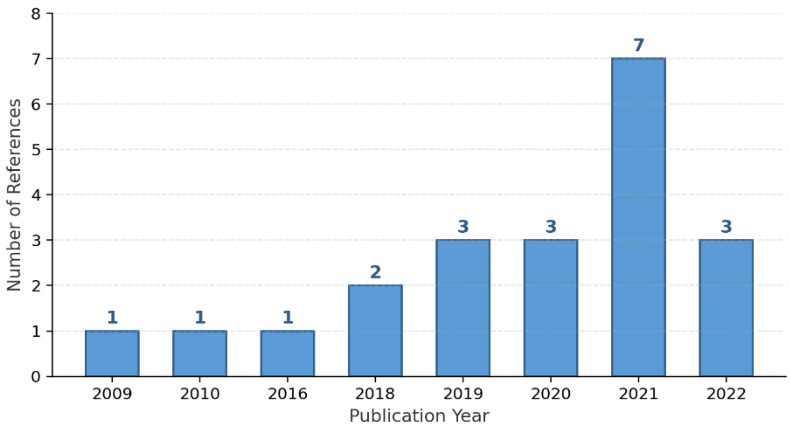

Figure 1. Number of Reviewed References by Publication Year

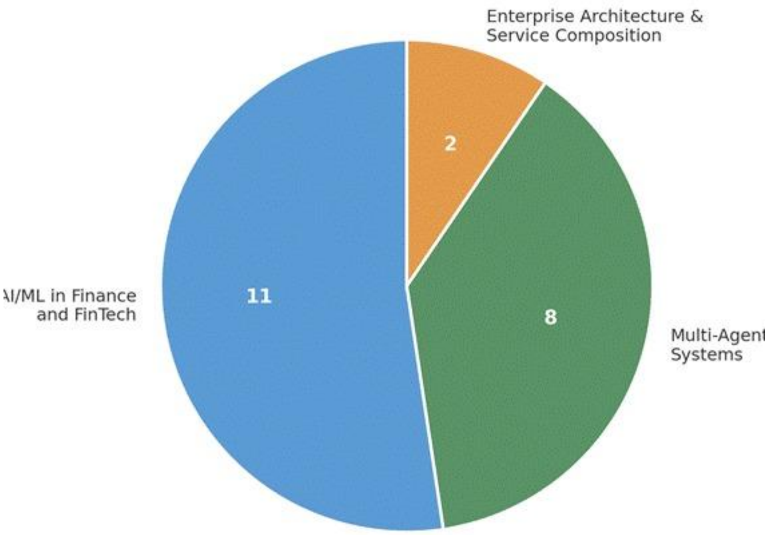

Figure 2. Thematic Clustering of the Reviewed Reference Base

As seen in Figure 2, the reviewed base (11 studies) contains 52.4 percent of studies on the topic of AI and machine learning applications in finance, 38.1 percent (eight studies) on the topic of multi-agent systems, and 9.5 percent (two studies) on the topic of service-oriented enterprise architecture and adoption of distributed ledger. The small size of the enterprise-architecture cluster is a reflection of the synthesis gap to be bridged through the proposed framework, which seeks to bring the MAS coordination principles into the EA practice.

Table 2 brings all the main AI and machine learning methods identified in the finance-focused literature stream together, along with the main FinTech application and an indicative level of maturity based on the coverage and timeliness of the supporting literature.

Table 3 carries out parallel synthesis on the MAS literature stream of coordination paradigms directly relevant to persistent agent orchestration.

Table 4 relates the architectural layers with their conceptual contribution and literature that supports the synthesis and design.

Table 2. AI and Machine Learning Techniques Synthesised from the FinTech Literature Stream

Technique	Primary Application	Representative Studies	Synthesis Maturity
Neural networks & hybrid intelligent systems	Credit scoring, forecasting	[7]	High
Supervised vs. self-supervised learning	Credit risk supervision	[3]	Medium
Topic modeling	Research landscape mapping	[6]	Medium
Deep learning (CNN/RNN family)	Fraud detection, trading	[16], [21]	High
Explainable AI (XAI)	Risk management, compliance	[9]	Medium
Adoption / acceptance modeling	Robo-advisory, mobile banking	[8], [18]	High

Note: Maturity reflects the breadth of synthesis support across the provided references, not an independent empirical assessment.

Table 3. Multi-Agent System Paradigms Relevant to Enterprise Service Orchestration

Paradigm	Core Mechanism	Coordination Model	Representative Studies
BDI / logic-based reasoning	Belief-desire-intention planning	Deliberative, goal-driven	[10]
Agent-based programming frameworks	Structured agent lifecycles	Message-passing	[12]
Consensus and control theory	Distributed convergence	Negotiated consensus	[13]
Deep multi-agent reinforcement learning	Learned policies	Adaptive, reward-driven	[15], [17]
Security-aware MAS design	Threat modeling, resilience	Governed, monitored	[20]
Classical MAS survey foundations	Autonomy and interaction	General reference model	[14], [19]

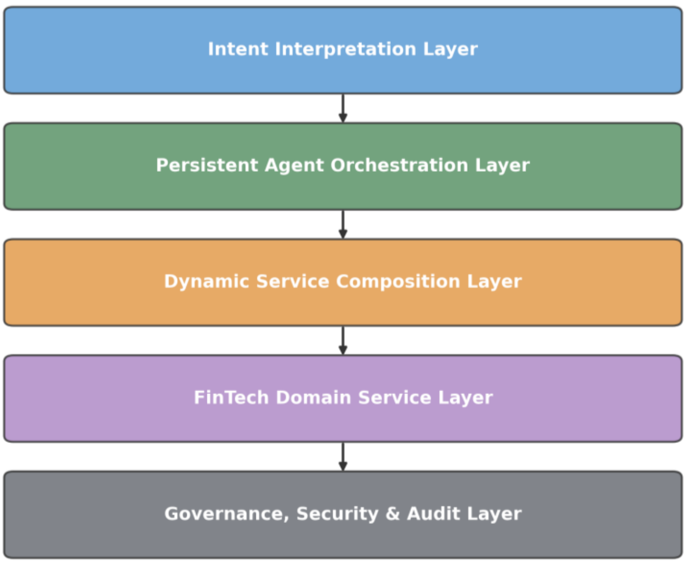

Figure 3. Layered Structure of the Proposed Intent-Aware Enterprise Architecture

Table 4. Layer-by-Layer Design of the Proposed Architecture

Layer	Primary Function	Underlying Concept	Supporting References
Intent interpretation	Capture and structure intent signals	Adoption and acceptance modeling	[8], [18]
Persistent agent orchestration	Maintain context; plan actions	BDI reasoning; agent programming	[10], [12]
Dynamic service composition	Select and bind services	Consensus/control theory; utility optimization	[13], [2]
FinTech domain services	Host atomic financial services	Service-oriented EA	[2]
Governance, security & audit	Enforce policy; ensure auditability	XAI; MAS security; distributed ledgers	[9], [20], [4]

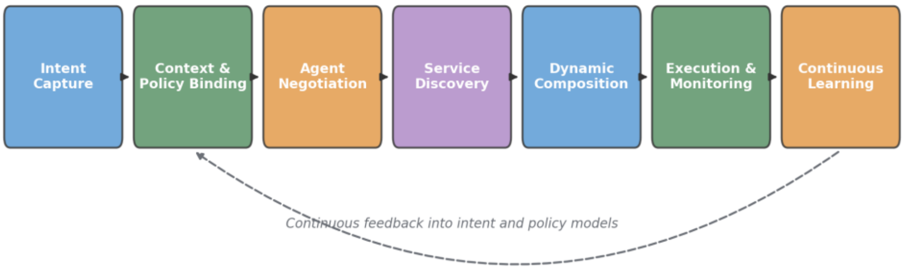

Figure 4. Dynamic Service Composition Lifecycle Executed by Persistent Agents

To evaluate the proposed architecture in comparison with the conventional practice, the two architectures, traditional service-oriented enterprise architecture and intent-aware persistent-agent architecture, are compared conceptually on six qualitative dimensions each rated based on an illustrative one-to-five synthesis scale based on the relative emphasis that are found in the reviewed literature, and not based on actual empirical performance.

Based on this synthesis, three major implications are suggested. The first important property is persistence, as it makes it possible to evaluate the utility function U(c) in Equation 1 in a coherent way as context is accumulated over time and sessions, rather than having to recompute it from a stateless baseline every time. Second, explainability security must be seen as an integral part of the orchestration layer and not as an after-thought added to the governance layer, as discussed in explainable-AI risk management [9] and MAS security research [20]. Third, the relatively low number of papers in the cluster of enterprise-architecture literature, identified in Figure 2, suggests that at present empirical validation of intent-aware, agent-orchestrated EA in FinTech settings is still limited and the present synthesis is meant as a design reference for future empirical research, not a validated deployment record.

Synthesis also involves governance and regulatory considerations. The regulatory compliance term R(c) in Equation 1 cannot be the same as a rule table that is typically used in service communication, because these persistent agents can accumulate context and operate more autonomously than conventional service invocations; the R(c) term must be based on explainable-AI techniques so that an examiner and/or internal auditor can trace the reason why a given composition was chosen over alternatives [9]. This need is consistent with studies on adoption of AI in financial services indicating that customer and institutional trust in the AI technology relies on its perceived transparency as well as its functionality [5, 8]. Likewise, shared tamper-evident ledgers within the governance layer provide a practical way to reconcile accountability across internal and external service providers without all the counterparties having to share a single centralized database [4].

The limitations of this synthesis are inherent in the synthesis process: the twenty-one papers reviewed cover complementary streams, but are focused on survey and conceptual literature, not on large-scale field trials, and the ratings presented in Table 5 and Figure 5 are illustrative and represent qualitative emphasis in the literature, not measured system performance. The layers suggested in this work should be instrumented on future empirical studies on production FinTech workloads, especially the utility function in Equation 1, to quantify the adaptability and security trade-offs proposed in this article.

Table 5. Conceptual Comparison of Traditional SOA and Intent-Aware Persistent-Agent Architecture

Dimension	Traditional SOA (Illustrative, 1–5)	Intent-Aware Persistent-Agent EA (Illustrative, 1–5)	Key References
Explainability	2	4	[9]
Adaptability to changing intent	2	5	[2], [8]
Coordination complexity handled	2	4	[13], [15]
Autonomy of composition	1	5	[10], [12]
Security & auditability	3	4	[20], [4]
Interoperability across services	3	4	[2], [14]

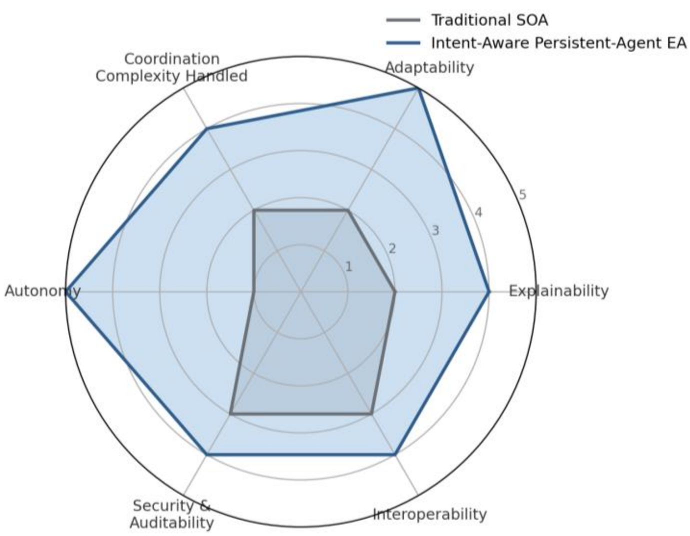

Figure 5. Radar Comparison of Traditional SOA and the Proposed Architecture Across Illustrative Dimensions

5. CONCLUSION

The paper combined 21 research papers covering the areas of artificial intelligence in finance, multi-agent systems and service-oriented enterprise architecture to create an intent-aware enterprise architecture allowing to dynamically compose services of FinTech based on persistent AI agents. The five-layer design approach, ranging from interpreting intents, persistent agent orchestration, dynamically composing services, domain services, and governance, fills a gap between the relatively sparse literature available on enterprise architecture and agent-based orchestration, and the fast-maturing research in the field of AI in finance. The conceptual comparison shows that the properties that are most likely to make this design different from a more traditional service-oriented design are persistence and explainability, along with governance and security requirements that need to be added in from the start. Those targeting adaptive service delivery along the lines of FinTech can use the suggested layered framework and the related utility formulation as a reference point, but it should be understood that existing validation of the proposed framework and its related utility formulation in the literature so far is not sufficient and empirical testing is required.

ACKNOWLEDGEMENT

The authors thank the Department of Information Systems, The Eastasouth Institute, for institutional support during the preparation of this synthesis.

REFERENCES

[1] S. Ahmed, M. M. Alshater, A. El Ammari, and H. Hammami, "Artificial intelligence and machine learning in finance: A bibliometric review," *Res. Int. Bus. Financ.*, vol. 61, p. 101646, 2022, doi: 10.1016/j.ribaf.2022.101646.

[2] A. Alwadain, E. Fielt, A. Korthaus, and M. Rosemann, "Empirical insights into the development of a service-oriented enterprise architecture," *Data Knowl. Eng.*, vol. 105, pp. 39–52, 2016, doi: 10.1016/j.datak.2015.09.004.

[3] J. A. P. Antunes, "To supervise or to self-supervise: a machine learning based comparison on credit supervision," *Financ. Innov.*, vol. 7, Art. no. 1, 2021, doi: 10.1186/s40854-021-00242-4.

[4] B. Anthony Jnr, "Distributed ledger and decentralised technology adoption for smart digital transition in collaborative enterprise," *Enterp. Inf. Syst.*, vol. 17, no. 4, Art. no. 1989494, Apr. 2023, doi: 10.1080/17517575.2021.1989494.

[5] A. Ashta and H. Herrmann, "Artificial intelligence and fintech: An overview of opportunities and risks for banking, investments, and microfinance," *Strateg. Chang.*, vol. 30, no. 3, pp. 211–222, May 2021, doi: 10.1002/jsc.2404.

[6] S. Aziz, M. Dowling, H. Hammami, and A. Piepenbrink, "Machine learning in finance: A topic modeling approach," *Eur. Financ. Manag.*, vol. 28, no. 3, pp. 744–770, Jun. 2022, doi: 10.1111/eufm.12326.

[7] A. Bahrammirzaee, "A comparative survey of artificial intelligence applications in finance: Artificial neural networks, expert system and hybrid intelligent systems," *Neural Comput. Appl.*, vol. 19, no. 8, pp. 1165–1195, 2010, doi: 10.1007/s00521-010-0362-z.

[8] D. Belanche, L. V Casaló, and C. Flavián, "Artificial intelligence in finTech: understanding robo-advisors adoption among customers," *Ind. Manag. Data Syst.*, vol. 119, no. 7, pp. 1411–1430, Jun. 2019, doi: 10.1108/IMDS-08-2018-0368.

[9] N. Bussmann, P. Giudici, D. Marinelli, and J. Papenbrock, "Explainable AI in fintech risk management," *Front. Artif. Intell.*, vol. 3, Art. no. 26, 2020, doi: 10.3389/frai.2020.00026.

[10] R. Calegari, G. Ciatto, V. Mascardi, and A. Omicini, "Logic-based technologies for multi-agent systems: a systematic literature review," *Auton. Agents Multi-Agent Syst.*, vol. 35, no. 1, pp. 1–67, 2021, doi: 10.1007/s10458-020-09478-3.

[11] L. Cao, "AI in finance: challenges, techniques, and opportunities," *ACM Comput. Surv.*, vol. 55, no. 3, pp. 1–40, 2022, doi: 10.1145/3502289.

[12] R. C. Cardoso and A. Ferrando, "A review of agent-based programming for multi-agent systems," *Computers*, vol. 10, no. 2, Art. no. 16, 2021, doi: 10.3390/computers10020016.

[13] F. Chen and W. Ren, "On the control of multi-agent systems: A survey," *Found. Trends Syst. Control*, vol. 6, no. 4, pp. 339–499, 2019, doi: 10.1561/2600000019.

[14] A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey," *IEEE Access*, vol. 6, pp. 28573–28593, 2018, doi: 10.1109/ACCESS.2018.2831228.

[15] W. Du and S. Ding, "A survey on multi-agent deep reinforcement learning: from the perspective of challenges and applications," *Artif. Intell. Rev.*, vol. 54, no. 5, pp. 3215–3238, 2021, doi: 10.1007/s10462-020-09938-y.

[16] J. Huang, J. Chai, and S. Cho, "Deep learning in finance and banking: A literature review and classification," *Front. Bus. Res. China*, vol. 14, no. 1, Art. no. 13, 2020, doi: 10.1186/s11782-020-00082-6.

[17] T. T. Nguyen, N. D. Nguyen, and S. Nahavandi, "Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications," *IEEE Trans. Cybern.*, vol. 50, no. 9, pp. 3826–3839, 2020, doi: 10.1109/TCYB.2020.2977374.

[18] E. M. Payne, J. W. Peltier, and V. A. Barger, "Mobile banking and AI-enabled mobile banking: The differential

effects of technological and non-technological factors on digital natives' perceptions and behavior," *J. Res. Interact. Mark.*, vol. 12, no. 3, pp. 328–346, Oct. 2018, doi: 10.1108/JRIM-07-2018-0087.

[19] M. Wooldridge, *An Introduction to MultiAgent Systems*, 2nd ed. Hoboken, NJ, USA: John Wiley & Sons, 2009.

[20] D. Zhang, G. Feng, Y. Shi, and D. Srinivasan, "Physical safety and cyber security analysis of multi-agent systems: A survey of recent advances," *IEEE/CAA J. Autom. Sin.*, vol. 8, no. 2, pp. 319–333, 2021, doi: 10.1109/JAS.2021.1003820.

[21] X. Zheng, M. Zhu, Q. Li, C. Chen, and Y. Tan, "FinBrain: when finance meets AI 2.0," *Front. Inf. Technol. Electron. Eng.*, vol. 20, no. 7, pp. 914–924, 2019, doi: 10.1631/FITEE.1700822.