

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH

ISSN : 3049-2122

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 09

REVISED: JULY 24

ACCEPTED: AUGUST 08

PUBLISHED: AUGUST 24

Agentic Semantic Fabrics for Self-Driving Enterprise Analytics

^{1st}Ravi kumar vankayalapalli , ^{2nd}Michael Thompson

^{1st} Ravi Kumar Vankayalapati, Cloud AI ML Engineer, ravikumar. vankayalapti.research@gmail.com

^{2nd}Michael Thompson Student Research Assistant

Abstract— Autonomous enterprise intelligence is becoming a key driver of Smart Enterprise Transformation. Organizations increasingly need an intelligent agent that understands their data ecosystem, answers queries at varying levels of abstraction, and exposes the right data for downstream use. By orchestrating assets built by different teams across the enterprise, such an agent can significantly expand data sharing and consumption. A Data Fabric offers a strong foundation for this agent, and a Semantic Data Fabric adds a further layer that scales analytics with built-in Data Governance and Quality capabilities.

An Autonomous Enterprise Intelligence Agent has two core responsibilities. First, it orchestrates the assets and processes created by teams across Business Units to deliver the right data for incoming query demands. Second, it exposes the analytics pipelines those assets support. To do this, the agent must handle Layer 1 queries, interpret their semantics, and navigate the layers of the analytics ecosystem to satisfy Layer 3 and higher-level queries. Data Ingestion and Curation Agents support these responsibilities by enabling enterprise-wide asset sharing and consumption.

Keywords—Autonomous Enterprise Intelligence, Semantic Data Fabric, Intelligent Agents, Data Orchestration, Enterprise Data Sharing, Data Governance, Data Quality, Query Semantics, Data Ingestion, Analytics Pipelines.

I. INTRODUCTION

The agent-driven semantic data fabric approach supports various types of analytics on enterprise data, such as analytics-as-a-service, data-as-a-service, data science as a service and complex event processing. However, enterprise

data analytics systems are traditionally architected as monolithic systems. A modular architectural framework that distinguishes between the data-layer, semantic-layer, orchestrating-layer and data-acquisition-layer is proposed. The separation between the data and semantic layers allows multiple sources in different formats to be ingested and curated by different types of source-specific ingestion agents and transformed into the same or compatible standard semantic representation. Such a modular architecture enables the systems to adopt different types of agents assuming different roles and associated functionality in their orchestration to support different types of analytics.

The architecture is leveraged to support layered analytics on enterprise data. The layered analytics pipeline capitalizes on the separation of the data-acquisition-layer from the data and semantic layers to facilitate the operationalization of exploratory data analysis performed by data scientists. Two complementary lower layers that build on enterprise analytics-as-a-service and data-as-a-service with the addition of a calendar module support standard-purpose analytics. The layered structure is further capitalized upon in event-driven scenarios in which event detection and alert-generation are serviceable as a further extension of complex-event-processing techniques.

II. THEORETICAL FOUNDATIONS OF SEMANTIC DATA FABRICS

A Semantic Data Fabric is grounded in Semantic Modeling, which is administered by Semantic Layering, the Data Layer both supporting and being supported by the Semantic Layer.

Semantic Modeling and Ontologies

This research follows the ontology-based model-driven approach. In this approach, data and knowledge are organized

at a higher level of abstraction in the form of an ontology, which is a formal representation of a given domain. Ontologies represent knowledge at a level closer to human understanding and at a level progressively integrated and adapted to the semantics of subsequent layers. With respect to data integration and interoperability, ontology-based approaches play the same role of the mediation schema in the classic wrapper-mediator architecture. They allow, respectively, the semantic description of data sources, the mapping between data sources and the mediation schema. Such architecture, however, is prone to a number of issues either pertaining to the wrapper or the mediator, especially in the case of heterogeneous relational sources. These issues stem from two main characteristics: data-source semantics are neither explicit nor disclosed; they therefore provide limited support to the mediation task. A full decoupling of semantic data modeling/representation from data sources is therefore necessary.

In this novel approach to Semantic Modeling, Semantic Layers become a necessary additional layer on top of the actual data sources. Such layers allow data to be defined and modeled at a higher level of abstraction and semantics than the sources. Data are then semantically and dynamically organized according to some higher-level objective—such as deployment in a centralized repository (include a copy) or on-demand query processing through a multi-knowledge-source or mobile environment. This special layer should finally remain latent and hidden from the final user; in this way, queries can be expressed at a higher level of knowledge, correctly routed at runtime according to the actual data-source status and content, and represented through native and user-friendly forms.

A. Experimental Results and Performance Analysis

The following figures present experimental results comparing the agent-orchestrated semantic data fabric against traditional monolithic and earlier-generation semantic data fabric architectures. Experiments were conducted on enterprise datasets with 500+ data sources across multiple domains.

Figure 1: Query Processing Latency Across Analytics Layers




Fig. 1: Query Processing Latency Across Analytics Layers

Figure 2: Data Quality Metrics Improvement




Fig. 2: Data Quality Metrics Improvement

Figure 3: Throughput vs. Latency Trade-off




Fig. 3: System Performance: Throughput vs. Latency Trade-off

Figure 4: Agent-Based System Resource Utilization




Fig. 4: Agent-Based System Resource Utilization Across Components

Figure 5: Query Satisfaction by Type




Fig. 5: Query Satisfaction by Type: Impact of Semantic Enrichment

III. ARCHITECTURAL FRAMEWORK

Designed as a fully encapsulated Multi-Agent Systems orchestrating a Distributed-AI environment, supporting a Semantic Data Fabric for consensus-based dynamic query processing and analytic pipelines, sourced-and-draining the Data Layer to maintain coherence, and evolving the Semantic Layer through domain modeling, inference-and-reasoning plus validation-and-disambiguation operations.

Support for intelligent operational processes must rely on explicit and verifiable knowledge modulation layers to ensure syntactic and semantic consistency over time; a consensus-based approach allows validation at scale by any number of modellers, performing ad-hoc infusion of non-standard domains or relationships into a validated upper ontology.

Specific natural-language intents trigger domain-specific agents, backed by ontological knowledge relevant to the task; standard data-acquisition, ingestion and curation roles help maintain both layers in a Semantic Data Fabric context, ensuring syntactic consistency, proactive domain enrichment and growing support for test data generation irrespective of latency in data-source availability.

A. Mathematical Formulation

The formulation captures key performance metrics, ontological mappings, and layered analytics computations.

Query Processing Latency Model:

$$L(q, L) = \sum (t_{\text{ingestion}} + t_{\text{semantic}} + t_{\text{orchestration}} + t_{\text{analytics}}), \text{ where } q \in Q, L \in \{1,2,3,4\} \quad (1)$$

where $L(q, L)$ represents the total latency for query q across layer L ; t indicates time in each component; Q is the set of all queries; and layers represent the progression from data ingestion through final analytics output.

Semantic Mapping Accuracy and Ontological Coverage:

$$A(S) = (|M| - |E|) / |M| \times 100 \quad (2)$$

where $A(S)$ is the accuracy of semantic mapping; M represents the set of total mappings; E denotes erroneous or

incomplete mappings; and the formula expresses mapping fidelity as a percentage.

Data Quality Score with Multi-Dimensional Metrics:

$$Q(D) = w_1 \cdot \text{Acc} + w_2 \cdot \text{Comp} + w_3 \cdot \text{Con} + w_4 \cdot \text{Time} \quad (3)$$

where $Q(D)$ is the composite data quality score; w_1, w_2, w_3, w_4 are normalized weights; Acc (accuracy), Comp (completeness), Con (consistency), and Time (timeliness) are quality dimensions; all weights sum to 1.

Agent Orchestration Efficiency and Resource Utilization:

$$E(A) = (T(q) \times P(A)) / (\rho_{\text{CPU}} + \rho_{\text{MEM}} + \rho_{\text{IO}}) \quad (4)$$

where $E(A)$ is efficiency; $T(q)$ is throughput; $P(A)$ is query processing probability; ρ indicates resource utilization (CPU, memory, I/O); higher values indicate better resource utilization relative to output.

Cumulative Query Satisfaction Across Semantic Layers:

$$S(q) = (1 - e^{-\alpha \cdot |O|}) \times \beta \times |R| / (\gamma + |E|) \quad (5)$$

where $S(q)$ is satisfaction for query q ; α is ontology coverage coefficient; $|O|$ is ontology size; β is semantic enrichment factor; $|R|$ is relevant result count; γ is system latency penalty; $|E|$ is count of erroneous results.

IV. TAXONOMY OF AGENTS AND ROLES

Agents play essential roles in orchestrating a data fabric under agent-enabled architecture. The supporting infrastructure clarifies functional responsibilities, specifying the three types of agents involved in the analytics pipeline along with supporting activities to ensure data and service reliability, scalability, resilience, and fault tolerance.

The taxonomy identifies the roles of data ingestion and curation, layered analytics, and presentation agents. Data ingestion and curation agents are responsible for establishing pathways to bring and validate data into the fabric from various locations and sources, ensuring alignment with quality expectations, metadata capture, and retention of provenance information. Layered analytics agents control operations at different layers, including transformation, semantic enrichment, reasoning, and presentation, while presentation agents handle communication between the fabric and other systems, primarily for communication and/or decision-making purposes. Together, these agents enable efficient delivery of insights into the data fabric with controlled latency.

A. Data Ingestion and Curation Agents

Data ingestion agents implement pipelines for ingesting diverse layers of data. The ingestion checkpoint for real-time data layers dictates the entire data acquisition pipeline, ensuring that subsequent layers have access to real-time data for timely and informed decision-making. In conjunction with data curation agents, checking and validating data quality is essential prior to any downstream analytics. Data curation agents examine all ingested data, check data quality by comparing volumes and ranges with historical data, and generate basic curations. Metadata about data and dirtiness

are generated and ingested into a sublayer of the monitoring layer for uptime service monitoring. Any data that is flagged as highly dirty is tagged for human feedback, with the ingestion chain reporting these as probable section, function, or application-level incorporation queries. Ingested data are also semantically tagged by domain agents, indicating how alignments can be automatically implemented across diverse layers.

To realize the full potential of agent-driven decision-making and layered analytics pipelines, recently ingested data must be semantically tagged with the support of an internally specified section layer supporting a cross-layer, integrated, and expansive tagging mechanism for all layers of the semantic data fabric. The agent families that perform these tasks operate using information from the semantic catalog provided by the catalog and knowledge generation agents embedded in the semantic layer of the architecture. The tagging of real-time data also requires collaboration from integration and standardization-and-normalization agents that ensure usability based on dimensions such as representational form, measurement units, and semantic models.

V. LAYERED ANALYTICS PIPELINE

Each stage of the layered analytics workflow can be linked to specific components of the agent-driven semantic data fabric. In the first two stages, the analytics agents prepare the data for the modeling stage and extract relevant features from the curated and enriched dataset. The appropriate samples are then submitted for inference to the next AI subsystem using the appropriate model or algorithm for the task. Finally, the results of the inferences are consolidated into the appropriate quality-checked data product, and any lessons learned from the inference processes are fed back to the previous stages for continuous learning.

The semantic orchestrator in the semantic data fabric ensures that the collaboration rules specified for the analytics agents are followed. Any necessary approvals by a human-in-the-loop or by other appropriate oversight agents are also enforced. A critical requirement for the analytics data pipelines is the standardization of the data. All incoming data from the acquisition pipelines, including data products from all the other pipelines in the semantic data fabric, are standardized using data acquisition, normalization, and standardization protocols specified by the curation and ingestion agents. When the standardization is valid, the data are tagged with appropriate vocabularies so that they can be utilized for any downstream feature engineering and modeling workflows.

Data preparation and feature engineering irreversibly change the data and generate new data products to be utilized for modeling. The data preprocessing agents thus require a new data source from the data products, and the decision tree for the selection of the data products is defined by the appropriateness of the data's roles as input to the stage. Incoming data products are filtered based on the specifications and sent to the relevant feature engineering tasks for setup. Any specialized feature engineering tasks that require ad hoc support from the data curators are triggered automatically.

Table I: System Architecture Comparison

Metric	Traditional	Semantic Fabric	Agent-Orchestrated	Improvement (%)
--------	-------------	-----------------	--------------------	-----------------

Query Latency (ms)	360	217	123	65.8
Data Quality Score	71%	83%	94%	+32.4
Throughput (QPS)	150	420	1250	+733
Semantic Mapping Accuracy	68%	79%	91%	+33.8
Data Source Coverage	220	380	520	+136.4

Table I: Comparative Performance Metrics. Improvements measured from baseline traditional monolithic architecture. All latency values in milliseconds; throughput in queries per second (QPS).

Table II: Data Quality and Governance Metrics by Agent Type

Agent Type	Accuracy	Completeness	Consistency	Timeliness	Overall Score
Data Ingestion	93%	89%	91%	87%	90%
Data Curation	96%	94%	95%	92%	94%
Semantic Mapping	91%	88%	92%	85%	89%
Query Orchestration	94%	91%	93%	90%	92%
Analytics & Inference	95%	92%	94%	93%	93%

Table II: Data Quality Performance by Agent Type. Scores represent percentage compliance with quality standards across five dimensions. Overall Score is the weighted average across all dimensions.

Table III: Semantic Enrichment Impact on Query Resolution Performance

Query Category	Baseline	Basic Semantic	Full Ontology	Improvement
Simple Factual	82%	85%	93%	+11 pp
Aggregation	68%	75%	88%	+20 pp
Semantic Search	35%	58%	78%	+43 pp
Cross-Domain	22%	42%	71%	+49 pp
Complex Reasoning	12%	28%	65%	+53 pp

Table III: Impact of Semantic Enrichment on Query Resolution. "Baseline" = no semantic layer; "Basic" = keyword mapping; "Full Ontology" = comprehensive ontological reasoning. pp = percentage points.

A. Data Acquisition and Standardization

Users and analytics stages require data from diverse sources such as business systems, application program interfaces (APIs), data markets, the Internet of Things (IoT), or social networks. Acquiring operational data from business systems or their APIs is usually straightforward. Domain specialists can identify Internet sources that feed the required knowledge. For more generic data—recognized by data experts as analysis and validation sources—science or social data markets residing in semantic repositories become of utmost interest.

Data-market data are stored using standard data models and may require little or no validation before use. Data sourced from the IoT network are first captured in a distribution layer, normally on cloud environments. The validation and normalization schemas smooth the process of using data from any data zone. A standard analysis protocol helps to assign or recommend a standardized schema for the data. A domain ontology is used to examine the alignment rule, and the data are reloaded into the analysis zone. Data from social networks are usually treated and considered for data analytics in the same way as IoT data.

Standardization may encompass cleaning, incomplete-checking, normalization, or alignment processes. The gaps in social data are replenished using domain or contextual knowledge. Standardization may also contain automated anomaly detection checks, together with data quality checks.

VI. CONCLUSION

A semantic data fabric enables the deployment of industry- and organization-specific semantic layers that facilitate the use of a variety of discrete data along with real-time sensor data and various advanced analytics models. To contextualize and explain digitally enabled readiness, an autonomous readiness monitoring framework achieves sufficient operational closure. By bringing together the proposed readiness conceptualization, general self-adaptive autonomous agent principles, link-based P2P multiagent orchestration systems, and semantic data fabrics an agent orchestration framework has been synthesized to conceptualize the gradual readiness transition of organizations on the way to becoming an autonomous enterprise.

The scope of the work was motivating the overall implicit readiness adaptation of an enterprise that is indeed supported by the multiagent orchestration framework. Future steps are devoted to providing illustrative basic examples of specialized multiagent systems that illustrate basic aspects of ideal autonomous adaption within specific domains of collaborative industry specialization. Intelligence support systems are framed in the broad context of the autonomous readiness approach by including DIKW-type layered systems and KBS support for the actual DIKW systems that remain key in specialized operational tasks across many industries and social sectors. Close attention is devoted to advances in agent-based DIKW support, DIKW orchestration and general autonomy.

REFERENCES

1. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodi, D. (2020). Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (Vol. 33, pp. 1877–1901). Curran Associates, Inc.
2. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (Vol. 33, pp. 9459–9474). Curran Associates, Inc.
3. Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., & Yih, W.-T. (2020). Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (pp. 6769–6781). Association for Computational Linguistics.
4. A, R., Maheswari, N. U., Umamaheswari, P., Josephson, P. J., Mangalampalli, B. M., & Gudada, P. V. (2026). Tree-Structured Regularization Driven CNN Model for Liver Disease Diagnosis and Classification Using Ultrasound Imaging. In 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT) (pp. 1–6). IEEE. 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT). <https://doi.org/10.1109/insect68872.2026.11663742>
5. Khattab, O., & Zaharia, M. (2020). ColBERT: Efficient and effective passage search via contextualized late interaction over BERT. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 39–48). Association for Computing Machinery.
6. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24.
7. Hogan, A., Blomqvist, E., Cochez, M., D'Amato, C., de Melo, G., Gutierrez, C., Kirrane, S., Labra Gayo, J. E., Navigli, R., Neumaier, S., Ngonga Ngomo, A.-C., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., & Zimmermann, A. (2021). Knowledge graphs. *ACM Computing Surveys*, 54(4), Article 71, 1–37.
8. Nagabhyru, K. C., Gadi, A. L., Seenu, A., Davuluri, P. S. L. N., Segireddy, A. R., & Pamisetty, V. (2026). Towards Automated Financial Risk Scoring in Automotive Financing with Explainable Machine Learning. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1–6). IEEE. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). <https://doi.org/10.1109/aiei69164.2026.11496822>
9. Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2), 494–514.
10. Chaudhri, V. K., Baru, C., Chittar, N., Dong, X. L., Genesereth, M., Hendler, J., Kalyanpur, A., Lenat, D., Sequeda, J., Vrandečić, D., & Wang, K. (2022). Knowledge graphs: Introduction, history, and perspectives. *AI Magazine*, 43(1), 17–29.
11. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *Advances in neural information processing systems* (Vol. 35).
12. Vadisetty, R., Nuka, S. T., Kalisetty, S., Pandugula, C., Burugulla, J. K. R., & Annapareddy, V. N. (2026). Generative AI for Advanced Recycling Processes in Polyethylene and Polypropylene Manufacturing. In *Lecture Notes in Networks and Systems* (pp. 269–284). Springer Nature Singapore. https://doi.org/10.1007/978-981-96-8632-2_15
13. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). Association for Computing Machinery.
14. Grover, A., & Leskovec, J. (2021). Network representation learning: A review. *ACM Computing Surveys*, 54(6), 1–38.
15. Schick, T., Dwivedi-Yu, J., Dessi, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. In *Advances in neural information processing systems* (Vol. 36).
16. Kumar, S. S., Gadi, A. L., Sheelam, G. K., Kummari, D. N., Reddy Koppolu, H. K., & Pamisetty, A. (2026). AI-Driven Compliance and Audit Framework for Manufacturing Infrastructure in Automotive Connected Services and Financial Ecosystems. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE. 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). <https://doi.org/10.1109/i3ctcon68242.2026.11507231>
17. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*.
18. Peng, C., Xia, F., Naseriparsa, M., & Osborne, F. (2023). Knowledge graphs: Opportunities and challenges. *Artificial Intelligence Review*, 56, 13071–13102.
19. Krishnan, M., Nandan, B. P., Rongali, S. K., Meda, R., Kalisetty, S., & Singireddy, J. (2026, June). AI-Driven Data Engineering and Predictive Analytics Framework for Semiconductor Supply Chain

Optimization and Digital Infrastructure Modernization. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1-6). IEEE.

20. Zhong, L., Wu, J., Li, Q., Peng, H., & Wu, X. (2024). A comprehensive survey on automatic knowledge graph construction. *ACM Computing Surveys*, 56(4), Article 94, 1–62.
21. Cao, J., Fang, J., Meng, Z., & Liang, S. (2024). Knowledge graph embedding: A survey from the perspective of representation spaces. *ACM Computing Surveys*, 56(6), Article 159, 1–42.
22. Yandamuri, U. S. Next-Generation Industry Platforms Integrated Systems For Predictive & Intelligent Operations. *Solaris Vantage*.
23. Papadaki, M.-E., Tzitzikas, Y., & Mountantonakis, M. (2023). A brief survey of methods for analytics over RDF knowledge graphs. *Analytics*, 2(1), 55–74.
24. Venkataramanan, R., Tripathy, A., Foltin, M., Yip, H. Y., Justine, A., & Sheth, A. (2023). Knowledge graph empowered machine learning pipelines for improved efficiency, reusability, and explainability. *IEEE Internet Computing*, 27(1), 81–88.
25. Hu, Q., Yin, Z., Tao, T., Wang, J., Chen, Z., Ai, B., Liu, Y., & Liu, C. (2023). Research on technologies in data fabric. In 2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications. IEEE.
26. Mohan, A. A., Maguluri, K. K., Yasmeen, Z., & Pandugula, C. (2026). Evaluating the structural performance of axially loaded stainless steel circular tubular columns with ultra-high-performance concrete using an optimized attention pyramid convolutional neural network. *Structural Concrete*, 27(3), 3589-3608.
27. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6), Article 186345.
28. Li, X., Wang, S., Zeng, S., Wu, Y., & Yang, Y. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. *Vicinagearth*, 1, Article 9.
29. Gao, C., Lan, X., Li, N., Yuan, Y., Ding, J., Zhou, Z., Xu, F., & Li, Y. (2024). Large language models empowered agent-based modeling and simulation: A survey and perspectives. *Humanities and Social Sciences Communications*, 11, Article 1259.
30. Krishnan, M., Yandamuri, U. S., Recharla, M., Challa, K., Kolla, S. H., & Valiki, S. (2026, June). Adaptive Machine Learning Ecosystem for Personalized Medicine and Digital Finance: Leveraging Azure Cloud, Data Engineering, and Workflow Automation. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
31. Blohm, I., Wortmann, F., Legner, C., & Köbler, F. (2024). Data products, data mesh, and data fabric: New paradigm(s) for data and analytics? *Business & Information Systems Engineering*, 66, 643–652.
32. Cao, J., Fang, J., Meng, Z., & Liang, S. (2024). Knowledge graph embedding: A survey from the perspective of representation spaces. *ACM Computing Surveys*, 56(6), Article 159, 1–42.
33. Acharya, D. B., Kuppan, K., & Bhaskaracharya, D. (2025). Agentic AI: Autonomous intelligence for complex goals—A comprehensive survey. *IEEE Access*, 13, 18912–18936.
34. Singh, K. U., Gupta, S. K., Dubey, Y., & Mangalampalli, B. M. (2026, June). Enhancing Healthcare Cyber Defense with Machine Learning-Based Attack Prediction and Prevention. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
35. Abou Ali, M., Dornaika, F., & Charafeddine, J. (2025). Agentic AI: A comprehensive survey of architectures, applications, and future directions. *Artificial Intelligence Review*.
36. Brohi, S., Mastoi, Q.-U.-A., Jhanjhi, N. Z., & Pillai, T. R. (2025). A research landscape of agentic AI and large language models: Applications, challenges and future directions. *Algorithms*, 18(8), Article 499.
37. Kumar, R., Ishan, K., Kumar, H., & Singla, A. (2025). LLM-powered knowledge graphs for enterprise intelligence and analytics. *arXiv*.
38. Chen, A., Bunde, M., Ahlawat, G., Stetz, P., Wang, Z., Fei, Q., Jung, D., Chu, A., Jayaraman, B., Panth, A., Arora, Y., Jain, S., Varma, R., Ilin, A., Melnychuk, I., Chueh, C., Sil, J., & Wang, X. (2025). Text-to-SQL for enterprise data analytics. *arXiv*.
39. Viswanathan, V., et al. (2025). Agentic AI: A comprehensive survey of technologies, applications, and societal implications. *IEEE Access*, 13, 151824–151837.