

Reasoning-Driven Contextual Retrieval

Anumandla Mukesh

Independent Researcher
mukeshanumand@gmail.com

Abstract

Knowledge-Centric Enterprise Intelligence combines autonomous reasoning with contextual information retrieval to support decision-making that is transparent and evidence-based. Autonomous reasoning automates portions of enterprise decision-making, particularly for problems that recur, must be resolved in real time or at short intervals, or follow predictable patterns. Realizing its predictive potential requires reducing reliance on manual configuration, so that decisions are made in ways that account for uncertainty in the underlying causal models while remaining interpretable within the context of the system that produced them. While autonomous decision support shares similarities with case-based reasoning, its distinguishing focus is on managing the decision-making process itself, rather than merely generating a final recommendation or organizing support for human decision-makers. This work examines the goals, capabilities, limitations, and evaluation metrics of autonomous reasoning in business settings, comparing three paradigms: rule-based, probabilistic, and hybrid approaches.

Keywords: Enterprise Intelligence; Knowledge Management; Autonomous Reasoning; Contextual Information Retrieval; Ontological Foundations; Data Layer and Knowledge Layer Separation; Integration Methodologies Knowledge-Centric Enterprise.

1. Introduction

Knowledge-Centric Enterprise Intelligence integrates autonomous reasoning and contextual information retrieval to enhance decision-making with transparent, evidence-based processes. Enterprise Intelligence makes specific decisions based on situational circumstances, beyond purely predictive techniques. Advanced enterprise architectures share a structural dependency with Knowledge Graph solutions. Similarities with Knowledge Graph management mechanisms and contextual information retrieval are evident. Both approaches focus on dynamic databases and the interplay among the data management, knowledge management, decision support, and knowledge representation layers.

Table 1. Components of the Knowledge-Centric Enterprise Intelligence Framework

Component	Primary Function	Input	Output	Enterprise Benefit
Data Layer	Store structured enterprise data	Enterprise databases	Raw operational data	High-speed data access
Knowledge Layer	Maintain semantic knowledge graph	Processed enterprise data	Ontological knowledge	Semantic understanding

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 17

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 25



Component	Primary Function	Input	Output	Enterprise Benefit
Autonomous Reasoning Engine	Infer decisions using knowledge	Knowledge graph	Decision recommendations	Intelligent automation
Contextual Information Retrieval	Retrieve relevant contextual evidence	User query + context	Supporting information	Improved decision accuracy
Governance Module	Monitor policies and compliance	Enterprise policies	Compliance validation	Regulatory adherence
Provenance Manager	Track data lineage	Data transactions	Audit trail	Trustworthy AI

Autonomous Reasoning is contrasted with predictive reasoning. Goals, capabilities, and drawbacks of reasoning paradigms applicable to the enterprise decision-making context are summarized. Rule-based, probabilistic, and hybrid reasoning paradigms are examined. Complementary evaluation criteria across three distinct qualities required for Autonomous Reasoning—autonomy, explainability, and reliability—are defined.

A. Autonomous Reasoning Paradigms

- Knowledge and Vision Evaluation: Evaluation of the capability of complex systems to generate and update transparent Knowledge Bases that capture the Knowledge and Vision of the enterprise on any topic of interest

- Semantic Reasoning Evaluation: Analysis of the Novelty and Impact of Inferences made by a Semantic Reasoning Component that can independently reason and detect hidden valid knowledge not present in the Knowledge Base that could be either true or false

Other evaluations guarantee adequate performance, namely:

- Explanation Evaluation: Evaluation of the capability of a complex intelligent system to explain all its decisions and recommendations

- Privacy Evaluation: Evaluation of the Privacy Guarantees provided by the intelligent system and its components

Table 2. Comparison of Autonomous Reasoning Paradigms

Reasoning Paradigm	Decision Logic	Explainability	Handling Uncertainty	Computational Complexity	Enterprise Suitability
Rule-Based	IF-THEN Rules	Very High	Low	Low	Compliance Automation
Probabilistic	Bayesian Inference	Medium	Very High	Medium	Risk Prediction
Hybrid	Rules + Probabilistic Models	High	High	High	Enterprise Intelligence
Knowledge Graph Reasoning	Semantic Relationships	High	Medium	Medium	Knowledge Discovery

Reasoning Paradigm	Decision Logic	Explainability	Handling Uncertainty	Computational Complexity	Enterprise Suitability
Agentic Reasoning	Autonomous Multi-step Planning	High	High	High	Enterprise Automation

2. Foundations of Knowledge-Centric Enterprise Intelligence

Knowledge-Centric Enterprise Intelligence integrates autonomous reasoning and contextual information retrieval to enhance decision-making with transparent, evidence-based processes.

The Shaping Tomorrow’s World initiative aims to create an autonomous and contextual information-centric enterprise intelligence. Such systems support decision-making or action selection and reduce associated workloads. Operation is supported by evidence- augmented information retrieval. Exploring the autonomous reasoning aspect identifies automation goals and capabilities, considered primarily from an enterprise perspective. Next, the capabilities and limitations of reasoning paradigms are summarized and compared—rule-based, probabilistic, or hybrid systems. To facilitate meaningful comparisons and informed technology selection, metrics for autonomy, explainability, and reliability are defined.

Table 3. Data Layer and Knowledge Layer Comparison

Feature	Data Layer	Knowledge Layer
Storage Type	Relational Database	Knowledge Graph

Feature	Data Layer	Knowledge Layer
Data Representation	Tables	Entities & Relationships
Query Language	SQL	SPARQL
Reasoning Support	Limited	Extensive
Semantic Awareness	Low	High
Scalability	High	High
Update Strategy	Transaction-Based	Knowledge Synchronization
Primary Purpose	Data Storage	Knowledge Representation

Modern enterprises amass extensive data stores over time, but current approaches do not account for this depth of knowledge. Instead, proponents often assume and exploit a conceptual knowledge structure. Retaining this data access provides potential for augmented information retrieval from reference or contextually matched data. Data workloads can be increased by retrieval operations. Full integration of information-retrieval workloads permits action selection while preserving contextual coherence. The integration of provenance tracking and associated auditing directly aligns enterprise workload distribution with regulatory considerations.

A. Knowledge Representation and Ontologies

Data interpretation and sense-making in complex, dynamic worlds requires extra information beyond the available raw or pre-processed data. Knowledge represents the content of that extra information, which pervades and influences every aspect of intelligent reasoning, learning, and decision behaviour. While critical to intelligent

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 17

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 25

ISSN: 3067-1140



behaviour, it is notably one of the weakest elements in today's cognitively-inspired systems. Two requirements, among others, for a more convincing treatment of knowledge are 1) a proper degree of realism, with embedding ontological prior knowledge that finances richer, more reliable, faster learning; and 2) support for knowledge evaluation and validation, which is especially pertinent for large-scale, crowd-sourced knowledge bases.

With this foundation, the specific aspects of knowledge representation and ontologies that require treatment within Knowledge-Centric Enterprise Intelligence (KCEI)—simplified demonstrations of the benefits of knowledge-driven, rather than data-driven, reasoning as applied to EIs—are carefully outlined and motivated. These aspects include the design criteria underpinning ontologies, the specification of domain ontologies using conceptual schemas, the construction of domain-specific taxonomies and other kinds of ontologies, the principles governing the alignment of these specific ontologies with one another, and the advantages of leveraging upper ontologies to facilitate semantic interoperability among the data sources of the enterprise and the reasonable controllers thereof.

Table 4. Evaluation Metrics for Knowledge-Centric Enterprise Intelligence

Metric	Description	Measurement	Expected Objective
Autonomy Score	Degree of automated decision making	Percentage	Maximize
Explainability Score	Transparency of reasoning	Explanation coverage (%)	Maximize
Reliability	Correct decision rate	Accuracy (%)	High

Metric	Description	Measurement	Expected Objective
Retrieval Precision	Relevant context retrieval	Precision	High
Retrieval Recall	Information coverage	Recall	High
Semantic Consistency	Knowledge graph correctness	Consistency Index	High
Governance Compliance	Policy adherence	Compliance Rate	100%
Provenance Completeness	Lineage availability	Metadata Coverage	Maximize

3. Architectural Frameworks

To achieve Knowledge-Centric Enterprise Intelligence in a seamless manner it is beneficial to clearly separate the data layer of an information system from the knowledge layer. Such Architectural Frameworks are common in data lakes and knowledge graphs. In these architectures, the information stored in the data layer constitutes a Data Store frequently implemented in a traditional relational DBMS. The information stored in the Knowledge Store, usually an information store implementing a Knowledge Graph, is produced by a set of processes or procedures. These processes or procedures consume information from the Data Store and produce information for the Knowledge Store. These two layers are loosely coupled making it possible to operate and manage these layers independently. Communication is usually achieved source-producing SQL query against the Data Store and SPARQL query against the Knowledge Store. Since, due to the heterogeneity of the data sources composing the Data Store, information staled in the Data Store may become



inconsistent in short time, it is important to provide all the information necessary to keep the Knowledge Store consistent and up-to-date.

Mathematical Formulas:

1. Enterprise Intelligence Decision Function

$$D_t = f(K_t, C_t, R_t, G_t)$$

Where D_t is the enterprise decision, K_t is knowledge state, C_t is contextual information, R_t is reasoning output, and G_t is governance constraint.

2. Contextual Retrieval Score

$$CR(q, d) = \alpha \cdot Sim(q, d) + \beta \cdot Rel(d, C) + \gamma \cdot Trust(d)$$

Where q is query, d is document, C is context, and $Trust(d)$ is source reliability.

3. Autonomous Reasoning Output

$$R_t = \arg \max_{a \in A} P(a | K_t, C_t, E_t)$$

Where a is an action, K_t is knowledge base, C_t is context, and E_t is evidence.

4. Knowledge Graph Consistency

$$KC = 1 - \frac{N_{conflict}}{N_{total}}$$

Where $N_{conflict}$ is the number of inconsistent triples and N_{total} is the total number of knowledge triples.

5. Knowledge Layer Update Function

$$K_{t+1} = K_t + \Delta D_t + \Delta P_t$$

Where ΔD_t represents new data updates and ΔP_t represents provenance-enriched updates.

6. Provenance Reliability Score

$$PRS(x) = \frac{V_x + T_x + A_x}{3}$$

Where V_x is validity, T_x is traceability, and A_x is auditability of data item x .

7. Explainability Index

$$EI = \frac{N_{explained}}{N_{decisions}}$$

Where $N_{explained}$ is the number of decisions with explanations and $N_{decisions}$ is the total number of decisions.

8. Autonomy Level

$$AL = \frac{D_{auto}}{D_{total}}$$

Where D_{auto} is the number of autonomously executed decisions and D_{total} is the total number of decisions.

9. Retrieval-Augmented Reasoning Score

$$RAR = \lambda R_t + (1 - \lambda) CR_t$$

Where R_t is reasoning confidence, CR_t is contextual retrieval score, and λ controls reasoning–retrieval balance.

10. Decision Reliability

$$DR = Accuracy \times Trust \times Consistency$$

Where *Accuracy*, *Trust*, and *Consistency* measure the correctness, source reliability, and knowledge coherence of enterprise decisions.

Knowing the inconsistency of the information inside the Data Store, updating the information in the Knowledge Store deriving from the Data Store become crucial. Latency may be a problem in Knowledge Graph

Enrichment since processes may take time to produce the enriched Knowledge Graph. For this a caching mechanism can be implemented, keeping the enriched Knowledge Graph in memory. In this case during the caching time the un-enriched .

A. Data Layer and Knowledge Layer Separation

Intelligent agents rely on a combination of different sources of knowledge to answer queries in real time. The Knowledge-Centric Enterprise Intelligence approach separates commonly used data stores from expensive knowledge stores and determines a strategy to use both types of storage transparently. Imitating an enterprise environment, two storage layers are defined: a traditional data layer, which follows the relational data model, and a knowledge layer, where content is represented using an ontological model. These two storage mechanisms are persistently synchronized and interact with each other through specific operations. Such a division allows an efficient management of both types of storages, taking advantage of fast data access while still allowing expressive reasoning and inference.




Fig 1: Reasoning paradigm comparison — Bar chart contrasting rule-based, probabilistic, and hybrid systems across five key dimensions (autonomy, explainability, reliability, adaptability, scalability).

While both data stores and knowledge stores are used in parallel, the data layer usually answers most queries. This is possible because the files stored in the data layer are

normally accessed many times during a query. Therefore, the strategy consists of caching in the data layer all data required by the knowledge layer, either by proxying the queries or by preloading the data. When the query that requires the knowledge store does not return invalidation messages, this query response is cached in the data layer. When the amount of cached information in the data layer increases, a purge mechanism is executed to remove unused elements, it is also possible to track which elements of the knowledge layer are being used by the data layer, and therefore make the purge mechanism smarter.




Fig 2: KCEI architectural framework — Structural diagram of the two-layer architecture with the strict data/knowledge layer separation and caching middleware in between.

4. Methodologies for Integration

Knowledge-Centric Enterprise Intelligence capabilities are applied within enterprise systems, in particular the features of autonomous reasoning, to enable and improve information retrieval tasks when supported by known and trusted information sources. These features are enabled through the application of specific methodologies tailored for integration-level concerns, such as the construction of the underlying information ecosystem or the enforcement of data governance. Although the emphasis is placed on autonomous reasoning, aspects of contextual information retrieval and their interaction with knowledge

creation and exploitation functions are noted in some of the presented methodologies.




Fig 3: **ADE metric radar** — Radar chart of the four Autonomous Decision-making Evaluation sub-metrics, comparing target capability against a typical system baseline.

Implementation of these integrations begins with a focus on data governance and information provenance, two aspects that open and close the integration loop. Formal policies and guidelines ensure that data remains in the expected conditions throughout its lifecycle and that all resulting knowledge usage and modifications can be reliably traced, thus supporting the complex background knowledge indispensable for the formal verification of enterprise decisions and actions during subsequent provenance-based information retrieval operations.

5. Conclusion

Progressively maturing knowledge-centric enterprise intelligence that integrates autonomous reasoning with contextual information retrieval is emerging. Autonomous reasoning aims to support complex decision-making processes by inferring conclusions from sparse information with minimal to no human intervention. Emphasis has shifted from the cognitive capacities of such systems to the cognitive and explicit aspects of their operation. Contextual information retrieval aims to enhance, enrich, or explain the decisions made by autonomous reasoning through the targeted retrieval of supplementary

information stored in structured and semi-structured databases rather than unstructured data sources.

The performance of autonomous reasoning is commonly evaluated against a set of metrics defined by the deployed reasoning paradigm, such as accuracy and reliability, explainability, and responsiveness. A knowledge-based system designed for enterprise environments should fulfil the following requirements: operation at varying levels of autonomy; transparency in its reasoning processes; a complete description of its assumptions and data at the time of decision-making; and a clear delineation of when and why supplementary information is retrieved.

References

1. Bonifacio, L. H., Abonizio, H. Q., Fadaee, M., & Nogueira, R. (2022). InPars: Data augmentation for information retrieval using large language models. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2345–2358.
2. Gao, L., Dai, Z., & Callan, J. (2021). COIL: Revisit exact lexical match in information retrieval with contextualized inverted list. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3030–3042.
3. Pranitha Potturi. (2023). Cloud-Driven Business Intelligence Modernization for Scalable Enterprise Data Integration and Decision Support. *International Journal of Intelligent Systems and Applications in Engineering*, 11(3s), 387 –. Retrieved from <https://www.ijisae.org/index.php/IJISAE/article/view/8505>
4. Gao, L., Ma, X., Lin, J., & Callan, J. (2023). Precise zero-shot dense retrieval without relevance labels. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 1762–1777.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 17

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 25



5. Izacard, G., Lewis, P., Lomeli, M., Hosseini, L., Petroni, F., Schick, T., Dwivedi-Yu, J., Joulin, A., Riedel, S., & Grave, E. (2023). Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24, 1–43.
6. Li, X., Zhu, Y., Liu, S., Ju, J., Qu, Y., & Cheng, G. (2023). DyRRen: A dynamic retriever-reranker-generator model for numerical reasoning over tabular and textual data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11), 13139–13147.
7. Sasidhar Reddy Mondeddula. (2023). Natural Language Processing for Intelligent Automation of Financial Documents and Banking Operations. *International Journal of Intelligent Systems and Applications in Engineering*, 11(8s), 605–612. Retrieved from <https://www.ijisae.org/index.php/IJISAE/article/view/8527>
8. Qu, Y., Ding, Y., Liu, J., Liu, K., Ren, R., Zhao, W. X., Dong, D., Wu, H., & Wang, H. (2021). RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 5835–5847.
9. Ram, O., Shachaf, G., Levy, O., Berant, J., & Globerson, A. (2022). Learning to retrieve passages without supervision. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2687–2700.
10. Ren, R., Lv, S., Qu, Y., Liu, J., Zhao, W. X., She, Q., Wu, H., Wang, H., & Wen, J.-R. (2021). PAIR: Leveraging passage-centric similarity relation for improving dense passage retrieval. *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2173–2183.
11. Mattaparthi, R. (2023). Deep Learning-Driven Combustion Anomaly Detection in Diesel Powertrains: A Multi-Sensor Fusion Approach for Real-Time ECM Adaptation. *International Journal of Intelligent Systems and Applications in Engineering*, 11, 1084.
12. Santhanam, K., Khattab, O., Saad-Falcon, J., Potts, C., & Zaharia, M. (2022). ColBERTv2: Effective and efficient retrieval via lightweight late interaction. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3715–3734.
13. Sachan, D. S., Lewis, M., Yogatama, D., Zettlemoyer, L., Pineau, J., & Zaheer, M. (2023). Questions are all you need to train a dense passage retriever. *Transactions of the Association for Computational Linguistics*, 11, 600–616.
14. Trivedi, H., Balasubramanian, N., Khot, T., & Sabharwal, A. (2023). Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 10014–10037.
15. Kolla, S. K., & Reddy, V. A. R. (2023). Deep Learning Architectures For Multimodal Medical Data Integration. *South Eastern European Journal of Public Health*, 248-260.
16. Arora, S., Lewis, P., Fan, A., Kahn, J., & Ré, C. (2023). Reasoning over public and private data in retrieval-based systems. *Transactions of the Association for Computational Linguistics*, 11, 902–921.
17. Wang, L., Yang, N., & Wei, F. (2023). Query2doc: Query expansion with large language models. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9414–9423.
18. Dai, Z., Zhao, V. Y., Ma, J., Luan, Y., Ni, J., Lu, J., Bakalov, A., Guu, K., Hall, K. B., & Chang, M.-W. (2022). Promptagator: Few-shot dense retrieval from 8 examples. *arXiv preprint arXiv:2209.11755*.
19. Gao, L., Ma, X., Lin, J., & Callan, J. (2022). Precise zero-shot dense retrieval without relevance labels. *arXiv preprint arXiv:2212.10496*.
20. Formal, T., Piwowarski, B., & Clinchant, S. (2021). SPLADE: Sparse lexical and expansion model for first stage ranking. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2288–2292.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 17

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 25



21. Ren, R., Qu, Y., Zhao, W. X., Liu, J., She, Q., Wu, H., Wang, H., & Wen, J.-R. (2021). RocketQA v2: A joint training method for dense passage retrieval and passage reranking. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 1767–1777.
22. Ram, O., Shachaf, G., Levy, O., Berant, J., & Globerson, A. (2022). Learning to retrieve passages without supervision. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2687–2700.
23. Loganathan, R. (2022). Converging Security Architecture and Compliance Management in Enterprise Data Center Ecosystems: A Unified Control Framework. *International Journal of Scientific Research and Modern Technology*, 1(12), 295–312.
24. Santhanam, K., Khattab, O., Saad-Falcon, J., Potts, C., & Zaharia, M. (2022). ColBERTv2: Effective and efficient retrieval via lightweight late interaction. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3715–3734.
25. Izacard, G., Lewis, P., Lomeli, M., Hosseini, L., Petroni, F., Schick, T., Dwivedi-Yu, J., Joulin, A., Riedel, S., & Grave, E. (2023). Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24, 1–43.
26. Asai, A., Wu, Z., Wang, Y., Sil, A., & Hajishirzi, H. (2023). Self-RAG: Self-reflective retrieval augmented generation. *NeurIPS 2023 Workshop on Instruction*, 1–12.
27. Arora, S., Lewis, P., Fan, A., Kahn, J., & Ré, C. (2023). Reasoning over public and private data in retrieval-based systems. *Transactions of the Association for Computational Linguistics*, 11, 902–921.
28. Wang, L., Yang, N., & Wei, F. (2023). Query2doc: Query expansion with large language models. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9414–9423.
29. Gao, L., Dai, Z., & Callan, J. (2021). COIL: Revisit exact lexical match in information retrieval with contextualized inverted list. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3030–3042.
30. Bharath Penaka. (2023). AI-Driven Predictive Analytics for Early Detection of High-Risk Patients in Cloud-Based Healthcare Systems . *International Journal of Medical Toxicology and Legal Medicine*, 26(3 and 4), 118–124. Retrieved from <https://ijmtlm.org/index.php/journal/article/view/1506>
31. Qu, Y., Ding, Y., Liu, J., Liu, K., Ren, R., Zhao, W. X., Dong, D., Wu, H., & Wang, H. (2021). RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 5835–5847.