

Agentic Business Intelligence Architecture for Autonomous Data Exploration, Insight Generation, and Strategic Decision Orchestration

Pranitha Potturi

¹Sr. Application Developer, 10009-0002-2673-809X, pranithapotturi@gmail.com

Abstract: Business Intelligence (BI) is critical for supporting tactical, operational, and strategic decision-making in enterprises. Traditionally, BI solutions have focused on creating data use dictionaries that allow decision leaders to collect and reference those data narratives that are most suited to their decision making. Contemporary digital business ecosystems generate immense data streams, and the creation of dictionaries for these streams is impractical and insufficient. Autonomous BI that allows autonomous data exploration and narrative generation is therefore a necessity. To meet this need, the Decision-Orchestration Adaptive Business Intelligence (DOABi) architecture enables the formation of autonomous decision-supporting BI nodes. These nodes connect to existing data streams with little or no dependence on human or administrative processing. Decisions needing support are inquired through BI agent entities that serve the interest of the decision-maker. A formal exploration model specifies the autonomy and supporting mechanisms of such service providers. These service BI agents operate at three cognitive levels and can ingest and explore data under a controlled-adaptive mechanism. The architecture is supported by an illustrative live enterprise case study and is further strengthened by literature analysis and industrial, academic, and entrepreneurship expert opinions. Data exploration is not new to data science. Making discovery easier, faster, and more fully realized is the question of the day. Meyer et al. [2018] highlight the shortcomings of traditional BI systems, which have focused on creating use dictionaries for data resources. Data streams generated by information systems, IoT devices, and social media, along with imagery and video captured by drones and mobile devices, evolve too rapidly for such dictionaries to be a fitting solution. The exploration of such uncharted data sources is better left to intelligent agents and the AI behind them. In a framework for autonomous business decision orchestration, autonomous exploration is a major enabler.

Keywords: Data exploration, agent-based modeling, business intelligence, exploratory query generation, data insights generation, information autonomy

1. Introduction

The Business Intelligence (BI) architecture enables the autonomous Business Data Analysts to explore the organization BI environment offered by the Architecture and discover insights to support their business operations. Such autonomous insights exploration can be achieved with a foundational architecture that automates the ingestion and integration of relevant external sources and that supports the autonomous data exploration. An Agentic BI Architecture supports the automatic ingestion and integration of external relevant data sources that can complement companies' internal data and enable a proper autonomous exploration.

An Agentic Business Intelligence Architecture enables Business Data Analysts to interactively explore the data space supported by the joint capability of traditional Business Intelligence environments and AI models. The Business Data Analyst enjoys complete autonomy in choosing the attitude towards exploration, from a more cautious approach, driven mainly by queries and dashboards, to a more exploratory attitude, relying mostly on natural language interfaces and exploratory paths. The Architecture also offers autonomy at the Cognitive Autonomy Control Level, allowing the system to choose the attitude towards exploration along the interaction process.



Table I. Cognitive Autonomy Levels in the DOABi Architecture

Level	Designation	Core Capability	Human Role
L1	Reactive	Responds to explicit requests: query answering, monitoring, why-questions.	Directive
L2	Semi-Autonomous	Suggests tasks and actions on request; assists when the system flags a need.	Supervisory
L3	Autonomous Exploration	Explores data without human input; infers insights and candidate hypotheses.	Consultative
L4	Decision Orchestration	Detects patterns, forms hypotheses, and reaches decisions guiding managerial action.	Oversight-only

2. The Concept of Agentic BI in Modern Enterprise

In the 1970s and 80s, during the times when business intelligence began to emerge as a major topic in the research and business community, computer systems were at a rudimentary stage, and analytical capabilities were limited. Today, technology advancements are extremely rapid, and the scientific community has also improved its data exploration and analysis capabilities thanks to natural language processing and transformative AI technologies. An autonomously acting data discovery and analysis agent is within technical capabilities. Such an agent autonomously explores the data, constructs models, makes forecasts, generates insights and recommendations, and orchestrates the related actions. Support for making strategic decisions offers tremendous business value, and the first commercial implementations of these ideas are already available. Automating simple tasks has been widely adopted, and the capability spectrum of agents is gradually increasing. In the future, research will focus on higher levels of cognitive autonomy.

The term “Agentic Business Intelligence” refers to an integration framework for rapid and autonomous data exploration, insight generation, and decision orchestration. Business intelligence systems should exploit the benefits of agentic frameworks to create an autonomous agent exploring the enterprise data assets and assembling forecasts, recommendations, and action plans. Agentic BI integrates the principles of self-configuring, self-tuning, and self-healing systems, enhancing information flow and synergy among business information framework components.

2.1 Table II. Comparative Positioning of DOABi

Dimension	Traditional BI	Agentic BI (general)	DOABi (proposed)
Query initiation	User-driven	Mixed user/agent	Agent-driven, event-triggered
Data dictionary dependence	High	Reduced	Minimal to none
Exploration scope	Predefined dashboards	Guided exploration	Open-ended, adaptive exploration
Decision role	Advisory only	Advisory / suggestive	Advisory through deciding (L1-L4)
External data ingestion	Manual / scheduled ETL	Semi-automated	Continuous, relevance-scored

3. Architectural Foundations

Despite rapid progress in sensing, data generation, and storage capacities, the majority of BI users do not fully exploit the available data due to skills shortages and the level of effort required to find valuable insights. Business decisions supported by BI are more novel when users can directly explore and query data themselves or by using intelligent assistants, with both skills and time constraints considered. Increased autonomy through automatic data

exploration without requiring explicit user input has the potential to reveal new facts and generate insights beyond user-driven exploration. The candidate architecture consists of free-standing, cognitively autonomous, self-interested agents responsible for automatic and periodic data exploration and automated insight generation; it ends when a query-based decision-support intervention is triggered.

Business events fulfil the contexts of causality, consequence, intentionality, and time related to the triggering of literary information-seeking and confidence-exploration queries. Continuous-loop or exploratory BI systems seek breakthroughs, not specific answers; neither model explanatory data exploration into complex systems, and neither considers the approximation of portfolio theory potentially lowering the total investment risk without compromising expected return. Progress thus far also discusses novel concepts related to the architectural foundations of Agentic BI's cognitive autonomy, particularly operationalizing the theoretical formalization of technology-assisted auto-exploratory decision support.

Mathematical Formulas:

$$C(t) = \operatorname{argmax}_{k \in \{1, 2, 3, 4\}} [w_k \cdot \Phi_k(R, S, A, D)] \quad (1)$$

Eq. (1) — Cognitive autonomy level selection: the active level $C(t)$ is the level k whose weighted response function Φ_k over reactive (R), semi-autonomous (S), autonomous-exploration (A), and decision (D) signals is maximal (Section IV).

$$Rel(d) = \alpha \cdot Sim(d, D_{exist}) + \beta \cdot Freq(d) + \gamma \cdot Nov(d) \quad (2)$$

Eq. (2) — Dataset relevance score used during ingestion and integration: a weighted combination of similarity to existing datasets, access frequency, and novelty (Section III-A).

$$U(q) = E[Insight(q)] - \lambda \cdot Cost(q) \quad (3)$$

Eq. (3) — Expected exploration utility of a candidate query q , trading off expected insight value against exploration cost (Section V-A).

$$Decide(q) = 1 \text{ if } U(q) \geq \theta_c, \text{ else } Explore(q) \quad (4)$$

Eq. (4) — Decision-orchestration rule: the agent commits to a decision once utility $U(q)$ clears confidence threshold θ_c ; otherwise it continues exploring (Section IV, L4).

A. Data Ingestion and Integration

Systems for ingesting and integrating new data should possess cognitive competencies on par with those exhibited by individuals occupied with or supervising this task. Cognitive agents in human organizations often act to satisfy itching data needs without imposing additional burden on the organization's data team. The interlocutor can freely ingest new data with whatever effort it requires. Upon ingesting the new dataset, the system determines a large set of questions concerning the new data while also increasing existing datasets' relevance scores. Given these two cognitive methods, the ingestion-and-integration mechanism needs to be designed to automatically query the new dataset for frequently-less-frequent-outside-frequent-occasional-outlier patterns, in addition to the newly-formed combinations with existing datasets.

Cognitive and autonomous needs for exploration and recommendation of integration partners should also be fulfilled. When a key decision area is about to

become unstable, an organization's peers or competitors that also possess some share of overall relevance should be suggested to allow strategic decision-making outside the organization. Recommendations for new data sources to be ingested, developed, or integrated should follow the same out-of-used-categories-of-all-datasets system logic: occasional patterns with datasets that are presently far from the cognitive fringes. Although these systems run independently and do not need control for ordinary operation, there are still cases where final conclusions require assistance from traditional BI systems.

4. Autonomy and Control Mechanisms

The agentic BI architecture, supporting autonomous operation, requires advanced management mechanisms. Control management must address the autonomy degree, security and access rights, as well as interaction with traditional BI systems. Four levels of cognitive autonomy are proposed: Reactive systems respond to explicit requests, including query answering, monitoring, and answering why-questions. Task assignment and action suggestion are outputs of semi-autonomous systems, available when a system needs assistance. The higher level of cognitive autonomy enables BI systems to venture into the unknown, autonomously explore the data without human input, and infer insights. The final level elevates BI systems into a deciding role, allowing them not only to look at the data, detect patterns, and create hypotheses, but also to reach decisions based on them, guiding managerial actions and remedies.

The control management assures the functional characteristics and security of the exploration and discovery system's operation. These capabilities comprise content and metadata protection, control of access to the data and knowledge domains, monitoring of exploration and discovery paths, and the formulation of alerts in conditions assessed as critical or demanding intervention.




Fig. 1. Conceptual illustration of the inverse relationship between human dependence and system autonomy across the four DOABi cognitive levels (Section IV-A). Values are illustrative, not measured.

A. Cognitive Autonomy Levels

Agentic Business Intelligence Architecture for Autonomous Data Exploration, Insight Generation, and Strategic Decision Orchestration—present a rigorous, evidence-based, formal analysis with clear argumentation and scholarly references.

Data explorability is the first layer of cognitive autonomy, thus the simplest, and is of the declarative character. Autonomous explorability of data means that it is possible to ask a question about the data and receive a complete or partial answer without human intervention. There are two key related aspects: question formulation and query execution. Natural language facilities can assist with question formulation, enabling interrogative statements to be represented in a machine-computable demonstrative form. The resulting support system can be regarded as a task-oriented natural language interface. The

second aspect is query execution, the optimal approach to which depends on the nature of the questions. A complete, nonrecursive plan can be derived from knowledge of the data structure, including data locations.

Data exploration on top of data explorability is the second layer of cognitive autonomy. It is still in the declarative mode for questions that can be formulated with no more than a single query, but exploratory exploration mainly identifies interesting regions in the data space that are candidates for follow-up questions. In broader terms, the aim is to perform the uncertain task of identifying empirically interesting regions of the data space. The results can guide the formulation of interesting-query candidates and be interpreted by a person as hints, guidelines, or recommendations for further exploration.




Fig. 2. Conceptual capability comparison of DOABi against traditional BI and general agentic BI across six qualitative dimensions (1=Low, 5=High), reflecting the architectural discussion in Sections II–V.

5. Data Exploration in Practice

Operationalizing the workforce as agentic actors also enables their data exploration processes to cater to the need for autonomous self-service BI. The goal is to facilitate autonomous exploration of large multidimensional data sets without the necessity of SQL query declaration while assuring sufficient control over query formulation and execution without the operational burden of query management. Autonomous, adaptive, self-service exploration combines the merit of BI-driven systems with user-friendly interfaces.

The system should be capable of automatically generating SQL queries that enable traversing through a large multidimensional data cube. Such abilities mimic data exploration by users equipped with OLAP tools. Intelligent ornamentation of an unrestricted natural language interface to a relational database can make even basic users of the system comfortable. Further, the empowered user should be able to specify their data exploration wishes as natural language query patterns and the system should be able to automatically explore the data set and determine query answers that reflect their wishes as in a data mining process. This is achieved by a synergy of techniques from intelligent agents, automated natural language processing, data mining, and OLAP.

A. Autonomous Querying and Exploration Strategies

The Digital Enterprise Architecture Model (DEAM) allows for the independent and autonomous implementation of the agentic business intelligence function. The concept of autonomous data exploration (ADE) is proposed as an example of applying the agent-based paradigm to solve specific tasks within the agentic business

intelligence architecture.

The goal of agent-based business intelligence architecture is to provide an environment for autonomous discovery, retrieval, exploration, integration, aggregation, and management of data. An autonomous data exploration strategy enables the automated generation of analytical queries based on predefined heuristics, triggering the exploration of different databases linked to the business intelligence architecture and the analysis of tables, metadata, data characteristics, and relationships. Agent cognitive emotions complement existing machine-learning and neural network-based methods in addressing business-relevant queries via autonomous systems that learn from experience within the business context.

6. Conclusion

No matter how a business operates, agentic Business Intelligence solutions can help discover hidden key information and insights. Through agentic Business Intelligence, even the smallest business can find the specific knowledge it requires from the vast amounts of data available on the internet. Due to the sheer volume of information available, it is impossible for a human user to manually search for it. Knowledge is not limited to databases controlled by the business. Unstructured data found in various places can enrich the contents of a Business Intelligence system and discover hidden business opportunities or strategies. With agentic Business Intelligence solutions, a strategy is generated and executed automatically.

The developed architecture can incorporate external data not controlled by the business: the infrastructure scans data sources on the internet, checks the relevance of the new information for the business, enhances the current knowledge reservoir if necessary, formulates corresponding queries, sends them, and interprets the answers. If needed, the structure automatically explores the data using procedures known from Artificial Intelligence. Natural Language Processing techniques can discharge relevance checks and an overall pre-filtering of the amassed knowledge to reduce the set to be explored. Using natural language, businesses can use their Business Intelligence system as an advisor, explorer, and knowledge generator, and let the skeleton act completely independently.

REFERENCES

1. 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, 1259.
2. Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). Large language model based multi-agents: A survey of progress and challenges. *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 8048–8057.
3. Handler, A., Larsen, K. R., & Hackathorn, R. (2024). Large language models present new questions for decision support. *International Journal of Information Management*, 79, 102811.
4. Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, J., Zhang, C., Wang, Z., Yau, S. K. S., Lin, Z., Zhou, L., Ran, C., Xiao, L., Wu, C., & Schmidhuber, J. (2024). MetaGPT: Meta programming for a multi-agent collaborative framework. *Proceedings of the Twelfth International Conference on Learning Representations*.
5. Krishnamurthy, A., Harris, K., Foster, D. J., Zhang, C., & Slivkins, A. (2024). Can large language models explore in-context? *Advances in Neural Information Processing Systems*, 37.
6. Kromidha, E., & Davison, R. M. (2024). Generative AI-augmented decision-making for business information systems. In R. M. Davison & D. Kreps (Eds.), *Human choice and computers: 16th IFIP International Conference on Human Choice and Computers, HCC 2024* (pp. 46–55). Springer.
7. Qian, C., Liu, W., Liu, H., Chen, N., Dang, Y., Li, J., Yang, C., Chen, W., Su, Y., Cong, X., Xu, J., Li, D., Liu, Z., & Sun, M. (2024). ChatDev: Communicative agents for software development. *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, 15174–15186.
8. 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, 186345.
9. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Awadallah, A. H., White, R. W., Burger, D., & Wang, C. (2024). AutoGen: Enabling next-gen LLM applications via multi-agent conversation. *Proceedings of the First Conference on Language Modeling*.
10. Wu, T., Li, J., Bao, J., & Liu, Q. (2024). ProcessCarbonAgent: A large language models-empowered autonomous agent for decision-making in manufacturing carbon emission management. *Journal of Manufacturing Systems*, 76, 429–442.
11. Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I.,

- Brooks, L., Buhalis, D., Carter, L., et al. (2023). "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642.
12. Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66, 111–126.
 13. Jarrahi, M. H., Lutz, C., & Newlands, G. (2022). Artificial intelligence, human intelligence and hybrid intelligence based on mutual augmentation. *Big Data & Society*, 9(2).
 14. Mikalef, P., Conboy, K., Lundström, J. E., & Popovič, A. (2022). Thinking responsibly about responsible AI and "the dark side" of AI. *European Journal of Information Systems*, 31(3), 257–268.
 15. Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210.
 16. Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66–83.
 17. Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637–643.
 18. Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586.
 19. Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of big data: Evolution, challenges and research agenda. *International Journal of Information Management*, 48, 63–71.
 20. Mikalef, P., Pappas, I. O., Krogstie, J., & Giannakos, M. (2018). Big data analytics capabilities: A systematic literature review and research agenda. *Information Systems and e-Business Management*, 16(3), 547–578.
 21. Mikalef, P., Boura, M., Lekakos, G., & Krogstie, J. (2019). Big data analytics capabilities and innovation: The mediating role of dynamic capabilities and moderating effect of the environment. *British Journal of Management*, 30(2), 272–298.
 22. Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. *Journal of Business Research*, 70, 356–365.
 23. Akter, S., Wamba, S. F., Gunasekaran, A., Dubey, R., & Childe, S. J. (2016). How to improve firm performance using big data analytics capability and business strategy alignment? *International Journal of Production Economics*, 182, 113–131.
 24. Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165–1188.
 25. Elbashir, M. Z., Collier, P. A., & Davern, M. J. (2008). Measuring the effects of business intelligence systems: The relationship between business process and organizational performance. *International Journal of Accounting Information Systems*, 9(3), 135–153.
 26. Popovič, A., Hackney, R., Coelho, P. S., & Jaklič, J. (2012). Towards business intelligence systems success: Effects of maturity and culture on analytical decision making. *Decision Support Systems*, 54(1), 729–739.
 27. Trieu, V. H. (2017). Getting value from business intelligence systems: A review and research agenda. *Decision Support Systems*, 93, 111–124.
 28. Wieder, B., & Ossimitz, M. L. (2015). The impact of business intelligence on the quality of decision making: A mediation model. *Procedia Computer Science*, 64, 1163–1171.
 29. Cao, G., Duan, Y., & Li, G. (2015). Linking business analytics to decision making effectiveness: A path model analysis. *IEEE Transactions on Engineering Management*, 62(3), 384–395.
 30. Shollo, A., & Galliers, R. D. (2016). Towards an understanding of the role of business intelligence systems in organisational knowing. *Information Systems Journal*, 26(4), 339–367.