P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

A Distributed Execution Framework for Enterprise Decision Systems

Katarzyna Nowak
Asst Professor

Abstract

Enterprise governance of Decision Intelligence enables organizations to make decisions at scale, responsibly, in a world full of both challenges and opportunities. The adoption of Distributed AI Execution Frameworks has made such scalable decision intelligence achievable in practice; however, governing these distributed frameworks properly remains an open problem. Governance mechanisms are essential for mitigating the associated risks and ensuring that accountability and transparency requirements are met. Policy-driven decision governance addresses the specification, enforcement, compliance, override, and audit of decision processes — enabling enterprises to realize the full potential of this approach.

Decision Intelligence (DI) reflects a growing effort among organizations to address the many factors that shape decision-making in a holistic way. As a metadiscipline, DI brings together Data-Driven Decision-Making and Machine Learning-Based Decision-Making, ensuring that neither happens in isolation but instead forms part of a coherent, well-considered strategy. Enterprise governance of DI thus enables decision-making that is scalable, transparent, auditable, and responsible — a capability sorely needed in a world that demands swift, competent decisions at every level of society.

Keywords :Decision intelligence, enterprise governance, distributed AI, execution frameworks, orchestration, policy governance, scalability ,Scalable Decision Intelligence,Enterprise Governance,Distributed Artificial Intelligence,AI Execution Frameworks,Intelligent Decision Support Systems,Distributed Machine Learning,Enterprise AI Architecture,Governance Automation,Adaptive Decision Analytics,Explainable and Trustworthy AI.

1. Introduction

Decision intelligence is maturing into a full-fledged discipline, encompassing methodical decision analysis, the explicit specification of decision-making objectives including quality assurance, the assessment of data quality and provenance, considering bias and distortion, an exploration of alternative decision strategies and their combination, and ultimately the responsible deployment of decision-making systems, without losing sight of their effects on enterprise governance and risk management. A formal and comprehensive treatment of these concepts in the context of enterprise activities is lacking. Critical aspects include the quality and provenance of data informing decisions, the evaluation of past decisions and their alignment with business objectives, a suitable theoretical basis, and support for external governance at the enterprise level.

Decision Intelligence Function

$$DI = f(DQ, AI, G)$$

Where DQ= data quality, AI= artificial intelligence model, and G= governance control.

Decision intelligence examines the quality and governance of decision processes based on supporting data. Artificial intelligence—including machine learning, deep

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

learning, and planning—is now augmenting or substituting human decision making in many enterprise application areas. Regulatory scrutiny of AI systems is increasing, and large enterprises are deploying comprehensive governance frameworks. Policy-based execution mechanisms are necessary for the transparent and auditable use of AI components in production systems

Table 1. Core Concepts of the Proposed Framework

Component	Short Description
Decision Intelligence	Supports structured and responsible enterprise decision-making
Distributed AI Execution	Executes AI tasks across scalable distributed systems
Enterprise Governance	Ensures decisions follow policies, risk limits, and accountability
Policy Governance	Defines, enforces, audits, and controls decision rules
Explainable AI	Improves transparency and trust in AI-driven decisions

2. Foundations of Decision Intelligence in Enterprises

Decision Intelligence (DI) describes the application of machine learning and AI technologies to the decision-making processes within organizations. While academia has focused on technical aspects of various branches of AI/ML individually and problems such as safety, security, or explainability, enterprises must view these as an interconnected whole forming a Decision Intelligence ecosystem. Such an ecosystem must ensure sufficient data quality and provenance, measure the effect of decisions on key business metrics, align the actual decision-making activity and outcomes with governance objectives, and detect and allow mitigation of decision-making risks across the full breadth of the enterprise. Furthermore, such alignments and governance controls must equally support the consumption of machine-generated insights in human decision-making, the end-to-end orchestration of AI/ML pipelines, and the modular implementation of AI/ML as components within enterprise Solution Architectures.

Governance Score

$$GS = R + A + T$$

Where R= risk control, A= accountability, and T= transparency.

Distributed AI execution frameworks are a promising technology for transforming AI from a technological prejudice into an organizational capability through the modular but comprehensive and scalable realization of DI ecosystems. A DI ecosystem consists of four key aspects: the primitives needed to implement the full breadth of the Decision Intelligence ecosystem, the coordination model employed to realize the primitives and their interaction patterns, the metadata needed to support feature discovery, execution coordination, and communication and the architectural considerations required to integrate deployment within the enterprise IT ecosystem in a manner that properly supports the broader organizational objectives.

Table 2. Governance Requirements for Scalable AI

Requirement	Purpose
Accountability	Identifies responsibility for decisions and outcomes

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

Requirement	Purpose
Transparency	Makes decision processes understandable
Traceability	Tracks decision history for audits
Compliance	Ensures alignment with legal and regulatory rules
Risk Control	Prevents unacceptable enterprise risk exposure

3. Distributed AI Execution Frameworks: Concepts and Architectures

Distributed AI execution frameworks combine lower execution overhead for individual tasks with fault-tolerance, low-latency, and very large-scale potential by breaking a computational process into a set of elementary tasks and binding them together with an arbitrary shape. These frameworks are typically used to automate machine learning-based decisions, which are often made with reduced or no human input because of the scale, frequency, or context of the decision instances. However, such large-scale automation is usually perceived as risky by organizations because of the typical lack of governance controls in the AI execution environment, with decisions being considered black boxes. The potential impact of decision execution or service quality often depends on both decision quality and IT service quality; the former is typically modeled as a black box.

Decision Quality

$$DQ = \frac{C + P + A}{3}$$

Where C= consistency, P= precision, and A= accuracy.

DECISION INTELLIGENCE IN ENTERPRISES examines how organizations are accountable for their decisions at the right time, place, and context. Five research topics are derived: data quality and provenance for decision-making; introducing a measurement framework for risk quantification; aligning decision impact with enterprise governance objectives; investigating foundations of decision quality; and providing an account of decision-making process quality.

Table 3. Distributed AI Execution Elements

Element	Role in Framework
Tasks	Break large decision processes into smaller units
Events	Trigger execution and communication between components
Policies	Control execution behavior and governance rules
Data Streams	Carry input and output data across systems
Metadata Catalogues	Support task discovery and execution coordination

A. Core primitives and orchestration models

The core primitives supporting the execution of a distributed artificial intelligence execution framework are constituent tasks, events, policies governing non-functional properties of the execution, along with data streams flowing into and out of the framework. Execution is typically supported through concepts such as task orchestration and choreography, where orchestration is defined via an explicit controller and choreography involves direct interaction among operational components.

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

Risk-Aware Decision Output

$$D_o = M(x) - R_s$$

Where $M(x)$= model prediction and R_s= risk score.

Distributed execution across a substantially larger pool of inexpensive computational resources makes the solution more scalable and fault-tolerant. Solutions typically deploy workload strategies that leverage this property building fault-tolerant bespoke engines with appropriate latency and service-quality considerations.

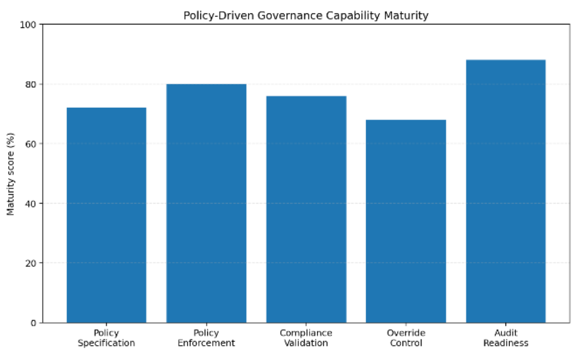

4. Governance Mechanisms for Scalable AI

Meeting risk appetite, compliance, and transparency requirements of enterprises is indispensable for enabling scalable AI execution. Risk management policies specify the level of risk deemed acceptable for an enterprise, sector, or region, and define processes to manage unacceptable risk exposures. Decision accountability relates to the attribution of responsibility for decisions, both ex-ante (establishing who was authorized to decide on a given matter) and ex-post (attribution of the outcome of the decision). Traceability supports evaluating the quality of decisions and consequences, enabling inquiry and institution corrections and learning from past mistakes. Finally, transparent decision-making procedures and decision-rationale disclosure are fundamental to enhance public trust in enterprises.

Policy Compliance

$$PC = \frac{N_c}{N_t}$$

Where N_c= compliant decisions and N_t= total decisions.

Governance issues are attracted once the scale of operations subsidizes the use of AI systems that are not individually but collectively understandable. Increasingly, budgets are allocated for computing resources, to develop and deploy large AI models (LLMs) serviced through APIs (Wordt et al. 2023). Incorporating sound governance practice will minimize reputational damage from occasional well-publicized blunders, and prevent incidents with potentially high-impact on individuals or society at large.

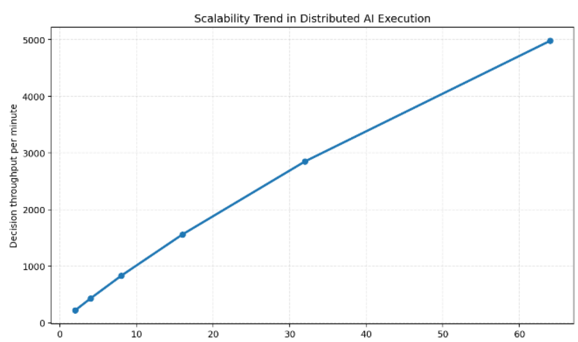

A. Policy-driven decision governance

Governance objectives for policy-driven decision intelligence encompass several parameters, including risk management for potential adverse decision repercussions, accountability for decision actions and outcomes, external transparency enabling scrutiny and oversight, and internal traceability supporting audits and investigations. Although the orientation towards clearly defined governance merits and objectives should favour a predominantly top-down architecture for distributed execution frameworks, it is

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

acknowledged that the underlying patterns employed may also require bottom-up policy-choreographed architectures such as those catering to electric power grid control.

Distributed Execution Efficiency

$$E_d = \frac{T_s}{T_d}$$

Where T_s= single-system execution time and T_d= distributed execution time.

Enforcement of governance policies governing decision-making processes generally assumes the presence of certain preconditions.2 Decision processes must be defined as fulfilment functions that return the actions to undertake for different decision input parameter observations. These may or may not also validate the decisions against policy checks. The decision functions must ensure compliance with applicable laws and regulations. Systems involved in the management of sensitive personal data must act in accordance with privacy protection policies, especially those addressing data usage consent. In determining the actual action to undertake for a decision, decision process functions may internally overload the necessary controls to comply with other applicable regulations.

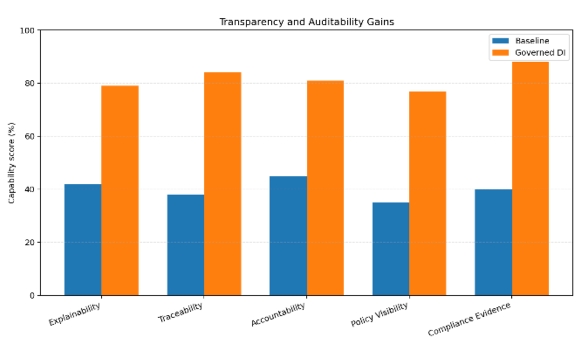

5. Conclusion

A synthesis of the decision intelligence enterprise governance foundations and investigation of policy-driven decision governance demonstrate the significance of a policy-driven approach to the management of decision intelligence and the deployment of AI technologies at scale in large organizations. The foundation of decision intelligence in enterprises outlines the major elements of the policy-driven governance framework. The description of distributed execution models provides a formal view of distributed AI execution frameworks, including properties, scalability patterns, core primitives, interoperability requirements, and modular architecture models. The most salient aspects of decision governance for scalable AI development and use remain a focus of attention.

Scalability Index

$$SI = \frac{P_n}{P_1}$$

Where P_n= performance with $n$nodes and P_1= performance with one node.

Enterprise governance is centered around shaping and contributing to the public and enterprise values enshrined in a policy document. At a more general level, governance deals with the rules of the game within which organizations operate. These rules encompass not just high-level objectives and ambitions but accountabilities, roles, main initiatives, boundaries, minefields, and how risks—not to mention the way the rules of the game and their significance change—are managed. In environments where artificial intelligence is deployed at scale, supporting these governance objectives is crucial. Nevertheless, managing these aspects represent main challenges for organizations following a more policy-driven development approach. Scalable decision processes revolve around risk-based aspects, management of legal mandates such as data protection regulations, and ensuring that decisions produced are aligned with the organization's objectives or preferred values.

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

Acknowledgment

Research in Decision Intelligence for Enterprise Governance in Scalable Distributed AI Implementation receives support from the European Unions Horizon 2020 Programme (Grant Number 101017091) and funds from the Foundation for Polish Science. The European Commissions Innovation Programme announces the results of the latest Open Call for SMEs under the H2020 Horizon 2020 Programme. The call invited small and micro enterprises from the Energy, Transport, and European Green Deal areas interested in the Automation Pilot to submit proposals supporting the European Commissions goal of achieving a climate-neutral Europe by 2050. The call was jointly managed and implemented by the Automation and KDT JU ecosystems participants. The Polish National Centre for Research and Development announces the results of the 16th Competition of the BRIDGE Programme. The BRIDGE Programme, which is co-financed by the European Union from the European Regional Development Fund, accommodates international cooperation of Polish enterprises and their partners within the framework of public calls for proposals in the area of applied research, development works, and demonstration projects that contribute to building and developing the Polish economy based on modern technologies. The BRIDGE Programme supports the cooperation of Polish enterprises with research units or groups of research units from the United States of America, the Republic of Korea, Canada, and Israel.

References

1. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60.
2. Rajbhandari, S., Rasley, J., Ruwase, O., & He, Y. (2020). ZeRO: Memory optimizations toward training trillion parameter models. Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 1–16.
3. Lepikhin, D., Lee, H., Xu, Y., Chen, D., Fırat, O., Huang, Y., Krikun, M., Shazeer, N., & Chen, Z. (2021). GShard: Scaling giant models with conditional computation and automatic sharding. International Conference on Learning Representations.
4. Garapati, R. S., Aitha, A. R., Yandamuri, U. S., Gottimukkala, V. R. R., Nagubandi, A. R., & Kolla, S. H. (2026, March). Cloud-Native Orchestration of Multi-Counterparty Derivatives and Collateral in Manufacturing Enterprises via AI-Assisted Financial Audit Engines. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
5. Fedus, W., Zoph, B., & Shazeer, N. (2022). Switch Transformers: Scaling to trillion parameter models with simple and efficient sparsity. Journal of Machine Learning Research, 23(120), 1–39.
6. Zheng, L., Li, Z., Zhang, H., Zhuang, Y., Chen, Z., Huang, Y., Wang, Y., Xu, Y., Zhuo, D., Xing, E. P., Gonzalez, J. E., & Stoica, I. (2022). Alpa: Automating inter- and intra-operator parallelism for

P-ISSN: 3050-9572 E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

distributed deep learning. Proceedings of the 16th USENIX Symposium on Operating Systems Design and Implementation, 559–578.

7. Lai, F., Dai, Y., Singapuram, S. S., Liu, J., Zhu, X., Madhyastha, H. V., & Chowdhury, M. (2022). FedScale: Benchmarking model and system performance of federated learning at scale. Proceedings of the 39th International Conference on Machine Learning, 162, 11814–11827.
8. Elkordy, A. R., Ezzeldin, Y. H., Han, S., Sharma, S., He, C., Mehrotra, S., & Avestimehr, S. (2023). Federated analytics: A survey. APSIPA Transactions on Signal and Information Processing, 12(1).
9. Segireddy, A. R., Davuluri, P. N., Mangala, N., & Nagubandi, A. R. (2026, April). Self-Healing Cloud-Native Quorum Systems: AI-Enabled DevOps for Zero-Downtime Financial Payments. In 2026 International Conference on NextGen Data Science and Analytics (ICNDSA) (pp. 1-9). IEEE.
10. Dean, P., & Porter, B. (2021). The design space of emergent scheduling for distributed execution frameworks. Proceedings of the Symposium on Software Engineering for Adaptive and Self-Managing Systems, 186–195.
11. Thamsen, L. (2020). Mary, Hugo, and Hugo*: Learning to schedule distributed data-parallel processing jobs on shared clusters. Concurrency and Computation: Practice and Experience, 32(20), e5823.
12. Bergui, M., Najah, S., & Nikolov, N. S. (2021). A survey on bandwidth-aware geo-distributed frameworks for big-data analytics. Journal of Big Data, 8, Article 40.
13. Li, J., Wang, Z., & Zhang, X. (2021). Octopus-DF: Unified DataFrame-based cross-platform data analytic system. Parallel Computing, 105, 102879.
14. Widanage, C., et al. (2024). Supercharging distributed computing environments for high-performance data engineering. Frontiers in High Performance Computing, 2.
15. Isukapalli, S., & Srirama, S. N. (2024). A systematic survey on fault-tolerant solutions for distributed data analytics: Taxonomy, comparison, and future directions. Computer Science Review, 53, 100660.
16. Zhang, C., Li, Y., & Li, X. (2020). Distributed deep learning with communication-efficient optimization. IEEE Transactions on Parallel and Distributed Systems, 31(10), 2384–2397.
17. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N.,

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, M. J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

18. Aitha, A. R. (2026). Explainable Agentic AI Framework for Automated Insurance Fraud Detection and Predictive Risk Intelligence. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 877-888.
19. Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., He, B., & Li, J. (2020). A survey on federated learning systems: Vision, hype and reality for data privacy and protection. IEEE Transactions on Knowledge and Data Engineering.
20. Bonawitz, K., Kairouz, P., McMahan, H. B., Rambhatla, S., et al. (2022). Towards federated learning at scale: System design. Proceedings of Machine Learning and Systems, 4, 374–388.
21. Verbraeken, J., Wolting, M., Katzy, J., Kloppenburg, J., Verbelen, T., & Rellermeyer, J. S. (2020). A survey on distributed machine learning. ACM Computing Surveys, 53(2), Article 30.
22. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2020). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 7(4), 3164–3178.
23. Satyanarayanan, M. (2020). The emergence of edge computing. Computer, 53(1), 30–39.
24. Inala, R., Garapati, R. S., Aitha, A. R., Komaragiri, V. B., Gottimukkala, V. R. R., Recharla, M., ... & Varri, D. B. S. (2026). U.S. Patent Application No. 19/389,108.
25. Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2020). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proceedings of the IEEE, 107(8), 1738–1762.
26. Zhou, X., Chen, H., & Wang, X. (2021). Distributed machine learning systems: A survey of architectures, algorithms, and applications. IEEE Access, 9, 109897–109916.
27. Ben-Nun, T., & Hoefler, T. (2020). Demystifying parallel and distributed deep learning: An in-depth concurrency analysis. ACM Computing Surveys, 52(4), Article 65.
28. Sergeev, A., & Del Balso, M. (2020). Horovod: Fast and easy distributed deep learning in TensorFlow. Journal of

P-ISSN: 3050-9572
E-ISSN: 3050-9580
EUROPEAN DATA SCIENCE JOURNAL
Volume: 04 Issue: 03
RECEIVED: JULY 07 REVISED: JULY 26 ACCEPTED: AUGUST 05 PUBLISHED: AUGUST 17

Machine Learning Research, 21(79), 1–10.
29. Peng, Y., Zhu, Y., Bao, Y., Chen, Y., Wu, C., Guo, C., & Lin, Q. (2020). A generic communication-efficient parallelization method for distributed deep learning. Proceedings of the 2020 USENIX Annual Technical Conference, 1–14.
30. Huang, C.-C., Jin, G., & Dai, J. (2020). SwapAdvisor: Pushing deep learning beyond the GPU memory limit. Proceedings of the 27th International Conference on Architectural Support for Programming Languages and Operating Systems, 367–380.
31. Zhuang, Y., Zheng, L., Zhao, Z., Li, Z., & Stoica, I. (2023). On optimizing the communication of model parallelism. Proceedings of Machine Learning and Systems, 5, 1–14.
32. Yandamuri, U. S. (2026). Scalable Cloud-Based Intelligent Decision Systems Leveraging AI and Big Data for Industry-Specific Optimization. Minnesota Journal of Business Law and Entrepreneurship, (1), 584-601.
33. Li, Z., Zheng, L., Zhong, Y., Liu, V., Sheng, Y., Jin, X., Huang, Y., Chen, Z., Zhang, H., Gonzalez, J. E., & Stoica, I. (2023). AlpaServe: Statistical multiplexing with model parallelism for deep learning serving. Proceedings of the 17th USENIX Symposium on Operating Systems Design and Implementation, 663–679.
34. Widanage, C., Perera, D., et al. (2024). Supercharging distributed computing environments for high-performance data engineering. Frontiers in High Performance Computing, 2.
35. Isukapalli, S., & Srirama, S. N. (2024). A systematic survey on fault-tolerant solutions for distributed data analytics: Taxonomy, comparison, and future directions. Computer Science Review, 53, 100660.