Enactive Artificial Intelligence: A Decision-Centric Architecture for Complex Systems

Zuo-Jun (Max) Shen **Yuan Qu** **Pujun Zhang**
Anbang Liu **Yunhao Liang**

The University of Hong Kong, Hong Kong SAR, China

maxshen@hku.hk yuanqu@hku.hk pjzhang@hku.hk
anbang@hku.hk yunhao8@hku.hk

July 23, 2026

Abstract

Industrial artificial intelligence increasingly combines language models, agents, knowledge graphs, optimization, and digital twins. These technologies are often deployed as separate capabilities, leaving a gap between organizational intent and reliable action in operational environments. This perspective article synthesizes adjacent research and introduces *Enactive AI* as a conceptual framework for enterprise reasoning, site-level decision support, and execution feedback. Three complementary roles organize the framework. An *Organizational World* represents goals, policies, business context, and evidence. A *Site World* represents decision-relevant operational conditions and possible interventions. *Schema Intelligence* maintains a shared structure across the two perspectives. The discussion positions Enactive relative to predictive world models, enterprise semantics, operations research, agentic AI, and digital twins. It then considers design principles, application domains, evaluation perspectives, governance requirements, and open research questions. Implementation details remain outside the scope of this article. The paper specifies conceptual roles and research boundaries without disclosing proprietary representations, orchestration policies, control logic, thresholds, or deployment parameters. It makes no claim of measured deployment benefit and identifies empirical validation as the central next step.

Keywords: Enactive AI; enterprise artificial intelligence; world models; industrial decision support

1 Introduction

Industrial decisions connect multiple levels of an enterprise. A question about delivery reliability, resource use, or operational risk may begin with business commitments and end with changes to schedules, logistics, maintenance, or site operations. The relevant information is distributed across enterprise systems, operational platforms, documents, local models, and human expertise. No single source normally contains the complete decision context.

Recent artificial-intelligence (AI) technologies address important parts of this problem. Language models interpret questions and interact with tools (Yao et al., 2023). Knowledge graphs stabilize entities and relations (W3C OWL Working Group, 2012). Operations research provides objectives, constraints, and decision methods (Elmachtoub and Grigas, 2022). Digital twins connect computational models with physical systems (Kritzinger et al., 2018). Agent frameworks coordinate specialized capabilities (Wu et al., 2023). The central challenge is not the absence of individual methods. It is the absence of a shared decision perspective that relates organizational meaning, operational context, possible action, and feedback.

World-model research offers a useful starting point because it studies representations that support prediction and planning. Enterprises, however, require more than a model of physical dynamics. Decisions are also shaped by policies, commitments, authority, evidence quality, competing objectives, and organizational interpretation. A technically feasible intervention may still be inappropriate because its business meaning, permissions, or consequences are unclear.

We use the term *Enactive World Model* for a conceptual framework in which an enterprise representation is evaluated by how well it supports responsible action and learning from consequences. This name draws on enactive accounts of cognition while retaining a deliberately limited engineering meaning (Varela et al., 1991; Froese and Ziemke, 2009; Rafiee and Sutton, 2026). Organism-like autonomy is not attributed to an enterprise system. Instead, the framework emphasizes the recurring connection among interpretation, decision support, operational action, and feedback.

This paper makes three contributions.

1. First, the paper synthesizes research on world models, enterprise semantics, operations research, agentic AI, and digital twins around a common industrial decision problem.
2. Second, a high-level framework assigns complementary roles to an Organizational World Model, a Site World Model, and Schema Intelligence.
3. Third, the public research agenda covers evaluation, governance, organizational adoption, and the limits of autonomy without prescribing a proprietary implementation.

The remainder of the paper reviews the relevant foundations and presents the conceptual framework. Later sections discuss the cross-layer decision cycle, applications, evaluation, governance, limitations, and research opportunities.

2 Introduction

Industrial decisions connect multiple levels of an enterprise. A question about delivery reliability, resource use, or operational risk may begin with business commitments and end with changes to schedules, logistics, maintenance, or site operations. The relevant information is distributed across enterprise systems, operational platforms, documents, local models, and human expertise. No single source normally contains the complete decision context.

Recent artificial-intelligence (AI) technologies address important parts of this problem. Language models interpret questions and interact with tools (Yao et al., 2023). Knowledge graphs stabilize entities and relations (W3C OWL Working Group, 2012). Operations research provides objectives, constraints, and decision methods (Elmachtoub and Grigas, 2022). Digital twins connect computational models with physical systems (Kritzinger et al., 2018). Agent frameworks coordinate specialized capabilities (Wu et al., 2023). The central challenge is not the absence of individual methods. It is the absence of a shared decision perspective that relates organizational meaning, operational context, possible action, and feedback.

World-model research offers a useful starting point because it studies representations that support prediction and planning. Enterprises, however, require more than a model of physical dynamics. Decisions are also shaped by policies, commitments, authority, evidence quality, competing objectives, and organizational interpretation. A technically feasible intervention may still be inappropriate because its business meaning, permissions, or consequences are unclear.

We use the term *Enactive AI* for a conceptual framework in which an enterprise representation is evaluated by how well it supports responsible action and learning from consequences. This name draws on enactive accounts of cognition while retaining a deliberately limited engineering meaning (Varela et al., 1991; Froese and Ziemke, 2009; Rafiee and Sutton, 2026). Organism-like autonomy is not attributed to an enterprise system. Instead, the framework emphasizes the recurring connection among interpretation, decision support, operational action, and feedback.

This paper makes three contributions.

1. First, the paper synthesizes research on world models, enterprise semantics, operations research, agentic AI, and digital twins around a common industrial decision problem.
2. Second, a high-level framework assigns complementary roles to an Organizational World , a Site World , and Schema Intelligence.
3. Third, the public research agenda covers evaluation, governance, organizational adoption, and the limits of autonomy without prescribing a proprietary implementation.

The remainder of the paper reviews the relevant foundations and presents the conceptual framework. Later sections discuss the cross-layer decision cycle, applications, evaluation, governance, limitations, and research opportunities.

3 Foundations and Research Landscape

3.1 Predictive World Models and Enaction

Model-based decision systems use representations of an environment to assess possible actions before committing them. Dyna combined learning, planning, and action within one architecture (Sutton, 1991). Neural world models later learned compact internal representations for prediction and policy development (Ha and Schmidhuber, 2018). MuZero showed that planning can rely on a learned representation without reconstructing every observable detail (Schrittwieser et al., 2020). More recent architectural proposals emphasize hierarchical prediction across multiple levels of abstraction (LeCun, 2022).

This literature establishes the value of action-oriented representation. Industrial systems broaden the representational problem. Relevant conditions may include physical resources and process state, but also contractual commitments, organizational priorities, permissions, and the source of evidence. The decision context therefore spans both operational and institutional realities.

Enaction contributes the idea that cognition is shaped through an ongoing history of perception and action (Varela et al., 1991). For enterprise systems, this idea is most useful as a design orientation rather than a claim of biological equivalence. A representation should retain distinctions that matter for decision and should be revised when outcomes reveal that its prior interpretation was incomplete. This interpretation aligns the framework with feedback-oriented engineering while acknowledging the stronger autonomy requirements associated with philosophical enactivism (Froese and Ziemke, 2009; Rafiee and Sutton, 2026).

3.2 Enterprise Semantics, Operations Research, and Digital Twins

Knowledge graphs and ontologies provide stable identities and explicit relations across heterogeneous information sources (W3C OWL Working Group, 2012; Haller et al., 2017). They help reconcile the multiple names, definitions, and contexts associated with enterprise objects. Their importance extends beyond retrieval: semantic consistency influences whether a decision system is addressing the intended object, metric, policy, or time horizon.

Operations research contributes a complementary perspective. It makes objectives, trade-offs, and constraints explicit and provides disciplined methods for comparing alternatives. Decision-focused learning also shows that predictive accuracy and decision quality are not identical objectives (Donti et al., 2017; Elmachtoub and Grigas, 2022). These methods are powerful once the decision problem has been framed, yet enterprise problem framing often depends on contextual and organizational interpretation that lies outside a conventional mathematical model.

Digital twins connect representations with physical assets and processes. Manufacturing research distinguishes digital models, digital shadows, and twins with bidirectional data exchange (Kritzinger et al., 2018). ISO 23247 provides a reference architecture for manufacturing digital twins and recognizes multiple forms of model composition (International Organization for Standardization, 2021, 2026). Digital twins can supply state, simulation, and monitoring capabilities, but they do not by themselves determine which organizational objective should govern a decision.

Recent enterprise-world-model research further narrows the open problem. The Business World Model represents business entities, dynamics, objectives, constraints, and possible actions for counterfactual reasoning (Pang and Sayama, 2026). Shared semantic worlds support coordination among enterprise agents (Mantsivoda and Gavrilina, 2026). Runtime discovery agents infer evolving enterprise processes when relevant rules are accessible (Nair et al., 2026). Digital-Twin Markov decision processes have also represented enterprise-agent behavior for offline learning and context engineering (Yang et al., 2026). These contributions indicate growing interest in enterprise-level world representations, while leaving room for a broader account of how organizational and site perspectives remain connected.

3.3 Language-Model Agents and Process Orchestration

Language-model agents combine reasoning with actions and observations (Yao et al., 2023). Multi-agent frameworks distribute work among specialized roles and maintain shared task context (Wu et al., 2023). Grounded planning research has also linked language-model proposals with estimates of executable skills (Ichter et al., 2023). These approaches expand the range of tasks that AI systems can address, especially when interpretation, retrieval, analysis, and tool use must be combined.

In industrial settings, orchestration is only one part of the problem. Agents also require consistent definitions, evidence boundaries, organizational accountability, and an understanding of operational consequences. Enactive AI therefore treats agentic AI as one family of capabilities operating within a broader decision framework rather than as the framework itself.

3.4 Synthesis and Positioning

Table 1 summarizes the different emphases of adjacent approaches. The comparison is conceptual and does not imply performance ranking. Enactive AI depends on several of these approaches as components.

Table 1: Conceptual positioning relative to adjacent approaches. "Partial" indicates that coverage depends on the implementation or is not the approach's primary focus.

Approach	**Enterprise meaning**	**Operational context**	**Decision methods**	**Governance**	**Feedback**
Language model or agent (Yao et al., 2023)	Partial	Partial	Partial	Partial	Partial
Knowledge graph or ontology (W3C OWL Working Group, 2012; Haller et al., 2017)	Strong	Limited	Limited	Limited	Partial
Operations-research model (Elmachtoub and Grigas, 2022)	Input dep.	Partial	Strong	Model dep.	Partial
Digital twin (Kritzinger et al., 2018; International Organization for Standardization, 2021)	Partial	Strong	Partial	Partial	Site level
Enterprise world representation (Pang and Sayama, 2026; Mantsivoda and Gavrilina, 2026)	Strong	Partial	Strong	Partial	Partial
Enactive AI	**Integrated**	**Integrated**	**Integrated**	**Explicit concern**	**Organizational and site**

The open research question is how to maintain a coherent connection between organizational interpretation and operational decision support without collapsing the two into a single undifferentiated representation. Enactive AI addresses this question at the level of roles, information relationships, and governance responsibilities.

4 Enactive AI as a Conceptual Framework

4.1 High-Level View

Enactive AI separates three concerns. Organizational context explains why a decision matters to the enterprise. Site context describes the operational setting in which intervention may occur. Schema Intelligence maintains enough shared structure for the two perspectives to remain mutually intelligible. This separation allows each layer to evolve with its own methods while preserving a common decision context.

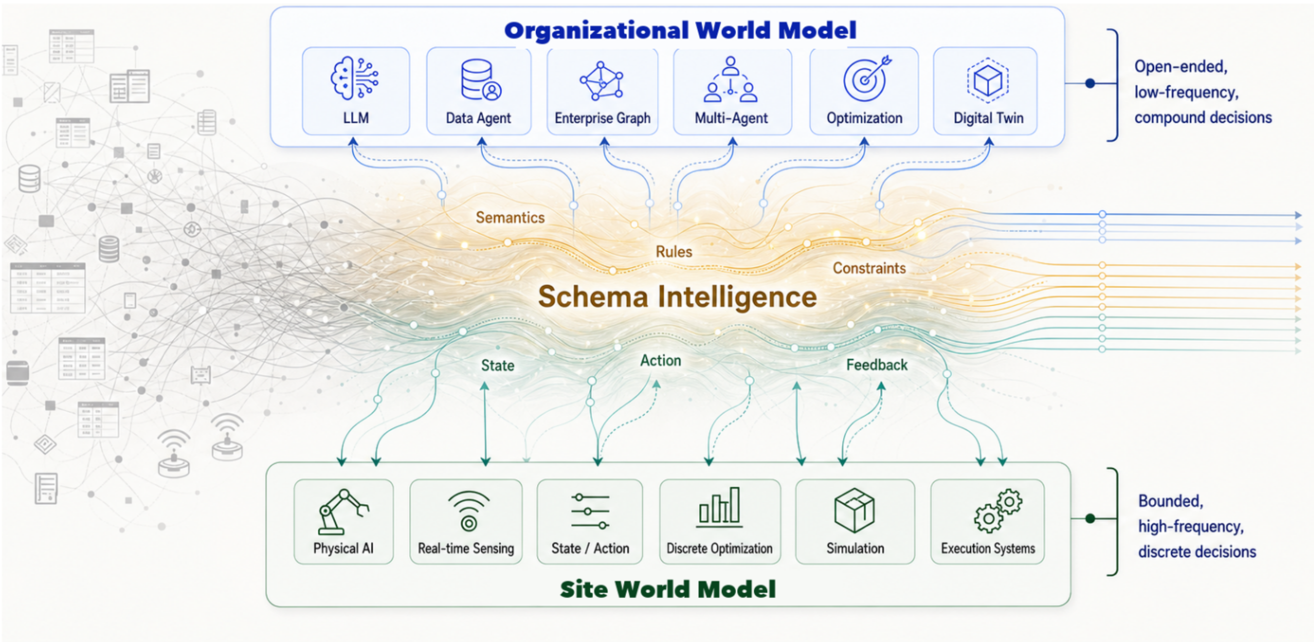

Figure 1: High-level view of Enactive AI. Organizational reasoning and site-level decision support remain distinct. Schema Intelligence connects semantics, rules, constraints, state, action, and feedback at a conceptual level.

The framework follows five general principles.

- **Decision grounding:** represent information in relation to consequential enterprise decisions rather than aiming for an exhaustive digital replica.
- **Layer separation:** distinguish organizational meaning from operational conditions while maintaining a clear connection between them.
- **Responsible action:** treat feasibility, safety, authority, and evidence as explicit concerns before operational release.
- **Progressive autonomy:** match the level of automation to context familiarity, consequence severity, reversibility, and human oversight.
- **Learning from outcomes:** use execution and human feedback to improve both operational understanding and organizational interpretation.

These principles describe a family of possible systems. They do not prescribe a universal data model, agent topology, solver, control policy, or deployment architecture.

4.2 Organizational World

An Organizational World provides the enterprise-facing perspective. It brings together the business context needed to interpret a decision: relevant objects and relationships, goals, policies, evidence, commitments, and organizational responsibilities. Its purpose is not to reproduce every enterprise database or document. Instead, it creates a coherent view of what the organization is trying to achieve and why particular outcomes matter.

At a high level, this model supports four activities. It frames open managerial questions, organizes relevant context, connects analytical capabilities, and communicates decision-oriented outputs. Depending on the setting, those capabilities may include language models, knowledge graphs, forecasting, simulation, optimization, or human expertise. The framework remains neutral about how they are orchestrated.

It also provides an accountability boundary. Decision context should remain traceable to responsible owners and evidence sources, especially when objectives conflict or policies change. This emphasis distinguishes enterprise decision intelligence from unrestricted task automation.

4.3 Schema Intelligence

Schema Intelligence is the relationship-structuring layer of the Enactive AI. It maintains conceptual consistency across organizational and site perspectives. Its role includes reconciling context, relating business meaning to operational conditions, and supporting the exchange of decision-relevant information. The framework does not assume that one static enterprise ontology can serve every decision equally well. Different decisions may require different levels of detail and different connections among concepts.

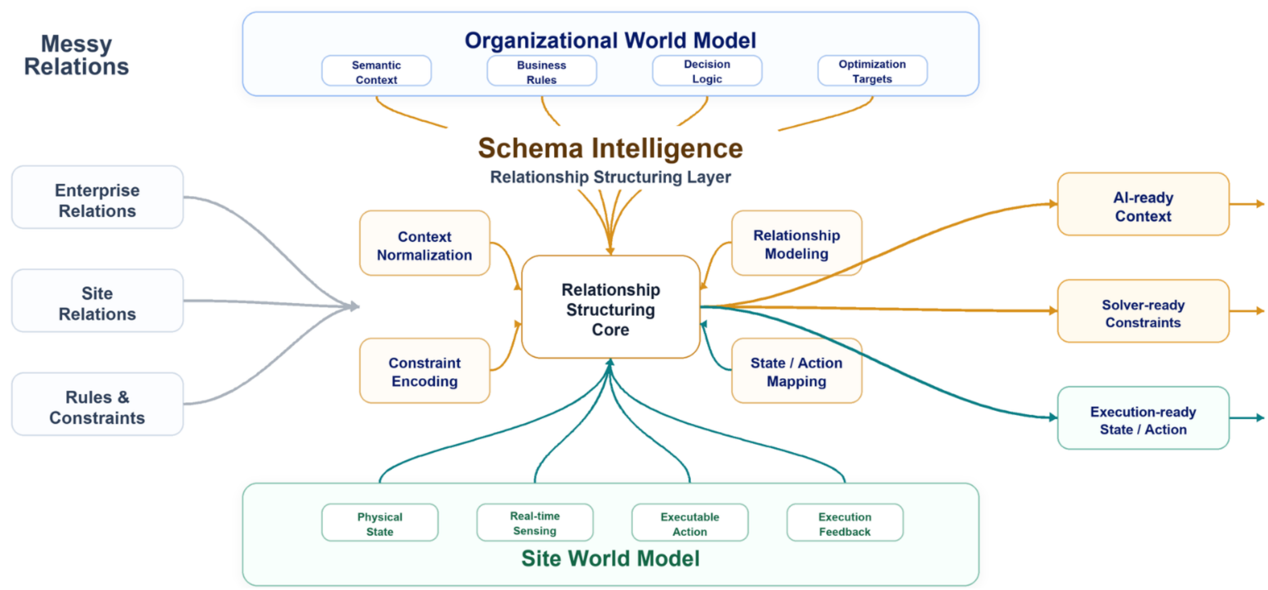

Figure 2: A general view of Schema Intelligence. It provides a relationship-structuring layer between organizational and site perspectives. The figure emphasizes conceptual roles rather than an implementation specification.

Three broad outputs are useful for understanding this role. First, AI-facing context gives reasoning systems a coherent interpretation of the decision. Second, analytical context makes objectives and constraints understandable to modeling tools. Third, execution-facing context connects approved intent with operational systems. These are conceptual interfaces; their technical realization may vary across organizations and use cases.

4.4 Site World

The SiteWorld provides the operational perspective. It represents the conditions, resources, constraints, and possible interventions relevant to a site-level decision. A site may be a factory, warehouse, logistics network, port, campus, transport system, or other cyber-physical environment. The model may draw on sensing, simulation, optimization, rules, and existing execution platforms.

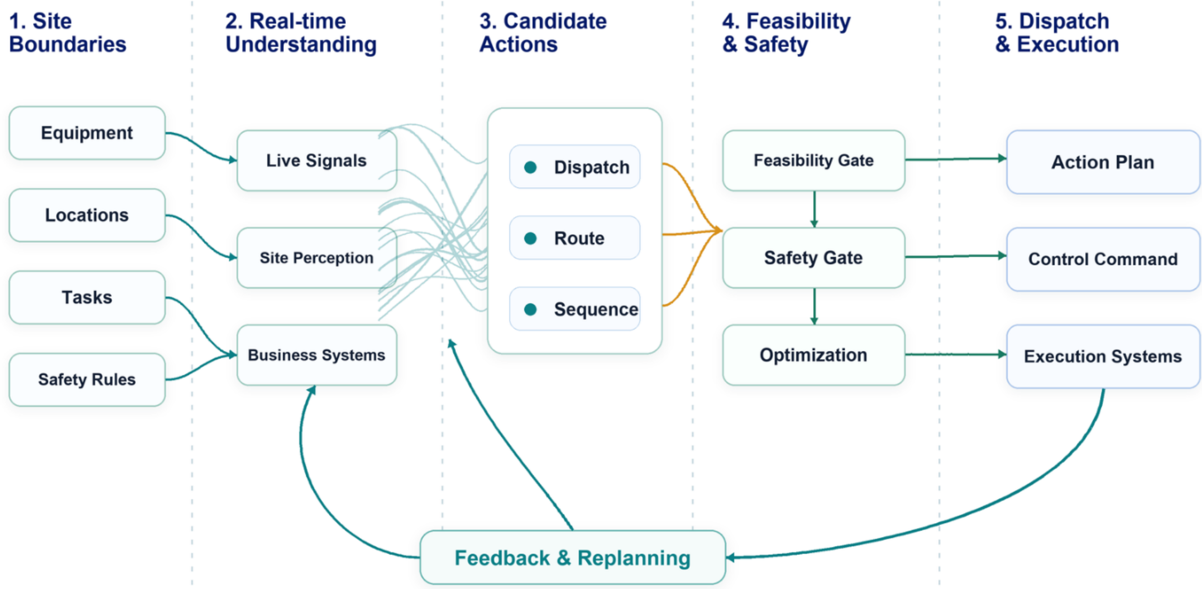

Figure 3: A general view of site-level decision support. The Site World connects operational boundaries with possible actions and safeguards. Execution and feedback complete the high-level cycle.

The Site World is not intended to replace deterministic controllers, transactional systems, or established safety mechanisms. Its role is to support higher-level decisions about what should be considered, coordinated, or reviewed. Existing operational systems retain their specialized responsibilities.

5 A Cross-Layer Decision Cycle

The framework can be understood through a recurring cycle with three broad movements. Organizations may implement this deliberately conceptual pattern with different technologies and governance arrangements.

5.1 From Organizational Intent to Decision Context

A decision begins with an organizational concern such as service reliability, resource use, risk exposure, or operational resilience. Organizational context clarifies the business meaning of that concern. Schema Intelligence then connects this interpretation with the operational concepts required for site-level analysis. This produces shared context rather than an automatically executable instruction. Business and operational participants can then reason about the same decision.

5.2 From Operational Context to Governed Action

The Site World examines the current setting and possible interventions at an appropriate level of abstraction. Candidate actions may be evaluated through rules, human judgment, optimization, simulation, or combinations of these methods. Consequential actions should remain subject to relevant feasibility, safety, authority, and evidence requirements. The framework does not require a particular release mechanism; it requires that release responsibility be explicit and auditable.

5.3 From Outcomes to Organizational Learning

Execution produces more than a physical outcome. It also reveals whether the enterprise interpreted the situation correctly, whether the selected models were adequate, and whether the consequences matched expectations. Feedback should therefore update both operational understanding and organizational context. This cross-level learning loop is central to the enactive orientation of the framework.

6 Application Perspectives

Enactive AI is relevant where business objectives and operational decisions are tightly coupled. The following domains illustrate the breadth of the research problem without specifying deployment procedures.

6.1 Manufacturing Operations

Manufacturing decisions connect orders, priorities, capacity, quality, maintenance, workforce considerations, and process constraints. An Organizational World can clarify commitments and priorities, while a Site World can support understanding of current production conditions. Schema Intelligence provides the shared vocabulary needed to relate business consequences to operational alternatives.

6.2 Supply Chains and Logistics

Supply-chain decisions span demand, inventory, transportation, supplier conditions, contractual obligations, and service expectations. The framework can help connect network-level concerns with local planning and execution environments. Its value lies in maintaining cross-level coherence when information and responsibility are distributed across organizations and systems.

6.3 Infrastructure and Complex Sites

Ports, campuses, transport systems, energy facilities, and other complex sites combine physical constraints with institutional rules and multiple stakeholders. In these settings, decision intelligence must account for safety, authority, coordination, and changing operational conditions. Enactive AI offers a common language for studying these interactions without assuming that one model or control system should govern the entire environment.

7 Evaluation Perspectives and Open Questions

Because Enactive AI is a conceptual framework, evaluation should examine whether the framework improves the quality and governability of decision processes rather than only the accuracy of an individual model. Table 2 summarizes five complementary perspectives.

Table 2: Evaluation perspectives for future empirical research.

Perspective	**Central question**
Semantic alignment	Do organizational and operational participants interpret the decision consistently?
Decision relevance	Does the representation retain the context needed to compare meaningful alternatives?
Operational reliability	Are recommendations compatible with the operational environment and its safeguards?
Governance and accountability	Can responsible participants understand, review, and intervene in consequential decisions?
Learning and adaptation	Does feedback improve subsequent interpretation and decision support without eroding control?

Future research should compare the framework with simpler alternatives in recorded, simulated, or sandboxed settings before considering live autonomy. One question is when a cross-layer representation provides material value. Others concern which decisions require shared context, how human expertise should enter the loop, and when coordination costs exceed the benefit.

Another research direction concerns representation sufficiency. A useful world model should omit irrelevant detail while preserving information that changes interpretation or decision. Determining that boundary is likely to depend on the decision domain, consequence severity, and organizational environment. Comparative studies across industries would help establish when reusable patterns exist and when local design remains unavoidable.

Evaluation should also report failures and abstentions. A framework for consequential decisions cannot be assessed only through successful examples. Research should examine misinterpretation, infeasible recommendations, delayed escalation, brittle transfer across sites, and disagreements between human and automated assessments.

8 Managerial and Governance Implications

8.1 Begin with a Consequential Decision, Not a Universal Enterprise Model

Organizations should start with a bounded decision whose owner, consequences, and operational context are understood. This approach limits unnecessary integration and helps reveal which organizational and site distinctions actually matter. Expansion should follow demonstrated decision needs rather than the ambition to model the entire enterprise at once.

8.2 Keep Organizational Meaning Visible

Managers should require AI initiatives to state the objective, evidence basis, responsibility, and expected operational consequence of a recommendation. Technical capability alone does not establish decision legitimacy. Organizational meaning should remain visible even when specialized models and agents perform much of the analysis.

8.3 Treat Autonomy as a Graduated Organizational Choice

The appropriate level of autonomy varies across decisions and conditions. Familiar, reversible, and well-understood settings may support greater automation. Novel, consequential, or weakly evidenced settings require stronger human involvement. Organizations should treat autonomy as a governed portfolio of choices rather than a single property of the system.

8.4 Preserve Independent Operational Safeguards

Enactive AI should complement rather than displace established safety, control, and transactional systems. Independent safeguards, clear authority boundaries, and recovery procedures remain necessary. Explanations and audit records should support review, but they should not be mistaken for proof that an action is safe or correct.

9 Limitations and Research Agenda

Enactive AI is presented here as a conceptual perspective rather than a complete technical specification. It does not define a universal schema language, agent architecture, optimization method, safety mechanism, or control interface. Such abstraction is intentional: implementations will vary across industries, organizations, and consequence levels, and some design elements may remain proprietary.

The framework also lacks empirical validation in its present form. Future work must test whether the proposed separation of roles improves decisions, coordination, or accountability relative to simpler architectures. Such studies should include negative results and settings in which the framework adds unnecessary complexity.

Five research themes appear especially important.

1. **Representation:** identify reusable concepts for linking organizational and operational perspectives without imposing one universal ontology.
2. **Human-AI organization:** study how responsibility, expertise, and escalation should be distributed among managers, operators, agents, and analytical tools.
3. **Evaluation:** develop benchmarks and field-study designs that capture decision quality, reliability, governance, and adaptation together.
4. **Transfer and change:** examine how world representations remain useful as policies, processes, equipment, and organizational priorities evolve.
5. **Responsible autonomy:** characterize the conditions under which automation should expand, remain bounded, or return control to human decision makers.

The Enactive AI concept should therefore be treated as a research program. Its usefulness will depend on whether future studies can show that cross-layer coherence improves consequential industrial decisions without creating excessive complexity or weakening established controls.

10 Conclusion

Industrial AI is moving from isolated analytical tools toward systems that interpret context, coordinate models, and influence operations. This transition creates a need for representations that connect organizational meaning with operational reality. Enactive AI addresses that need through three conceptual roles: the Organizational World , the Site World , and Schema Intelligence.

This synthesis brings together insights from predictive world models, enterprise semantics, operations research, agentic AI, and digital twins. Its central claim is modest but consequential: industrial decision intelligence requires a coherent relationship among intent, context, action, and feedback. Implementation remains intentionally open. Empirical research must determine where this perspective improves decision processes, where simpler systems are sufficient, and how responsible autonomy can be maintained as capabilities evolve.

References

Priya L. Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3fc2c60b5782f641f76bcefc39fb2392-Abstract.html.

Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022. doi: 10.1287/mnsc.2020.3922.

Tom Froese and Tom Ziemke. Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, 173(3–4):466–500, 2009. doi: 10.1016/j.artint.2008.12.001.

David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/2de5d16682c3c35007e4e92982f1a2ba-Abstract.html.

Armin Haller, Krzysztof Janowicz, Simon Cox, Danh Le Phuoc, Kerry Taylor, and Maxime Lefrançois. Semantic sensor network ontology. W3C Recommendation OGC 16-079, World Wide Web Consortium and Open Geospatial Consortium, October 2017. URL https://www.w3.org/TR/2017/REC-vocab-ssn-20171019/.

Brian Ichter, Anthony Brohan, Yevgen Chebotar, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Proceedings of the 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 287–318, 2023. URL https://proceedings.mlr.press/v205/ichter23a.html.

International Organization for Standardization. ISO 23247-2:2021 automation systems and integration—digital twin framework for manufacturing—part 2: Reference architecture. Technical Report ISO 23247-2:2021, International Organization for Standardization, 2021. URL https://www.iso.org/standard/78743.html.

International Organization for Standardization. ISO 23247-6:2026 automation systems and integration—digital twin framework for manufacturing—part 6: Digital twin composition. Technical Report ISO 23247-6:2026, International Organization for Standardization, 2026. URL https://www.iso.org/standard/87426.html.

Werner Kritzinger, Matthias Karner, Georg Traar, Jan Henjes, and Wilfried Sihn. Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11):1016–1022, 2018. doi: 10.1016/j.ifacol.2018.08.474.

Yann LeCun. A path towards autonomous machine intelligence, 2022. URL https://openreview.net/pdf?id=BZ5a1r-kVsf. Version 0.9.2, position paper.

Andrei Mantsivoda and Darya Gavrilina. Semantic modeling for world-centered architectures. *arXiv preprint arXiv:2604.01359*, 2026. doi: 10.48550/arXiv.2604.01359. URL https://arxiv.org/abs/2604.01359.

Jishnu Sethumadhavan Nair, Patrice Bechard, Rishabh Maheshwary, Surajit Dasgupta, Sravan Ramachandran, Aakash Bhagat, Shruthan Radhakrishna, Pulkit Pattnaik, Johan Obando-Ceron, Shiva Krishna Reddy Malay, Sagar Davasam, Seganrasan Subramanian, Vipul Mittal, Sridhar Krishna Nemala, Christopher Pal, Srinivas Sunkara, and Sai Rajeswar. Do enterprise systems need learned world models? the importance of context to infer dynamics. *arXiv preprint arXiv:2605.12178*, 2026. doi: 10.48550/arXiv.2605.12178. URL https://arxiv.org/abs/2605.12178.

Cecil Pang and Hiroki Sayama. Business world model. *arXiv preprint arXiv:2606.10044*, 2026. URL https://arxiv.org/abs/2606.10044.

Banafsheh Rafiee and Richard S. Sutton. Toward enactive artificial intelligence. *arXiv preprint arXiv:2605.24238*, 2026. URL https://arxiv.org/abs/2605.24238.

Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588:604–609, 2020. doi: 10.1038/s41586-020-03051-4.

Richard S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, 2(4):160–163, 1991. doi: 10.1145/122344.122377.

Francisco J. Varela, Evan Thompson, and Eleanor Rosch. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, MA, 1991. ISBN 9780262220422.

W3C OWL Working Group. OWL 2 Web Ontology Language: Document overview, second edition. W3c recommendation, World Wide Web Consortium, December 2012. URL https://www.w3.org/TR/2012/REC-owl2-overview-20121211/.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W. White, Doug Burger, and Chi Wang. AutoGen: Enabling next-gen llm applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*, 2023. URL https://arxiv.org/abs/2308.08155.

Xi Yang, Aurélie Lozano, Naoki Abe, Bhavya, Saurabh Jha, Noah Zheutlin, Rohan R. Arora, Yu Deng, and Daby M. Sow. A context engineering framework for improving enterprise AI agents based on Digital-Twin MDP. *arXiv preprint arXiv:2603.22083*, 2026. doi: 10.48550/arXiv.2603.22083. URL https://arxiv.org/abs/2603.22083.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=WE_vluYUL-X.