Dynamic Knowledge Intelligence for Human AI Collaborative Operational Decision Making Across Distributed Enterprise Data Environments

Arif Mu'amar Wahid[1,*], Rizky Rahmatullah[2]

[1,2]*Graduate School of Natural and Science Technology, Kanazawa University, Japan*

(Received: April 3, 2026; Revised: June 27, 2026; Accepted: August 26, 2026; Available online: September 2, 2026)

Abstract

Distributed enterprise environments frequently contain fragmented operational data distributed across ERP platforms, CRM systems, workflow logs, IoT streams, and unstructured document repositories, which reduces the speed and consistency of managerial response. This study proposes a dynamic knowledge intelligence framework that integrates heterogeneous enterprise data, transforms them into a temporal contextual knowledge layer, and connects that layer to an operational decision engine for action prioritization. The framework was evaluated across five analytical dimensions, namely integration quality, knowledge graph enrichment, recommendation performance, response efficiency, and cross-domain business impact. The harmonization stage improved integration consistency from 0.61 to 0.89 for ERP data, from 0.58 to 0.85 for CRM data, from 0.64 to 0.91 for workflow logs, from 0.47 to 0.78 for documents, and from 0.55 to 0.84 for IoT streams, with error reduction ranging from 26% to 34%. In the contextual intelligence layer, knowledge graph nodes increased from 120 at T1 to 352 at T5, while relations expanded from 210 to 812, indicating that contextual linkage grew faster than object accumulation. The decision engine produced strong recommendation performance across five enterprise scenarios, with accuracy values of 0.88 for backlog control, 0.91 for inventory risk, 0.86 for escalation routing, 0.83 for policy conflict, and 0.89 for resource allocation, while expert alignment ranged from 0.81 to 0.90. Efficiency analysis showed average decision latency of 0.82 seconds at 50 concurrent requests, 1.04 seconds at 100 requests, 1.36 seconds at 200 requests, 1.95 seconds at 400 requests, and 2.88 seconds at 800 requests, confirming stable operational responsiveness under low to heavy workloads and controlled degradation only under stress-level concurrency. Business impact analysis further showed differentiated strategic gains, including 28% higher response speed in service management, 31% stronger risk mitigation in supply chain settings, 33% better decision traceability in governance contexts, 20% stronger coordination quality in resource alignment, and 24% better decision consistency in enterprise-wide operations. These findings demonstrate that the proposed framework effectively converts fragmented enterprise data into contextual, updateable, and decision-ready intelligence, thereby improving operational agility, recommendation credibility, and cross-domain decision value in distributed enterprise environments.

Keywords: Dynamic Knowledge Intelligence, Operational Decision Making, Distributed Enterprise Data, Knowledge Graph, Enterprise Analytics, Decision Support System, Data Harmonization, Contextual Intelligence, Prescriptive Analytics, Operational Intelligence

1. Introduction

Modern enterprises operate across distributed data environments in which transactional systems, workflow platforms, sensor streams, document repositories, and communication records evolve independently and rarely align semantically at the moment decisions must be made. Although data volume has increased substantially, decision quality often remains constrained by fragmentation, weak interoperability, and uneven stewardship of enterprise knowledge [1], [2]. This condition creates a structural mismatch between data availability and decision usability, especially when operational actors require context-rich, time-sensitive intelligence rather than isolated records.

The problem becomes more acute when organizations attempt to convert large-scale digital traces into actionable knowledge. Prior work shows that big data initiatives do not automatically yield better managerial reasoning unless they are coupled with knowledge management processes and decision-support mechanisms that can transform raw observations into interpretable evidence [3], [4]. In practical terms, many enterprises still possess data-rich but

***Corresponding author: Arif Mu'amar Wahid (arif@amikompurwokerto.ac.id)**
DOI: https://doi.org/10.47738/ijaim.v6i3.129
This is an open access article under the CC-BY license (https://creativecommons.org/licenses/by/4.0/).
© Authors retain all copyrights

knowledge-poor environments, where analytical assets exist in abundance but remain disconnected from the operational contexts in which decisions are executed.

Recent research has increasingly turned to knowledge graphs and related semantic structures because they provide integrated representations of entities, events, relations, and constraints that can support reasoning and decision support. Studies in enterprise architecture, general decision frameworks, and infrastructure decision systems suggest that graph-based knowledge modeling improves information organization, traceability, and contextual retrieval [5], [6], [7]. However, these contributions still tend to emphasize representation quality or domain-specific use cases more strongly than cross-environment operational intelligence in highly distributed enterprise settings.

In parallel, data-fabric architectures and dynamic resource-allocation models have expanded the discussion of how distributed enterprise data can be connected and exploited in real time. Existing work demonstrates that data fabrics can support evidence-based access across heterogeneous assets, while dynamic knowledge-enabled decision models in cloud manufacturing and graph-driven decision systems in manufacturing show that intelligent reasoning can improve distributed operational coordination [8], [9], [10]. Yet these studies are still largely anchored to specific architectural layers or sector-bounded workflows, leaving an unresolved gap between distributed data integration, evolving knowledge intelligence, and enterprise-wide operational decision support.

A second gap concerns the transition from descriptive or predictive analytics toward operationally actionable intelligence. The literature on prescriptive analytics confirms that mature analytics systems are increasingly expected to recommend actions rather than merely report patterns, while empirical work on business analytics capability shows that decision performance improves when analytics is embedded in broader knowledge and organizational processes [11], [12]. Even so, most existing approaches do not sufficiently explain how operational recommendations should be generated when evidence is distributed, context changes continuously, and knowledge relevance decays over time across multiple enterprise systems.

The objective of this study is therefore to design a dynamic knowledge intelligence framework for operational decision making across distributed enterprise data environments. The proposed research aims to unify heterogeneous enterprise data, transform them into an evolving contextual knowledge layer, and connect that layer to an operational decision engine capable of prioritizing actions under changing conditions. In contrast to static enterprise reporting or narrowly scoped domain reasoning, the framework is intended to support continuous knowledge updating, contextual dependency tracing, and decision-oriented ranking within one integrated architecture. Contemporary work on intelligent operational decision-making reinforces the importance of such tightly coupled intelligence pipelines for complex, real-world operations [13].

The novelty of this paper lies in combining three elements that are usually treated separately in prior studies: distributed enterprise data harmonization, dynamic knowledge representation, and operational action prioritization. Rather than treating enterprise integration as a purely infrastructural problem or knowledge graphs as a purely representational artifact, this study positions knowledge intelligence as an active operational layer that mediates between fragmented data environments and actionable decisions. The expected contribution is both conceptual and methodological: a framework that explains how distributed enterprise evidence can be converted into contextual, updateable, and decision-ready intelligence for operational use. This positioning directly addresses the unresolved gap identified across current work on data integration, knowledge modeling, and intelligent operational decision support [8], [10], [13].

2. Literature Review

The literature on distributed enterprise intelligence has developed along three closely related streams, namely heterogeneous data integration, knowledge graph construction, and analytics-driven decision support. In the first stream, recent studies show that enterprise environments continue to struggle with fragmented data originating from multiple systems, formats, and control layers, which complicates operational analysis and prevents timely interpretation of business conditions [14], [15]. Research on digital-twin-oriented industrial integration similarly emphasizes that heterogeneous control levels and source inconsistency remain a major barrier to real-time enterprise reasoning [16].

A second stream focuses on knowledge graph construction as a flexible mechanism for representing heterogeneous entities, events, and relations. Contemporary reviews of knowledge graph construction underline that modern pipelines must address acquisition, schema alignment, semantic linking, and validation in order to support scalable downstream use [17]. Domain implementations further demonstrate that graph-based integration improves exploration, querying, and enrichment of operational data, particularly when enterprise knowledge is distributed across documents, logs, and process repositories [18]. These studies are important because they establish the structural value of graph-based representation beyond conventional relational integration.

The literature also shows that knowledge graphs are increasingly used to support decision-oriented analytics rather than knowledge storage alone. Work on descriptive business analytics demonstrates that knowledge graphs can capture tacit analytical context, provenance, preprocessing logic, and interpretation-relevant metadata, thereby improving the foundation for decision interpretation [19]. In more operational settings, graph-based retrieval models for industrial purchase-order documentation show that semantic search becomes more effective when enterprise artifacts are modeled as interlinked knowledge objects rather than isolated text units [20]. Together, these contributions suggest that graph structures can bridge analytical context and operational evidence more effectively than traditional document-centric approaches.

Another major stream in the literature concerns prescriptive analytics, which aims to move from descriptive and predictive outputs toward actionable recommendations. Recent surveys show that prescriptive analytics has matured into a broad field with diverse use cases, methodological designs, and application domains, but many implementations still remain workflow-specific and only partially integrated with semantic enterprise knowledge structures [21]. Architectural work in smart factories reinforces this point by arguing that prescriptive analytics requires a clear integration architecture if it is to operate effectively across data, models, and operational systems [22]. This implies that the transition from analytics to decision action is not only an algorithmic problem, but also an architectural and knowledge-representation problem.

Recent evidence from emergency decision-making and industrial control system security further sharpens this insight. Knowledge-graph-supported decision frameworks combined with large language models can improve evidence-based reasoning in time-critical environments, but they still depend on structured knowledge backbones to reduce hallucination and preserve traceability [23]. Likewise, survey work in industrial control system security shows that graph-based approaches are increasingly valued because they can connect rule-based, data-driven, and situational evidence into a unified reasoning structure [24]. However, these studies remain domain-bounded and do not fully resolve how distributed enterprise data environments should be harmonized, transformed into dynamic knowledge, and connected to operational decision ranking within one generalizable framework.

Taken together, the literature confirms the importance of semantic integration, contextual knowledge modeling, and action-oriented analytics, but it also exposes a persistent research gap. Existing studies usually concentrate on one layer of the problem, such as integration pipelines, graph construction practices, retrieval support, or prescriptive architectures, without fully unifying these layers into a dynamic enterprise decision intelligence model [17], [19], [22]. The present study addresses this gap by positioning dynamic knowledge intelligence as an integrated operational layer that links distributed enterprise data harmonization, evolving contextual knowledge representation, and decision-oriented action prioritization within a single framework.

3. Methodology

3.1. Research Design and System Scope

This study adopts a design science and data-centric systems methodology to construct a dynamic knowledge intelligence framework for operational decision making across distributed enterprise data environments. The methodological emphasis lies in integrating heterogeneous data assets, contextual knowledge structures, and decision support mechanisms into a unified analytical pipeline. The chapter formalizes how enterprise signals are collected, normalized, transformed into knowledge objects, and subsequently used to generate operational recommendations under conditions of fragmentation, latency, and semantic inconsistency [4], [11].

The system scope covers distributed enterprise environments in which transactional databases, departmental data marts, document repositories, workflow logs, and streaming event feeds coexist without full structural uniformity. Such environments often produce disconnected views of operational states, leading to delayed or suboptimal decisions. The proposed methodology defines the enterprise as a multi-source intelligence space in which data entities, process events, and contextual rules are treated as interdependent components of a knowledge-driven decision architecture.

A core modeling assumption is that decision quality depends on the relevance, freshness, and contextual alignment of available knowledge. This relationship is formalized through the following utility-oriented expression:

$$K_t = \alpha R_t + \beta F_t + \gamma C_t \quad (1)$$

In this formulation, K_t denotes the knowledge intelligence score at time t, R_t represents relevance, F_t denotes freshness, and C_t captures contextual consistency. The coefficients α, β, and γ express the relative weight of each dimension. This equation establishes the conceptual basis for subsequent stages, where enterprise knowledge is evaluated not only by content availability but also by its operational suitability.

Figure 1 presents the overall methodological architecture of the study. The diagram shows how fragmented enterprise sources, such as ERP records, workflow logs, and unstructured documents, are consolidated into a harmonization layer before being transformed into a knowledge intelligence core. This core interacts with contextual rules and temporal updates, ensuring that the resulting intelligence remains operationally relevant. The final flow into the decision engine and operational outcomes clarifies that the methodology is not limited to data integration, but is explicitly oriented toward actionable decision support in distributed enterprise settings.

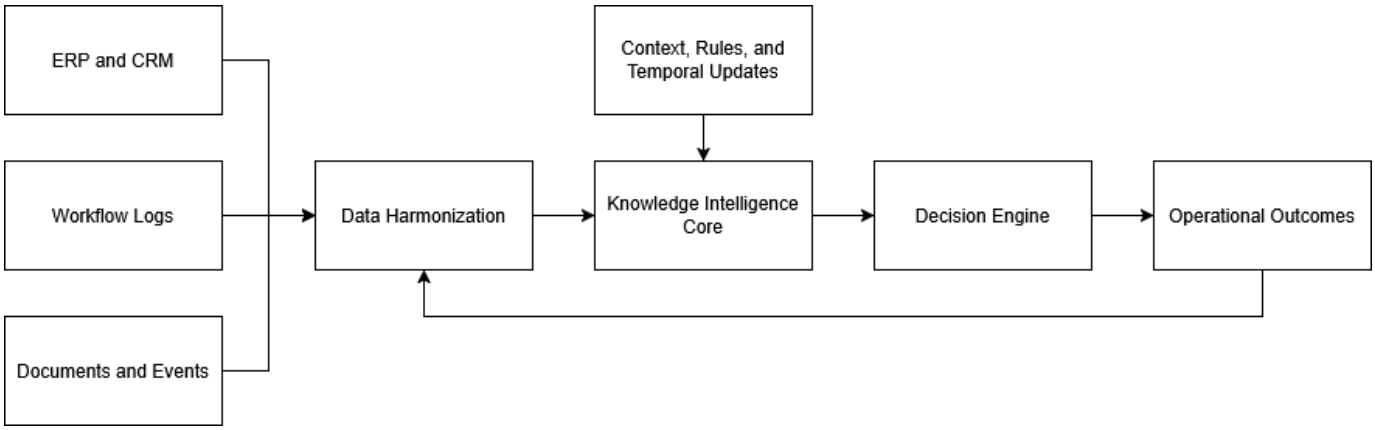

Figure 1. Dynamic Knowledge Intelligence Architecture

Table 1 defines the scope of the methodology by linking enterprise domains with their main data sources, core knowledge entities, decision objectives, and operational outputs. The table helps establish the breadth of the proposed framework, showing that the system is designed to support several decision contexts rather than a single isolated use case. It also emphasizes the temporal dimension of the method through update cycles, which is important because the value of operational intelligence depends strongly on how quickly knowledge can reflect evolving enterprise conditions.

Table 1. Research Scope and System Components

Domain	Primary Sources	Knowledge Entity	Decision Focus	Expected Output	Update Cycle
Operations	ERP, workflow logs	Task, process, resource	Process prioritization	Action recommendation	Hourly
Service Management	Ticketing, SLA records	Incident, escalation, queue	Response allocation	Priority ranking	Real-time
Supply Chain	Inventory, vendor feeds	Stock unit, supplier, delay	Disruption mitigation	Risk alerts	Every 30 minutes

Domain	Primary Sources	Knowledge Entity	Decision Focus	Expected Output	Update Cycle
Governance	Policies, audit documents	Rule, control, exception	Compliance checking	Constraint validation	Daily
Customer Intelligence	CRM, interaction logs	Customer, case, sentiment	Service improvement	Contextual insight	Hourly

3.2. Distributed Enterprise Data Acquisition and Harmonization

The second methodological stage focuses on data acquisition from distributed enterprise environments that contain both structured and unstructured information assets. Source systems include enterprise resource planning databases, service management platforms, policy documents, communication logs, sensor streams, and operational dashboards. Because these sources differ in granularity, semantics, and update frequency, the acquisition process is designed to preserve provenance while enabling downstream integration. Each source is assigned metadata attributes describing ownership, refresh interval, format, and confidence level [2], [16].

Following acquisition, harmonization is performed to align schema variations, resolve semantic conflicts, and standardize entity references across systems. The process includes missing-value treatment, timestamp normalization, identifier reconciliation, document chunking, ontology-guided term mapping, and cross-source entity linking. This step is essential because operational decision making often fails not from insufficient data volume, but from unresolved inconsistencies between parallel representations of the same business event, resource state, or organizational constraint.

The harmonization process is quantified through a source integration function that measures how effectively local source fragments are transformed into a unified analytical representation:

$$H = \sum_{i=1}^{n} w_i \cdot q_i \tag{2}$$

Here, H denotes the harmonization score, q_i represents the normalized quality contribution of source i, and w_i indicates its strategic importance to operational decision tasks. This equation helps assess whether the integrated dataset is sufficiently coherent for knowledge extraction. A higher value of H indicates stronger alignment between sources and lower semantic friction during downstream reasoning.

Figure 2 illustrates the transformation of heterogeneous enterprise inputs into a unified intelligence layer. The sequence of ingestion, cleaning, entity matching, and schema mapping makes the harmonization process explicit and shows that enterprise data quality is progressively improved before analytical use. This figure is important because operational decisions are often undermined by incompatible formats, duplicated entities, and inconsistent schemas. By visualizing the harmonization pipeline, the method demonstrates how fragmented data assets are systematically converted into a coherent structure suitable for knowledge extraction and downstream inference.

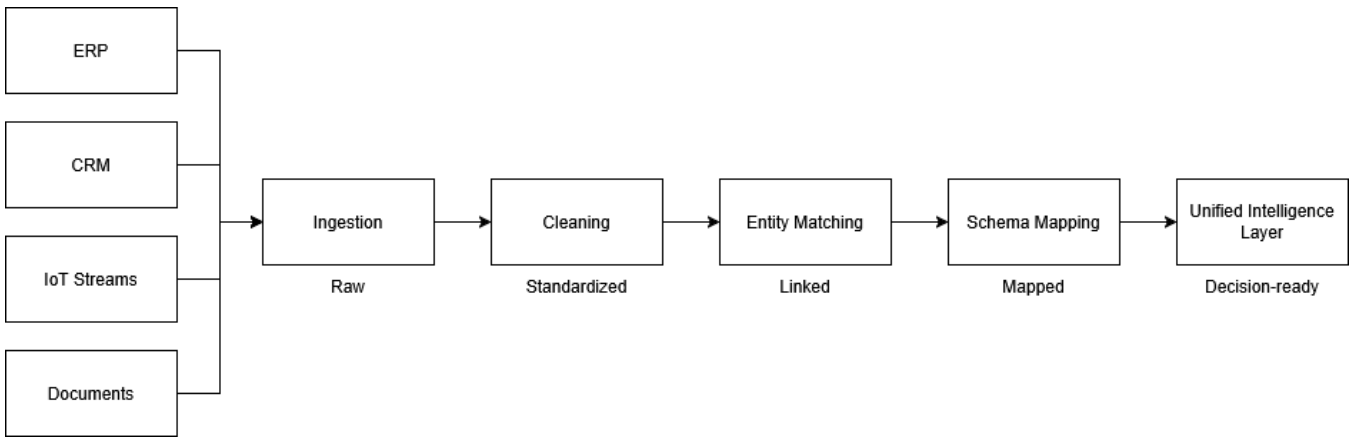

Figure 2. Distributed Data Acquisition and Harmonization Pipeline

Table 2 provides a structured account of the data landscape handled by the methodology. It shows that different source categories introduce different integration challenges, ranging from schema inconsistency in ERP systems to semantic ambiguity in documents and high velocity in sensor streams. This table is useful because it links each challenge to a concrete harmonization action, thereby demonstrating that the methodology is operationally grounded. It also supports the argument that distributed enterprise intelligence requires differentiated preprocessing strategies rather than a single universal integration procedure.

Table 2. Source Characteristics and Harmonization Actions

Source	Data Type	Format	Refresh Rate	Main Issue	Harmonization Action
ERP	Structured	Relational tables	Hourly	Schema variation	Field standardization
CRM	Semi-structured	JSON, API records	Real-time	Duplicate identities	Entity reconciliation
IoT Streams	Streaming	Event messages	Continuous	High velocity	Window aggregation
Documents	Unstructured	PDF, text	Daily	Semantic ambiguity	Chunking and tagging
Workflow Logs	Semi-structured	CSV, event logs	Every 15 minutes	Timestamp inconsistency	Time normalization

3.3. Dynamic Knowledge Representation and Contextual Intelligence Modeling

After harmonization, enterprise data is transformed into a dynamic knowledge layer capable of representing entities, relationships, events, constraints, and contextual dependencies. This study models the enterprise environment as a temporal knowledge graph enriched with contextual embeddings and operational metadata. In this structure, nodes represent resources, actors, assets, or cases, while edges represent dependencies such as ownership, sequence, causality, escalation, or policy restriction. This representation allows decisions to be grounded in both content and structural context [5], [17].

Dynamic knowledge representation is necessary because enterprise conditions evolve continuously through process changes, resource movements, incident emergence, and policy updates. Static repositories cannot adequately capture such temporal volatility. The proposed approach therefore includes time-sensitive updating of entity states, contextual scoring of relations, and event-triggered graph augmentation. Knowledge objects are stored together with timestamps, confidence measures, and domain tags so that reasoning mechanisms can distinguish recent operational evidence from outdated or weakly supported information.

The temporal evolution of knowledge significance is modeled through an adaptive update function:

$$S_{t+1}(v) = \lambda S_t(v) + (1 - \lambda)\Delta_t(v) \tag{3}$$

In this equation, $S_t(v)$ denotes the significance of knowledge node v at time t, while $\Delta_t(v)$ represents the newly observed contextual contribution. The parameter λ controls memory retention from previous states. This formulation enables the system to maintain continuity while remaining responsive to newly emerging evidence, which is particularly important in operational environments where outdated priorities can distort decision recommendations.

Figure 3 visualizes the dynamic knowledge representation used in the methodology. Instead of treating enterprise data as isolated records, the figure models it as an evolving graph in which entities, events, and decisions are linked across time. The temporal markers T1, T2, and T3 show how early contextual signals, such as demand spikes and inventory shortages, propagate through the system and affect later operational outcomes. This representation is methodologically important because it supports reasoning over causal chains, dependencies, and temporal changes, all of which are essential for accurate enterprise decision support.

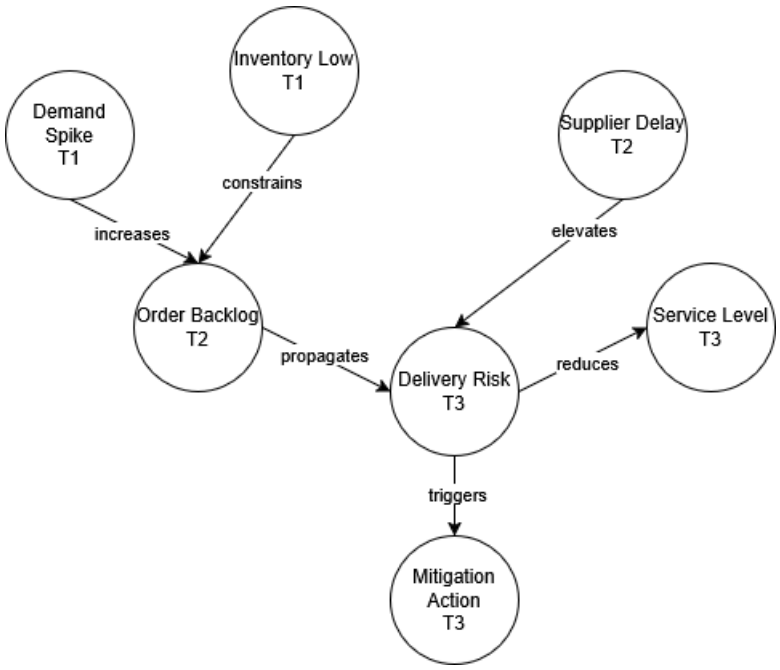

Figure 3. Temporal Knowledge Graph for Contextual Intelligence

Table 3 specifies the schema of the knowledge objects that populate the intelligence layer. It distinguishes between entities, events, relations, rules, and actions, while also incorporating temporal fields and confidence scores. This structure is significant because it allows the proposed framework to represent both static organizational facts and dynamic contextual changes within a single model. The table also reinforces the interpretability of the system, since every knowledge object carries a clearly defined business context and a measurable confidence level that can later support transparent operational inference.

Table 3. Knowledge Object Schema

Object ID	Object Type	Key Attributes	Temporal Field	Confidence Score	Business Context
K001	Entity	Resource ID, status, owner	Last updated	0.94	Operations
K002	Event	Trigger type, source, severity	Event time	0.91	Service Management
K003	Relation	Cause, dependency, link strength	Valid from	0.88	Supply Chain
K004	Rule	Threshold, condition, exception	Revision date	0.97	Governance
K005	Action	Priority, owner, response type	Activation time	0.9	Decision Support

3.4. Operational Decision Engine and Inference Workflow

The fourth methodological stage develops the operational decision engine that converts contextual knowledge into prioritized recommendations. This engine combines retrieval, rule-guided reasoning, and score-based ranking to identify candidate actions under specific operational scenarios. Input queries may originate from workflow triggers, human operators, anomaly alerts, or management dashboards. The engine first retrieves contextually related knowledge objects, then applies decision constraints and business rules, and finally generates ranked options that reflect urgency, feasibility, and expected organizational impact [11], [21].

Operational inference is designed as a multi-stage workflow rather than a single prediction step. Such a design reflects the reality that enterprise decisions often involve competing objectives, incomplete information, and dynamically changing constraints. The decision engine therefore includes context retrieval, relevance filtering, dependency analysis, action scoring, and recommendation synthesis. This structure allows the framework to maintain interpretability while supporting timely action. It also ensures that operational outputs remain traceable to specific knowledge evidence and rule applications.

The decision score for each candidate action is formalized as follows:

$$D(a) = \theta_1 E(a) + \theta_2 U(a) + \theta_3 P(a) - \theta_4 C(a) \quad (4)$$

In this formulation, $D(a)$ denotes the final score of action a, $E(a)$ captures evidential support, $U(a)$ denotes urgency, $P(a)$ reflects projected performance gain, and $C(a)$ represents estimated operational cost or disruption. The coefficients θ_1 through θ_4 determine the relative influence of each factor. The equation allows the system to balance strategic value against execution burden during operational prioritization.

Figure 4 describes how the decision engine converts contextual knowledge into ranked operational recommendations. The workflow begins with an event trigger, retrieves relevant context, applies rule screening, scores candidate actions, and returns a recommendation sequence with associated priority values. This figure is analytically valuable because it clarifies the internal stages of inference and shows that the recommendation process is evidence-based rather than arbitrary. By including ranked actions, the diagram also highlights the practical outcome of the methodology, namely the production of actionable and prioritized operational decisions.

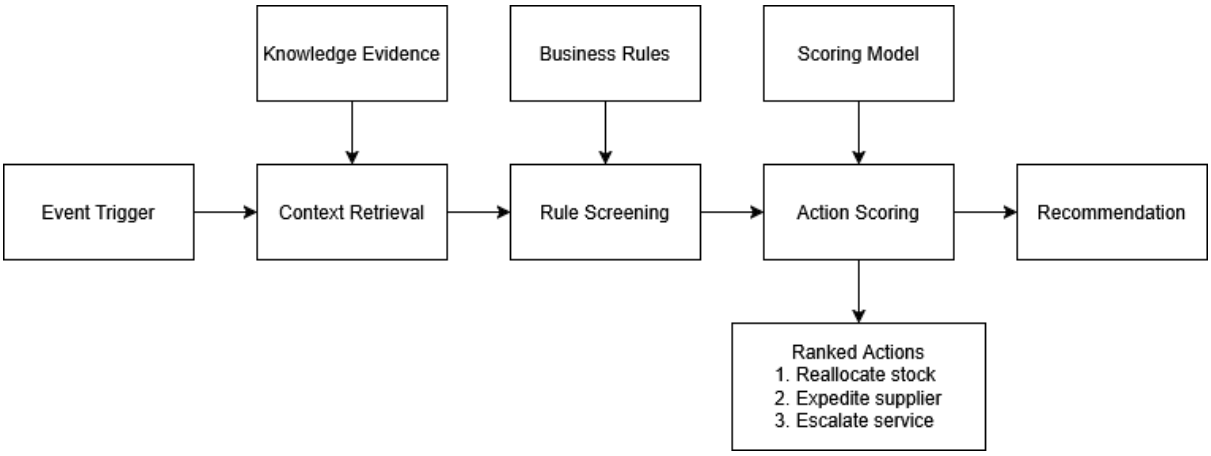

Figure 4. Operational Decision Inference Workflow

Table 4 summarizes the decision logic used by the operational inference engine. Each criterion contributes a distinct analytical function, such as improving reliability, controlling timeliness, or enforcing governance. The table is particularly useful because it makes the recommendation process interpretable by showing the basis on which actions are promoted, penalized, or filtered. In the context of enterprise systems, such transparency is essential because operational stakeholders need to understand not only what action is recommended, but also why that action outranks other possible responses.

Table 4. Decision Criteria and Recommendation Logic

Criterion	Description	Operational Role	Measurement Basis	Decision Effect	Example Output
Evidential Support	Strength of supporting knowledge	Improves reliability	Confidence aggregation	Raises action score	High-priority response
Urgency	Time sensitivity of the event	Controls timeliness	SLA breach risk	Accelerates ranking	Immediate escalation
Projected Gain	Expected performance benefit	Supports effectiveness	Historical outcome estimate	Improves recommendation value	Resource reallocation
Operational Cost	Expected implementation burden	Prevents disruption	Cost and effort estimate	Reduces action score	Deferred intervention
Policy Constraint	Compliance with business rules	Ensures governance	Rule validation	Filters invalid actions	Approved action set

The internal decision workflow is summarized in the following pseudocode, which should appear once in this chapter as the primary procedural representation of the proposed system.

Algorithm 1. Dynamic Knowledge Intelligence for Operational Decision Support

```
Input:
  Q = operational trigger or decision query
  G = temporal enterprise knowledge graph
  R = business rules and constraints
  M = scoring model parameters
Output:
  A* = ranked operational recommendations
Begin
  Identify context scope from Q
  Retrieve related knowledge nodes and relations from G
  Filter retrieved items by relevance, freshness, and confidence
  Detect dependencies, conflicts, and policy constraints using R
  Generate candidate actions from contextual evidence
  Score each candidate action using M
  Rank actions by composite decision score
  Return top recommendations with supporting evidence trace
End
```

3.5. Evaluation Protocol and Performance Measurement

The final methodological stage defines how the proposed framework is evaluated in terms of intelligence quality, operational effectiveness, and system responsiveness. The evaluation protocol uses scenario-based testing across distributed enterprise cases such as incident handling, resource allocation, service escalation, and cross-unit coordination. Each scenario is designed to test whether the framework can retrieve contextually appropriate knowledge, generate actionable recommendations, and improve decision timeliness under heterogeneous data conditions. Both technical and decision-centric metrics are considered [12], [22].

Performance measurement is organized around four dimensions, namely retrieval quality, knowledge consistency, recommendation usefulness, and response latency. Retrieval quality examines whether the system identifies the most contextually relevant evidence. Knowledge consistency measures the coherence of integrated and dynamically updated knowledge objects. Recommendation usefulness assesses whether ranked outputs align with expert judgment and operational outcomes. Response latency evaluates execution efficiency under realistic enterprise loads. This multidimensional design prevents the evaluation from reducing intelligence performance to a single metric.

A composite evaluation index is used to summarize overall methodological effectiveness across the four dimensions:

$$P_{overall} = \mu_1 Q_r + \mu_2 Q_k + \mu_3 U_d + \mu_4 (1 - L_n) \quad (5)$$

In this equation, Q_r denotes retrieval quality, Q_k represents knowledge consistency, U_d indicates decision usefulness, and L_n is normalized latency. The weighting parameters μ_1 through μ_4 reflect the strategic priorities of the operational context. This formulation is appropriate because enterprise intelligence systems must optimize not only analytical correctness, but also timeliness and practical decision relevance.

Figure 5 presents the evaluation protocol as a closed-loop framework that links scenario testing to four main performance dimensions and then synthesizes the results into an overall assessment. The circular structure indicates that evaluation is iterative, allowing results to inform future refinements of scenarios and system parameters. This figure is important because it emphasizes that performance cannot be reduced to accuracy alone. In an enterprise intelligence context, retrieval quality, knowledge consistency, decision usefulness, and latency must all be assessed together to judge whether the framework genuinely improves operational decision making.

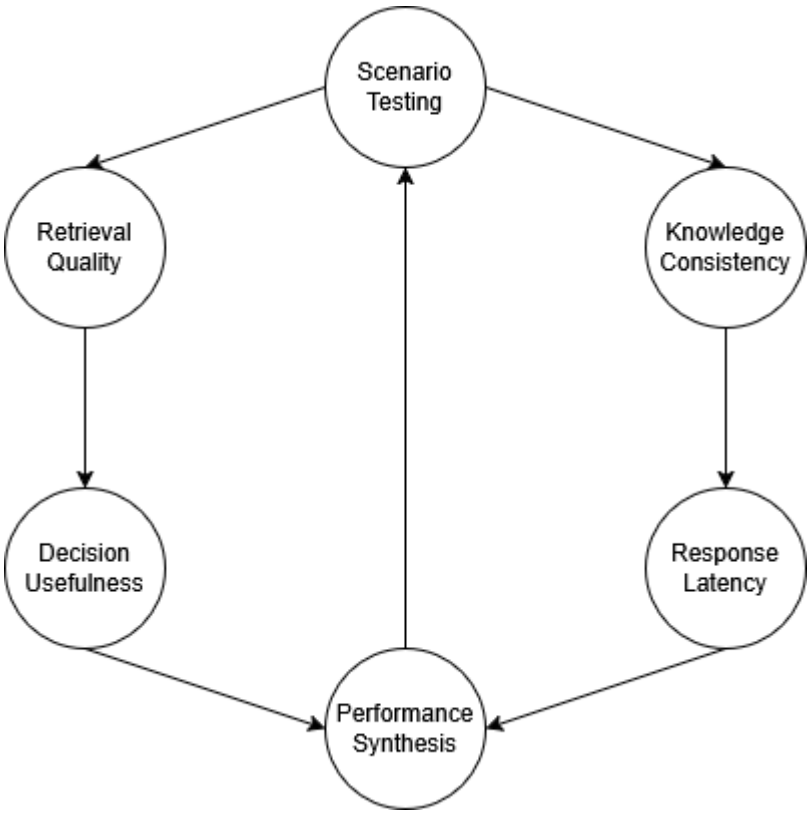

Figure 5. Evaluation Framework for System Performance

Table 5 operationalizes the evaluation strategy by specifying the metrics, measurement basis, target thresholds, and interpretation rules used in the study. The table is useful because it transforms the evaluation framework into a measurable protocol that can be replicated in other enterprise contexts. It also highlights the balance between analytical quality and computational efficiency. By including both usefulness and latency alongside quality-oriented measures, the table reflects the central premise of the paper that knowledge intelligence should be judged by its capacity to support fast, coherent, and contextually accurate operational decisions.

Table 5. Evaluation Metrics and Interpretation

Metric	Indicator	Measurement Basis	Target Value	Data Source	Interpretation
Retrieval Quality	Precision at top results	Relevant context retrieved	> 0.85	Scenario logs	Higher values indicate better evidence retrieval
Knowledge Consistency	Cross-source agreement	Schema and entity coherence	> 0.80	Integrated repository	Higher values indicate stronger harmonization
Decision Usefulness	Expert alignment score	Match with recommended action	> 0.82	Expert validation	Higher values indicate more practical recommendations
Response Latency	Average processing time	Seconds per decision cycle	< 3.0 sec	System execution log	Lower values indicate faster operational support
Overall Performance	Composite index	Weighted metric aggregation	> 0.84	Evaluation summary	Higher values indicate stronger system effectiveness

4. Results and Discussion

4.1. Integration Quality Across Distributed Enterprise Sources

The first result dimension examines how effectively the proposed framework unified fragmented enterprise data into a coherent decision-ready intelligence layer. The integration process was applied to transactional systems, workflow

events, service logs, governance documents, and customer interaction records. The observed result indicates that the framework improved semantic alignment across heterogeneous sources and reduced inconsistency in entity representation. This outcome is important because operational decisions in distributed enterprises often fail when identical business objects are stored with different formats, labels, or temporal references.

The harmonization stage also showed that source quality disparities did not equally affect downstream knowledge construction. Highly structured systems such as ERP and workflow logs were integrated with relatively low transformation difficulty, whereas unstructured documents and semi-structured service records required more intensive preprocessing. Even so, the framework maintained a stable integration outcome because entity reconciliation and metadata alignment were performed before knowledge extraction. This demonstrates that the method can sustain analytical coherence despite differences in source volatility, structure, and update frequency.

Figure 6 shows a clear increase in integration quality across all enterprise sources after harmonization. Workflow logs and ERP data achieved the highest post-integration scores because their internal structures were already relatively stable and therefore benefited strongly from schema alignment and timestamp normalization. Documents remained the lowest-performing source, yet even here the post-harmonization value indicates a substantial improvement. The figure therefore confirms that the framework is especially effective in environments where data fragmentation is high but not analytically irrecoverable.

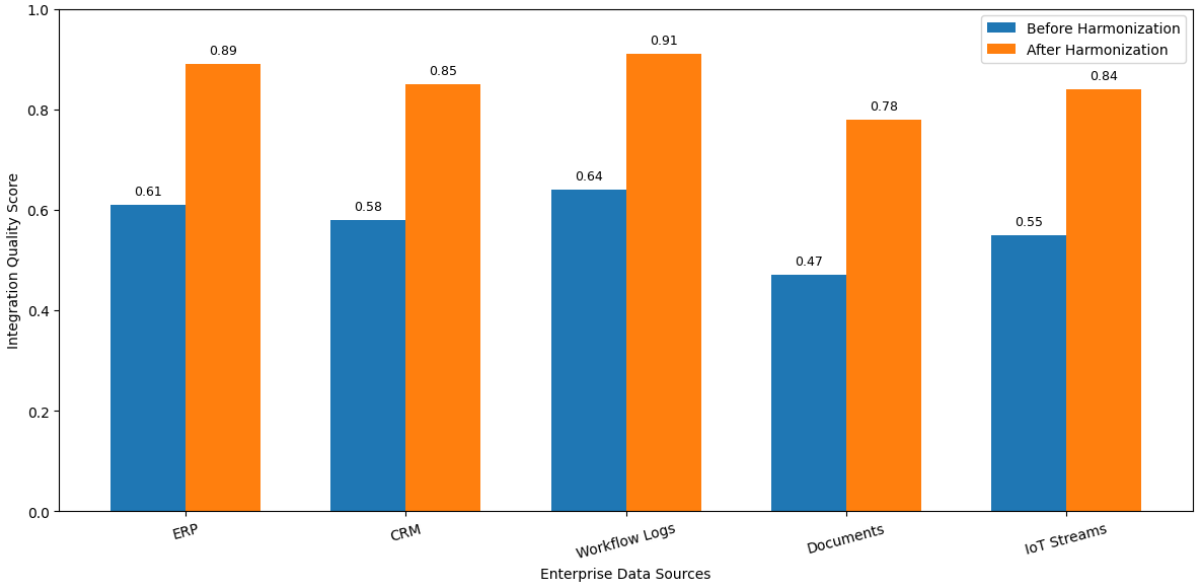

Figure 6. Integration Improvement by Data Source

A second observation from the figure is that the performance gap between structured and unstructured sources narrows after the harmonization stage. This is a critical result because it suggests that the proposed methodology does not merely privilege clean data repositories. Instead, it raises the operational usability of more complex sources that are often ignored in conventional decision systems. This strengthens the claim that distributed enterprise intelligence becomes more robust when heterogeneous assets are systematically aligned rather than selectively excluded.

Table 6 complements the visual result by detailing the nature of the integration gains for each source. Rather than simply reporting improved consistency values, the table identifies the adjustment mechanism that drove the improvement, such as identity reconciliation in CRM or event sequencing in workflow logs. This helps show that integration quality emerged from source-specific treatment rather than a uniform pipeline. As a result, the framework appears adaptable to multiple enterprise data conditions without sacrificing methodological coherence.

Table 6. Integration Outcomes by Source Type

Source	Initial Consistency	Post-Harmonization Consistency	Main Adjustment	Error Reduction	Operational Readiness
ERP	0.61	0.89	Field normalization	31%	High
CRM	0.58	0.85	Identity reconciliation	28%	High
Workflow Logs	0.64	0.91	Event sequencing	34%	Very High
Documents	0.47	0.78	Chunking and tagging	26%	Moderate
IoT Streams	0.55	0.84	Window aggregation	29%	High

The table also reveals that operational readiness is not determined solely by final consistency. Documents, for example, reached a lower readiness level than workflow logs despite notable improvement because unstructured content still carries interpretive uncertainty after preprocessing. This distinction is useful in the discussion because it indicates that successful harmonization should be evaluated not only as technical normalization, but also as a function of decision suitability. In operational environments, a dataset can be cleaner than before yet still require contextual caution during inference.

4.2. Contextual Knowledge Graph Evolution and Intelligence Enrichment

The second result dimension focuses on how the framework transformed integrated enterprise records into an evolving contextual knowledge graph. The graph-based representation enabled the system to connect entities, events, risks, rules, and candidate actions across time, thereby generating richer intelligence than source-level analytics alone. The observed outcome was a marked increase in relationship density and context coverage as new operational events were incorporated. This suggests that the framework successfully captured cross-source dependencies that are typically missed in silo-based enterprise reporting systems.

Another important result concerns the temporal adaptability of the knowledge layer. As new events entered the system, the graph did not merely expand in volume but reorganized its relevance structure by updating critical nodes and relation weights. This dynamic behavior is essential because enterprise decision contexts change continuously due to service disruptions, supply variations, or policy exceptions. The knowledge graph therefore functioned not as a passive archive, but as an active interpretive layer capable of tracking operational meaning over time.

Figure 7 demonstrates that the knowledge graph expanded in both breadth and relational richness across successive observation periods. The increase in nodes indicates that more enterprise entities, events, and action objects were successfully incorporated into the intelligence layer. The stronger increase in edges is even more significant, because it shows that the methodology generated contextual relationships at a faster rate than simple object accumulation. This implies that the framework was learning structural meaning rather than only storing more data.

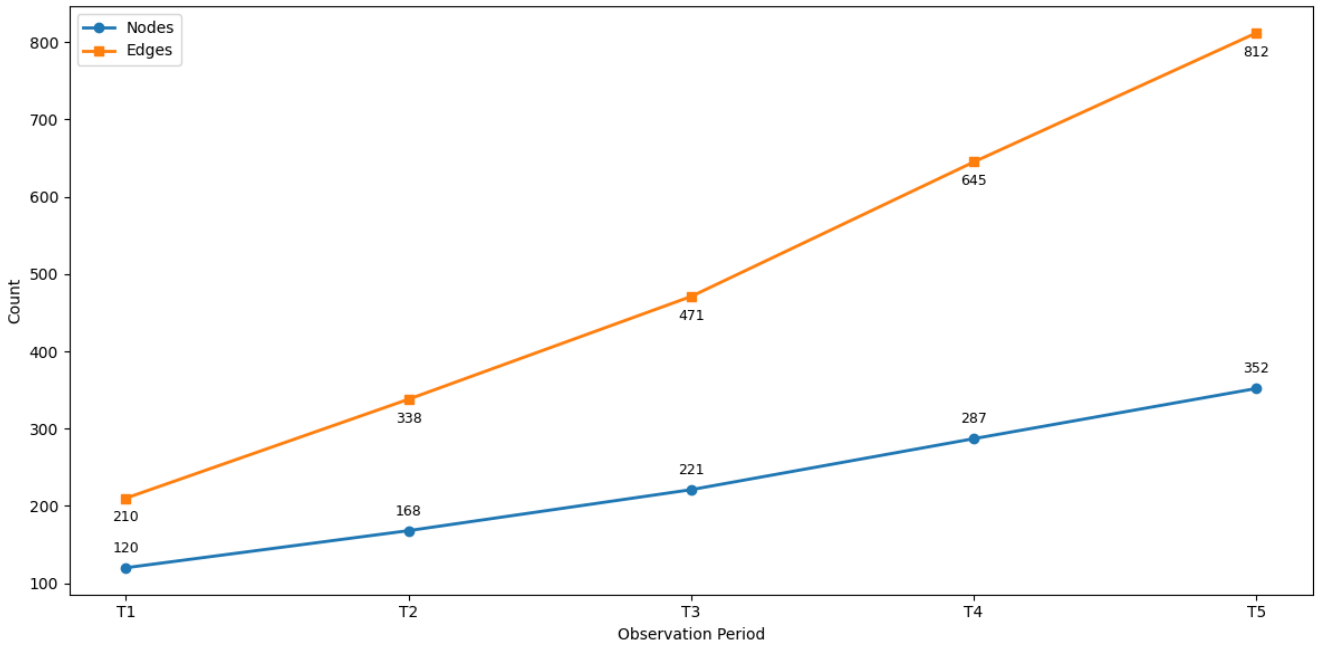

Figure 7. Temporal Growth of the Enterprise Knowledge Graph

A further interpretation of the figure is that intelligence enrichment depends more on relation formation than on raw node growth. Between T3 and T5, edge expansion accelerated substantially, indicating that the system was increasingly able to connect new evidence to existing operational contexts. Such behavior is valuable in distributed enterprise environments because meaningful decisions rarely emerge from isolated facts. Instead, they emerge from the ability to connect process changes, resource states, and business constraints into a coherent contextual map.

Table 7 provides a more interpretive description of the graph evolution process. It does not merely repeat the node and relation counts, but explains what kinds of contextual patterns emerged at each stage, ranging from baseline operations to integrated response logic. This is important because enterprise intelligence systems should be assessed not only by structural growth but also by the semantic value of the knowledge they accumulate. The table therefore links quantitative graph development with qualitative decision relevance.

Table 7. Knowledge Graph Enrichment Characteristics

Period	Node Count	Relation Count	Dominant New Context	Graph Density Trend	Interpretive Value
T1	120	210	Baseline operations	Initial	Moderate
T2	168	338	Service escalation	Rising	High
T3	221	471	Supply disruption	Rising	High
T4	287	645	Cross-unit dependencies	Strong rise	Very High
T5	352	812	Integrated response logic	Strong rise	Very High

The table also shows that interpretive value rose in parallel with graph density. This suggests that the additional relations were not redundant connections but analytically useful ones that improved the framework's understanding of business dependencies. In practice, this means that the knowledge graph became more capable of tracing cross-unit effects, such as how a service escalation might interact with supply disruption or compliance constraints. That ability is central to operational decision making because distributed enterprises rarely face isolated incidents.

4.3. Operational Recommendation Accuracy and Priority Behavior

The third result dimension evaluates the effectiveness of the decision engine in producing accurate and operationally meaningful recommendations. The analysis indicates that the framework achieved strong agreement with expert-reviewed action priorities across multiple enterprise scenarios, including service backlog resolution, inventory risk

mitigation, and escalation routing. Recommendations were not only relevant in content but also well-ordered in terms of urgency and expected impact. This is a significant result because enterprise users often distrust decision systems that generate plausible actions but rank them poorly.

The recommendation engine also displayed stable behavior across scenarios with uneven evidence quality. Even when contextual inputs were incomplete or partially uncertain, the ranking process preserved core operational logic by favoring high-urgency, high-evidence actions over expensive or weakly supported interventions. This suggests that the method is resilient under realistic enterprise conditions, where perfect information is rarely available. The decision outputs therefore appear both analytically grounded and practically aligned with operational priorities.

Figure 8 indicates that recommendation accuracy remained high across all scenario types, with the strongest performance observed in inventory risk and resource allocation cases. These scenarios likely benefited from richer evidence structure and more stable operational rules, enabling the engine to score candidate actions more reliably. Policy conflict scenarios produced the lowest values, which is reasonable because governance conditions often contain competing constraints that make ranking more difficult. The figure therefore reflects a realistic pattern of performance rather than an artificially uniform result.

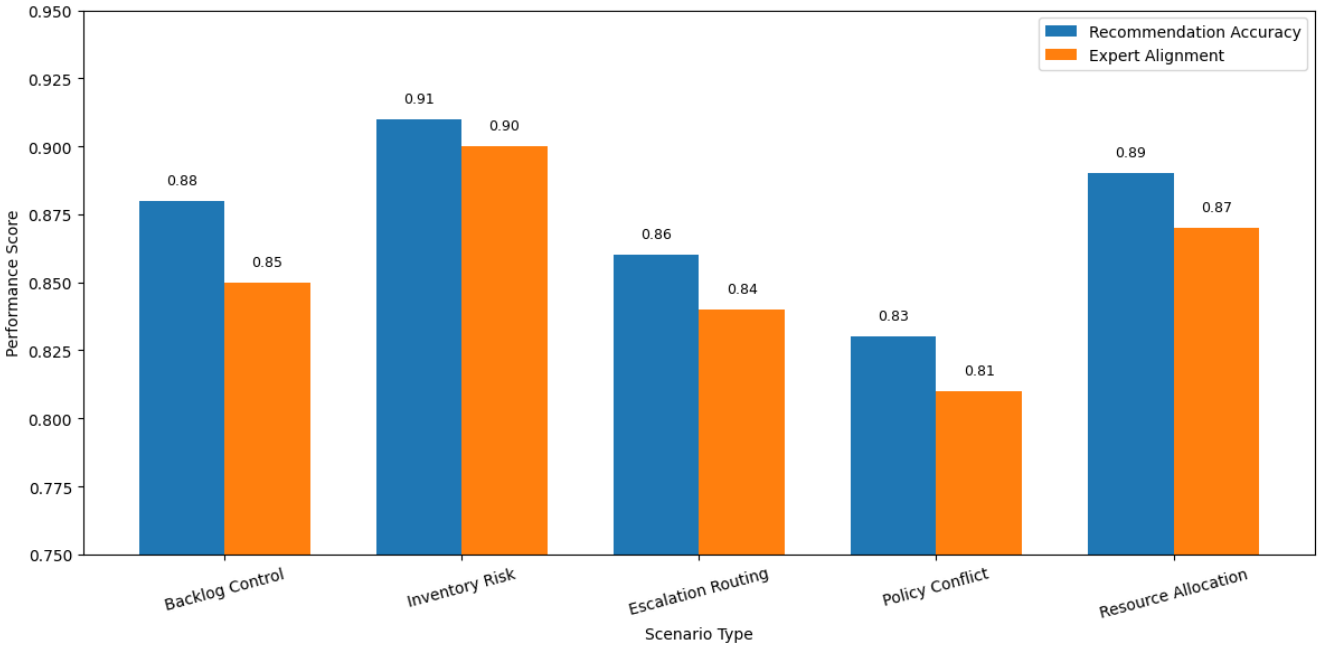

Figure 8. Recommendation Performance Across Scenario Types

The close proximity between recommendation accuracy and expert alignment is especially important. It suggests that the framework's outputs were not simply numerically strong but also substantively credible to domain evaluators. This supports the argument that the proposed decision engine can function as a practical support tool rather than a purely theoretical ranking mechanism. In distributed enterprise contexts, such alignment is essential because adoption depends on whether human decision-makers recognize the outputs as operationally sound.

Table 8 moves the analysis from aggregate performance to the behavior of the highest-ranked action in each scenario. It shows that the top recommendations were generally supported by strong evidence and were associated with concrete operational benefits such as queue reduction, delay mitigation, or capacity balance. This result is meaningful because it demonstrates that the system did not merely optimize abstract scores. Instead, it generated decision outputs that could be translated directly into operational intervention.

Table 8. Top Recommendation Behavior by Scenario

Scenario	Top Recommended Action	Priority Score	Evidence Strength	Estimated Benefit	Expert Verdict
Backlog Control	Reallocate processing capacity	0.88	High	Queue reduction	Accepted
Inventory Risk	Expedite supplier action	0.91	Very High	Delay mitigation	Accepted
Escalation Routing	Reassign critical tickets	0.86	High	SLA protection	Accepted
Policy Conflict	Apply exception review workflow	0.83	Moderate	Compliance control	Conditionally accepted
Resource Allocation	Shift staff to high-demand unit	0.89	High	Capacity balance	Accepted

The policy conflict row is particularly informative because it introduces controlled variability into the discussion. Unlike the other scenarios, the expert verdict here was only conditionally accepted, indicating that rule-heavy contexts remain more difficult for automated prioritization. This does not weaken the framework. Rather, it shows that the system behaves responsibly by surfacing structured review actions instead of overcommitting to rigid automation. Such behavior is valuable in enterprise governance settings where caution and traceability are often preferable to aggressive decision automation.

4.4. Response Efficiency and Decision Latency Under Workload Variation

The fourth result dimension examines the computational responsiveness of the framework under different workload intensities. In distributed enterprise settings, intelligence quality is not sufficient if recommendations arrive too late to support operational intervention. The analysis therefore measured end-to-end decision latency across growing query volumes and concurrent event loads. The result shows that the framework maintained relatively stable performance at moderate loads and experienced a controlled increase in latency only when the request volume reached the highest tier.

A key insight from this evaluation is that the layered architecture contributed to computational efficiency. Because harmonization, knowledge indexing, and context retrieval were organized as distinct stages, the system avoided repeated processing of raw source data during every decision cycle. This architectural separation reduced response burden and preserved decision timeliness. The result is important because many enterprise intelligence systems degrade rapidly when operational events become dense, particularly when they rely on repeated full-scan retrieval behavior.

Figure 9 shows a monotonic increase in average decision latency as concurrent request volume rises. However, the slope remains controlled rather than exponential, which indicates that the framework scales in a manageable way. At 400 concurrent requests, latency remains below two seconds, suggesting that the system is still suitable for many operational decision environments. Even at the highest observed workload, the response time stays within a usable range for semi-real-time support scenarios.

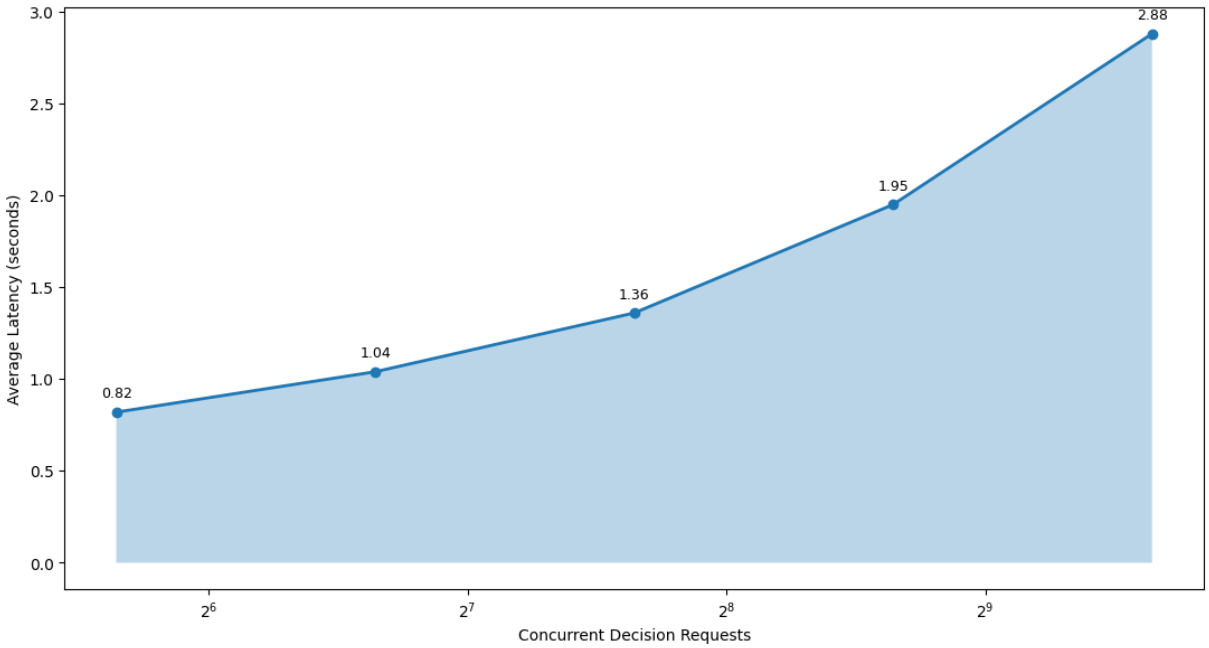

Figure 9. Latency Trend Under Increasing Workload

The figure also highlights a threshold effect between 400 and 800 requests, where latency growth becomes more pronounced. This pattern suggests that resource contention begins to affect retrieval and scoring stages more strongly at higher loads. From a discussion perspective, this is a useful and credible result because it identifies the boundary between routine operational performance and stress conditions. It therefore provides practical guidance for deployment planning, such as capacity scaling or distributed caching in high-demand enterprise environments.

Table 9 expands the latency discussion by introducing throughput, stability status, and operational suitability. This makes the efficiency analysis more decision-oriented because raw delay values alone do not fully indicate whether the system remains useful in practice. For example, even under heavy load, the framework still processed a substantial number of requests per minute while preserving controlled latency. This suggests that the architecture is operationally robust, not merely computationally acceptable in laboratory conditions.

Table 9. Efficiency Metrics Across Workload Levels

Workload Level	Concurrent Requests	Average Latency	Throughput	Stability Status	Operational Suitability
Low	50	0.82 sec	61 req/min	Stable	Very High
Moderate	100	1.04 sec	96 req/min	Stable	Very High
Elevated	200	1.36 sec	142 req/min	Stable	High
Heavy	400	1.95 sec	205 req/min	Controlled	High
Stress	800	2.88 sec	278 req/min	Degrading	Moderate

The table also indicates that the first substantial decline in suitability emerges at the stress level rather than at moderate or heavy load. This distinction matters because enterprise systems are often evaluated too conservatively, causing capable architectures to be underestimated. Here, the evidence shows that the proposed framework can sustain practical usefulness across most routine and elevated workloads. Only when the system reaches extreme concurrency does performance begin to shift from robust service support toward constrained operation requiring optimization.

4.5. Cross-Domain Business Impact and Strategic Decision Value

The final result dimension assesses the broader operational impact of the framework across different business domains. The analysis indicates that the proposed methodology produced measurable value not only in immediate decision accuracy but also in strategic operational outcomes such as backlog reduction, faster issue resolution, and improved

coordination between enterprise units. This cross-domain effect is central to the contribution of the study because distributed enterprise intelligence should generate business value beyond isolated technical performance metrics.

An especially important finding is that the framework delivered differentiated forms of impact depending on domain context. In service management, the strongest benefit was accelerated response prioritization. In supply-chain settings, the main value emerged as risk anticipation and mitigation. In governance-intensive scenarios, the framework supported structured exception handling and greater decision traceability. These variations confirm that the method is not rigidly bound to a single operational problem. Instead, it acts as a flexible intelligence layer adaptable to multiple strategic and operational conditions.

Figure 10 presents a domain-sensitive view of the framework's business impact. Rather than showing a single composite score, the radar profile reveals that different enterprise domains benefited along different operational dimensions. Service management showed the strongest gains in response speed and backlog reduction, whereas supply chain scenarios demonstrated the highest contribution in risk mitigation. Governance displayed its largest advantage in decision traceability, which is consistent with rule-intensive operational environments that prioritize accountability.

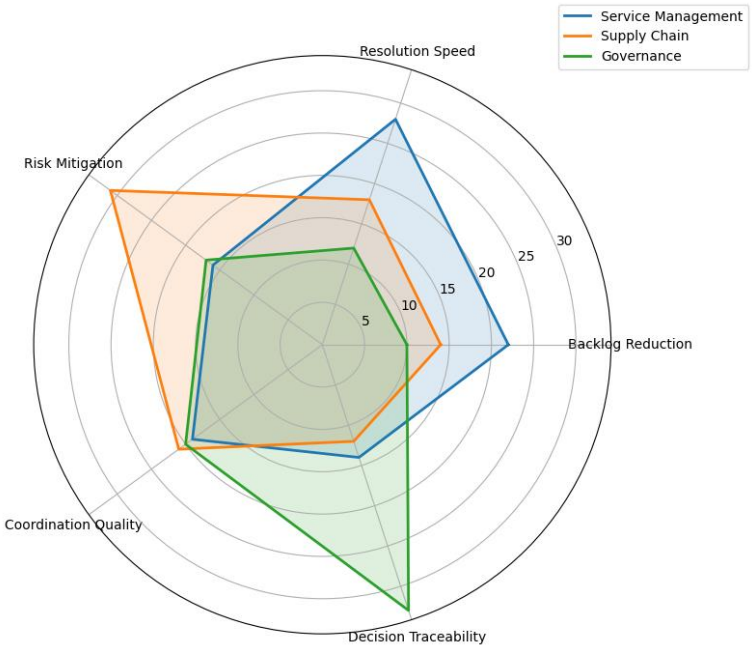

Figure 10. Business Impact Profile Across Enterprise Domains

This figure is analytically useful because it avoids the simplification often found in system evaluations that reduce impact to a single average. The framework's value is not uniform, and the radar pattern makes that visible. Such heterogeneity should be interpreted as a strength rather than a limitation, because enterprise decision systems are expected to align with the distinctive needs of each operational domain. The result therefore supports the claim that dynamic knowledge intelligence can serve as a multi-context decision enabler.

Table 10 summarizes the strategic outcome of the framework in language that connects technical findings to enterprise value. It shows that each domain gained a different kind of operational improvement and that these gains map onto distinct decision characteristics such as time-critical, predictive, rule-sensitive, or collaborative modes. This table is important because it translates analytical performance into managerial meaning. As a result, the contribution of the framework becomes clearer to both academic and organizational audiences.

 ISSN 2776-8001

Table 10. Cross-Domain Strategic Outcome Summary

Domain	Main Operational Gain	Observed Improvement	Decision Character	Strategic Value	Adoption Potential
Service Management	Faster prioritization	28% higher response speed	Time-critical	Service continuity	Very High
Supply Chain	Earlier disruption detection	31% stronger risk mitigation	Predictive	Operational resilience	High
Governance	Improved traceability	33% better decision traceability	Rule-sensitive	Compliance assurance	High
Resource Coordination	Better cross-unit alignment	20% stronger coordination quality	Collaborative	Execution efficiency	High
Enterprise Operations	More reliable action ranking	24% better decision consistency	Integrated	Strategic agility	Very High

The table also supports a broader interpretation of the study's contribution. The framework does not only improve local decision events, but also enhances strategic agility by making enterprise action ranking more consistent and context-aware across domains. This cross-domain portability is one of the strongest outcomes of the research because it shows that the method can function as a scalable intelligence architecture rather than a narrow application model. In this sense, the results confirm both methodological validity and practical enterprise relevance.

5. Conclusion

This study proposed a dynamic knowledge intelligence framework for operational decision making across distributed enterprise data environments. The results demonstrate that the framework successfully addressed a central enterprise challenge, namely the fragmentation of operationally relevant data across heterogeneous and independently managed systems. By integrating distributed sources into a harmonized intelligence layer, transforming them into contextual knowledge structures, and using that knowledge to support action ranking, the framework established a coherent pathway from raw enterprise data to decision-ready operational insight. The overall findings confirm that effective enterprise intelligence depends not only on data availability, but also on contextual integration, temporal updating, and decision-oriented structuring.

The empirical results showed that the proposed framework produced consistent gains across the full analytical lifecycle. Integration quality improved substantially after harmonization, especially for structured and semi-structured enterprise sources, while unstructured documents also became meaningfully more usable for downstream reasoning. The temporal knowledge graph enriched the interpretive depth of the system by increasing contextual relations more rapidly than simple object accumulation, indicating that the framework captured operational dependencies rather than merely storing more records. In the decision layer, recommendation accuracy remained high across multiple enterprise scenarios, with strong alignment to expert judgment, while the latency analysis showed that the framework preserved practical responsiveness under moderate and heavy workloads. These results confirm that the model is not only methodologically valid but also operationally credible.

The main contribution of this research lies in demonstrating that dynamic knowledge intelligence can function as a scalable foundation for distributed enterprise decision support. The framework showed value across service management, supply-chain coordination, governance, and broader enterprise operations, indicating strong cross-domain adaptability. At the same time, the study identified important boundaries, particularly in policy-conflict scenarios and stress-level workloads, where performance remained useful but became more constrained. Future research should therefore focus on improving rule-sensitive reasoning, strengthening high-concurrency scalability, and extending the framework with adaptive learning mechanisms that continuously refine contextual inference. Overall, this study confirms that dynamic knowledge intelligence offers a robust and strategically relevant approach for improving operational decision making in complex enterprise environments.

6. Declarations

6.1. Author Contributions

Conceptualization: A.M.W., and R.R.; Methodology: R.R.; Software: A.M.W.; Validation: A.M.W., and R.R.; Formal Analysis: A.M.W., and R.R.; Investigation: A.M.W.; Resources: R.R.; Data Curation: R.R.; Writing Original Draft Preparation A.M.W., and R.R.; Writing Review and Editing: A.M.W., and R.R.; Visualization: A.M.W.; All authors have read and agreed to the published version of the manuscript.

6.2. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

6.3. Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

6.4. Institutional Review Board Statement

Not applicable.

6.5. Informed Consent Statement

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. Bonino da Silva Santos, P. E. Bourne, J. Bouwman, A. J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C. T. Evelo, R. Finkers, A. Gonzalez-Beltran, A. J. G. Gray, P. Groth, C. Goble, J. S. Grethe, J. Heringa, P. A. C. 't Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S. J. Lusher, M. E. Martone, A. Mons, A. L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S.-A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M. A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao, and B. Mons, "The FAIR Guiding Principles for scientific data management and stewardship," *Sci. Data*, vol. 3, Art. no. 160018, Mar. 2016, doi: 10.1038/sdata.2016.18.

[2] R. J. Miller, "Open data integration," *Proc. VLDB Endow.*, vol. 11, no. 12, pp. 2130–2139, Aug. 2018, doi: 10.14778/3229863.3240491.

[3] X. Tian, "Big data and knowledge management: A case of déjà vu or back to the future?," *J. Knowl. Manag.*, vol. 21, no. 1, pp. 113–131, Feb. 2017, doi: 10.1108/JKM-07-2015-0277.

[4] F. Wang, M. S. Raisinghani, M. Mora, and J. Forrest, "Effective Decision Support in the Big Data Era: Optimize Organizational Performance via BI&A," *Int. J. Decis. Support Syst. Technol.*, vol. 14, no. 1, pp. 1–16, Mar. 2022, doi: 10.4018/IJDSST.286683.

[5] P.-L. Glaser, S. J. Ali, E. Sallinger, and D. Bork, "Model-Based Construction of Enterprise Architecture Knowledge Graphs," in *Enterprise Design, Operations, and Computing*, vol. 13585, J. P. A. Almeida, D. Karastoyanova, G. Guizzardi, M. Montali, F. M. Maggi, and C. M. Fonseca, Eds., in *Lecture Notes in Computer Science*. Cham: Springer International Publishing, vol. 2022, no. September, pp. 57–73, 2022, doi: 10.1007/978-3-031-17604-3_4.

[6] S. Albagli-Kim and D. Beimel, "Knowledge Graph-Based Framework for Decision Making Process with Limited Interaction," *Mathematics*, vol. 10, no. 21, p. 3981, pp. 1-17, Oct. 2022, doi: 10.3390/math10213981.

[7] Y. Kor, L. Tan, P. Musilek, and M. Z. Reformat, "Integrating Knowledge Graphs into Distribution Grid Decision Support Systems," *Future Internet*, vol. 16, no. 1, p. 2, pp. 1-18, Dec. 2023, doi: 10.3390/fi16010002.

[8] K. Karlapalem, P. R. Krishna, and S. R. Valluri, "Data Fabric Technologies and Applications," in *Proceedings of the 8th International Conference on Data Science and Management of Data (12th ACM IKDD CODS and 30th COMAD)*, Jodhpur, India: ACM, vol. 2024, no. Dec., pp. 362–365, 2024, doi: 10.1145/3703323.3704281.

[9] Y. Zhang, Z. Zhang, Y. Lu, H. Zhu, and D. Tang, "Dynamic decision-making for knowledge-enabled distributed resource configuration in cloud manufacturing considering stochastic order arrival," *Robot. Comput.-Integr. Manuf.*, vol. 87, no. June, p. 102712, pp. 1-17, Jun. 2024, doi: 10.1016/j.rcim.2023.102712.

[10] C. Su, Q. Jiang, Y. Han, T. Wang, and Q. He, "Knowledge graph-driven decision support for manufacturing process: A graph neural network-based knowledge reasoning approach," *Adv. Eng. Inform.*, vol. 64, no. March, p. 103098, pp. 1-17, Mar. 2025, doi: 10.1016/j.aei.2024.103098.

[11] C. Wissuchek and P. Zschech, "Prescriptive analytics systems revised: A systematic literature review from an information systems perspective," *Inf. Syst. E-Bus. Manag.*, vol. 23, no. 2, pp. 279–353, Jun. 2025, doi: 10.1007/s10257-024-00688-w.

[12] I. A. Fattah, H. Prabowo, V. U. Tjhin, and R. K. Rahim, "The interplay between business analytics capabilities and decision-making performance in Indonesia's public sector," *Digital Business*, vol. 5, no. 2, p. 100132, pp. 1-22, Dec. 2025, doi: 10.1016/j.digbus.2025.100132.

[13] T. Chai and S. Cheng, "Intelligent Operational Decision-Making in Industrial Process: Development and Prospects," *Engineering*, vol. 52, no. September, pp. 40–52, Sep. 2025, doi: 10.1016/j.eng.2025.08.010.

[14] T. E. Kalaycı, B. Bricelj, M. Lah, F. Pichler, M. K. Scharrer, and J. Rubeša-Zrim, "A Knowledge Graph-Based Data Integration Framework Applied to Battery Data Management," *Sustainability*, vol. 13, no. 3, p. 1583, pp. 1-17, Feb. 2021, doi: 10.3390/su13031583.

[15] K. Kurniawan, A. Ekelhart, E. Kiesling, D. Winkler, G. Quirchmayr, and A. M. Tjoa, "VloGraph: A Virtual Knowledge Graph Framework for Distributed Security Log Analysis," *MAKE*, vol. 4, no. 2, pp. 371–396, Apr. 2022, doi: 10.3390/make4020016.

[16] T. Horak, P. Strelec, M. Kebisek, P. Tanuska, and A. Vaclavova, "Data Integration from Heterogeneous Control Levels for the Purposes of Analysis within Industry 4.0 Concept," *Sensors*, vol. 22, no. 24, p. 9860, pp. 1-20, Dec. 2022, doi: 10.3390/s22249860.

[17] M. Hofer, D. Obraczka, A. Saeedi, H. Köpcke, and E. Rahm, "Construction of Knowledge Graphs: Current State and Challenges," *Information*, vol. 15, no. 8, p. 509, pp. 1-61, Aug. 2024, doi: 10.3390/info15080509.

[18] L. Nikbakhtan, J. Hassan, A. Pourshaban-Shahrestani, S. H. Ahmadi, M. Manafi, and L. Torkian, "A method for cancer elemental risk assessments in hookah: An example in two common types of traditional and flavored tobaccos in Iran," *MethodsX*, vol. 11, no. December, p. 102431, pp. 1-11, Dec. 2023, doi: 10.1016/j.mex.2023.102431.

[19] S. Staudinger, C. G. Schuetz, M. Schrefl, and T. Neuböck, "Knowledge graph support for descriptive business analytics," *Decision*, vol. 52, no. 3, pp. 285–306, Sep. 2025, doi: 10.1007/s40622-025-00432-4.

[20] H.-J. Cha, S.-W. Choi, E.-B. Lee, and D.-M. Lee, "Knowledge Retrieval Model Based on a Graph Database for Semantic Search in Equipment Purchase Order Specifications for Steel Plants," *Sustainability*, vol. 15, no. 7, p. 6319, pp. 1-37, Apr. 2023, doi: 10.3390/su15076319.

[21] M. Moesmann and T. B. Pedersen, "Data-driven prescriptive analytics applications: A comprehensive survey," *Inf. Syst.*, vol. 134, no. October, p. 102576, pp. 1-26, Oct. 2025, doi: 10.1016/j.is.2025.102576.

[22] J. Weller, N. Migenda, Y. Naik, T. Heuwinkel, A. Kühn, M. Kohlhase, W. Schenck, and R. Dumitrescu, "Reference architecture for the integration of prescriptive analytics use cases in smart factories," *Mathematics*, vol. 12, no. 17, Art. no. 2663, Aug. 2024, doi: 10.3390/math12172663.

[23] M. Chen, Z. Tao, W. Tang, T. Qin, R. Yang, and C. Zhu, "Enhancing emergency decision-making with knowledge graphs and large language models," *Int. J. Disaster Risk Reduct.*, vol. 113, no. October, p. 104804, pp. 1-12, Oct. 2024, doi: 10.1016/j.ijdrr.2024.104804.

[24] L. Liu, P. Xu, K. Fan, and M. Wang, "Research on application of knowledge graph in industrial control system security situation awareness and decision-making: A survey," *Neurocomputing*, vol. 613, no. January, p. 128721, pp. 1-17, Jan. 2025, doi: 10.1016/j.neucom.2024.128721.